INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 12 Control in Petrochemicals and Polymers

CHAPTER 12

Loss-in-Weight Additive Feeders and the Holdup Their Scales Cannot Weigh

Pages 282–304

Edited by John D. Hedengren

How to cite this chapter. "Loss-in-Weight Additive Feeders and the Holdup Their Scales Cannot Weigh," in *Control in Petrochemicals and Polymers*, John D. Hedengren, Ed., vol. 12 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 282–304.

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

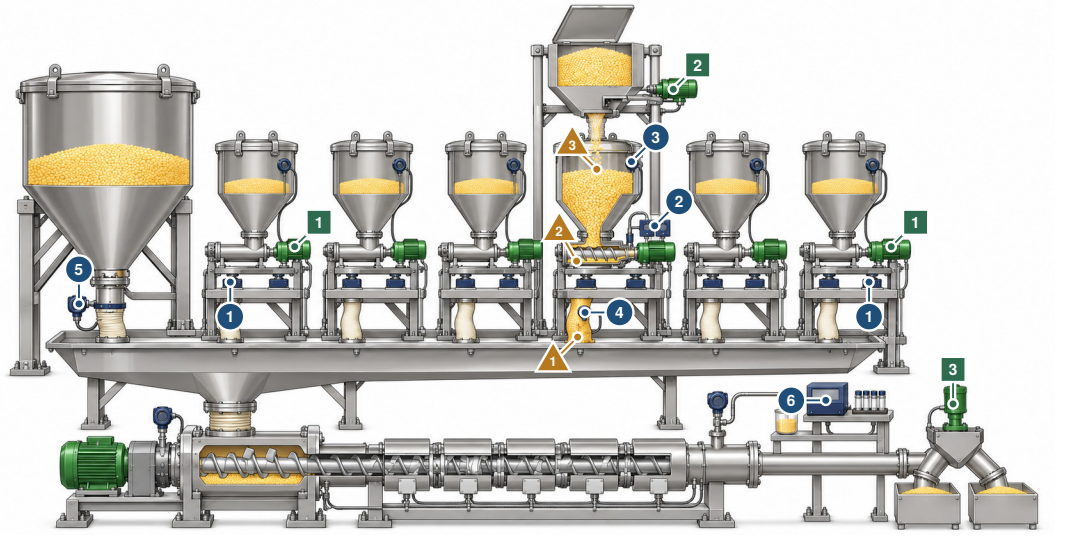
The app for this chapter
https://apmonitor.com/apps/weigh-feeder/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v12-petrochemicals-polymers/#weigh-feeder

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

CHAPTER 12

Loss-in-Weight Additive Feeders and the Holdup Their Scales Cannot Weigh



The controller moves 1 speed of each feeder's screw 2 refill valve above the hopper 3 diverter to prime or to the diversion bin

The controller reads 1 load cells under each feeder 2 screw speed and motor torque 3 hopper pressure 4 sock high-level switch 5 resin flow 6 laboratory assay of a sample, hours late

What disturbs it 1 powder held in the sock below the scale 2 buildup where the screw discharges 3 a refill, which can pressurize the hopper

Notice that each discharge sock hangs below its feeder's load cells, so powder held in it has left the weighed feeder without reaching the extruder. Illustration, not to scale.

IN PLAIN TERMS **A Scale That Cannot See the Last Step**

The main idea. Plastic pellets get small doses of powders that protect the plastic and help it work. Each powder sits in a hopper standing on a scale. A screw turns faster or slower so that the weight drops at the right rate, even when the screw starts to deliver less on each turn.

The challenge. The scale weighs only the hopper. Powder can pile up in the soft sleeve below it, so the pellets run short while every screen reads normal, and the laboratory finds out hours later.

Why it matters. A scale that reads right does not prove that the pellets are right.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in Section 12.1, the History box in Section 12.2, the Try it boxes in Sections 12.3 and 12.4, the History box in Section 12.5, the Did you know? box in Section 12.6, and Key points at the end of the Conclusion. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 12.6 with Figure 12.3, Section 12.7, and Section 12.8 with its Open challenge box.

The control question A loss-in-weight feeder holds the flow its scale weighs, but the product sees only what leaves the discharge sock below the scale and a laboratory reports it hours late. Which faults does the weighing loop correct, which can it not see, and what do a laboratory trim, a moving horizon estimator and a predictive controller add?

Learning objectives

- Convert a recipe in ppm of finished compound into feeder flows for neat and masterbatch ingredients, and account for the carrier in the resin rate.
- Compute the speed a feed-factor loss demands and the time a sock restriction takes to trip the high-level switch.
- Explain why the sock holdup is not observable from hopper weight, and name the measurements that reveal it.
- Decide when a delayed assay may move the laboratory trim, judged against its own collection interval and lot.
- Compare ratio control, MPC, the simple estimate and the MHE on matched runs, and state where each changes the product.

Prerequisites Sections 1.4.5 (grade changes) and 1.4.6 (laboratory bias update), 1.6.2 (velocity-form PID), 1.6.3 (MPC and MHE) and 1.6.4 (scenarios and the exported file).

The app https://apmonitor.com/apps/weigh-feeder/

See also Volume 5, Chapter 6 for loss-in-weight feeders with refill feedforward ahead of a tablet press, and Volume 5, Chapter 7 for diversion timed by residence. Volume 3, Chapter 6 for a dose computed from an upstream measurement, with the slow downstream measurement demoted to a ratio trim.

12.1 Introduction

A specification of 1,000 ppm of a phosphite antioxidant is one kilogram per tonne of compound, so a compounding line making 1,000 kg/h of pellets needs 1 kg/h of that powder, or 5.6 g every twenty seconds. The feeder that delivers it is usually a loss-in-weight (LIW) screw feeder, a hopper, motor and screw standing on load cells whose flow is taken as the rate at which the weighed assembly loses mass [1, 2]. The same principle meters active ingredients and excipients into continuous tablet lines [3, 4], the subject of Volume 5, Chapter 6, and the additive packages it meters into polyolefins enhance the polymer's properties and prolong its life [5]. A published polypropylene composition carries 0.1 % of a phosphite antioxidant, 0.05 % of a hindered phenolic antioxidant, 0.11 % of calcium stearate and 0.2 % of a sorbitol clarifier, compounded in a twin-screw extruder [6], and the simulator's grades sit at those magnitudes.

> ✦ DID YOU KNOW? **Fractions of a Percent**
>
> A 2018 review of additives in plastics gives typical contents of 0.05 to 3 % by weight for stabilizers and antioxidants, and 0.1 to 3 % for slip agents. Phenolic antioxidants are used in low amounts and phosphites in higher ones, and polyolefins take the lowest amounts [5].

This simulator targets a single belief, that a feeder holding its weighed flow on target holds the product on specification. The scale weighs the hopper, the motor and the screw. Whatever leaves the screw and stays in the flexible discharge sock below it has left the weighed system without reaching the product. In case 4 of the app the sock narrows to a fifth of its healthy capacity. The phosphite feeder's LIW reading then averages 1.00 kg/h from 15 to 30 min and still reads 1.03 kg/h when a level switch trips at 36.6 min, while the phosphite at the sample tap sits at 600 ppm against 1,000 ppm. The class of problem is general. A flow measured upstream of an unmeasured inventory, a product property seen only by a laboratory, and a model updated from infrequent assays describe polymer reactors [7] as well as feeders, and grade changes on such plants are judged by the off-specification material they leave [8], as Section 1.4.5 sets out and Chapter 8 measures on a fluidized bed.

However, the feeder literature is mostly pharmaceutical and treats the feed rate as the controlled quantity, with refill transients, feed-factor prediction and the coordination of several feeders as its subjects [4, 9, 10], while polymer quality control treats the reactor rather than the additive lanes. This chapter measures what ratio control, a laboratory trim, a short moving horizon estimator (MHE) and a supervisory model predictive controller (MPC) each add to the product, on a browser simulator that records each lane's weighed flow, its delivered flow, the off-scale holdup and the product concentration together. The balances, the proportional-integral (PI) loop, the MHE [11, 12] and the MPC [13] are standard methods. The results are measurements on one teaching plant, and three of them are negative findings, that the off-scale holdup is not observable from hopper weight, that the MHE is quieter than the simple estimate but not faster, and that the MPC's measured product benefit is confined to grade transitions. The model's total mass balance closes to 10^{-13} of the input, the recipe solver reports infeasible ingredient sets instead of clipping them, the feeder MHE solves a five-minute window in 0.2 ms, and no controller or estimator crosses the measurement boundary to read the simulator's truth. Ten cases run from recipe arithmetic to a graded twelve-hour shift.

Section 12.2 follows the domain, from gravimetric feeding to delayed laboratory assays. Section 12.3 builds the line, its instruments, the recipe and the laboratory. Section 12.4 works out by hand what the weighing boundary implies for the two restrictions and for washout. Section 12.5 describes the ratio and LIW loops, the two estimators, the laboratory trim, the supervisor and the MPC. Section 12.6 measures the ten cases and compares the four control architectures on matched runs. Section 12.7 names the hardware and the practice behind the model. Sections 12.8 and 12.9 state the limitations and the conclusions, and the Exercises put the cases to work.

12.2 Weighing Powder, Dosing Additives and Waiting for the Laboratory

A screw feeder delivers a volume of powder per revolution, so its mass flow moves with bulk density, hopper level and anything that changes how the screw fills. Gravimetric feeding answers this by weighing. In the loss-in-weight arrangement the hopper, the drive and the screw stand on load cells, the controller differentiates the net weight to obtain the flow, and it trims the screw speed to hold a setpoint [1]. The weight is a derivative measurement, so the flow estimate needs a window and is only as good as the scale beneath it. A hopper cannot be weighed while it is being refilled, which forces the feeder to run on its learned speed until the weight settles, and refills have been shown to cause measurable feed-rate deviations [9]. Characterization methods for these feeders measure their feed-rate variability and the feed factor, the delivery per revolution, as the hopper empties [2], and the feed factor has been predicted from the properties of the powder [10].

> ⧖ HISTORY **A Feeder Standing on a Scale, 1945**
>
> In a patent application filed in March 1945, a weigh hopper discharged through a star feeder driven by a variable-speed motor. The hopper, the feeder, the motor and the other accessories all stood on a scale platform as tare weight, so the scale showed the net weight of material left in the hopper [14]. A later application described how such systems held a rate. A counterpoise, moved along the scale by a motor-driven lead screw, set a weight that fell at the desired rate, and any unbalance of the beam adjusted the flow. The reservoir had to be refilled by hand and the scale balanced again after every emptying, and a change in density could set the flow oscillating about its target [15].

Continuous pharmaceutical manufacturing made the feeder the first unit of a line whose composition matters at every instant [3]. One study treated a set of LIW feeders as a control system in its own right [4], and a hybrid of MPC and PID control was implemented on a direct compaction pilot plant [16]. Polymer compounding has the same structure at larger scale and with sparser measurement. Additives of these kinds are found throughout plastic products [5], and assays of a stabilizer in a polymer commonly begin by extracting the additive from the matrix [17], so the product's composition is known from a laboratory rather than from an online instrument.

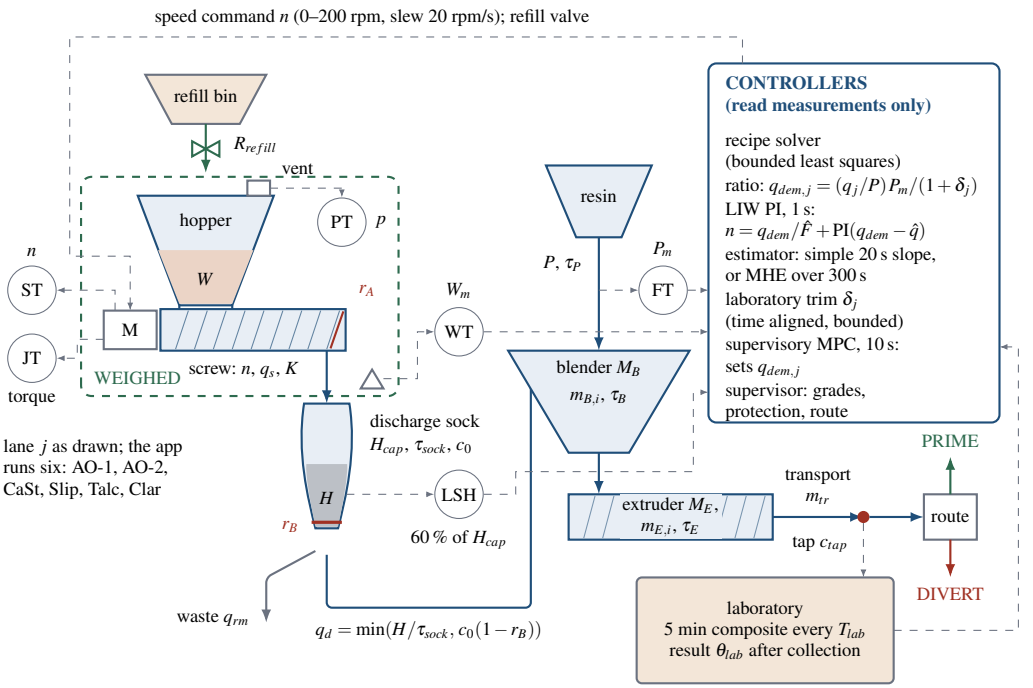


Figure 12.1 Additive compounding line as modeled. One lane is drawn in full: the hopper weight W, the screw at speed n with its flow q_s, the feed-factor loss r_A of a restriction that backs into the screw, the vent and hopper pressure p, and the refill R_{refill} lie inside the dashed weighing boundary, while the sock holdup H, its restriction r_B and its delivered flow q_d lie outside it. The resin P, the blender (M_B, τ_B), the extruder (M_E, τ_E), the transport m_{tr} to the tap c_{tap}, the laboratory and the diverter follow. Instruments: WT net weight W_m, ST speed, JT torque, PT pressure, LSH sock high level, FT resin flow. Not to scale.

Infrequent laboratory measurements were brought into polymer control through models whose parameters are updated each time a result arrives [7], the bias update of Section 1.4.6, and dynamic optimization of grade changes showed how much off-specification polymer a transition policy costs [8]. Moving horizon estimation, a least-squares fit of a model to a sliding window of data with constraints on states and disturbances [11, 12] (Section 1.6.3), is the estimation counterpart of the receding-horizon controllers that became standard in the process industries [13]. A robust penalty on the residuals in the sense of Huber [18] limits the influence of a single bad weight reading on the estimate. The simulator puts these pieces on one compounding line and asks which of them the product notices.

12.3 The Compounding Line as Modeled

The simulated line makes 1,000 kg/h of finished polyolefin compound, selectable from 500 to 2,000 kg/h, as shown in Figure 12.1. Base resin enters through its own gravimetric feeder, the master stream. Six additive lanes, one per active, each comprise a weighed hopper, motor and screw on load cells, a refill bin above, and a discharge sock below the weighing boundary. The lanes carry a phenolic antioxidant of the AO-1010 type (AO-1), a phosphite antioxidant of the AO-168 type (AO-2), calcium stearate (CaSt), an erucamide slip agent, a fine talc antiblock and a sorbitol clarifier (Clar). Everything discharges into a blender of 20 kg that feeds an extruder of 30 kg, and 16.7 kg of product downstream a sample tap feeds the laboratory before a diverter sends the product to prime or to a diversion bin.

Table 12.1 Variables by Role for the AO-2 Lane at the Design Point (PP-A grade)

Variable	Role	Design	Range
$c_{tap,i}$	CV, product at tap (lab only)	1,000 ppm	band ±10 %
q_{dem}	CV of the LIW loop, flow	1.000 kg/h	0–2.5 kg/h
n	MV, auger speed	80 rpm	0–200 rpm
$q_{dem,j}$	MV of the MPC, lane demand	1.000 kg/h	0–0.95$\hat{F}n_{max}$
Q	DV, production setpoint	1,000 kg/h	500–2,000 kg/h
r_A	DV, feed-factor loss	0	0–60 %
r_B	DV, sock restriction	0	0–100 %
s	DV, scale span error	0	±10 %
x_{MB}, c_{res}	DV, potency, resin AO-1	5 %, 0	3–6, 0–800 ppm
W	state, hopper	1.5 kg	0–3 kg
H	state, sock (unmeasured)	2.8 g	0–300 g
a, p	states, aeration, pressure	0, 0	0–0.06, Pa
P	state, resin flow	997.5 kg/h	—
M_B, M_E	states, mixing masses	20, 30 kg	—
W_m, torque	measurements	1.5 kg, 25 %	—

Table 12.1 assigns each variable its role. Per lane five states are integrated, the hopper weight W, the sock holdup H, the feed-factor loss r_A, the aeration a after a refill and the hopper pressure p. The line adds the resin flow P, the blender and extruder masses with six active masses each, and ledgers of resin and refill in and of product and waste out, so that conservation can be checked to round-off. The model integrates in seconds, and flows written in kg/h are divided by 3,600 where they enter a derivative. Concentrations are mass fractions of the finished compound, displayed in mass ppm, with $1\,\% = 10{,}000\,\text{ppm}$.

12.3.1 Hopper, Screw and Two Restrictions

The weighed hopper loses what the screw takes and gains what a refill adds, as shown in Equation (12.1).

$$\begin{aligned} \frac{\mathrm{d}W}{\mathrm{d}t} &= R_{refill} - \frac{q_s}{3600} \\ q_s &= K\,n\,f(W)\,(1-a)\,(1-r_A)\,\mathrm{acc}(H) \\ f(W) &= \left[1+\beta\left(\frac{W}{W_{max}}-\frac{1}{2}\right)\right]\min\left(1,\frac{W}{0.1\,W_{max}}\right) \end{aligned} \tag{12.1}$$

where K is the commissioned feed factor in kg/h per rpm, n the auger speed, $\beta = 0.08$ the hopper-level effect (±4 % from empty to full), and the second factor of f starves the screw below 10 % of the hopper. The lanes are sized so that the design rate $q_{des} = 80K$ sits at 80 rpm, with $K = 0.00625$, 0.0125, 0.0125, 0.0125, 0.0625 and 0.025 kg/h/rpm for AO-1, AO-2, CaSt, slip, talc and clarifier. The AO-2 value delivers 1 kg/h at 80 rpm, the example that opens the chapter. The feed-factor loss r_A represents a restriction that backs into the screw, so it is inside the weighing boundary, and it ramps toward its disturbance value at a buildup rate of 25 % per 30 min and clears with a 5 s time constant. The backpressure factor $\mathrm{acc}(H)$ is 1 below 70 % of the sock capacity H_{cap} and falls by a smoothstep to 0 at H_{cap}, so a full sock stops the screw instead of losing mass.

12.3.2 The Discharge Sock Below the Scale

The sock receives the screw flow and releases it with an availability-limited outflow, as shown in Equation (12.2).

$$\begin{aligned} \frac{\mathrm{d}H}{\mathrm{d}t} &= \frac{q_s - q_d - q_{rm}}{3600} \\ q_d &= \min\left(\frac{3600\,H}{\tau_{sock}},\; c_0\,(1-r_B)\right), \qquad c_0 = 3\,q_{des} \end{aligned} \tag{12.2}$$

where $\tau_{sock} = 10\,\mathrm{s}$ is the transit time of a healthy sock, c_0 its capacity, and r_B the sock restriction, with 1 a full plug. The outflow can never exceed what is in the sock, so H stays nonnegative when material enters an empty sock. Cleaning locks the lane out for 300 s, sets $q_d = 0$ and sends the holdup to waste at $q_{rm} = 3600H/30\,\mathrm{s}$, while knocking the sock clear removes r_B so that its contents fall into the blender at the healthy capacity c_0. Neither action deletes mass. The microfeeder socks hold $H_{cap} = 0.30\,\mathrm{kg}$ and the talc and masterbatch socks 2 kg.

12.3.3 Refill, Aeration and Hopper Pressure

The feeder's own sequence opens the refill valve when the measured weight falls below 25 % of the hopper and closes it at 85 % or after 40 s, with $R_{refill} = 0.6W_{max}/20\,\mathrm{s}$. Fresh powder is aerated and the displaced air must leave through the vent filter, as shown in Equation (12.3).

$$\frac{\mathrm{d}a}{\mathrm{d}t} = \frac{0.06\,R_{refill}}{0.6\,W_{max}} - \frac{a}{\tau_{aer}}, \qquad \frac{\mathrm{d}p}{\mathrm{d}t} = u_{refill} - \frac{p}{\tau_v} \tag{12.3}$$

where u_{refill} is 1 while the valve is open and 0 otherwise, p is in pascals, $\tau_{aer} = 120\,\mathrm{s}$, a full refill costs up to 6 % of the delivery per revolution, and the vent time constant τ_v is 2 s when healthy and 90 s with the filter plugged. A plugged filter therefore leaves about 18 Pa in the hopper at the end of a 20 s refill. The term represents hopper pressure reaching the load cells through the flexible connections, a declared teaching relation.

> **TRY IT IN THE APP** **6 · Refill With a Plugged Vent Filter**
>
> Run 15 min on the Feeder view and watch the AO-2 speed after the refill that starts at 220 s. With the timer alone the speed falls to 65.1 rpm and AO-2 at the tap dips to 911 ppm. Set Weight validity to Timer + pressure, press Reset and run again: the speed stays at 75.2 rpm or above and the tap holds 995 to 1,021 ppm.

12.3.4 What the Instruments Read

Nothing reaches the controllers except the instruments of Equation (12.4), each read once per second.

$$\begin{aligned} W_m &= (1+s)\,(W + \alpha H) + k_p\,p + v_W \\ \tau_q &= 15 + 25\,\frac{n}{200} + 40\,\frac{H}{H_{cap}} + 20\,r_A + d_\tau + v_\tau \\ \mathrm{LSH} &= [H > 0.6\,H_{cap}] \end{aligned} \tag{12.4}$$

where s is a span error, $\alpha = 0$ the share of the sock carried by the scale (a commissioning value), $k_p = 0.5\,\mathrm{g/Pa}$ the pressure coupling, v_W colored noise of 0.2 g (1 g on the talc and masterbatch lanes) with a 1 s correlation time, τ_q the motor torque in percent with a slow drift d_τ of 2 % over 2 h and noise of 1.5 %, and LSH a discrete sock high-level switch. The drive's speed, the hopper pressure (σ 0.2 Pa), the refill valve, the maintenance log of cleanings and knocks, the resin flow (σ 0.3 %) and the blender level (σ 0.2 kg) complete the packet. The sock holdup itself is measured by nothing.

12.3.5 Resin, Blender, Extruder and Tap

Ingredient j carries A_{ij} kilograms of active i per kilogram, 1 for a neat additive, 0.05 for the slip masterbatch whose other 0.95 is carrier polymer, and 1/3 and 2/3 for the optional antioxidant premix. The blender and the extruder are level-regulated ideal mixers, as shown in Equation (12.5).

$$\begin{aligned} Q_{in} &= P + \sum_j q_{d,j}, \qquad \dot{m}_{in,i} = P\,c_{res,i} + \sum_j A_{ij}\,q_{d,j} \\ \frac{\mathrm{d}m_{B,i}}{\mathrm{d}t} &= \frac{\dot{m}_{in,i}}{3600} - \frac{m_{B,i}}{\tau_B}, \qquad \frac{\mathrm{d}m_{E,i}}{\mathrm{d}t} = \frac{m_{B,i}}{\tau_B} - \frac{m_{E,i}}{\tau_E} \\ c_{tap}(t) &= c_E\left(M_{prod}^{-1}(M_{prod}(t) - m_{tr})\right) \end{aligned} \tag{12.5}$$

where the total masses follow the same pair with Q_{in}, $\tau_B = 72\,\text{s}$ and $\tau_E = 108\,\text{s}$ correspond to 20 and 30 kg at 1,000 kg/h, and the tap reads the extruder outlet concentration of the product that left $m_{tr} = 16.7\,\text{kg}$ earlier, so its delay follows the throughput. The resin follows its command with $\tau_P = 15\,\text{s}$, and its command moves at most 20 kg/h per second.

12.3.6 The Recipe and Its Feasibility

For a finished-compound demand Q, targets c_i and a declared resin background $c_{res,i}$, the solver seeks nonnegative ingredient flows that satisfy Equation (12.6).

$$\begin{aligned} P + \sum_j q_j &= Q, \qquad q_j \geq 0 \\ P\,c_{res,i} + \sum_j A_{ij}\,q_j &= c_i\,Q \quad \forall i \end{aligned} \tag{12.6}$$

Eliminating P leaves one row per active, and the rows are scaled by each active's tolerance and solved as a bounded least-squares problem with Goldfarb and Idnani's dual active-set method [19]. The recipe is feasible only when every residual is below $10^{-9}Q$. Otherwise the solver names the active and the reason, a premix whose fixed ratio differs from the targets or a resin that already carries more than the target, and never clips one ingredient on its own. For PE-B at 1,000 kg/h it returns 0.5, 1, 0.5, 1 and 5 kg/h of neat actives and 992 kg/h of resin. With a 5 % slip masterbatch the slip lane carries 20 kg/h, of which 19 kg/h is carrier, and the resin falls to 973 kg/h. Both agree with hand arithmetic to 3×10^{-12} kg/h.

12.3.7 The Laboratory

A five-minute composite is drawn from the tap every $T_{lab} = 3\,\text{h}$ (2 to 4 h) from 1 h, weighted by the product mass that passes it, and reported $\theta_{lab} = 45\,\text{min}$ (20 to 90 min) after collection ends. Each species carries a relative noise of 3 % for the chromatographic assays, 4 % for calcium stearate assayed as calcium and 5 % for talc assayed as ash, with a 10 ppm floor and a 20 ppm reporting limit. A scheduled event can lose a sample or multiply the next result by 1.35. Each record keeps its collection interval and the grade requested when collection began, so a late result belongs to its own lot.

12.3.8 Parameters, Initialization and Numerics

The values of Table 12.2 are teaching assumptions sized for the line's operating envelope, except where a source is named. The line starts lined out on the controller's recipe, with the declared potency and declared resin background, while the product already carries the true composition of those flows. Hoppers start half full, socks at their healthy holdup $q\tau_{sock}$, and the blender and extruder at their steady masses. With this start the derivatives of the sock, blender and extruder states are zero, and a hidden fault reaches the controllers only through a measurement. The engine of Section 1.6.1 takes fourth-order Runge–Kutta steps of 1 s, runs the controllers every second, the MPC every 10 s and each MHE every 5 s, and logs a row every 10 s. Refills, the sock-full mark, cleaning, sampling and reports are events at step boundaries with historian marks, and the derivative function is pure.

12.3.9 Assumptions

The actives are conserved through the extruder and the assay measures the intact parent compound, so a low phosphite result is always a dosing error here, which on a real line it need not be. The blender and extruder are ideal mixers with constant residence, with no segregation and no axial dispersion in the transport. Calcium stearate is assayed as calcium and talc as ash through exact mappings with larger noise. A grade change between polypropylene and polyethylene requires a resin purge that is not simulated, and the grade selector refuses it. The torque relation, the pressure coupling, the sock capacity and the bands are declared teaching relations, not supplier correlations or standards, and a real grade's bands come from its customer specification and its assay validation.

Table 12.2 Parameters (V visible, A adjustable, H hidden truth)

Symbol	Meaning	Value	
K	feed factor, AO-2	0.0125 kg/h/rpm	V
W_{max}	hopper, micro (talc, MB)	3 (20, 25) kg	V
β, starve	level effect, starving	0.08, 10 %	V
τ_{aer}, a_{max}	aeration after refill	120 s, 0.06	V
τ_{sock}	sock transit	10 s	A
c_0	healthy sock capacity	$3q_{des}$	A
H_{cap}	sock capacity, micro (large)	0.30 (2) kg	V
$\dot{r}_A$	buildup rate	25 %/30 min	A
α	sock load on scale	0	A
k_p	pressure coupling	0.5 g/Pa	A
τ_v	vent, healthy (plugged)	2 (90) s	A
σ_W	load-cell noise, micro (large)	0.2 (1) g	V
τ_P	resin feeder	15 s	A
τ_B, τ_E	blender, extruder residence	72, 108 s	A
m_{tr}	product, extruder to tap	16.7 kg	A
T_{lab}, θ_{lab}	sample interval, turnaround	3 h, 45 min	A
σ_{rel}, σ_0	assay noise, floor	3 %, 10 ppm	A
bands	nonzero target, zero target	$\pm$10 %, $\leq$ 50 ppm	V
x_{MB}	masterbatch potency	5 %	H
c_{res}	resin AO-1 background	0 ppm	H

12.4 What the Weighing Boundary Implies

The two restrictions sit on opposite sides of the scale and have opposite consequences. A feed-factor loss r_A removes delivery per revolution from the weighed flow, so the scale sees it and speed restores it. At a half-full hopper the AO-2 lane needs $1/(0.0125 \times 0.75) = 106.67$ rpm for 1 kg/h after a 25 % loss, 33.3 % more than the healthy 80 rpm, and the 200 rpm limit allows losses up to 60 %. A sock restriction r_B removes nothing from the weighed flow until the sock pushes back. Its capacity $3(1 - r_B)q_{des}$ falls below the demand once r_B exceeds two thirds, and no speed can then deliver the demand, while a full plug delivers nothing at any speed. At 200 rpm with the sock plugged the delivered flow after an hour is 0.0000 kg/h and the sock holds 299 g of its 300 g.

The partial sock restriction of case 4 can be followed by hand. The sock starts at $1\,\text{kg/h} \times 10\,\text{s} = 2.8\,\text{g}$ and fills at $1 - 0.6 = 0.4\,\text{kg/h}$, so it reaches the switch point of 180 g after 26.6 min, at $t = 36.6$ min. The engine trips the switch at 36.62 min. Over the 15 to 30 min window the LIW reading averages 0.9998 kg/h while the sock delivers 0.600 kg/h, and the phosphite at the tap is 608 ppm at 20 min. Figure 12.2(a) shows the three flows parting at 10 min and joining again only after the sock passes 70 % of its capacity, when backpressure throttles the screw. Released, the holdup is a slug. With the speed fixed and the sock restricted 80 % from 600 s, the screw still leaves the hopper at 0.98 to 0.99 kg/h from 1,200 to 2,400 s while 0.600 kg/h is delivered, the sock holds 247.9 g when it is knocked clear at 3,990 s, and the phosphite at the tap peaks at 2,871 ppm.

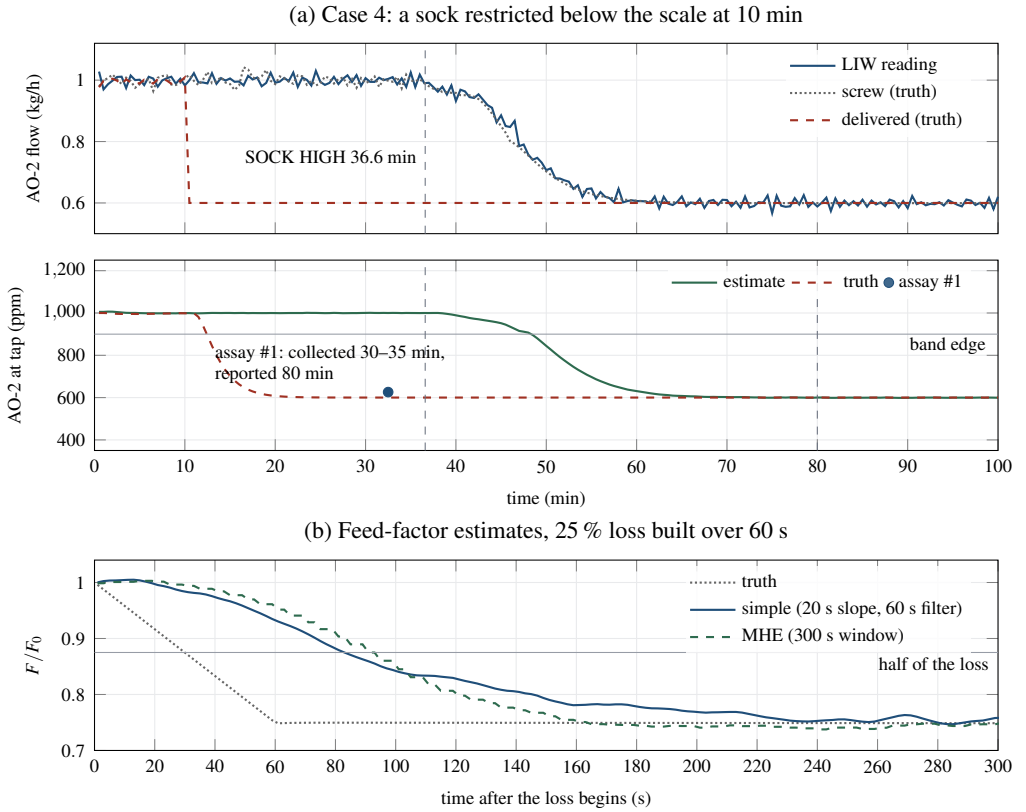


Figure 12.2 Two blind spots of the weighing loop. (a) Case 4 as shipped: the LIW reading follows the screw while the delivered flow and the phosphite at the tap fall, and the estimate stays near target until backpressure throttles the screw and pulls it to 599 ppm against a true 600 ppm by 100 min. The assay reported at 80 min is accepted as evidence of holdup and leaves the trim at zero. (b) A 25 % feed-factor loss on the AO-2 lane in Manual: the MHE is smoother but reaches half the loss later than the simple estimate.

> **TRY IT IN THE APP** **4 · Hidden Holdup: The Scale Says All Is Well**
>
> Load the scenario, keep the Line view and press Run. Watch the AO-2 feeder's card: its LIW flow keeps reading about 1 kg/h. Now watch AO-2 at the tap, the dashed true line. It is below 900 ppm from 12.4 min, and nothing on the feeder changes until the sock high-level switch trips at 36.6 min.

The line's residence decides when a recipe change reaches the product. The blender, extruder and transport hold 66.7 kg, which is 240 s at 1,000 kg/h, so the scheduled feed switch leads a grade change by that time. Removal is slower than the mean residence suggests. After the clarifier feeder stops, its concentration passes three lags of 10, 72 and 108 s, and integrating them by hand gives 520.5 s to fall by the factor of 40 from 2,000 ppm to 50 ppm. With the 60 s transport the clarifier should leave the band 9.67 min after the feed stops at 26 min, at 35.67 min, and the engine reports 35.7 min.

12.5 Control and Estimation as Implemented

Every controller, estimator and supervisor reads only the instrument packet of Equation (12.4), the released laboratory results, the controls and its own state. The exported file marks the simulator's

truth columns with the prefix `truth_`, so a controller replayed from the file cannot read them. Each control execution, once per second, reads the instruments, forms the recipe, updates the estimator and the laboratory trim, runs the supervisor, solves the MPC if it is due, and updates the six LIW loops, in that order.

12.5.1 Ratio Control and the LIW Loop

The recipe flows of Equation (12.6) at the production setpoint, computed with the declared ingredient matrix and the declared resin background, become ratios to the resin, and each lane's demand is that ratio times the measured resin flow P_m, divided by the laboratory trim, as shown in Equation (12.7).

$$\begin{aligned} q_{dem,j} &= \frac{q_j}{P}\,\frac{P_m}{1+\delta_j} \\ n &= \frac{q_{dem}}{\hat{F}} + K_c\,e + \frac{K_c}{\tau_I}\int e\,dt, \qquad e = 100\,\frac{q_{dem}-\hat{q}}{q_{des}} \end{aligned} \tag{12.7}$$

where $\hat{F}$ is the estimated feed factor, the feedforward inverse, $\hat{q}$ the estimated flow, and the error is in percent of the lane's design rate so that one tuning reads the same on every lane. The shipped tuning is $K_c = 0.2$ rpm per percent and $\tau_I = 60$ s. The PI is the velocity-form PID of Section 1.6.2 with no derivative action, a 0 to 200 rpm range and a slew limit of 20 rpm/s, reduced to 0.5 rpm/s for 5 min after a cleaning. The feedforward enters as the loop's moving bias. While the weight is invalid, during a refill, the 10 s settling after it, and optionally while the hopper pressure is outside 1.5 Pa, the lane runs volumetrically at $q_{dem}/\hat{F}$ and the PI tracks that speed, so the return to gravimetric operation is bumpless.

> ⧖ HISTORY **A Load Cell, a Derivative and a Hold for the Refill, 1967**
>
> An application filed in 1967 replaced the moving weight with electronics. A transducer, which could be a load cell built on strain gauges, turned the weight of the reservoir and its contents into a voltage. A circuit differentiated it with respect to time, and a controller compared that mass flow with the desired rate and adjusted the outflow. Differentiation magnifies anything that changes faster than the weight, so frequencies above about one cycle per second were filtered out to remove the vibration of nearby machinery. Two level settings opened and closed a refill valve. While the weight rose, its derivative would have driven the controller to the wrong flow, so the controller held its last output or switched to manual until the refill ended [15]. The volumetric mode described above does the same.

12.5.2 The Simple Estimate

The classical LIW reading is the least-squares slope of the last 20 s of net weight, valid only when no sample in the window is masked. Its feed factor is that flow divided by the speed, bounded to 0.2 to 1.5 times K, filtered with a 60 s time constant and updated only while valid and while the lane is not in alarm.

12.5.3 The Feeder MHE

The moving horizon estimator works at the scale of the feeder, not of the product: it replaced an eight-hour composition window considered first. A feeder empties gradually and fills repeatedly, its flow is the change of weight over time, and the estimator's job is to detect a change of delivery per revolution against speed for the controller, so its horizon is about five minutes. Over the last $N = 300$ s of 1 s samples it chooses the starting weight W_0, one feed factor F_b per 10 s block and one

free mass J_g per refill that ends in the window, as shown in Equation (12.8).

$$\begin{aligned} \min \sum_{t\in\mathcal{V}} \rho_H\Big(\frac{W_{m,t}-W_t}{\sigma_W}\Big) + \frac{(W_0-\bar{W}_0)^2}{\sigma_{W0}^2} + \frac{(F_1-\bar{F})^2}{\sigma_{F0}^2} \\ + \sum_b \frac{(F_{b+1}-F_b)^2}{\sigma_q^2} + \varepsilon \sum_g J_g^2 \\ W_t = W_0 - \sum_{\tau<t} \frac{F_{b(\tau)}\, n_\tau}{3600} + \sum_{g\le t} J_g \\ 0.2K \le F_b \le 1.5K \end{aligned} \tag{12.8}$$

where $\mathcal{V}$ holds the valid samples, so refill and settling samples are masked, ρ_H is the Huber penalty with threshold 2, $\bar{W}_0$ and $\bar{F}$ are the previous window's trajectory at this window's start, $\sigma_{W0} = 2\,\mathrm{g}$, $\sigma_{F0} = 0.05K$, the random-walk step is $\sigma_q = 0.002K$ per block, and ε is a small regularization that keeps the problem strictly convex. The estimate is carried across a refill by the random walk and is never reset. The weights are linear in the decisions, so each window is a dense convex quadratic program of 31 to 33 variables and 60 bounds, solved by the same dual active-set method [19]. The Huber penalty is applied by iteratively reweighted least squares, two solves per update with weights $\min(1, 2\sigma_W/|r|)$ from the first. A solve is rejected when its status is anything but optimal, when a value is infinite or when a block leaves its bounds, and the lane then keeps its previous estimate, as it also does for a window with fewer than 30 valid samples. The lane's flow is then $\hat{q} = F_B\, n$ from the last block.

Three things are deliberately outside the estimator. A constant scale offset is absorbed by W_0, but a span error or drift cannot be told from a feed-factor change by the weight alone and is left to the laboratory trim. Laboratory assays do not enter the MHE. The sock holdup is not observable from the hopper weight at all, because the weight records what the screw removed and the holdup is what happened to it afterward. Its witnesses are the torque, the high-level switch and the laboratory.

🕮 **Pause.** Try Exercise 12.1 before reading on.

The same problem written for the shared estimator of the framework, with one noise vector per one-second step, has 602 decisions and 903 rows and costs 161 ms per update at $N = 300$. The block structure and one mass per refill reduce it to the 31 to 33 decisions above, which cost 0.205 ms per solve in Node, about 785 times less.

12.5.4 The Product Between Assays and the Laboratory Trim

Between assays the product concentration is a model driven by the estimated lane flows, each times its declared composition and its trim $1+\delta_j$, plus the measured resin times the declared background, passed through the declared sock, blender and extruder lags at the measured throughput and the 16.7 kg transport. A released result is used once, against this model's prediction averaged over its own collection interval, never against the present state, as shown in Equation (12.9). It is the bias update of Equation (1.11) written as a gain on each lane's delivery and time-aligned to the collection interval.

$$\delta_j \leftarrow \delta_j + \lambda\Big(\frac{y_{lab}-b}{\hat{y}_{raw}(t_s,t_e)-b} - 1 - \delta_j\Big), \qquad |\delta_j| \le 0.2 \tag{12.9}$$

where $\lambda = 0.5$, b is the background's share and $\hat{y}_{raw}$ the untrimmed prediction for the interval $[t_s, t_e]$. A result within 1.5 assay standard deviations of the present trim moves nothing. A result more than 25 % from its prediction is held as a possible outlier unless the next result confirms it on the same side or the lane's own diagnostics, the switch or torque above the healthy relation, were raised after its collection began. The estimate's standard deviation combines the assay noise, an aging term of 2 % of target per hour since the last accepted assay, and a holdup term of 30 % of target per unit of torque and switch evidence, the switch alone counting one half.

🕮 **Pause.** Try Exercise 12.2 before reading on.

12.5.5 Supervisor

The supervisor keeps three grade states apart from the laboratory, namely old grade, transition, and new grade predicted on specification once the predicted product has held every band for 10 min, with laboratory confirmation only from an accepted result collected after that instant. With Preview on, the feed recipe switches ahead of a scheduled change by the declared holdup over the measured throughput, 240 s at 1,000 kg/h. Protection acts per lane without waiting for the laboratory. AT LIMIT means a speed at 98 % of maximum with the flow more than 10 % short for 30 s, SOCK HIGH is the switch, and PLUGGED is the switch with the flow under half its demand for 30 s or with the torque above 75 %, or AT LIMIT with either. Under SOCK HIGH or PLUGGED the speed may not rise above its present value, because more speed only fills the sock. Auto routing diverts during a transition, while a lane reads SOCK HIGH or PLUGGED, for 15 min after a knock recorded in the maintenance log, while the recipe is infeasible and while the predicted product is outside a band.

12.5.6 Supervisory MPC

The MPC replaces the ratio demands every 10 s in Auto and leaves the LIW loops underneath it. It has the general form of Section 1.6.3. One small problem per lane uses the declared sock, blender and extruder lags at the measured throughput, with the estimator's own lag states as the initial condition and the trim as the gain, as shown in Equation (12.10).

$$\begin{aligned} \min_{u} \; & \sum_{k} \left\| \frac{y_k - r_k}{tol} \right\|^2 + S \sum \left(\frac{\Delta u}{q_{des}} \right)^2 \\ & + R \sum \left(\frac{u - u_{ratio}}{q_{des}} \right)^2 + \rho \|\sigma\|^2 \\ & 0 \le u \le 0.95 \, \hat{F} \, n_{max}, \qquad |\Delta u| \le q_{des} \end{aligned} \tag{12.10}$$

where y_k is the predicted concentration of each active the lane carries, r_k the destination grade when that material reaches the tap (Preview on), $S = 2$, $R = 0.02$ and u_{ratio} the ratio demand. The bands are soft at $0.8\,tol$ about target with a price $\rho = 50/tol^2$ judged every ninth step with one slack per judged step. The horizon is 15 min, 90 steps, with twelve graduated move blocks per lane, four of one step each and the rest lengthening through 2, 2, 3, 4, 6, 9, then 15, to a final 45. While the supervisor holds a lane, u may not exceed its present flow. The lane takes the ratio demand for a decision whenever the plan's status is not optimal, a move is infinite or a move leaves its box. The MPC never reads a fault, the true sock or an unannounced event, and its capacity is what the estimated feed factor allows, so it cannot overcome a plug.

12.5.7 Plant, Estimator and Controller Models

Three models are kept apart in the app, the simulator's model of Section 12.3, the estimator's declared lags and compositions, and the MPC's three-state linear model per lane. Neither of the last two contains the sock holdup, the fill-level effect, aeration or the true potency. Twin runs that differ only in a hidden truth confirm the boundary. With a 25 % feed-factor loss on AO-2 the first measured difference appears at 2 s and the first speed difference at 5 s. With an 80 % sock restriction the torque differs at 2 s, but no speed differs until 1,602 s, one second after the high-level switch trips and the supervisor holds the lane. A resin background of 300 ppm AO-1 in case 1 and a masterbatch potency of 4.0 % on the shift as shipped change no speed in two hours. A source scan finds no read of the plant state, a disturbance, a run parameter, the case script, the scorecard or a laboratory record's hidden flags in the five controller-side files.

Table 12.3 The Ten Cases

#	Setup	Lesson
1	PE-A, slip neat or as 5 % masterbatch, 1 h	carrier displaces resin
2	PP-A, Manual, production 1,000 → 1,500 kg/h at 10 min	ratio, not speed
3	AO-2 feed factor −25 % over 30 min from 10 min	inside the scale
4	AO-2 sock restricted 80 % at 10 min, sample at 30 min	outside the scale
5	AO-2 sock plugged at 10 min, clean or knock	infeasible, slug
6	AO-2 refill at 220 s with the vent plugged	weight validity
7	PE-A → PE-B at 0.5 h, back at 2.5 h	transition mass
8	PP-B → PP-A at 0.5 h	clarifier washout
9	PP-A → PP-B at 1 h, samples every 2 h, one lost, one botched	late lots
10	12 h shift: two grade changes, 4.3 % masterbatch, AO-2 loss 20 %	challenge

Table 12.4 Measured Outcomes at the Default Seed

#	Result
4	Tap below 900 ppm at 12.4 min, switch and first diversion at 36.6 min. 403 kg off-spec prime. Assay #1 626 ppm against 1,000 ppm predicted, accepted as holdup evidence at 80 min, trim unchanged
5	Switch and diversion at 20.7 min, PLUGGED at 27 min, speed held at 77.8 rpm. Cleaned at 50 min: 293 g to waste, in band at 71.3 min, 148 kg off-spec prime. Knocked: peak 2,902 ppm, 169 kg off-spec prime on Auto, 903 kg on Prime
6	Timer only: speed down to 65.1 rpm, tap down to 911 ppm (MHE: 68.2, 929). Timer and pressure: 281 s volumetric, speed ≥ 75.2 rpm, tap 995–1,021 ppm
7	Transition mass ratio + Preview 410 kg (1.6, 3.0 min). MPC 361 kg (0.7, 1.0 min). Ratio without Preview 541 kg (5.5, 6.9 min)
8	Clarifier above 50 ppm until 35.7 min, 262 kg transition, predicted on spec at 45.8 min. Without Preview 39.7 min. MPC 34.3 min, 239 kg
9	#1 (PP-A lot, 0.90–0.98 h) reported at 1.73 h, clarifier 104 ppm against 89 ppm predicted. #2 lost. #3 reads 1,392 ppm AO-2 against 1,000, rejected. No trim moved
10	Shipped 0 kg on spec, 9,000 kg transition. Suggested (Lab trim and Preview on) 9,574 kg on spec, 450 kg transition, 1,896 kg off-spec prime before the first assay

12.6 Ten Cases, Measured

The ten cases are listed in Table 12.3. Every number in this section comes from headless runs of the engine at each case's default seed and shipped settings, or from three-seed means where so stated.

12.6.1 Recipes, Rates and the Two Restrictions

The measured outcomes of the cases at their default seed are collected in Table 12.4. In case 1 the slip lane moves from 0.5 kg/h neat to 10 kg/h as masterbatch, the resin falls from 995 to 985.5 kg/h, and the slip at the tap stays at 495 to 501 ppm and 498 to 501 ppm. Asked for 1,500 ppm of AO-2 the 1:2 premix is reported infeasible, its closest point being 577 and 1,154 ppm. In case 2 fixed speeds leave AO-2 at 660 ppm by 30 min and Auto routing diverts 1,176 kg of the hour, while ratio control takes the lane to 118.8 rpm with a worst error of 0.07 of a tolerance. In case 3 the speed averages 104.6 rpm between 50 and 90 min (100.6 to 108.5 rpm as the hopper cycles), the tap stays at 997 to 1,003 ppm, and the simple and MHE feed factors end at 0.00941 and 0.00944 kg/h/rpm against a true 0.00947 kg/h/rpm. The same buildup in Manual leaves AO-2 at 764 ppm by 90 min.

In case 4 the scale reads normal for the whole time the product is wrong. The phosphite at the tap is below 900 ppm from 12.4 min, the switch trips at 36.6 min with the LIW reading 1.03 kg/h, and Auto routing diverts on the switch in the same scan. By then 403 kg of off-specification product has

gone to prime, all of it made while every feeder screen read normal. The first assay, collected at 30 to 35 min and reported at 80 min, reads 626 ppm against a prediction of 1,000 ppm. It is accepted although it lies outside the 25 % gate, because the switch corroborates it, and for the same reason it is accepted as evidence of holdup and does not move the trim. The full-width view of this case at 45 min is shown in Figure 12.4. In case 5 the switch trips and diverts the product at 20.7 min, and the protection holds the speed at 77.8 rpm, so the lane never climbs toward 200 rpm. Cleaning sends the 293 g in the sock to waste. Knocking it clear with the route on Prime instead releases a slug that peaks at 2,902 ppm of phosphite, and the case ends with 903 kg of off-specification prime.

12.6.2 Refills, Transitions and Late Results

A plugged vent leaves 18 Pa in the hopper, which at 0.5 g/Pa is a 9 g weight error decaying with 90 s, and the derivative reads the decay as extra flow. With the timer alone the AO-2 speed dips to 65.1 rpm and the tap to 911 ppm, while the pressure check keeps the lane volumetric for 281 s and holds the tap at 995 to 1,021 ppm. The MHE with the timer alone dips less, to 68.2 rpm and 929 ppm. Grade changes in case 7 cost 410 kg of transition material with ratio control and Preview, 541 kg without Preview, and 361 kg with the MPC, which overfeeds within capacity on the way up and cuts early on the way down. In case 8 the clarifier feeder stops at 26 min, ahead of the 30 min change, and the clarifier stays above its 50 ppm limit until 35.7 min. No controller can remove it faster than the line washes out, because a feeder cannot deliver a negative flow. In case 9 the first result belongs to the PP-A lot although it arrives after the change to PP-B, the lost sample moves nothing, and the botched analysis is rejected.

12.6.3 The Four Architectures on Matched Runs

The comparison runs each case with ratio control or MPC, each with the simple estimate or the MHE, on identical measurements, schedules, fault injections and sample availability, at three paired seeds, as shown in Table 12.5. The conventional reference matches or beats the advanced options on the nominal, feed-factor-loss and holdup cases. On the nominal hour the mass-weighted error is 0.021 of a tolerance for ratio control with the simple estimate against 0.032 for MPC with the MHE, and on the feed-factor ramp it is 0.018 against 0.046. The MPC's product benefit is confined to grade transitions, 414 kg of transition material for ratio control against 370 and 365 kg for the two MPC variants, and to the delayed-laboratory case's transition, 219 against 192 kg. Over the twelve-hour shift with the suggested settings the four architectures make 9,576 kg (ratio, simple), 9,570 kg (ratio, MHE), 9,616 kg (MPC, simple) and 9,613 kg (MPC, MHE) on specification, within 0.5 % of each other. The MHE estimates the nominal feed factor with 0.28 % mean error against 0.47 % for the simple estimate, but it changes no product outcome measurably. Its visible effect is on the actuator. With the MPC, the AO-2 speed travels 895 rpm in the nominal hour against 1,463 rpm with the simple estimate, and 12,294 rpm against 20,428 rpm over the shift. Nothing beats the off-scale holdup, which no estimator can see: all four rows of case 4 end with 403 kg of off-specification prime, set by when the sock switch trips. Figure 12.3 draws the error and transition rows of Table 12.5.

> ✦ DID YOU KNOW? **The Plainest Pair Is the Most Accurate**
>
> From the simulator: in Table 12.5, ratio control with the simple estimate has the smallest error of the four architectures on the nominal hour, 0.021 of a tolerance, and on the feed-factor loss, 0.018. The MPC with the MHE scores 0.032 and 0.046 on the same runs.

12.6.4 How Long the Feeder MHE Should Look

The horizon was chosen by an open-loop test on the AO-2 lane in Manual, in which the feed factor falls 25 % over 60 s and detection is the first time an estimate falls past half the loss and stays there for 30 s, as shown in Table 12.6. A 60 s window detects in 64 s, faster than the simple estimate's 83 s, but with 0.88 % steady noise against 0.48 %. From 120 s up the noise falls to 0.36 to 0.38 % and

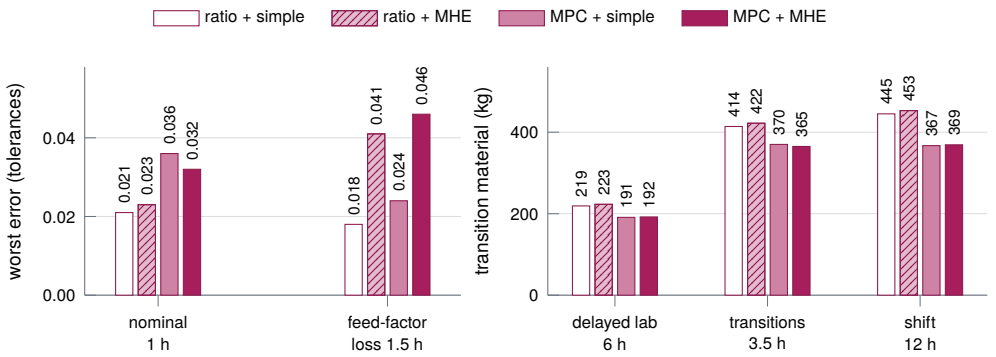


Figure 12.3 The error and transition rows of Table 12.5, three-seed means. Left, the worst active's mass-weighted error in tolerances. Right, the transition material in kilograms. Open and hatched bars are ratio control with the simple estimate and the MHE, medium and dark bars the MPC with the same two. Notice that ratio control with the simple estimate has the smallest errors, that the MPC's gain is confined to the transitions, and that the holdup and plug rows, not drawn, read 403 kg and 148 kg in every column.

Table 12.5 Four Architectures, Three Paired Seeds (Means)

Case, metric	R+S	R+M	M+S	M+M
nominal 1 h, error (tol)	0.021	0.023	0.036	0.032
nominal, AO-2 travel (rpm)	1,522	1,397	1,463	895
delayed lab 6 h, transition (kg)	219	223	191	192
feed-factor loss 1.5 h, error (tol)	0.018	0.041	0.024	0.046
holdup 2 h, off-spec prime (kg)	403	403	403	403
plug 2 h, off-spec prime (kg)	148	148	148	148
transitions 3.5 h, transition (kg)	414	422	370	365
shift 12 h, on spec (kg)	9,576	9,570	9,616	9,613
shift, transition (kg)	445	453	367	369
shift, AO-2 travel (rpm)	20,524	19,652	20,428	12,294

R ratio + PI, M MPC (first letter). S simple estimate, M MHE (second letter). Error is the worst active's mass-weighted error in tolerances. MPC: 0 fallbacks, 0.57 to 0.65 ms per decision for all lanes in Node.

Table 12.6 MHE Horizon Against the Simple Estimate (Three Seeds)

N (s)	noise MHE	noise simple	half-loss MHE	half-loss simple
60	0.88 %	0.48 %	64 s	83 s
120	0.38 %	0.48 %	87 s	83 s
300	0.36 %	0.48 %	91 s	83 s
600	0.36 %	0.48 %	91 s	83 s

detection slows to 87 to 91 s, and beyond 300 s nothing changes, because the random-walk step, not the window, sets the estimator's memory. At the shipped 300 s the MHE is quieter than the simple estimate but not faster, 91 s against 83 s, as shown in Figure 12.2(b). The 300 s window was kept because it spans a refill and its settling, up to 281 s of invalid weight in case 6, with valid data on both sides.

🕮 **Pause.** Try Exercise 12.3 before reading on.

12.6.5 Holdup Is Not a Delivery Loss

Following case 4 past its end tests how the laboratory trim treats a holdup. The trim corrects a lasting error in delivery, while the sock's shortfall is transient and leaves with the sock's contents. A result

taken while a lane's holdup alarms are raised is therefore accepted as evidence and leaves that lane's trim at zero, and a cleaning or a knock recorded in the maintenance log resets the trim. Cleaned at 90 min, the lane returns to its recipe flow and the phosphite at the tap averages 990 ppm between 110 and 170 min, with a minimum of 875 ppm in the restart and a maximum of 1,001 ppm. Knocked clear instead, it averages 1,000 ppm. At 150 min the estimate and the truth both read 1,000 ppm. By 3 h the off-specification prime is 403 kg after a cleaning and 407 kg after a knock, with the trim on or off. The app's suite checks the diversion at the switch, the unmoved trim and the 1,100 ppm band edge after the cleaning.

12.6.6 The Shift Challenge

The shift runs PE-A, PE-B from 3 h and PE-A again from 8 h, with a slip masterbatch that is really 4.3 % active and an AO-2 feed factor that falls 20 % from 1.5 h. Every scale reads correctly while the slip runs 14 % low, so as shipped, with Lab trim and Preview off, no product is on specification and all 9,000 kg made after the first change count as transition. The target is 9,200 kg on specification and at most 500 kg of transition by 12 h. With both settings on, the three seeds make 9,574 kg, 9,587 kg and 9,566 kg on specification with 450, 430 and 455 kg of transition, and all three pass. Lab trim alone makes 9,564 to 9,586 kg but 562 to 589 kg of transition and fails, and Preview alone makes none.

12.6.7 Verification

The total mass balance, inventories against the resin and refill in and the product and waste out, closes to 5.8×10^{-14} of the input over case 4 and to 8.9×10^{-14} through a sock restriction and a knock, and the app's own test holds it below 10^{-9} through refills, a plug, cleaning to waste and a knock. Halving the step to 0.5 s with the controllers still at 1 s moves case 7's transition mass from 409.72 to 409.31 kg, its recoveries from 93 and 180 s to 93 and 179.5 s and its peak from 5.853 to 5.869 tolerances, case 4's off-specification prime from 402.90 to 402.76 kg, and the slug's peak not at all at 2,871 ppm. The four-way benchmark reproduces the app's recorded file in every non-timing field of all 32 rows. In Node a simulated hour of the shift costs 0.86 s of wall time with ratio control and 1.05 s with MPC and MHE, and the browser plays the twelve-hour shift in 72 s at its top speed of 600 times real time. The app's own suite passes 53 of 53 checks.

12.7 Technology in Practice

What Measures It

Every lane's controlled flow comes from one measurement, the net weight of a hopper, its motor and its screw, read once per second as W_m in Equation (12.4). A real feeder obtains it from load cells under the weighing platform, and its controller divides each change of weight by the time between readings to obtain the flow [3]. The box below describes the cell and the forces that can pass for powder, and Figure 12.5 shows its principle.

> **TECHNOLOGY Sensor: Strain-Gauge Load Cells Under the Feeder**
>
> A load cell is a metal element of precisely known elasticity with strain gauges bonded to it. Under load the element compresses on a microscopic scale and the gauges change their resistance. A vessel usually stands on several cells whose signals are summed, because its weight is not shared equally among its supports. The cells read every force on the vessel: a pipe without a flexible coupling, vibration from nearby machinery, or a current through a cell from a welder grounded to the frame all look like a change in contents [20]. A flow taken from the change of weight is a net flow, so it cannot separate an outflow from a refill [20]. WT of the chapter stands for the summed cells, with 0.2 g of noise, a span error s and the hopper pressure reaching it at 0.5 g/Pa.

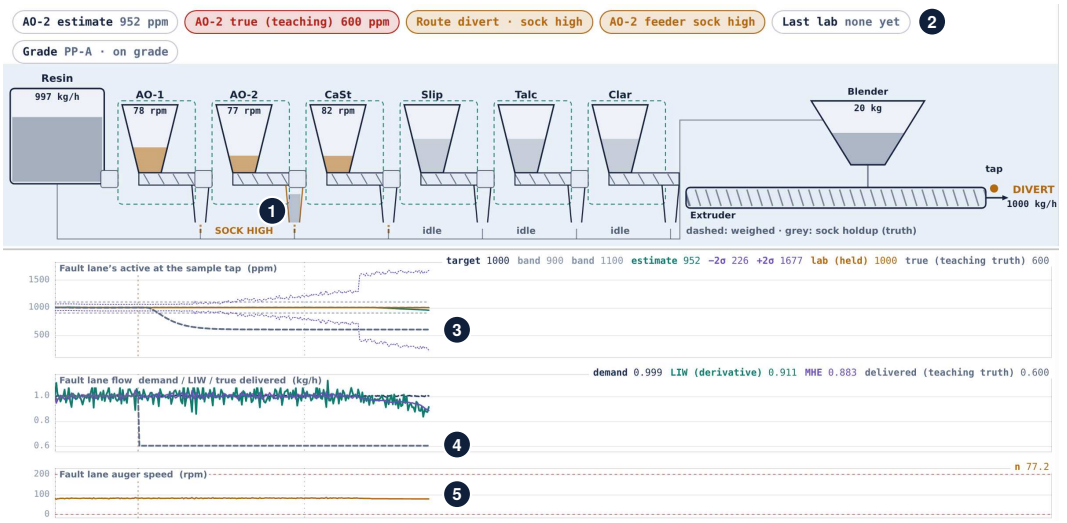


Figure 12.4 Hidden holdup of case 4 at $t = 45$ min on the Line view, cropped to its badges, its plant view and its first three trends. (1) The AO-2 lane reads SOCK HIGH with its sock drawn below the dashed weighing boundary. (2) The badges: the phosphite estimate is 952 ppm (± 725 ppm, two standard deviations) against a true 600 ppm and a target of 1,000 ppm, and the switch has sent the product to diversion. (3) The phosphite at the sample tap, the dashed true line below the estimate and its band. (4) The flow trend shows a demand of 0.999 kg/h, an LIW reading of 0.911 kg/h, an MHE flow of 0.883 kg/h and 0.600 kg/h delivered, (5) at 77.2 rpm.

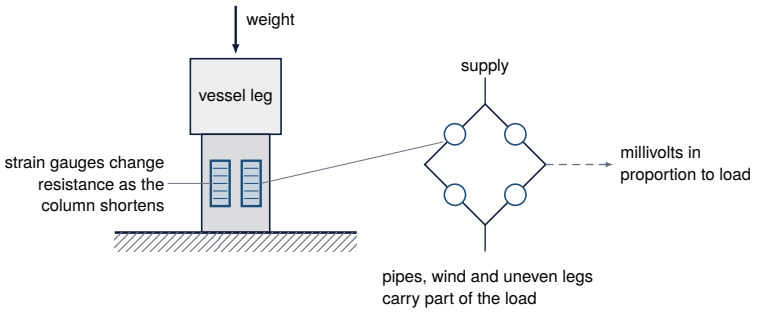


Figure 12.5 The strain-gauge load cell from the book's instrument library. Notice that the cell measures every force that reaches it through the leg. In the chapter's model the hopper pressure enters W_m this way, through the flexible connections, and the share α of the sock that the scale carries is set to 0, which is why the holdup of case 4 never reaches the weight.

What Moves It

The manipulated variable of each LIW loop is the auger speed n, and the torque JT is the lane's only continuous witness of a sock that pushes back. In a real feeder the speed is set by a motor drive turning a metering screw, as the box below explains. Neither the screw nor its drive can see the sock: a full sock throttles the screw through the backpressure factor $\mathrm{acc}(H)$, and the supervisor then forbids more speed.

TECHNOLOGY Final Element: A Metering Screw on a Variable-Speed Drive

A loss-in-weight feeder is a volumetric feeder, a hopper with a metering screw, mounted on a weighing platform under a gravimetric controller that corrects the screw speed whenever the weight lost per unit time departs from the set rate. Screws come single or twin, with various thread depths, as spirals or augers [3]. A variable-frequency drive sets the speed of an alternating-current motor through the frequency of its supply, and drives commonly limit speed, torque and acceleration [20]. A powder of low density can compact in the flights until the motor reaches its maximum torque and the feeder shuts down [3]. The chapter's screw runs from 0 to 200 rpm under a slew limit of 20 rpm/s, and JT reads its torque.

Where the Controller Runs

Figure 12.6 draws one lane of the chapter with hardware in place of the schematic's symbols. The gravimetric controller is part of the feeder itself, together with the volumetric feeder and the weighing platform [3]. The recipe, the ratio to the resin and the laboratory trim sit a level higher, in a control system whose supervisor sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [21]. In the app the six PI blocks run in velocity form every second, the MHE every 5 s and the MPC every 10 s, and the lane falls back to volumetric operation whenever its weight is invalid. The settings of that fallback are a matter of practice.

IN PRACTICE Choosing When to Refill

During a refill a loss-in-weight feeder runs volumetrically and is blind to density: the incoming powder can compress the bed below it or aerate it. The number of refills is therefore chosen so that the feed-rate deviation stays acceptable to the operations downstream. In the study a 2018 review reports, refilling a hopper at a low fill level such as 20 % gave larger deviations, while refilling at 60 % fed more accurately but demanded more frequent, smaller refills. A hopper designed for the feed rate and the powder also reduces the disturbance [3]. The chapter's lanes start a refill below 25 % of the hopper and stop at 85 %.

What Is Modernizing It

Three developments bear on this line. In-line Raman and near-infrared spectroscopy at the extruder reads the active content of the melt itself. In one study a model built on in-line Raman spectra saw a powder accumulation in the barrel that a soft sensor built on the feeders' data could not, because that sensor knew only what the feeders had dosed [22]. Near-infrared probes at a feeder outlet have been used to follow the bulk density of the powder and the screw speed that compensates for it [3]. A hybrid of predictive and PID control, informed by process analytical tools, has been implemented on a continuous direct-compaction pilot plant for tablets [16]. The first development is the one this chapter's holdup asks for, and the Open challenge box in Section 12.8 says why it is not yet routine.

12.8 Limitations

The equipment sizes, the sock capacity and backpressure, the torque relation, the pressure coupling and the bands are teaching values sized for one self-consistent line at 1,000 kg/h, not data from a plant or a supplier. The blender and extruder are ideal mixers, the actives are conserved, and the assays map exactly onto the actives with added noise, so the model cannot represent antioxidant consumption, segregation, axial dispersion or an ambiguous ash result. A polypropylene to polyethylene changeover is refused rather than modeled. The faults arrive as scheduled steps and ramps.

The MHE estimates only the feed factor and the weighed flow. It cannot see a span error, a potency error or the sock, and no confidence interval is claimed for it. The MPC plans each lane separately, has no resin-rate decision, and carries no stability guarantee beyond its linear prediction rebuilt every 10 s. The laboratory trim separates holdup from delivery only through the torque and the switch: a holdup that raises neither would still be learned as a lasting efficiency loss. Protection and routing are

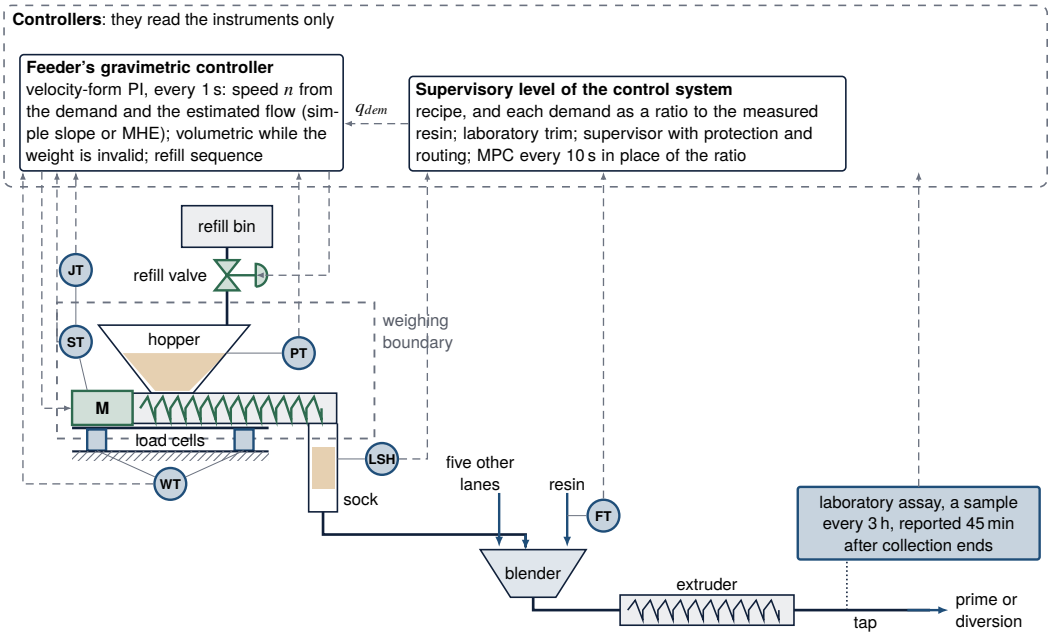


Figure 12.6 One additive lane with its hardware. Blue marks are sensing devices: the load cells WT under the weighing platform, the speed and torque of the drive, the hopper pressure, the sock high-level switch and the resin flow FT. Green symbols are final elements: the screw with its variable-speed drive and the refill valve. Dashed arrows are signals. Notice that the dashed weighing boundary encloses the hopper, the drive and the screw but not the sock, and that the only measurement of the product reaches the controllers from a laboratory, 45 min after each sample. The other five lanes and the diverter are left out. Not to scale, and not a piping diagram.

an educational supervisor, not certified interlocks, and a simulated pass is not evidence of compliance with any product specification.

> ⚠ **OPEN CHALLENGE** **Reading the Additive in the Melt**
>
> In-line spectroscopy could report an additive in the melt while the product is made, not hours later. Near-infrared and Raman measurements of molten polymer have been researched intensively, yet industry still applies them sparsely because their measurements are unpredictable and unstable [23]. In one study of polymer additive content in an extruder, a calibration built on near-infrared spectra from one time lost its accuracy on spectra collected at another [22]. Progress would be an in-line reading that keeps its calibration across days, grades and machines, so that a holdup like the sock of case 4 shows in the product long before the next assay.

12.9 Conclusion

A loss-in-weight feeder measures what leaves its hopper, and this simulator logs that reading beside the flow that reaches the blender, the holdup between them and the product at the tap. A restriction inside the weighing boundary is corrected by speed, from 80 to about 105 rpm for a 25 % loss, while one in the sock below it keeps the reading near 1 kg/h as the product falls to 600 ppm, and 403 kg of it reaches prime before the sock switch diverts it. On matched runs ratio control matches the advanced options except on grade transitions, where the MPC saves 44 to 49 kg of 414 kg, and the five-minute MHE estimates the feed factor with less error and less speed travel but reaches half a step loss in 91 s against 83 s. The holdup is unobservable from the weight, so the laboratory trim must not learn

from a result taken under the holdup alarms. The model is a teaching model with declared relations. Natural extensions are a measured sock load ($\alpha > 0$), a holdup estimate from the torque, antioxidant consumption in the extruder and a resin purge between polymer families. The simulator shows that a feeder's scale vouches for the hopper, never for the pellets.

KEY POINTS

- A loss-in-weight feeder weighs what leaves its hopper, never what reaches the pellets.
- A restriction inside the weighing boundary is corrected by speed. One in the sock below it keeps the reading near 1 kg/h while the product falls to 600 ppm.
- The holdup is not observable from the weight. Its witnesses are the torque, the high-level switch and the laboratory, so the trim must not learn from a result taken under the holdup alarms.
- On matched runs ratio control matches the advanced options except on grade transitions, where the MPC saves 44 to 49 kg of 414 kg.

Where else this applies. A dose set from an upstream measurement, with the slow laboratory result demoted to a trim, is also the coagulant feedforward of Volume 3, Chapter 6, and loss-in-weight feeders meter the dose ahead of the tablet press of Volume 5, Chapter 6.

References

[1] M. Hopkins, "Loss in weight feeder systems," *Measurement and Control*, vol. 39, no. 8, pp. 237–240, 2006.

[2] W. E. Engisch and F. J. Muzzio, "Method for characterization of loss-in-weight feeder equipment," *Powder Technology*, vol. 228, pp. 395–403, 2012.

[3] C. A. Blackshields and A. M. Crean, "Continuous powder feeding for pharmaceutical solid dosage form manufacture: a short review," *Pharmaceutical Development and Technology*, vol. 23, no. 6, pp. 554–560, 2018.

[4] J. Hanson, "Control of a system of loss-in-weight feeders for drug product continuous manufacturing," *Powder Technology*, vol. 331, pp. 236–243, 2018.

[5] J. N. Hahladakis, C. A. Velis, R. Weber, E. Iacovidou, and P. Purnell, "An overview of chemical additives present in plastics: Migration, release, fate and environmental impact during their use, disposal and recycling," *Journal of Hazardous Materials*, vol. 344, pp. 179–199, 2018.

[6] J. Pezzutti, A. Benito, G. Cassano, L. Roth, W. Schoene, H. Siebert, A. Winter, A. Dimeska, and V. Galiatsatos, "Supertransparent high impact strength random block copolymer," U.S. Patent 8,076,429, Dec. 13, 2011.

[7] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[8] ——, "Optimal grade transitions in a gas phase polyethylene reactor," *AIChE Journal*, vol. 38, no. 10, pp. 1564–1576, 1992.

[9] W. E. Engisch and F. J. Muzzio, "Feedrate deviations caused by hopper refill of loss-in-weight feeders," *Powder Technology*, vol. 283, pp. 389–400, 2015.

[10] F. Tahir, J. Palmer, J. Khoo, J. Holman, I. K. Yadav, G. Reynolds, E. Meehan, A. Mitchell, and G. Bajwa, "Development of feed factor prediction models for loss-in-weight powder feeders," *Powder Technology*, vol. 364, pp. 1025–1038, 2020.

[11] D. G. Robertson, J. H. Lee, and J. B. Rawlings, "A moving horizon-based approach for least-squares estimation," *AIChE Journal*, vol. 42, no. 8, pp. 2209–2224, 1996.

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

[13] 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.

[14] R. J. Harkenrider, "Apparatus for feeding material at a predetermined rate," U.S. Patent 2,544,155, filed 27 March 1945, granted 6 March 1951, 1951.

[15] H. H. Marhauer, "Mass flow controller," U.S. Patent 3,481,509, filed 18 October 1967, granted 2 December 1969, 1969.

[16] R. Singh, A. Sahay, K. M. Karry, F. Muzzio, M. Ierapetritou, and R. Ramachandran, "Implementation of an advanced hybrid MPC–PID control system using PAT tools into a direct compaction continuous pharmaceutical tablet manufacturing pilot plant," *International Journal of Pharmaceutics*, vol. 473, no. 1–2, pp. 38–54, 2014.

[17] H. J. Vandenburg, A. A. Clifford, K. D. Bartle, J. Carroll, I. Newton, L. M. Garden, J. R. Dean, and C. T. Costley, "Analytical extraction of additives from polymers," *The Analyst*, vol. 122, no. 9, pp. 101R–116R, 1997.

[18] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[19] 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.

[20] 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.

[21] 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.

[22] N. Munir, M. Nugent, D. Whitaker, and M. McAfee, "Machine learning for process monitoring and control of hot-melt extrusion: Current state of the art and future directions," *Pharmaceutics*, vol. 13, no. 9, p. 1432, 2021.

[23] S. Zhu, Z. Song, S. Shi, M. Wang, and G. Jin, "Fusion of near-infrared and raman spectroscopy for in-line measurement of component content of molten polymer blends," *Sensors*, vol. 19, no. 16, p. 3463, 2019.

12.10 Exercises

12.1 Concept check · 2 min

In case 4 the AO-2 discharge sock is restricted to a fifth of its healthy capacity, below the weighing boundary. Which measurements can reveal the holdup before a laboratory result does?

(a) The hopper weight, once the MHE replaces the simple 20-second slope

(b) The LIW derivative, once its window is lengthened from 20 to 60 seconds

(c) The auger speed, which climbs as the loop makes up the missing flow

(d) The motor torque and the sock high-level switch, read outside the weight

12.2 Concept check · 1 min

Assay #1 of case 4 reads 626 ppm of AO-2 against a prediction of 1,000 ppm, and the AO-2 sock switch was raised after its collection began. Why does the app accept the result without moving the AO-2 trim?

(a) The shortfall sits in the sock and leaves with it, so a trim would overdose after the repair

(b) The result is more than 25 % from its prediction, so it is held as a possible outlier

(c) The result arrives 45 min after collection, too late to describe the present product

(d) The MPC sets the AO-2 demand in Auto, so the laboratory trim is not used

12.3 Concept check · 2 min

Table 12.6 shows that lengthening the feeder MHE's window from 300 s to 600 s changes neither its noise nor its detection time. Why?

(a) The Huber penalty discards the older samples of a longer window as outliers

(b) The random-walk step of the feed factor, not the window, sets how far back the estimate remembers

(c) The solver stops after two reweighted solves, before the extra feed-factor blocks can move

(d) A refill falls in every 600 s window, and its samples are masked from the fit

12.4 Theory and calculation · 5 min

The PE-B grade at 1,000 kg/h of finished compound asks for 500 ppm of AO-1, 1,000 ppm of AO-2, 500 ppm of calcium stearate, 1,000 ppm of slip and 5,000 ppm of talc, and no clarifier. The resin carries none of them.

(a) With every additive fed neat, find each feeder's flow and the resin flow. *Checkpoint:* 0.5, 1, 0.5, 1 and 5 kg/h, and 992 kg/h of resin.

(b) The slip is fed instead as a 5 % masterbatch. Find the slip feeder's flow, the carrier it brings and the new resin flow. *Checkpoint:* 20 kg/h, 19 kg/h of carrier, 973 kg/h of resin.

(c) The masterbatch is really 4.3 % active while the recipe assumes 5 %. Find the slip in the product and say whether the ±10 % band holds it. *Checkpoint:* 860 ppm, outside the band.

12.5 Theory and calculation · 8 min

The AO-2 lane has the feed factor $K = 0.0125\,\mathrm{kg/h/rpm}$, a design rate of 1 kg/h at 80 rpm, a speed limit of 200 rpm, a sock of capacity $H_{cap} = 0.30\,\mathrm{kg}$ with transit time $\tau_{sock} = 10\,\mathrm{s}$ and outflow capacity $c_0 = 3q_{des}$, and a high-level switch at $0.6\,H_{cap}$. Take the hopper half full ($f = 1$) and no aeration.

(a) A buildup at the screw discharge removes 25 % of the delivery per revolution. Find the speed that restores 1 kg/h, and the largest loss the speed limit can make up. *Checkpoint:* 106.67 rpm; 60 %.

(b) The sock is restricted instead, with $r_B = 0.8$. Find its outflow capacity, and the smallest r_B at which no speed can deliver 1 kg/h. *Checkpoint:* 0.6 kg/h; $r_B > 2/3$.

(c) The screw keeps sending 1 kg/h into that sock from $t = 10\,\mathrm{min}$, starting from its healthy holdup. Find when the switch trips. *Checkpoint:* After 26.6 min, at $t = 36.6\,\mathrm{min}$.

12.6 Theory and calculation · 6 min

The blender holds 20 kg, the extruder 30 kg and the transport to the tap 16.7 kg, at 1,000 kg/h. A stopped feeder's concentration passes three first-order lags, the sock (10 s), the blender (72 s) and the extruder (108 s), and then the transport delay.

(a) Find the lead that Preview gives the feeders ahead of a scheduled grade change, and the transport delay alone. *Checkpoint:* 240 s; 60 s.

(b) The response of the three lags to a unit step down is $y(t) = \sum_i a_i e^{-t/\tau_i}$ with $a_i = \tau_i^2 / \prod_{j \neq i}(\tau_i - \tau_j)$. Find the three a_i, and show that $y = 1/40$ at $t = 520.5\,\mathrm{s}$. *Checkpoint:* $a_i = 0.0165$, -2.323 and 3.306; $y(520.5) = 0.0250$.

(c) The clarifier feeder of case 8 stops at 26 min from 2,000 ppm. Find when the clarifier at the tap falls below its 50 ppm limit. *Checkpoint:* 35.67 min.

12.7 App challenge · 9 min

Load scenario 3, "Gradual sock restriction: the speed rises a third". Despite its name, the fault is a buildup at the AO-2 screw discharge, inside the weighing boundary: the feed factor falls 25 % over 30 min from 10 min. Use the default seed and the speed 300× (300 simulated seconds per wall-clock second). The case pauses at 1.5 h. Export the CSV file, whose time column is in hours.

(a) Run the case as loaded, in Auto. Find the mean of `rpm_1` over the rows from 50 to 90 min, and the range of `truth_ppm_1` over the same rows. *Checkpoint:* 104.5 rpm, within 0.3 rpm; 997 to 1,004 ppm, within 1 ppm.

(b) Read `ff_simple_1`, `ff_mhe_1` and `truth_ff_1` in the last row. *Checkpoint:* 0.00941, 0.00944 and 0.00947 kg/h/rpm, within 0.00001.

(c) Press the reload button, switch to Manual before Run, and run to 1.5 h. Read `truth_ppm_1` in the last row. *Checkpoint:* 764 ppm, within 2 ppm.

Compare (a) with the 106.67 rpm of a half-full hopper.

12.8 App challenge · 8 min

Load scenario 4, "Hidden holdup: the scale says all is well", and change no setting. The AO-2 sock is restricted 80 % at 10 min. Use the default seed and the speed 300×. Run until the case pauses at 2 h and export the CSV file, whose time column is in hours.

(a) Find the mean of `liw_kgh_1` over the rows from 15 to 30 min, and read `truth_del_kgh_1`, `truth_ppm_1` and `est_ppm_1` in the row at 30 min. *Checkpoint:* 1.00 kg/h, within 0.01 kg/h; 0.600 kg/h, 600 ppm and 1,000 ppm, each within 0.5 %.

(b) Find the first row whose `route` reads `divert`, and its `route_why`. Read the Off-spec prime tile at 2 h. *Checkpoint:* 36.7 min, within 0.2 min, "sock high"; 403 kg.

(c) Find the first row in which `lab_ppm_1` differs from 1,000 ppm, its value, and `trim_1` in the last row. *Checkpoint:* 80.2 min, 626 ppm within 1 ppm; a trim of 0.

12.9 App challenge · 9 min

Load scenario 7, "PE-A ↔ PE-B on a known schedule". The grade changes to PE-B at 0.5 h and back to PE-A at 2.5 h. Use the default seed and the speed 600×, and run each case until it pauses at 3.5 h. Read the second mass and the last recovery on the Diverted · transition tile.

(a) Run the case as loaded, with Ratio + PI and Preview On. *Checkpoint:* 410 kg; last recovery 3.0 min.

(b) Set Preview to Off, press Reset and run again. *Checkpoint:* 541 kg; 6.9 min.

(c) Set Preview to On and Control structure to MPC, press Reset and run again. *Checkpoint:* 361 kg; 1.0 min.

Accept a mass within 2 kg. Explain what Preview and the MPC each change.

12.10 Application area · 4 min

A plugged vent filter leaves 18 Pa in a hopper at the end of a 20 s refill, and the pressure decays with a time constant of 90 s. The load cells read the pressure at 0.5 g/Pa, and the feeder treats the weight as valid again 10 s after the refill ends.

(a) Find the weight error at the end of the refill, and the extra flow the weight derivative reads 10 s after it. *Checkpoint:* 9 g; about 0.32 kg/h.

(b) Say what the AO-2 loop does with that reading, and what a plant can do about it. *Checkpoint:* It cuts the speed. Keep the weight invalid while the hopper pressure is out of its band, and keep the vent filter clean.

12.11 Application area · 3 min

In case 4, all four control architectures send 403 kg of off-specification product to prime, and the AO-2 at the tap falls below 900 ppm 24 min before the sock switch trips. A plant engineer is asked what would shorten that gap. Name two changes to the equipment or the instruments, and say why a sample requested at 10 min would not help.

12.12 Review the agent · 5 min

An AI assistant reviewing case 4 writes: "The simple LIW estimate is a 20 s slope of a noisy weight, which is why the restriction went unnoticed for 26 min. The MHE fits the last 5 min of weight against speed, and its feed-factor error on the nominal hour is 0.28 % against 0.47 %. Setting Estimator to MHE will detect the sock restriction sooner and cut the 403 kg of off-specification prime." Find the flaw and say how you would check it.