INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 8 Batch, Run-to-Run, and Hybrid Control

CHAPTER 8

TRISO Coating Reactor Simulator for Inferential Control of an Unmeasured Product

Pages 193–220

Edited by John D. Hedengren

How to cite this chapter. "TRISO Coating Reactor Simulator for Inferential Control of an Unmeasured Product," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 193–220.

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

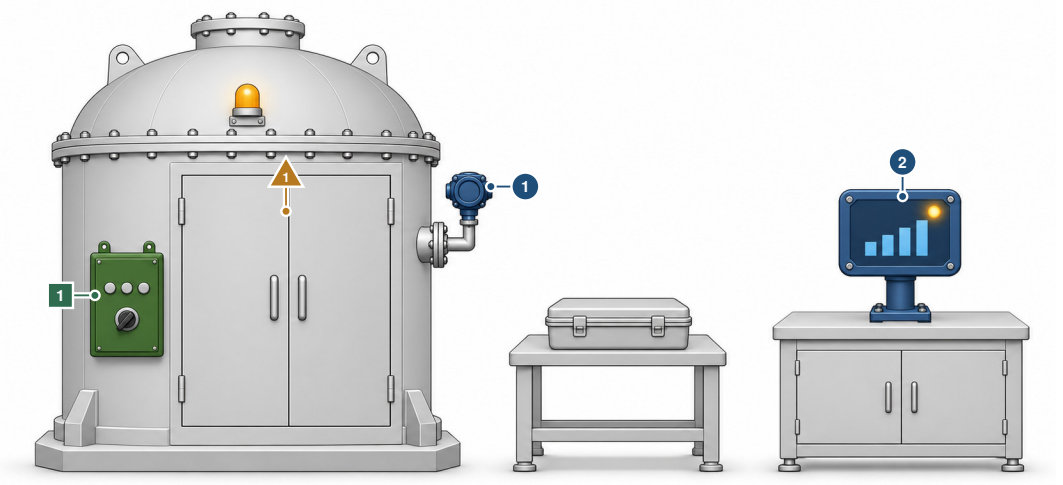
The app for this chapter
https://apmonitor.com/apps/triso-coating/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v08-batch-hybrid/#triso-coating

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

CHAPTER 8

TRISO Coating Reactor Simulator for Inferential Control of an Unmeasured Product



The controller moves 1 furnace power

The controller reads 1 bed temperature thermocouple 2 layer thickness, measured on a sample a week later

What disturbs it 1 coating efficiency nobody can see

Notice that the coating grows behind closed doors, so the thickness is read only on a sample at a separate laboratory station after the batch. Illustration, not to scale.

IN PLAIN TERMS **Coating Fuel Grains Nobody Can Watch**

The main idea. Nuclear fuel grains, each about the size of a grain of sand, get four thin coats inside a very hot furnace. No instrument can see a coat grow, so a computer model adds up how much should have landed and ends each coat when the sum reaches its target.

The challenge. If less of the gas sticks than the model assumes, the model still reaches its target on time while the real coats come out thin. Every reading in the control room looks normal.

Why it matters. Weighing the whole coater as it gains mass catches the missing coating and saves a batch of fuel that would otherwise be rejected.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history and the first fact in Section 8.2, the second fact and the two try-its in Section 8.6, and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 8.6), Technology in Practice (Section 8.8), Limitations (Section 8.9) and the open challenge

The control question No instrument reads a coating thickness while a TRISO fuel batch is in the furnace, so each of its four layers is ended on an integrated model. Which disturbances can such an endpoint defend against, and what does one real measurement, a load cell, add?

Learning objectives

- Compute the particle count, bed depth and minimum spouting velocity of a charge, and explain why the carrier flow must be scheduled.
- Explain why a timer, an inferential endpoint and an endpoint corrected by the bed mass give the same batch on a nominal plant and different ones under a disturbance.
- Invert a measured mass gain through the particle geometry into a layer thickness.
- Identify which errors a gain correction cannot see, and the check that catches them.
- Relate the fluidization setpoint to the spread of the silicon carbide layer on which a batch is accepted.

Prerequisites Sections 1.4.2 (guards that end a phase), 1.4.4 (endpoint constraints) and 1.6.2 (the PID of the furnace loop), and Chapter 6 for a hold ended on an observer's estimate.

The app https://apmonitor.com/apps/triso-coating/

See also Volume 5, Chapter 10 for a soft sensor corrected by a delayed laboratory result, and Volume 5, Chapter 11 for a product temperature no instrument reaches. Volume 18, Chapter 11 for inferential measurements on another nuclear-materials process, a waste vitrification melter.

8.1 Introduction

A block diagram teaches a student that a controller holding its measurement is controlling the process. On a fluidized-bed coater growing nuclear fuel that statement is false, and the falsehood is expensive. The bed temperature holds its setpoint to a mean error of 2.7 °C through the four layers, all three flows hold theirs, the bed spouts, the integrated thickness estimate climbs a smooth curve to exactly the right number, every layer ends on schedule, and the batch is 28.4 µm thin. The simulator of this chapter reproduces that run, and it reproduces it beside the nominal run so that the two are visibly identical on every trend the operator has. What separates them is a coating efficiency of 0.72 rather than 1.00, and no instrument in the coater carries that number.

The misconception the simulator corrects is the substitution of the measured variable for the variable of interest. A coating thickness cannot be measured through the wall of a 1500 °C graphite furnace, so the controlled variable is computed from an integrated deposition model driven by a thermocouple, three mass flow controllers and an estimate of how much particle surface is in the pot. The application integrates both models, the plant's seven-class truth and the control system's single-particle estimate, and the controller sees only the second. A transfer function between precursor flow and thickness cannot pose the problem, because the error that matters is a parameter inside the estimator rather than a term in the plant.

The class of problem is closed-loop control on an unmeasured quality variable through a soft sensor, and it is general. Reviews of soft-sensor practice place the technique wherever the economically important variable is unmeasurable, expensive to measure, or measurable only after a delay [1, 2]. Semiconductor film thickness is measured between wafers rather than during a deposition, so the loop that sets the process runs run to run rather than within run [3], and polymer melt index is inferred from reactor conditions between laboratory samples hours apart and corrected by the sample when it comes [4]. Pharmaceutical tablet coating, distillation composition from tray temperatures and kiln free lime have the same shape, and Volume 5, Chapter 10 corrects a psychrometric soft sensor of powder moisture with a delayed laboratory result. Two features recur in all of them. Model error and estimator error are indistinguishable from process error until an independent measurement arrives, so the design question is never how the model is tuned but which of its parameters can move without any instrument noticing. And the batch is accepted on the variance rather than the mean, because a population of a hundred thousand particles is only as good as its thin tail, which is why the acceptance criterion here is a standard deviation that nothing on the operating screen displays.

This chapter sets out a browser-hosted TRISO coating reactor in the detail needed to rebuild it, and shows on that plant that an endpoint scheme defends against exactly the disturbances its estimator can see and against nothing else. The analysis rests on three measured results. A silicon carbide step run 50 °C cold gives 26.2 µm on a timer and 35.0 µm on the integrated estimate for 21.6 min of extra time. A coating efficiency of 0.72 leaves the same estimate 28.4 µm thin with all four layers rejected and nothing moving on any trend, and a load cell closes that gap to 0.79 µm for 28 extra minutes. And the silicon carbide standard deviation rises from 0.61 to 3.53 µm as the spouting ratio falls from 2.20 to 1.15 with every mean thickness within 0.15 µm of target. The implementation is a 34-state plant carrying a seven-class particle ensemble, a ten-phase recipe sequencer, an instrument chain, a cruder estimator model, four endpoint architectures and fourteen scripted scenarios, the last of which withholds the one parameter the reader most needs. The deposition kinetics, the Mathur-Gishler correlation and the correction of a fast model by a slow measurement are all established methods.

The chapter is organized as follows. Section 8.2 recalls coated particle fuel and why its coaters became a control problem. Section 8.3 sets down the states, the constants and the instruments in enough detail for the plant to be reproduced, and closes with the assumptions. Section 8.4 gives the thickness estimator the control system runs and its correction by the load cell. Section 8.5 describes the temperature loop, the fluidization schedule, the recipe sequencer with its four endpoint architectures and the acceptance criteria. Section 8.6 reports the measured batches: the nominal run, the temperature loop, the fluidization window, the central result on coating efficiency, what the correction cannot see, variability, errors upstream of the instruments and the graded challenge. Section 8.8 describes the measurements, gas trains and control platforms of real coaters, Sections 8.7 and 8.9 discuss the structure of the result and state the limitations, Section 8.10 summarizes, and the exercises follow in Section 8.11.

8.2 Coated Particle Fuel and Its Coaters

Coated particle fuel began as a way to keep fission products inside a fuel element that had no cladding. The Dragon project, a cooperative European high-temperature reactor program operated at Winfrith, was built to test fuel for a helium-cooled reactor and had to solve the containment problem inside the fuel particle itself [5]. The pyrolytic-carbon-coated particle came out of that work, and the four-layer arrangement that followed placed a porous buffer against the kernel, a dense inner pyrocarbon over it, a silicon carbide shell over that and a dense outer pyrocarbon on the outside. The silicon carbide layer is the pressure boundary and the barrier to metallic fission products, and it is the layer that decides whether the fuel is any good [6].

Two reactor programs then ran the fuel at scale and produced the operating record the fabrication specifications were written from. The AVR pebble-bed reactor at Jülich raised its average hot-gas temperature to 950 °C and reported the fission-product release from coated-particle fuel under that duty [7], and Fort St. Vrain in Colorado ran prismatic coated-particle fuel as a commercial unit [6]. The German program that grew out of AVR converted to low-enriched uranium dioxide TRISO and, over two decades of irradiation and heating tests, established that particles made to its process released almost nothing below 1600 °C [8]. A direct comparison of the German and United States campaigns found the difference to be in fabrication rather than in the design of the particle, and identified the coating conditions, the kernel and the resulting layer properties as what separated fuel that performed from fuel that did not [9]. That result is why a coater is a control problem rather than a furnace with a timer on it.

The United States answered it with a qualification program rather than with a single recipe. The Advanced Gas Reactor fuel development and qualification campaign, run from Idaho National Laboratory with Oak Ridge National Laboratory, was laid out as a sequence of irradiations, AGR-1 through AGR-5/6/7, each preceded by a fabrication campaign whose product was characterized before it went into a capsule [10]. Its recipes, its coaters and the result of its first irradiation are set

out in the history box that follows. The layer thicknesses used here are the AGR-2 specification and the coater is a 75 mm laboratory unit of the size those campaigns ran development batches in.

> ⧖ HISTORY **2002 to 2018: Several Recipes Before One**
>
> A national program to qualify coated particle fuel for licensing began in 2002 [6]. Earlier campaigns in the United States had coated their particles under one nominal set of conditions, and the fuel often performed below what was required under irradiation. The first irradiation of the new program therefore carried a baseline recipe and three variants: a hotter inner pyrocarbon step, a richer inner pyrocarbon gas, and a silicon carbide step at a lower temperature with argon mixed into its hydrogen. Those particles came from a laboratory coater of 2 inches, and the later irradiations from a pilot-scale coater of 6 inches [11]. The first irradiation reached a peak burnup of 19.6 % fissions per initial metal atom with no failed particle detected among the roughly 300 000 irradiated [12].

Acceptance moved from the mean to the distribution over the same period, and for a mechanical reason. A batch is a hundred thousand independent pressure vessels, the internal gas pressure they hold is common to all of them, and the through-wall stress in each is set by its own silicon carbide thickness, so the probability that any particle fails is fixed by the thin tail of the thickness distribution rather than by its mean. Characterization for the AGR campaigns is correspondingly statistical, reporting a mean and a dispersion measured on a sample of particles from every coating batch [6], and a specification written that way is not defended by any loop that holds a mean. Fabrication has continued to be studied as a multiscale problem, coupling the bed hydrodynamics that decide how uniformly the particles see the gas to the chemistry that decides how fast the layer grows [13]. What has not changed since Dragon is that the growing layer cannot be seen. The batch is sectioned, polished and measured about a week after it leaves the coater, so the feedback loop that actually sets the process is a week long and everything inside the batch is open loop with respect to the property that matters. That is the situation inferential control exists for.

> ✦ DID YOU KNOW? **Ten Thousand Particles to a Pebble**
>
> The most common fuel pebble holds around ten thousand coated particles in a sphere 60 mm across. One fuel line built for a pebble-bed demonstration plant was designed for 300 000 spheres a year, and more than 500 000 had been made for its first core [6].

8.3 Coater Model

Coater CV-301 is a conical spouted bed of 75 mm diameter at the cylindrical section with an 8 mm gas inlet nozzle in the bottom of the cone, standing inside a graphite resistance furnace whose susceptor surrounds the vessel. The charge is 200 g of 500 µm uranium oxycarbide kernels, 279 865 particles presenting 0.220 m^2 of surface in a 7.16 mm bed. Carrier gas and precursor enter through the nozzle from three mass flow controllers, off-gas and soot leave through the freeboard and a filter, thermocouple TE-301 reads the bed, differential pressure cell PDT-301 reads across the bed, and load cell WT-301 carries the whole coater. Figure 8.1 shows the arrangement, the instruments and every symbol used in the equations that follow. The clock is in minutes, because a coating run is three hours long and the silicon carbide layer alone is an hour of it.

Table 8.1 classifies what the control system interacts with. The first entry is the one the application is about. A layer thickness is a controlled variable with no transmitter, so it appears only as the estimate $\hat{\delta}_L$, and the spouting ratio R is in the same position for a different reason. Every unmeasured disturbance listed is a quantity the estimator either assumes or reads through an instrument that does not see it. Thirty-four states are integrated, six for the plant and twenty-eight for the particle ensemble.

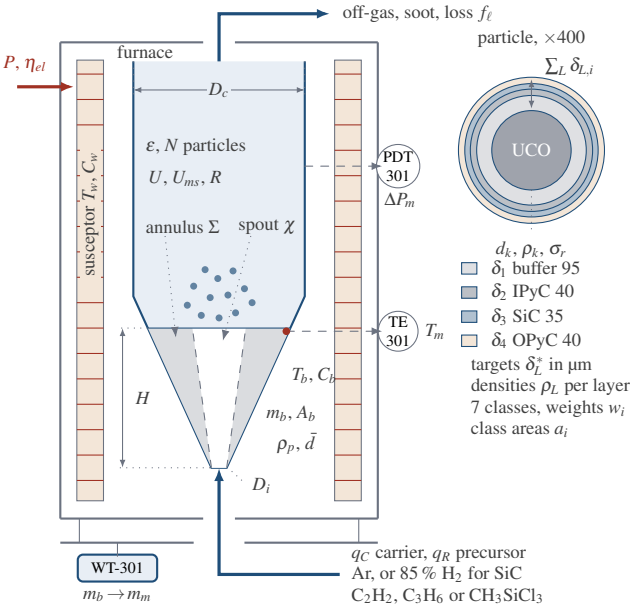


Figure 8.1 Conical spouted-bed coater CV-301 in cutaway inside its graphite furnace, with one particle magnified beside it. Carrier gas q_C and precursor q_R enter through the nozzle of diameter D_i and spout a bed of depth H in a column of diameter D_c, throwing a fountain of particles that rain back down the dense annulus. Superficial velocity U against the minimum spouting velocity U_{ms} sets the ratio R, which sets the circulation quality χ and the segregation strength Σ that shares the deposition flux unevenly among the seven size classes. Furnace power P heats the susceptor at T_w, which heats the bed at T_b. Only three instruments read the pot, thermocouple TE-301, differential pressure cell PDT-301 across the bed and load cell WT-301 under the whole coater, and none of them reports a thickness $\delta_{L,i}$.

Table 8.1 Classification of Variables. Nominal Values Are Those at the Start of the Buffer Layer, Where Every Scenario That Coats Begins.

Symbol	Role	Nominal	Range
$\hat{\delta}_L$	CV, inferred only	0	0–150 µm
$R = U/U_{ms}$	CV, inferred only	2.20	0.80–4.00 SP
T_b (TE-301)	CV, measured	1250 °C	1100–1650 °C SP
P	MV	2.04 kW	0–12 kW
q_C	MV	18.4 slpm	0–300 slpm
q_R	MV, recipe only	5.5 slpm	0–40 slpm
η_d coating efficiency	DV, unmeas.	1×	0.40–1.30×
ε_m charge error	DV, unmeas.	0 %	±20 %
γ_C, γ_R delivered	DV, unmeas.	1×	0.4–1.3, 0.5–1.4×
η_{el} element condition	DV, unmeas.	1×	0.35–1.10×
b_{TC} thermocouple bias	DV, unmeas.	0 °C	±80 °C
T_w, T_b	states	1275, 1250 °C	0–2600, 2400 °C
$\tilde{q}_C$, $\tilde{q}_R$	states	18.4, 5.5 slpm	0–300, 40 slpm
f_ℓ, a loss, agglomeration	states	0	0–0.80, 0–1.50
$\delta_{L,i}$, 4×7	states	0	m

8.3.1 The Particle Ensemble

The bed is represented by $N_c = 7$ size classes on a Gaussian grid, $z_i = -3, -2, \ldots, 3$, with weights and kernel diameters given by Equation (8.1).

$$w_i = \frac{\mathrm{e}^{-z_i^2/2}}{\sum_j \mathrm{e}^{-z_j^2/2}}, \qquad d_{k,i} = d_k\,(1 + z_i\,\sigma_r) \tag{8.1}$$

where d_k is the mean kernel diameter and σ_r the relative standard deviation of the kernel size distribution. Each class has a kernel mass $m_{k,i} = \rho_k \pi d_{k,i}^3/6$ and four coating thicknesses. Writing $r_{0,i} = d_{k,i}/2$ and $r_{L,i} = r_{L-1,i} + \delta_{L,i}$ for $L = 1 \ldots 4$, the class mass, diameter and surface area follow from Equation (8.2).

$$\begin{aligned} m_i &= m_{k,i} + \tfrac{4}{3}\pi \sum_{L=1}^{4} \rho_L \left(r_{L,i}^3 - r_{L-1,i}^3\right) \\ d_i &= 2 r_{4,i}, \qquad a_i = \pi d_i^2 \end{aligned} \tag{8.2}$$

where ρ_L is the density of layer L. The number of particles actually in the pot is set by the charge as it was weighed, by the weighing error and by the fraction carried over, as shown in Equation (8.3).

$$N = \frac{m_0\,(1 + \varepsilon_m/100)}{\bar{m}_k}\,(1 - f_\ell), \qquad \bar{m}_k = \sum_i w_i m_{k,i} \tag{8.3}$$

The bed mass, surface area, mean particle density and bed depth are the class averages multiplied by that count, as shown in Equation (8.4).

$$\begin{aligned} m_b &= N \sum_i w_i m_i, \qquad A_b = N \sum_i w_i a_i \\ \rho_p &= \frac{\sum_i w_i m_i}{\sum_i w_i \pi d_i^3/6}, \qquad H = \frac{m_b}{\rho_p (1 - \varepsilon) A_{col}} \end{aligned} \tag{8.4}$$

where ε is the loose-packed voidage and $A_{col} = \pi D_c^2/4$ is the cylindrical cross-section. A_b is the denominator of the whole deposition problem, and it runs from 0.220 to 0.745 m^2 through a batch.

8.3.2 Gas, Fluidization and the Spouting Ratio

The carrier is argon for the three carbon layers and 85 % hydrogen for silicon carbide, and its molar mass and molar heat capacity are mixed linearly in the hydrogen fraction f_{H_2}. The precursor is not a trace, because acetylene is nearly a quarter of the buffer gas and methyltrichlorosilane is 149 g/mol against hydrogen's 2, so it is carried explicitly in the mixture properties of Equation (8.5).

$$\begin{aligned} q_C &= \tilde{q}_C\,\gamma_C, \quad q_R = \tilde{q}_R\,\gamma_R, \quad q_T = q_C + q_R, \quad y_R = \frac{q_R}{q_T} \\ M_{mix} &= \frac{q_C M_c + q_R M_r}{q_T}, \qquad \rho_g = \frac{P_{atm} M_{mix}}{\mathscr{R}\,T}, \\ U &= \frac{q_T \times 10^{-3}\,(T/T_0)}{60\,A_{col}} \end{aligned} \tag{8.5}$$

where $\tilde{q}_C$ and $\tilde{q}_R$ are the flows the mass flow controllers indicate in standard liters per minute, γ_C and γ_R are the fractions actually delivered to the cone, M_c and M_r are the carrier and precursor molar masses, T is the bed temperature in kelvin, T_0 is 273.15 K, U is the superficial velocity in m/s and the factor T/T_0 converts standard liters to bed conditions. The minimum spouting velocity is the Mathur-Gishler correlation for a conical-cylindrical spouted bed [14], as shown in Equation (8.6).

$$U_{ms} = \frac{\bar{d}}{D_c}\left(\frac{D_i}{D_c}\right)^{1/3} \sqrt{\frac{2gH(\rho_p - \rho_g)}{\rho_g}}, \qquad R = \frac{U}{U_{ms}} \tag{8.6}$$

where $\bar{d} = \sum_i w_i d_i$ is the mean particle diameter, D_i the inlet nozzle diameter and R the spouting ratio. Equation (8.6) states the control problem in the hydrodynamics on its own. The requirement rises linearly with the particle diameter, with the square root of the bed depth and with the square root of the particle-to-gas density ratio, and the coating is moving all three. Below $R = 1$ the spout collapses, and above the carry-over ratio R_{carry} the fountain clears the freeboard.

The bed pressure drop is the one hydrodynamic measurement available, and Equation (8.7) shows why it is nearly useless for control.

$$\Delta P = \frac{f_s m_b g}{1000 A_{col}} \underbrace{\min(R^{0.9}, 1)}_{\text{packed branch}} \underbrace{\left[1 + \Delta p_{pk}\, \mathrm{e}^{-((R-1)/0.20)^2}\right]}_{\text{onset peak}} \tag{8.7}$$

where f_s is the fraction of the bed weight the gas carries when the bed is spouting, Δp_{pk} is the fractional overshoot at the onset of spouting, and the result is in kPa. Above $R = 1$ the first bracket is unity and the second decays, so the reading settles at about three quarters of the bed weight per unit area and then barely moves with flow. It says whether the bed is spouting. It says nothing about how hard, which is why the fluidization loop of Section 8.5 runs on a computed ratio instead. Circulation quality is a normalized distance from the fully developed spout, as shown in Equation (8.8).

$$\chi = \mathrm{clamp}\left(\frac{R - 1}{R_{good} - 1}, 0, 1\right) \tag{8.8}$$

where R_{good} is the ratio at which circulation is fully developed.

8.3.3 Deposition Kinetics

Deposition is two limits in series. The kinetic limit is Arrhenius in bed temperature and a power law in precursor mole fraction, which is the form fitted to pyrocarbon deposition from light hydrocarbons [15] and to silicon carbide from methyltrichlorosilane [16]. The supply limit is the precursor arriving, multiplied by the fraction of it that lands on a particle, divided by the total particle surface area. Equation (8.9) gives both and combines them.

$$\begin{aligned} s &= \mathrm{clamp}\left(\frac{y_R - y_{soot}}{w_{soot}}, 0, 1\right), \qquad \eta = \eta_L\, \eta_d\, (1 - k_s s) \\ J_{sup} &= \frac{\eta\, Y_L\, q_R}{A_b} \\ J_{kin} &= k_L \exp\left[-\frac{E_{a,L}}{\mathscr{R}}\left(\frac{1}{T} - \frac{1}{T_{ref,L}}\right)\right]\left(\frac{y_R}{y_{ref,L}}\right)^{n_L} \\ J &= \frac{J_{kin} J_{sup}}{J_{kin} + J_{sup}}, \qquad \dot{m}_{dep} = J A_b \end{aligned} \tag{8.9}$$

where J is the bed-average deposition flux in kg/(m^2 min), s is the soot fraction, y_{soot} is the hydrocarbon mole fraction at which carbon begins nucleating in the gas phase, w_{soot} is the width of that transition, k_s is the efficiency lost to fully developed sooting, η_L is the coating efficiency the recipe assumes for layer L, η_d is the coating-efficiency disturbance, Y_L is the mass of deposit a standard liter of precursor could make if every molecule of it landed on a particle, k_L is the kinetic flux at the reference condition, $E_{a,L}$ is the activation energy, $T_{ref,L}$ and $y_{ref,L}$ are the reference temperature and mole fraction, and n_L is the reaction order. A harmonic mean is the standard way to write a process that is kinetically controlled when it is cold and feed-limited when it is hot, and it is why a temperature excursion changes the rate at all. η_d is the disturbance the estimator has no variable for, because Section 8.4 uses η_L alone.

The bed-average flux is shared out unevenly among the classes. A fully developed spout turns the bed over every few seconds and every particle sees the same average gas, while a slack bed leaves

the annulus as a dead zone in which the heavier particles, which sink, spend proportionally longer. Equation (8.10) gives the strength of that effect and the per-class flux factors.

$$\begin{aligned} \Sigma &= \sigma_{min} + (\sigma_{max} - \sigma_{min})(1-\chi)^{3/2} + \sigma_s\, s \\ f_i &= \max\!\left(0.05,\; 1 - \Sigma\,\frac{m_i - \sum_j w_j m_j}{\sum_j w_j m_j}\right) \cdot \frac{\sum_j w_j a_j}{\sum_j w_j a_j f_j^{0}} \end{aligned} \tag{8.10}$$

where Σ is the segregation strength, σ_{min} and σ_{max} are its values at full circulation and at minimum spouting, σ_s is the extra segregation contributed by a sooting bed, and f_i^0 is the unnormalized factor. The final quotient normalizes the factors so that the area-weighted mean is unity, which leaves the bed depositing JA_b in total and lets the classes disagree only about who receives it. The active layer then grows by Equation (8.11).

$$\frac{\mathrm{d}\delta_{L,i}}{\mathrm{d}t} = \frac{J\, f_i}{\rho_L}, \qquad \frac{\mathrm{d}\delta_{K,i}}{\mathrm{d}t} = 0 \;\text{ for } K \neq L \tag{8.11}$$

Equation (8.11) is where poor fluidization becomes a standard deviation rather than a slower mean. The bed mean and spread reported to the scorecard are the class averages of Equation (8.12).

$$\bar{\delta}_L = \sum_i w_i \delta_{L,i}, \qquad \sigma_L = \left[\sum_i w_i \left(\delta_{L,i} - \bar{\delta}_L\right)^2\right]^{1/2} \tag{8.12}$$

8.3.4 Furnace, Bed Energy and Actuators

Power enters a graphite susceptor, crosses to the bed by conduction and radiation lumped into one temperature-dependent conductance, and leaves as radiation from the susceptor and as sensible heat in the gas. Equations (8.13) and (8.14) are the two energy balances.

$$UA = UA_0 \left(\frac{T_w + T_0}{1773.15\,\mathrm{K}}\right)^3 \tag{8.13}$$

$$\begin{aligned} C_w \frac{\mathrm{d}T_w}{\mathrm{d}t} &= 60\,P\,\eta_{el} - UA\,(T_w - T_b) - Q_{rad} \\ Q_{rad} &= k_{rad}\,\phi\left[(T_w{+}T_0)^4 - (T_a{+}T_0)^4\right] \\ C_b \frac{\mathrm{d}T_b}{\mathrm{d}t} &= UA\,(T_w - T_b) - C_g\,(T_b - T_g) + \Delta H_L\,\dot{m}_{dep} \end{aligned} \tag{8.14}$$

where T_w is the susceptor temperature, C_w and C_b are the susceptor and bed heat capacities in kJ/K, P is the furnace power in kW, the factor 60 converts it to kJ/min, η_{el} is the element condition, Q_{rad} is the susceptor's radiative loss with coefficient k_{rad} and multiplier ϕ, T_a is room temperature, T_g is the inlet gas temperature and ΔH_L is the heat released per kilogram of deposit. The gas heat capacity flow is $C_g = (q_T/V_m)\,C_{p,mol}/1000$ in kJ/(K min) with V_m the standard molar volume. The cube in (8.13) carries the radiative part of the wall-to-bed transfer, and the fourth power in Q_{rad} is the loss that makes the loop gain a function of temperature. Both balances carry kJ/min on both sides.

Pause. Try Exercise 8.2 before reading on.

The two gas trains are first-order lags on the commanded flows, as shown in Equation (8.15), and the carry-over and agglomeration states are given by Equation (8.16).

$$\tau_{MFC}\frac{\mathrm{d}\tilde{q}_C}{\mathrm{d}t} = u_C - \tilde{q}_C, \qquad \tau_{MFC}\frac{\mathrm{d}\tilde{q}_R}{\mathrm{d}t} = u_R - \tilde{q}_R \tag{8.15}$$

$$\begin{aligned} \frac{\mathrm{d}f_\ell}{\mathrm{d}t} &= k_{el}\,\max(0,\, R - R_{carry})^{3/2}\,(1 - f_\ell) \\ \frac{\mathrm{d}a}{\mathrm{d}t} &= k_{ag}\,\max(0,\, 1-R)\begin{cases} 1 & \text{depositing} \\ 0.25 & \text{otherwise} \end{cases} \end{aligned} \tag{8.16}$$

Table 8.2 The Four-Layer Recipe. Targets Are the AGR-2 Specification. Every Entry in the Upper Block Is Adjustable From the Screen.

	Buffer	IPyC	SiC	OPyC
Deposit	porous PyC	dense PyC	SiC	dense PyC
Precursor	C_2H_2	C_2H_2/C_3H_6	CH_3SiCl_3	C_3H_6
Target δ_L^* (µm)	95	40	35	40
T_{sp} (°C)	1250	1300	1500	1300
q_R (slpm)	5.5	4.0	4.5	5.0
Hold time (min)	26.0	24.0	59.5	24.0
ρ_L (kg/m^3)	1050	1900	3200	1900
k_L (10^{-3} kg/(m^2 min))	13.00	6.300	3.050	5.100
$E_{a,L}$ (kJ/mol)	190	200	230	200
$T_{ref,L}$ (°C)	1250	1300	1500	1300
$y_{ref,L}$	0.240	0.0878	0.0698	0.0603
Y_L (10^{-3} kg/L)	1.0719	1.3399	1.7893	1.6078
η_L assumed; n_L	0.35; 1.0	0.50; 1.0	0.35; 0.8	0.50; 1.0
f_{H_2}; M_r (g/mol)	0; 26.0	0; 34.1	0.85; 149.5	0; 42.1

Particles that leave take their surface area out of Equation (8.3), so the ones that remain coat faster than the model expects. The agglomeration index declares the batch lost at $a = 1$.

8.3.5 Instruments

Three instruments and two flow indications are all the control system has. The thermocouple carries a bias, the load cell is filtered hard because the bed is bouncing on the gas, and both carry first-order-filtered Gaussian noise scaled so that the stationary standard deviation matches the specified value (0.8 °C and 0.5 g by default), as shown in Equation (8.17).

$$\begin{aligned} \tau_{TC}\frac{\mathrm{d}T_m}{\mathrm{d}t} &= T_b + b_{TC} + n_T(t) - T_m \\ \tau_{DP}\frac{\mathrm{d}\Delta P_m}{\mathrm{d}t} &= \Delta P - \Delta P_m \\ \tau_{cell}\frac{\mathrm{d}m_m}{\mathrm{d}t} &= m_b + n_m(t) - m_m \end{aligned} \tag{8.17}$$

where T_m, ΔP_m and m_m are the measured bed temperature, bed pressure drop and bed mass. Each filter is advanced at the control scan by the exact zero-order-hold update $x_k = \alpha x_{k-1} + (1-\alpha)u_k$ with $\alpha = \exp(-\Delta t/\tau)$, so no filter can go unstable at any sample period. The mass flow controllers report $\tilde{q}_C$ and $\tilde{q}_R$, which are the flows through the meters and not the flows into the cone.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.3.6 Parameters, Initialization and Numerics

The recipe is given in Table 8.2 and the remaining constants in Table 8.3. The pre-exponential factors k_L are fitted so that each layer takes the time a real run takes at the recipe's flow and temperature, while the activation energies, the reference temperatures and the stoichiometric yields are physical.

Every scenario opens at a genuine steady state. The thicknesses are set to the targets of any layer the chosen start phase has already passed, the bed temperature to the recipe value for the current phase, the precursor flow to the recipe value, and the carrier flow to the value that puts the bed at the spouting ratio setpoint. That flow has to be found iteratively, because the velocity required depends on the gas density, which depends on the composition, which depends on how much carrier the answer turns out to be. Six passes of Equation (8.6) put the opening ratio within a part in ten thousand of the setpoint, except at the start of the silicon carbide layer, where methyltrichlorosilane dominates the mixture mass and the residual is 0.07 %. The susceptor temperature is then set to the value that holds the bed, obtained by solving the steady form of Equations (8.13) and (8.14) in four passes, which

Table 8.3 Model Constants. Every Constant Except Room Temperature Is Adjustable in the Model Dialog. The Lower Block Sets the Segregation, Carry-Over, Agglomeration and Pressure-Drop Laws.

Symbol	Meaning	Value	Unit	Vis.
D_c, D_i	coater, nozzle diameter	0.075, 0.008	m	adj.
ε	loose-packed voidage	0.42	—	adj.
R_{good}, R_{carry}	circulation, carry-over	2.00, 3.00	—	adj.
m_0	charge as weighed	0.200	kg	adj.
d_k	mean kernel diameter	500	µm	adj.
σ_r	kernel size spread	0.024	—	adj.
ρ_k	UCO kernel density	10900	kg/m^3	adj.
y_{soot}, w_{soot}, k_s	soot onset, width, cost	.30, .18, .55	—	adj.
ΔH_{pyc}, ΔH_{sic}	heat of deposition	9450, 900	kJ/kg	adj.
C_w, C_b	susceptor, bed capacity	8.0, 0.40	kJ/K	adj.
UA_0	conductance, 1500 °C	1.50	kJ/K·min	adj.
k_{rad}	radiative loss	1.7004×10^{-11}	kJ/min·K^4	adj.
P_{max}	furnace rating	12.0	kW	adj.
T_g	inlet gas temperature	120	°C	adj.
τ_{TC}, τ_{MFC}, τ_{DP}	instrument lags	.05, .10, .15	min	adj.
τ_{cell}, τ_{gas}, τ_{obs}	cell, gas, trim	1.0, 1.2, 3.0	min	adj.
m_{gate}	correction noise gate	0.006	kg	adj.
t_{soak}, t_{cool}, t_{max}	soak, cool, abort	3.0, 20, 150	min	adj.
T_{band}	at-temperature band	8.0	°C	adj.
σ^*_{SiC}	SiC variability limit	2.0	µm	adj.
σ_{min}, σ_{max}, σ_s	segregation strengths	.30, 2.60, .90	—	adj.
k_{ag}, k_{el}	agglomeration, carry-over	0.80, 0.020	1/min	adj.
f_s, Δp_{pk}	bed weight, onset peak	0.72, 0.30	—	adj.
T_a	room temperature	30	°C	hidden

returns $T_w = 1275\,^{\circ}\mathrm{C}$ and $P = 2.04\,\mathrm{kW}$ at the buffer condition. Both initializations are checked by a startup test that evaluates every derivative at the opening state of scenario 2 and requires the largest relative value to be below 2×10^{-3}.

The integrator is classical fourth-order Runge-Kutta at a fixed 0.04 min (2.4 s) step, which is well inside the 6 s flow controller lag, the fastest dynamics it integrates. The control system declares and executes a scan every three physics steps, 0.12 min (7.2 s), which is close to what a batch coater's distributed control system runs. The instrument filters, including the 3 s thermocouple, are advanced at that scan, and the historian keeps five samples per simulated minute. Initialization does not advance the thickness estimator or its load-cell trim before plant time passes. Noise comes from the framework's seeded generator (Section 1.6.1), here with seed 20 260 906, so a noisy scenario replays identically. After every step the thicknesses are floored at zero, the flow states are clipped to their actuator spans, f_ℓ to 0.80 and a to 1.50.

8.3.7 Assumptions

The model is detailed where the control lesson lives and thin everywhere else, and each simplification fails somewhere.

Deposition is a lumped two-limit fit rather than a mechanism. There is no gas-phase chemistry, no intermediate species and no boundary-layer transport in Equation (8.9). The activation energies of 190 to 230 kJ/mol and the stoichiometric yields are physical, but the pre-exponential factors are fitted to reproduce layer durations, so the model cannot support a claim about a real machine's deposition rate.

Coating efficiency is one constant per layer times a disturbance. On a real coater it depends on temperature, gas velocity, wall condition and cone geometry and drifts within a run. Making it a constant is what allows the lesson to be posed cleanly, and it is also a simplification of exactly the quantity the lesson is about.

Layer anisotropy, density and crystallinity are absent. The optical anisotropy of the inner pyrocarbon matters as much as its thickness, because an anisotropic layer shrinks unevenly under irradiation and cracks [9]. Nothing here predicts whether a batch would survive a reactor.

Seven size classes on a Gaussian grid, not a population balance. There is no attrition, no breakage, no bimodal charge and no representation of the far tail of the distribution that a real qualification is written on [6]. The seven classes reproduce a mean, a spread and the direction in which poor fluidization moves them.

One correlation and a two-parameter circulation map for the hydrodynamics. Equations (8.6) to (8.10) carry no computational fluid dynamics, no explicit spout-annulus exchange and no cone-angle dependence, and the segregation coefficients are fitted rather than derived from a residence-time distribution [17]. The slump and carry-over flows are the right order of magnitude and the right trend, not design numbers.

The furnace is two capacities and a radiative loss. There is no gradient down the cone and no separate wall node in the deposition zone, so wall deposition, one of the real reasons coating efficiency is not what it should be, appears only inside the lumped efficiency of Equation (8.9).

The thermocouple reads the bed directly. A bias is offered as a disturbance, but on a real coater the well reads something between the gas and the graphite and a 40 °C discrepancy is ordinary, which through Equation (8.9) moves the kinetic rate by 30 to 47 %, depending on the layer and the sign.

Residual precursor swept out during a purge deposits nothing, although its molar mass is still carried in Equation (8.5) because dropping it at the changeover would halve the mixture density in one step. In reality a small amount of the last layer continues to grow through the purge.

The numerics are more repeatable than a control system. An identical trajectory on every machine is convenient and false.

8.4 The Thickness Estimator

Everything in this section is what the control system believes rather than what the plant is doing. The estimator is cruder than the plant in five separate ways. It carries one mean particle instead of seven classes, the charge as weighed rather than as charged, the nominal coating efficiency η_L rather than $\eta_L \eta_d$, no soot term at all, and the flows as indicated by the mass flow controllers rather than as delivered. Every one is a place the estimate can walk away from the truth.

8.4.1 One Mean Particle

From the four estimated thicknesses $\hat{\delta}_L$ the estimator builds a single particle and multiplies it by a particle count taken from the batch ticket, as shown in Equation (8.18).

$$
\begin{aligned}
\hat{n} &= \frac{m_0}{\bar{m}_k}, \qquad \hat{r}_0 = \frac{d_k}{2}, \qquad \hat{r}_L = \hat{r}_{L-1} + \hat{\delta}_L \\
\hat{m}_p &= \bar{m}_k + \tfrac{4}{3}\pi \sum_{L=1}^{4} \rho_L \left(\hat{r}_L^3 - \hat{r}_{L-1}^3 \right) \\
\hat{d} &= 2\hat{r}_4, \qquad \hat{A} = \hat{n}\pi\hat{d}^2, \qquad \hat{m}_b = \hat{n}\,\hat{m}_p
\end{aligned}
\tag{8.18}
$$

The estimated particle density and bed depth follow from $\hat{m}_p$ and $\hat{d}$ as in Equation (8.4). Comparing (8.18) with (8.3) locates the errors precisely. The charge error ε_m and the carry-over fraction f_ℓ are both absent from $\hat{n}$, and the class spread is absent from $\hat{A}$.

8.4.2 The Integrated Estimate

The estimator runs its own copy of Equation (8.9) on measured and indicated quantities, as shown in Equation (8.19).

$$\begin{aligned} \hat{y}_R &= \frac{\tilde{q}_R}{\tilde{q}_C + \tilde{q}_R}, \qquad \hat{J}_{sup} = \frac{\eta_L Y_L \tilde{q}_R}{\hat{A}} \\ \hat{J}_{kin} &= k_L \exp\left[-\frac{E_{a,L}}{\mathscr{R}}\left(\frac{1}{T_m} - \frac{1}{T_{ref,L}}\right)\right]\left(\frac{\hat{y}_R}{y_{ref,L}}\right)^{n_L} \\ \frac{\mathrm{d}\hat{\delta}_L^{r}}{\mathrm{d}t} &= \frac{1}{\rho_L}\frac{\hat{J}_{kin}\hat{J}_{sup}}{\hat{J}_{kin} + \hat{J}_{sup}}, \qquad \hat{\delta}_L = \hat{\delta}_{L,0} + \kappa\left(\hat{\delta}_L^{r} - \hat{\delta}_{L,0}^{r}\right) \end{aligned} \tag{8.19}$$

where $\hat{\delta}_L^{r}$ is the uncorrected integral, κ is the load-cell correction factor of Section 8.4.3, and the zero subscripts are the values at which the current layer began. Both integrals are kept, because the correction has to be formed against the uncorrected one. Equation (8.19) contains the measured bed temperature T_m, so a run that is 50 °C cold extends itself, and it contains no η_d, so a run whose coating efficiency is 0.72 does not. The quantity the application exists to display is the difference between the two integrals, summed over the batch so that the divergence reads as one curve rather than four, as shown in Equation (8.20).

$$G = \sum_{L=1}^{4} \hat{\delta}_L - \sum_{L=1}^{4} \bar{\delta}_L \tag{8.20}$$

G is the inferential gap. No instrument on a real coater produces it, and it is on the screen only because the application is the plant as well as the control system.

8.4.3 The Bed-Mass Correction

The bed goes from 200 g to 383 g over a nominal run, and that gain is a direct measurement of how much material landed on the particles. It is the only measurement in the coater that bears on the controlled variable. The correction inverts the estimator's own mass model to find the thickness that accounts for the measured gain, as shown in Equation (8.21).

$$\hat{m}_b\left(\hat{\delta}_1, \ldots, \delta_L^{m}, \ldots, \hat{\delta}_4\right) = \hat{m}_{b,0} + (m_m - m_{m,0}) \tag{8.21}$$

where δ_L^{m} is the unknown, $\hat{m}_{b,0}$ and $m_{m,0}$ are the modeled and measured bed masses at the start of the layer, and the left side is Equation (8.18). It is monotone in δ_L^{m}, so twenty passes of bisection on $[0, 4\delta_L^{*}]$ are exact to better than a nanometer. The correction factor is the ratio of Equation (8.22).

$$\kappa \leftarrow \mathrm{clamp}\left(\frac{\delta_L^{m} - \hat{\delta}_{L,0}}{\tilde{\delta}_L^{r} - \hat{\delta}_{L,0}^{r}},\ 0.25,\ 3\right), \qquad \tau_{cell}\frac{\mathrm{d}\tilde{\delta}_L^{r}}{\mathrm{d}t} = \hat{\delta}_L^{r} - \tilde{\delta}_L^{r} \tag{8.22}$$

Three details in Equation (8.22) are deliberate and all three are general. The measured mass is inverted through the geometry rather than used as a mass ratio, because mass is not linear in thickness and a shell on a 250 µm kernel radius grows by a factor of 1.38 by the time the buffer is finished. The denominator is the *uncorrected* integral, because against the corrected one the loop would be chasing its own tail and would converge on the square root of the error instead of the error. And the model is put through the same 1 min filter as the load cell, which is what $\tilde{\delta}_L^{r}$ is, because comparing a filtered measurement against an unfiltered model reads short by whatever arrived during the filter's time constant.

Two gates protect the correction. It stays at unity until the modeled mass gain exceeds $m_{gate} = 6\,\mathrm{g}$, since a two-gram measurement over a two-gram model is not information, and on the first gated scan it snaps to the ratio and is filtered with τ_{obs} after that so it does not spend the first minutes of

every layer walking away from one. At every changeover the finished layer is re-based on the cell by solving Equation (8.21) once more with $\hat{\delta}_{L,0}$ substituted for the accumulated estimate. That reading is taken three load-cell time constants into the following purge rather than at the instant the precursor is shut off, with nothing depositing and the filter caught up, which is also when an operator would read it. Because the correction compares a measured gain against a modeled gain, it can only see errors that change the gain, and Section 8.6 shows what that excludes.

8.5 Control System

Three actuators carry three very different kinds of loop. Furnace power holds bed temperature in an ordinary PID on a thermocouple. Carrier flow holds a spouting ratio that nobody measures. Precursor flow has no loop at all, and is the recipe's number for whichever step the sequencer is in.

8.5.1 Bed Temperature

TIC-301 is the incremental, or velocity, form of PID of Section 1.6.2, with the derivative taken on the filtered measurement rather than on the error, executed every $\Delta t_c = 0.12\,\mathrm{min}$. Each scan forms increments, as shown in Equation (8.23), and accumulates them into the furnace command, as shown in Equation (8.24). In the symbols of Chapter 1, s_k is the running sum S_k of Equation (1.15), c_k the output before limiting of Equation (1.16) and I_k the integral of Equation (1.17).

$$\begin{aligned} e_k &= T_{sp,k} - T_{m,k}, & \Delta u_k^P &= K_c\,(e_k - e_{k-1}) \\ \Delta u_k^I &= \frac{K_c}{\tau_I}\,e_k\,\Delta t_c, & \Delta u_k^D &= -K_c\,\tau_D\left(\dot{T}_{m,k}^f - \dot{T}_{m,k-1}^f\right) \end{aligned} \tag{8.23}$$

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D, & c_k &= s_k + I_{k-1} \\ u_k &= \mathrm{clamp}(c_k,\,0,\,P_{max}) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{8.24}$$

where T_m^f is the measurement passed through a first-order filter of time constant 0.5 min and $\dot{T}_{m,k}^f$ is its backward difference over one scan, s_k accumulates the proportional and derivative increments without limit, c_k is the power the algorithm demands, u_k is the power applied, and the integral I used at scan k was accumulated through scan $k-1$. I is seeded at the steady power of the opening state, so every scenario starts without a bump. The loop is reverse acting, because the process gain is positive and a bed above setpoint calls for less power. The derivative filter is advanced by the same zero-order-hold update as the instruments, so it is unconditionally stable at the 7.2 s scan. Anti-windup is the back-calculation term of Equation (8.24), with tracking time $T_t = \tau_I$, so while the command is clipped each scan removes $\Delta t_c/\tau_I = 3\,\%$ of the excess from I alone. Manual mode tracks the operator's power as Section 1.6.2 describes, so the transfer is bumpless in both directions. The default settings are $K_c = 0.130\,\mathrm{kW/^\circ C}$, $\tau_I = 4.0\,\mathrm{min}$ and $\tau_D = 0.60\,\mathrm{min}$, and the output carries no rate limit because a resistance furnace has none.

The setpoint is ramped rather than stepped, at a rate that is itself a recipe parameter, as shown in Equation (8.25).

$$T_{sp,k+1} = T_{sp,k} + \mathrm{clamp}\left(T_{tar} - T_{sp,k},\; -\dot{T}\,\Delta t_c,\; \dot{T}\,\Delta t_c\right) \tag{8.25}$$

with $\dot{T} = 15\,^\circ\mathrm{C/min}$ by default. The target T_{tar} belongs to the phase and is re-read every scan rather than latched at the changeover, so moving a recipe temperature while a step is running does what the operator expects.

8.5.2 Fluidization as a Schedule With a Trim

The spouting ratio has no transmitter, so FIC-301 holds the ratio computed from the indicated flows and the estimated bed. Because the relationship between flow and ratio is algebraic and known, the

honest structure is a schedule carrying the flow and feedback closing only the residual, as shown in Equation (8.26).

$$\begin{aligned} \hat{R} &= \frac{\hat{U}}{\hat{U}_{ms}\left(\hat{d}, \hat{H}, \hat{\rho}_p, \hat{\rho}_g\right)} \\ q_{C,ff}:\ \hat{R}\left(q_{C,ff}\right) &= R_{sp} \quad \text{by six passes of (8.6)} \\ u_C &= \text{clamp}\left(q_{C,ff} + \Delta q_C,\ 0,\ 300\,\text{slpm}\right) \end{aligned} \tag{8.26}$$

where $\hat{U}$ and $\hat{U}_{ms}$ are evaluated from $\tilde{q}_C$, $\tilde{q}_R$, the measured temperature T_m and the estimated bed of Equation (8.18), and Δq_C is the trim. The trim is a PI controller of the same form as (8.24) with $\tau_D = 0$, output limits of ± 120 slpm, a slew limit of 120 slpm/min applied after the clamp, back-calculation against the slewed output, and default settings $K_c = 3.0$ slpm per unit of ratio and $\tau_I = 1.5$ min. The final clamp of Equation (8.26) lies outside the trim, so its integral would not see it, although no shipped scenario reaches either limit of that clamp. Every quantity in (8.26) comes from upstream of the cone, so the loop is exact when the delivered flow and the estimated bed are right and arithmetically correct but physically false when they are not.

8.5.3 The Recipe Sequencer

A coating recipe is not a setpoint, it is a state machine. Ten phases run in order, heat-up, buffer, purge, inner pyrocarbon, purge, silicon carbide, purge, outer pyrocarbon, cool-down and complete, and every transition changes the temperature target, the carrier composition and the precursor at once. The carrier changeover and the precursor ramp are both first-order with time constant $\tau_{gas} = 1.2$ min rather than steps, and that is not decoration. Replacing argon with hydrogen moves U_{ms} by a factor of more than two through the ρ_g term of (8.6), and methyltrichlorosilane at 6 to 8 % of the stream carries 57 to 64 % of the mixture's mass, so cutting it in one step would more than halve the gas density within one scan and drop the bed out from under a flow controller that had done nothing wrong.

A heat-up or a purge ends when the bed is within T_{band} of a target the ramp has already reached and the phase has run longer than its soak or purge time, and a cool-down ends on its allowance. What ends a *deposition* phase, the guard of Section 1.4.2, is the design decision, and four architectures are selectable.

- MANUAL. Nothing advances on its own. The operator steps the recipe from the panel, and in Manual mode also holds all three actuators.
- TIMER. The phase ends at $t_{phase} \geq$ the hold time, which is how every coater built before about 1990 decided when to stop.
- INFERENTIAL. The phase ends at $\hat{\delta}_L \geq \delta_L^*$, with $\kappa \equiv 1$.
- INFERENTIAL + MASS. The same condition with κ from Equation (8.22) and the end-of-layer re-basing active.

No phase may run forever. A phase that has not met its condition after $t_{max} = 150$ min is abandoned with a latched fault, because a step that cannot finish is a fault rather than a reason to wait. The mechanism does select. Derating the elements to $\eta_{el} = 0.35$ leaves the bed unable to reach 1500 °C, the purge ahead of the silicon carbide layer is abandoned at 150 min, the layer itself is abandoned 150 min later with 11.4 µm on the particles, and the batch finishes at 422 min instead of 191. The same choice among a clock, a measured event and an estimate ends primary drying in Chapter 5 and each column loading in Chapter 11.

8.5.4 Acceptance

The batch is graded once, after it comes out, on the six criteria of Equation (8.27).

$$\begin{aligned} &\left|\bar{\delta}_1 - 95\right| \le 8, \quad \left|\bar{\delta}_2 - 40\right| \le 3, \quad \left|\bar{\delta}_3 - 35\right| \le 2, \\ &\left|\bar{\delta}_4 - 40\right| \le 3, \quad \sigma_3 \le 2.0\,\mu\text{m}, \quad f_\ell \le 0.02, \quad a < 1 \end{aligned} \tag{8.27}$$

The fifth condition is the one that decides most real batches and the only one with no corresponding number anywhere on the operating screen.

8.6 Results

Fourteen scenarios are arranged as a laboratory sequence, with the briefings, completion rules and challenge of Section 1.6.4. Scenario 1 coats the buffer by hand and scenario 2 step-tests the temperature loop. Scenarios 3 and 4 find the slump and carry-over walls on a fresh and on a coated bed. Scenarios 5 and 6 run the whole recipe on the timer and on the inferential endpoint. Scenarios 7, 8 and 10 to 13 each load one fault: a cold silicon carbide step, a low coating efficiency, a mis-weighed charge, excess acetylene, a slack bed and a carrier leak. Scenario 9 repeats scenario 8 with the load-cell correction on. Scenario 14 is the graded challenge. Each fixes the starting phase, the endpoint architecture, the tuning, the disturbances and any timed events. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt, as does an extra learner disturbance event when it fires. Every number below was produced by driving the shipped model headlessly at the physics step, exactly as the browser steps it, until the batch completed.

8.6.1 The Nominal Batch

The design point is 200 g of 500 µm kernels, 279 865 particles presenting 0.2199 m^2 in a 7.16 mm bed at 10 900 kg/m^3, with $U_{ms} = 0.2282$ m/s reached by 18.4 slpm of argon alongside 5.5 slpm of acetylene at a spouting ratio of 2.20, and 2.04 kW of furnace holding 1250 °C with the susceptor at 1275 °C. The opening buffer rate is 5.10 µm/min. The sequencer walks all ten phases, starting the buffer at 6.6 min, the inner pyrocarbon at 36.7 min, the silicon carbide at 74.3 min and the outer pyrocarbon at 147.3 min, and finishes at 191.4 min on 12.0 kWh. The particles leave 921 µm across, the bed weighs 383 g and presents 0.745 m^2, so the denominator of Equation (8.9) rose by a factor of 3.4 while the recipe ran.

The timer and the inferential endpoint produce the same batch on a nominal plant, which is worth stating as a result. The first two rows of Table 8.4 give 95.2, 39.8, 35.1 and 40.0 µm against 95, 40, 35 and 40 on the timer and 95.2, 40.1, 35.0 and 40.1 µm on the integrated estimate, both accepted, in 191.1 and 191.4 min. Over 210 µm of coating the gap of Equation (8.20) ends at 0.09 µm and never exceeds 0.35 µm, and that residual is not drift but the difference between the seven-class plant and the one-particle estimator. An endpoint scheme that changed the answer on a nominal plant would be a scheme with a defect in it.

8.6.2 The Temperature Loop at Both Ends of the Recipe

The bed temperature loop is the only conventional one on the machine, and its process gain is not a constant. Solving the steady form of Equations (8.13) and (8.14) over a 50 °C interval with the carrier flow rescheduled by Equation (8.26) and deposition active gives 285 °C/kW at the 1250 °C buffer condition and 132 °C/kW at the 1500 °C silicon carbide condition, a factor of 2.2, because the radiative loss Q_{rad} grows far faster than the transferred heat does. The dominant time constant moves the other way for the same reason, from $C_w/UA = 8.0$ min at the buffer condition to 4.7 min at the silicon carbide condition. One set of settings therefore has to work across a gain that more than doubles, and the default $K_c = 0.130$ kW/°C with $\tau_I = 4.0$ min closes a 50 °C setpoint move to 63 % in 3.1 min at either end. Furnace power peaks at 8.31 kW of the 12 kW rating during the ramp to 1500 °C, and derating the elements to $\eta_{el} = 0.66$ puts the command against that ceiling for 4.6 min while still completing the batch in 192 min.

Table 8.4 Measured Batch Outcomes. Thicknesses and the Inferential Gap G Are in µm. Scenarios 7 and 12 Start at the Purge Ahead of the Silicon Carbide Layer and Scenarios 11 and 13 at the Buffer, So Their Batch Times Are Not Comparable With a Full Run.

#	Case	Buffer	IPyC	SiC	OPyC	σ_{SiC}	G	t (min)	Verdict
5	Timer, nominal plant	95.2	39.8	35.1	40.0	0.53	−0.08	191.1	accepted
6	Inferential, nominal plant	95.2	40.1	35.0	40.1	0.52	0.09	191.4	accepted
7	Timer, SiC 50 °C cold	95.0	40.0	26.2	41.6	0.46	−0.04	117.7	rejected, SiC mean
7	Inferential, SiC 50 °C cold	95.0	40.0	35.0	40.1	0.61	0.01	139.4	accepted
8	Inferential, $\eta_d = 0.72$	79.7	34.9	32.0	35.4	0.52	28.35	191.4	rejected, all four means
9	Inferential + mass, $\eta_d = 0.72$	97.2	41.2	34.8	39.9	0.52	−0.79	218.9	accepted
10	Inferential + mass, $\varepsilon_m = +20\%$	83.8	36.7	31.6	36.7	0.50	22.00	185.8	rejected, all four means
11	Inferential, acetylene 12 slpm	79.3	41.9	36.0	41.3	0.49	11.82	170.3	rejected, buffer mean
12	Inferential, $R_{sp} = 1.30$	95.0	40.0	35.1	40.2	2.89	−0.04	100.8	rejected, SiC variability
13	Inferential, $\chi_C = 0.62$	105.2	44.6	40.9	44.0	0.96	−24.35	183.8	rejected, all four means
14	Challenge as loaded	77.3	46.2	35.2	52.9	0.54	−34.52	325.6	rejected, three means and loss
14	Challenge, suggested settings	96.5	40.0	34.3	40.1	0.51	−1.47	208.7	accepted

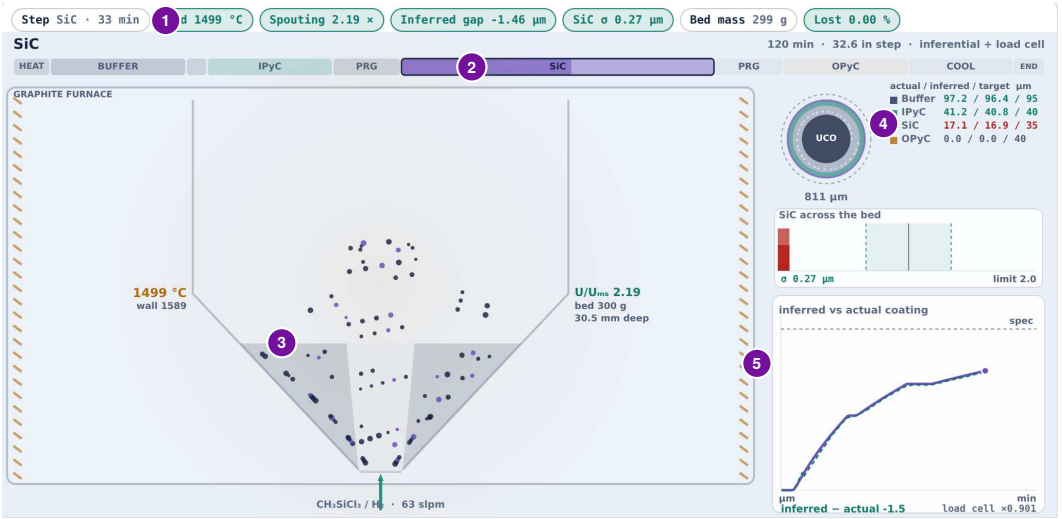


Figure 8.2 Scenario 9 at $t = 120$ min, thirty-three minutes into the silicon carbide layer, with the coating efficiency at 0.72 and the load-cell correction working (plant view cropped from the app's screen). The badges (1) read an inferential gap of −1.46 µm, which ends at −0.79 µm (Table 8.4), a spread of 0.27 µm, a load cell at 299 g against the 200 g charge and nothing lost. The recipe strip (2) shows the sequencer in the silicon carbide step, and the furnace view (3) shows the spout and the annulus. The layer table (4) reads the quantity the machine cannot measure beside the quantity it infers: 97.2 against 96.4 µm on the buffer, 41.2 against 40.8 on the inner pyrocarbon, and 17.1 against 16.9 µm of the 35 µm silicon carbide target so far, on a particle 811 µm across. The coating plot (5) carries the correction factor derived from the bed mass, 0.901, which is the mechanism this chapter argues for caught in the act. The trends of the full screen, cut here, read the bed at 1499.2 °C (thermocouple 1499.3) against a 1500 °C setpoint on 6.1 kW, and the spouting ratio at 2.201 inferred against 2.188 actual on 58.1 slpm of hydrogen carrier. The spread is far inside the 2.0 µm acceptance limit.

The gain that matters is the other one. In the exponential of Equation (8.9) at the silicon carbide condition, +50 °C multiplies the kinetic rate by 1.534 and −50 °C leaves 0.636 of it. Bed temperature is not a comfort variable here. It is a term inside an exponential in the quantity that has no instrument,

Table 8.5 The Fluidization Window at Four Points in the Recipe. Carrier Flows Are Those That Put the Bed at $R = 1$, at the Setpoint of 2.20 and at the Carry-Over Ratio of 3.00, With the Recipe Precursor Flowing.

	Buffer start	SiC start	SiC end	OPyC end
Carrier	Ar	85 % H_2	85 % H_2	Ar
T_b (°C)	1250	1500	1500	1300
$\bar{d}$ (µm)	500	770	840	920
H (mm)	7.16	26.1	33.9	44.6
ρ_p (kg/m^3)	10900	3988	3807	3352
m_b (g)	200.0	267.0	330.8	382.6
A_b (m^2)	0.220	0.521	0.621	0.744
U_{ms} (m/s)	0.228	0.601	0.789	0.565
q_C at slump (slpm)	5.9	11.0	17.1	20.9
q_C at $R = 2.20$ (slpm)	18.4	49.5	66.4	52.2
q_C at carry-over (slpm)	26.7	78.4	102.3	73.0

and scenario 7 prices it: the timer runs the layer for its 59.5 min either way and delivers 26.2 µm against a target of 35 with a tolerance of 2, while the inferential endpoint delivers 35.0 µm by running the layer for 81.1 min. That converts a quality failure into a schedule cost of 21.6 min.

8.6.3 Fluidization Through the Batch

Table 8.5 evaluates Equation (8.6) at four points in the recipe and reports the carrier flow at each wall of the operating window.

The minimum spouting velocity rises from 0.228 to 0.565 m/s between the first minute and the last, a factor of 2.47, and the carrier flow that holds the setpoint rises from 18.4 to 52.2 slpm. Three of the four terms in Equation (8.6) move. The particles grow 84 % in diameter and U_{ms} is linear in that, the bed becomes six times deeper and U_{ms} goes as its square root, and working the other way the particle density falls from 10900 to 3352 kg/m^3 as low-density carbon replaces uranium oxycarbide, which takes back a factor of 1.8. The requirement is not monotone, because the 85 % hydrogen carrier of the silicon carbide layer is five times lighter than argon at 7.71 g/mol against 39.95, and takes U_{ms} to 0.789 m/s at the end of that layer before the return to argon brings it back to 0.565. The carrier command peaks at 88.6 slpm in the purge ahead of that layer, where hydrogen has replaced argon and no precursor is in the stream to weigh the mixture down.

The consequence is that the flow which spouts the fresh charge slumps the finished bed. The window runs from 5.9 to 26.7 slpm on the fresh charge and from 20.9 to 73.0 slpm on the finished bed, overlapping only between 20.9 and 26.7 slpm. No fixed number in a recipe can hold the bed for three hours, which is why the fluidization loop has to be the schedule of Equation (8.26) and why it inherits every error the thickness estimator carries.

8.6.4 The Central Result

Scenario 8 sets $\eta_d = 0.72$ and changes nothing else. The recipe assumes that 35 % of the acetylene and 35 % of the methyltrichlorosilane land on a particle, and today 0.72 of that does. Every loop stays on setpoint, the batch takes the same 191.4 min as the nominal one, each layer ends on schedule, the estimate finishes at 95.2, 40.1, 35.0 and 40.1 µm exactly on target, and the particles carry 79.7, 34.9, 32.0 and 35.4 µm. The gap is 28.35 µm, 13 % of the coating, and all four layers are rejected. Nothing on any trend moved, because Equation (8.19) has no term in which η_d appears.

Table 8.6 Inferential Gap G and Silicon Carbide Mean Against Coating Efficiency, With and Without the Load-Cell Correction. Both Columns Use the Same Plant and the Same Recipe.

η_d	1.00	0.90	0.80	0.72	0.60	0.50
Inferential endpoint						
G (µm)	0.09	9.31	19.45	28.35	43.35	57.77
SiC mean (µm)	35.02	34.07	33.00	32.02	30.31	28.58
Batch (min)	191	191	191	191	191	191
Verdict	pass	pass	fail	fail	fail	fail
Inferential endpoint with bed-mass correction						
G (µm)	−0.86	−2.13	−1.08	−0.79	−1.44	−1.75
SiC mean (µm)	34.45	34.97	34.58	34.84	35.02	34.35
Batch (min)	191	199	207	219	236	255
Verdict	pass	pass	pass	pass	pass	pass

TRY IT IN THE APP **8 · The Estimate Lies, and Nothing on the Screen Looks Wrong**

Load the scenario and press Run. Watch the total coating trend: the estimate climbs to each target and every layer ends on schedule, in the same 191.4 min as the nominal batch. Then read the scorecard. The particles carry 79.7, 34.9, 32.0 and 35.4 µm, and all four layers are rejected.

Scenario 9 runs the identical plant with the load-cell correction of Section 8.4.3 switched on and obtains 97.2, 41.2, 34.8 and 39.9 µm, accepted, with a gap of −0.79 µm against 28.35, a reduction of 97 %, and a correction factor settling at $\kappa = 0.890$. The cost is 28 min, 218.9 min against 191.4, and 1.7 kWh of extra furnace energy, 13.8 against 12.1 kWh, because each layer now runs until the corrected estimate reaches target. Table 8.6 sweeps the disturbance to show the structure of the result rather than one instance of it.

The uncorrected gap is nearly linear in the efficiency shortfall and the batch time never moves, because the estimate reaches target on schedule whatever the plant does, as Figure 8.3 draws. The corrected gap stays inside 2.2 µm across the whole range while the batch time absorbs the disturbance, from 191 to 255 min. The corrected gap moves by about 1 µm with the load-cell and thermocouple noise seed, from −0.98 to +0.62 µm over six seeds at $\eta_d = 0.72$, so the differences between its entries are within that spread. That is the shape of a correctly built inferential scheme. An integrated model supplies resolution and speed, one real measurement supplies the level, and the schedule rather than the product carries the disturbance.

DID YOU KNOW? **The Clock Cannot Tell**

From the simulator: on the inferential endpoint every batch takes 191 min, whether the coating efficiency is 1.00 or 0.50, and the silicon carbide mean falls from 35.02 to 28.58 µm. With the bed-mass correction the mean stays near target and the batch stretches from 191 to 255 min (Table 8.6).

8.6.5 What the Correction Cannot See

Scenario 10 charges 20 % more kernels than the batch ticket records, with the correction switched on, and the batch is rejected on all four means at 83.8, 36.7, 31.6 and 36.7 µm with the estimate 22.00 µm high. The reason is structural. A mis-charge scales the particle count, and through Equations (8.3) and (8.4) it scales the surface area and the mass gain by the same factor, so the measured gain matches the modeled gain, κ stays near unity, and every thickness is reported 20 % high because it is divided by a count that came off a piece of paper. The measurement that would have caught it is the same load cell read differently. Its *absolute* value is 240 g against a ticket of 200, and in a run that starts from the bare charge the application writes that comparison as a latched warning at 0.04 min, the

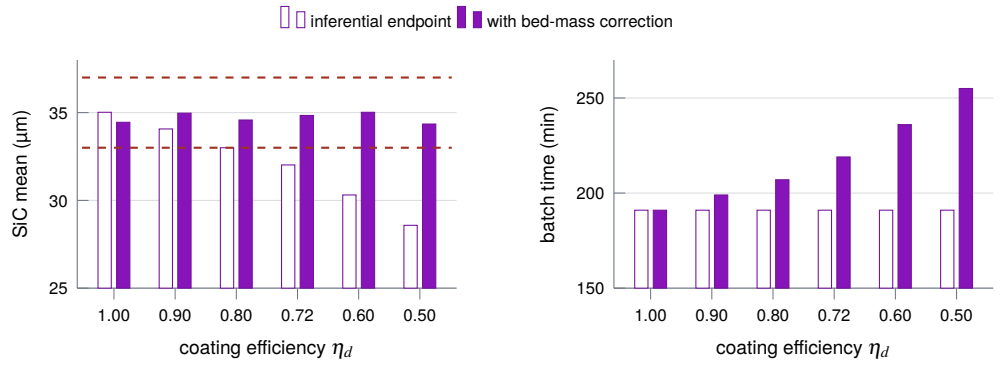


Figure 8.3 The efficiency sweep of Table 8.6. Without the correction (white bars) the batch always takes 191 min, the silicon carbide mean falls with the efficiency to the edge of the acceptance band (dashed) and below it, and every batch from 0.80 down fails. With it (solid bars) the mean stays near target and the batch time absorbs the disturbance.

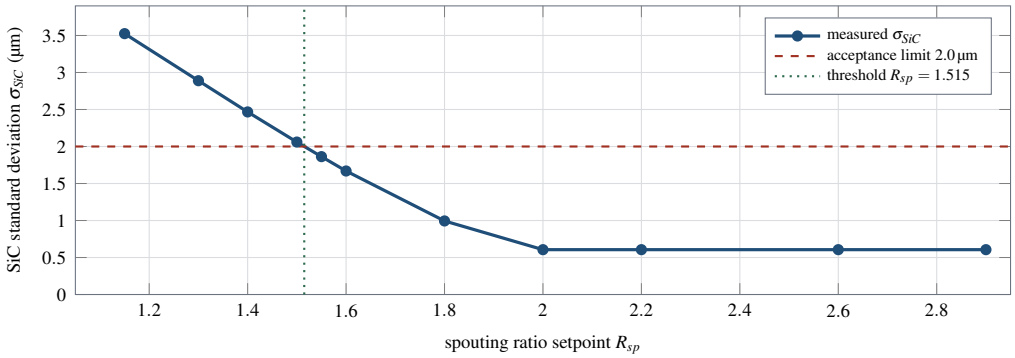


Figure 8.4 Silicon carbide layer variability against the fluidization setpoint, measured on the shipped model at setpoints on the slider grid. Every point is a silicon carbide layer whose mean thickness is within 0.15 µm of the 35 µm target, so the acceptance decision turns entirely on the width of the distribution. The spread saturates at 0.606 µm above $R_{good} = 2.00$, where circulation is fully developed. The threshold line comes from model runs at 1.51 and 1.52, between grid points.

first physics step. It saves the opening load-cell reading before deposition and compares it with the expected opening bed mass from the ticket and any completed layers. The same check therefore applies to cases that open partway through the recipe without mistaking deposited mass for a charge error. A correction built on a gain can only see errors that change the gain, so errors in the baseline have to be caught by looking at the baseline. Every soft sensor needs both, and the gross-error check is the cheaper of the two.

🕮 **Pause.** Try Exercise 8.3 before reading on.

8.6.6 Variability, and Why the Mean Is Not the Specification

Scenario 12 runs the silicon carbide layer at a spouting ratio of 1.30 instead of 2.20. The bed is still spouting, no fluidization alarm is raised, and the layer reaches 35.1 µm, dead on target. The batch is rejected, on a standard deviation of 2.89 µm against a limit of 2.0, with every mean on the scorecard inside tolerance. Figure 8.4 sweeps the setpoint.

The measured spreads are 3.53, 2.89, 2.47, 2.06, 1.67, 0.99 and 0.61 µm at setpoints of 1.15, 1.30, 1.40, 1.50, 1.60, 1.80 and 2.20, and the silicon carbide mean stays between 35.00 and 35.14 µm

across all seven. On the slider's 0.05 grid the batch is rejected at 1.50 and accepted at 1.55 with $\sigma_{SiC} = 1.86\,\mu m$, and model runs between those grid points put the crossing at 1.515, where the segregation strength of Equation (8.10) passes $\Sigma = 1.08$. Above $R_{good} = 2.00$ the curve is flat at 0.606 µm, because χ has saturated and the only remaining spread is the kernel size distribution feeding through Equation (8.2). The fluidization setpoint is therefore a product quality variable, and the operator's instinct to run the gas gently is the wrong instinct.

> **TRY IT IN THE APP** **12 · The Bed Goes Slack During the SiC Layer**
>
> Load the scenario, press Run and watch the histogram beside the magnified particle. The layer ends at 35.1 µm, on target, and the batch is rejected on a spread of 2.89 µm. Press Reset, set the spouting ratio setpoint to 1.55 before the layer starts and run again: the spread is 1.86 µm and the batch is accepted.

8.6.7 Errors Upstream of Every Instrument

Scenario 13 delivers 62 % of the indicated carrier to the cone. The inferred ratio of Equation (8.26) averages 2.200 over the whole run, exactly on setpoint, while the true ratio is 1.432 twenty minutes in. Nothing in the control system has done anything wrong, because the flow meter measures what passes through the flow meter and the loop divides that by an estimated minimum spouting velocity to obtain a number that is arithmetically correct and physically false. What it costs is the spread and the mean together. The buffer standard deviation rises from 1.95 to 7.97 µm because the bed circulates at less than half the intended vigor for the entire layer, and every layer comes out over target at 105.2, 44.6, 40.9 and 44.0 µm because a smaller total flow is a richer gas. Taking the delivery to 45 % before the run slumps the bed at 27.0 min, accumulates 72.3 min below minimum spouting and agglomerates the batch at 162.0 min. The differential pressure cell is the one instrument downstream of the leak. By Equation (8.7) it cannot report the ratio, but it collapses when the bed does, and that is the general point. An inferential measurement is only as good as its most upstream input, and the defense is an independent measurement further down, however crude.

Scenario 11 raises the acetylene from 5.5 to 12 slpm to speed the buffer up, which takes the hydrocarbon mole fraction to 33 % and past the soot onset of Equation (8.9). The companion run is the same case with the acetylene set back to 5.5 slpm before Run or carried through Reset. Both paths initialize at the selected flow and produce no soot. The layer does finish faster, in 11.7 min instead of 25.0, and it finishes at 79.3 µm instead of 95.1. 15.0 g of soot leaves through the vent and the estimate reads 95.3 µm, 16.0 µm high, because Equation (8.19) has no soot term. Precursor utilization needs its observation window stated. Read at the end of the buffer step it falls from 21.3 % to 16.4, and read over the whole batch it falls from 19.8 % to 18.5, because the three later layers are untouched. Two consequences are worse than the thickness. The buffer standard deviation rises from 1.95 to 4.11 µm through the $\sigma_s s$ term of Equation (8.10), and the error propagates, because the estimator now believes the particles are larger than they are and carries an oversized $\hat{A}$ into every later layer.

The fluidization trim works against a process gain that falls through the batch. Differentiating Equation (8.6) numerically gives 0.096 units of ratio per slpm on the fresh charge, 0.028 at the start of the silicon carbide layer and 0.038 on the finished bed, so a fixed trim gain is three times more aggressive at the beginning than in the middle. Sweeping K_c over a nominal batch, the carrier command travels 177 slpm at the default value of 3.0 and 180 slpm at 20, with the extra reversals coming from thermocouple noise rather than from the loop. At $K_c = 22$ the travel jumps to 1170 slpm (927 to 1247 over five noise seeds) and the peak inferred-ratio error from 0.19 to 0.39, and at $K_c = 30$, the top of the slider, the travel is 5258 slpm. The onset of sustained oscillation is therefore near $K_c = 22$, and the batch stays acceptable throughout the sweep because the layers are long compared with the cycle. The app's hint on this gain once put the onset at 8 and now puts it at 22.

8.6.8 The Graded Challenge

The challenge loads four faults at once, an unknown coating efficiency of 0.78, a timer endpoint, a spouting ratio setpoint of 3.30 and a deliberately weak temperature tuning, and asks for an accepted batch in under 240 min. As loaded it gives 77.3, 46.2, 35.2 and 52.9 µm in 325.6 min, rejected on three means and on particle loss, because the fountain is over the freeboard for the whole run. The suggested settings, the inferential endpoint with the bed-mass correction at a spouting ratio of 2.20 and the tuned temperature loop, give 96.5, 40.0, 34.3 and 40.1 µm with $\sigma_{SiC} = 0.51$ µm, accepted in 208.7 min. The challenge carries 3 °C of thermocouple noise, and over eight noise seeds the batch as loaded is rejected on the same four criteria in 317.7 to 329.4 min while the suggested settings are accepted in 208.1 to 210.8 min. Removing the three fixes one at a time separates what each is for. The plain inferential endpoint at a ratio of 3.30 is rejected on three means and on loss. Adding the bed-mass correction while leaving the ratio at 3.30 is still rejected, on the buffer, silicon carbide and outer pyrocarbon means at 134.3, 46.3 and 47.7 µm and on 27.3 % particle loss (27.3 to 34.7 % over the seeds), because carry-over removes surface area from Equation (8.3) faster than the correction can follow it. Fixing the ratio as well gives an accepted batch, but in 329.2 min, far past the deadline, because the weak tuning makes every ramp crawl, and tuning the loop brings the same batch to 208.7 min, 31.3 min inside the deadline whether the settings are made before Run or carried through Reset. Two habits follow. Fix the measurement before the controller, because no tuning recovers information the loop does not have, and separate the faults that cost quality from the faults that cost time, because they are answered by different things.

8.7 Discussion

The four architectures are not four controllers. Manual, timer and inferential are all open loop with respect to thickness, and what the fourth closes is not a thickness loop but a model-correction loop driven by a load cell. That is the honest shape of a great deal of industrial advanced control. The improvement comes from a measurement nobody had thought to use, wired into a model that was already there, and the load cell costs a few thousand dollars and converts a rejected batch of nuclear fuel into an accepted one. When an inferential scheme is presented for approval the useful question is never how the model is tuned. It is which parameter in the model can move without any instrument noticing, and what measurement would notice.

The economics of the plant force the structure. A coating run is three hours and the metallography that says whether it worked is a week of sectioning, polishing and image analysis on a sample of particles [6]. The feedback loop that actually sets the process is therefore a week long, and everything inside the batch is open loop with respect to the property that matters. Every in-situ measurement worth having on such a plant is a way of shortening that loop, and the ranking of candidates follows from Table 8.6 rather than from instrument catalogs. The load cell is first because it bears directly on deposited mass. An off-gas analyzer would be second, because unreacted precursor and soot are the difference between the two branches of Equation (8.9). A pyrometer down the freeboard would be third, because it addresses a quantity the thermocouple already reports imperfectly rather than one nothing reports.

The same equations support methods beyond the implemented loops. The estimator of Section 8.4 is an open-loop integration with one scalar correction, and the natural replacement is a moving horizon estimator on Equations (8.9) to (8.14) that reconciles the load cell, the thermocouple and the bed pressure drop at once and estimates η_d as a parameter rather than correcting its consequence. The endpoint decision is a terminal-constraint optimal control problem, minimizing batch time subject to four terminal thicknesses, a variance limit, the carry-over limit of Equation (8.16) and the furnace rating, and its solution would schedule temperature and precursor flow together rather than holding both at recipe values.

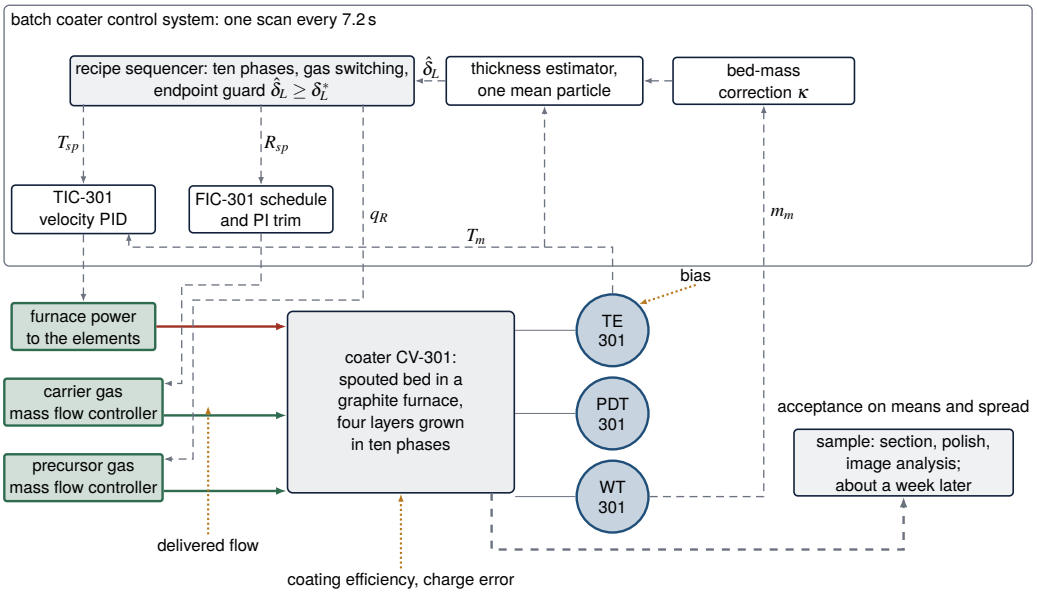


Figure 8.5 The chapter's loops on plant hardware. TIC-301 holds the bed temperature on the thermocouple through furnace power, FIC-301 schedules the carrier flow with a trim, and the sequencer sets the targets, switches the gases and ends each layer when the estimate reaches it. The load cell feeds the correction of the estimate. The thickness itself is read only on a sample, after the batch (dashed). Tags are the chapter's.

8.8 Technology in Practice

What Measures It

The coater of this chapter carries a thermocouple in the bed, a differential pressure cell across it and a load cell under the whole vessel, beside the flows its mass flow controllers indicate. None of them reads a thickness. In a qualification program the thicknesses are measured on a random sample of the coated particles, by sectioning, polishing and image analysis, and specifications are written on both the mean and the tail of each property [18]. In the model the thermocouple has a 3 s lag and an optional bias, and the load cell is filtered at 1 min. Those two signals are all the estimator and its correction use.

> **TECHNOLOGY** **Sensor: Layer Thickness on a Polished Section**
>
> The thickness this chapter controls is measured only after coating, on a random sample of particles. The particles are sectioned and polished to near their midplane, and optical micrographs of the sections are analyzed by automated image analysis. Kernel diameters are read from shadowgraphs [18]. The image processing finds each layer boundary and corrects for the distance of the cut from the midplane. The first method read 80 particles mounted in a fixed grid, and an automated stage later extended it to mounts of hundreds of particles in random positions [11]. Buffer and inner pyrocarbon densities come from hot samples of the bed, taken before the next layer is deposited [18]. In the model the true thicknesses $\delta_{L,i}$ are known only to the simulator, and the chapter places the measurement about a week after the batch.

What Moves It

Furnace power heats a graphite susceptor, and the gas enters through the nozzle. Acetylene deposits the low-density buffer, propylene mixed with acetylene and argon deposits the dense pyrocarbon layers, and methyltrichlorosilane in hydrogen deposits the silicon carbide, with or without argon [6].

The program's list of coating improvements adds argon dilution during the silicon carbide step, changes to the coater retort to raise yields, and a better vaporizer for methyltrichlorosilane that drew on computer chip manufacture [11]. In the chapter's recipe every changeover switches the carrier and the precursor together, as a first-order ramp rather than a step. Figure 8.5 draws the loops on that hardware.

> **TECHNOLOGY Final Element: Thermal Mass Flow Controllers**
>
> Among the coating improvements of the qualification program were mass flow controllers that set the gas flows during the deposition of each layer [11]. A thermal mass flowmeter carries the gas past a heated sensor and an unheated one, and the heat the flow takes away measures its mass flow. A small mass flow controller adds its own throttling valve and control electronics. The reading depends on the specific heat of the gas, so a calibration holds for one composition only, and a change between laminar and turbulent flow shifts it in some designs [19]. In the model each gas train is a first-order lag of 6 s, and the controllers report the flow through the meter, not the flow that reaches the cone.

Where the Controller Runs

A distributed control system places a supervisory level over local controllers that hold their process conditions at set points, and loops in industry run with cycle times from milliseconds to minutes [20]. The app's control system scans every 7.2 s, which the chapter takes as close to what a batch coater's distributed control system runs. Read as plant hardware, TIC-301 and the fluidization schedule FIC-301 are local regulatory loops. The recipe sequencer, with its endpoint guards, the thickness estimator and the bed-mass correction, belongs to the supervisory level that ends each phase and hands the next targets down.

> **IN PRACTICE Acceptance on a Confidence Interval**
>
> A coated particle specification limits the mean of each property and the tail of its population. For one program's silicon carbide layer the mean had to lie between 32 and 38 µm, and at most 1 % of the particles could be 28 µm or thinner, both shown at 95 % confidence from a random sample. A lot whose property lies close to a limit needs many more particles measured before it can be accepted. Quick checks such as coating thickness are often run on every batch, so that one bad batch does not spoil a composite of several [18]. This chapter writes the tail as a silicon carbide standard deviation of at most 2.0 µm.

What Is Modernizing It

The first development is automated characterization. Layer thicknesses are now measured on mounts of hundreds of particles by image analysis, and computed tomography counts the particles in a fuel compact in minutes, although its instruments are expensive and may not be found in industrial plants [11]. The second is a closer link between the bed and the layer. One study fitted three nozzle geometries to a laboratory coater, mapped their particle circulation by particle image velocimetry at room temperature, and found clear differences in the structure of the pyrocarbon each one deposited [21]. The third is modeling and data. Computational fluid dynamics can now track every particle in the bed during coating, machine learning is mining past coating work for the link between process and properties, and the field has asked for automation and for more checks during fabrication to lower the cost of quality control [22].

8.9 Limitations

The model is a demonstration of inferential control and not a coater design tool. With lumped kinetics, seven size classes and a correlation-based hydrodynamic treatment it cannot support a claim about a real machine's deposition rate, coating efficiency, particle size distribution or fluidization limits, and the slump and carry-over flows of Table 8.5 should be read as the right trend rather than as design numbers. Nothing here addresses whether a batch would survive irradiation, because layer

anisotropy, density and crystallinity are absent and those properties matter as much as thickness [9]. The population is too coarse for the far tail a real qualification is written on, so Figure 8.4 demonstrates a mechanism without predicting a failure fraction. The carry-over state is clamped at 80 %, so the challenge scenario as loaded reports a particle loss that is a bound rather than a measurement. Disturbances arrive as commanded steps rather than as the drifting behavior of a plant, and the coating efficiency the application turns on is a constant in the plant as well as in the estimator. Three elements are stylized on purpose. A coating efficiency that steps from 1.00 to 0.72 at the touch of a button is not something a coater does, but a step is the disturbance whose absence from every trend can be shown in one run. A charge weighed 20 % heavy is beyond what a real weighing error produces, although it separates a baseline error from a gain error cleanly. And the true thicknesses are displayed beside the estimated ones, which no coater can do, because the gap of Equation (8.20) is the object of the lesson, and covering that trace turns the simulator back into the real problem. Where the individual assumptions break is set out in Section 8.3.7.

> **OPEN CHALLENGE** **From the Mean to Every Particle**
>
> A retrospective of the qualification program asks that quality control move from bulk properties to particle-to-particle distributions, which would reduce the conservatism of a specification's mean. Buffer density is measured by interrupting the coating after the buffer and gives one average over many particles, and inner pyrocarbon density needs particles from interrupted runs or hot samples of the bed. Nondestructive methods that an industrial plant can deploy are wanted [11]. A workshop of the field asked for more checks during fabrication, such as particle diameter during coating [22]. Progress would be a measurement made inside the run, which would give an endpoint like this chapter's a second correcting signal beside the load cell.

8.10 Conclusion

This chapter has treated a TRISO coating reactor as an interactive laboratory for control of a quality variable that has no transmitter. The simulator integrates thirty-four states, a susceptor and a bed temperature, two mass flow controller lags, particle carry-over, an agglomeration index and one thickness per layer for each of seven particle size classes, with deposition written as the harmonic mean of an Arrhenius kinetic branch and a precursor supply branch divided by a surface area that rises from 0.220 to 0.745 m^2 through the batch. A cruder model, one mean particle at the charge as weighed and the nominal coating efficiency, is integrated alongside it and is the only thickness the control system sees. On a nominal plant the two end within 0.09 µm of each other over 210 µm of coating and the timer and the inferential endpoint deliver the same batch in 191 min. A silicon carbide step run 50 °C cold is inside the estimator's model, so the timer gives 26.2 µm and is rejected while the inferential endpoint gives 35.0 µm for 21.6 min of extra time. A coating efficiency of 0.72 is outside it and leaves the batch 28.4 µm thin with all four layers rejected, every setpoint held and every layer ending on schedule. A load cell closes that gap to 0.79 µm for 28 min, and stays inside 2.2 µm for efficiencies from 1.00 down to 0.50 while the batch time absorbs the disturbance from 191 to 255 min. The same correction is blind to a 20 % charge error, which scales the modeled and measured mass gain together and is visible only in the cell's absolute reading of 240 g against a 200 g ticket. Underneath all of it the minimum spouting velocity rises by a factor of 2.47 and the silicon carbide standard deviation rises from 0.61 to 3.53 µm as the fluidization setpoint falls from 2.20 to 1.15 with every mean thickness on target. Further work would add a moving horizon estimator on the same equations, an off-gas analyzer as a second correcting measurement and a population balance in place of the seven classes. The TRISO coating reactor is an accessible browser platform for learning what a soft sensor can and cannot defend.

KEY POINTS

- A controller can hold every measurement on its setpoint while the coating it never sees comes out thin.
- On a nominal plant the timer and the inferential endpoint give the same batch, and a cold silicon carbide step separates them.
- A coating efficiency of 0.72 leaves the batch 28.4 µm thin with all four layers rejected, and a load cell closes that gap to 0.79 µm for 28 min.
- The same correction is blind to a 20 % charge error, which only the cell's absolute reading reveals.
- The silicon carbide spread rises from 0.61 to 3.53 µm as the fluidization setpoint falls, with every mean on target.

Where else this applies. Film thickness on semiconductor wafers is measured between wafers, so its loop runs from run to run [3], as in the lithography line of Volume 14, Chapter 6.

References

[1] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[2] F. A. A. Souza, R. Araújo, and J. Mendes, "Review of soft sensor methods for regression applications," *Chemometrics and Intelligent Laboratory Systems*, vol. 152, pp. 69–79, 2016.

[3] T. F. Edgar, S. W. Butler, W. J. Campbell, C. Pfeiffer, C. Bode, S. B. Hwang, K. S. Balakrishnan, and J. Hahn, "Automatic control in microelectronics manufacturing: Practices, challenges, and possibilities," *Automatica*, vol. 36, no. 11, pp. 1567–1603, 2000.

[4] M. Ogawa, M. Ohshima, K. Morinaga, and F. Watanabe, "Quality inferential control of an industrial high density polyethylene process," *Journal of Process Control*, vol. 9, no. 1, pp. 51–59, 1999.

[5] S. Weintroub, "The Dragon high-temperature reactor project," *Nature*, vol. 197, no. 4864, pp. 240–241, 1963.

[6] P. A. Demkowicz, B. Liu, and J. D. Hunn, "Coated particle fuel: Historical perspectives and current progress," *Journal of Nuclear Materials*, vol. 515, pp. 434–450, 2019.

[7] H. Knüfer, "Preliminary operating experiences with the AVR at an average hot-gas temperature of 950 °C," *Nuclear Engineering and Design*, vol. 34, no. 1, pp. 73–83, 1975.

[8] H. Nickel, H. Nabielek, G. Pott, and A. W. Mehner, "Long time experience with the development of HTR fuel elements in Germany," *Nuclear Engineering and Design*, vol. 217, no. 1-2, pp. 141–151, 2002.

[9] D. A. Petti, J. Buongiorno, J. T. Maki, R. R. Hobbins, and G. K. Miller, "Key differences in the fabrication, irradiation and high temperature accident testing of US and German TRISO-coated particle fuel, and their implications on fuel performance," *Nuclear Engineering and Design*, vol. 222, no. 2-3, pp. 281–297, 2003.

[10] D. Petti, J. Maki, J. Hunn, P. Pappano, C. Barnes, J. Saurwein, S. Nagley, J. Kendall, and R. Hobbins, "The DOE advanced gas reactor fuel development and qualification program," *JOM*, vol. 62, no. 9, pp. 62–66, 2010.

[11] P. A. Demkowicz, D. A. Petti, J. D. Stempien, J. Palmer, W. F. Skerjanc, D. M. Scates, G. Hawkes, T. Gerczak, and G. Helmreich, "Advanced gas reactor fuel development and qualification program retrospective and lessons learned," Idaho National Laboratory, Idaho Falls, ID, Report INL/RTP-25-06381, Revision 0, 2026, oSTI 17429367.

[12] P. A. Demkowicz, J. D. Hunn, D. A. Petti, and R. N. Morris, "Key results from irradiation and post-irradiation examination of AGR-1 UCO TRISO fuel," *Nuclear Engineering and Design*, vol. 329, pp. 102–109, 2018.

[13] Z. Yan, L. Jiang, Y. Tian, R. Liu, Y. Shao, B. Liu, and M. Liu, "Multi-scale study of fluidized bed-chemical vapour deposition process in nuclear fuel coated particle fabrication for high-temperature gas-cooled reactor: A review," *The Canadian Journal of Chemical Engineering*, vol. 103, no. 4, pp. 1502–1517, 2025.

[14] K. B. Mathur and P. E. Gishler, "A technique for contacting gases with coarse solid particles," *AIChE Journal*, vol. 1, no. 2, pp. 157–164, 1955.

[15] Z. J. Hu and K. J. Hüttinger, "Mechanisms of carbon deposition—a kinetic approach," *Carbon*, vol. 40, no. 4, pp. 624–628, 2002.

[16] H. Sone, T. Kaneko, and N. Miyakawa, "In situ measurements and growth kinetics of silicon carbide chemical vapor deposition from methyltrichlorosilane," *Journal of Crystal Growth*, vol. 219, no. 3, pp. 245–252, 2000.

[17] M. Olazar, M. J. San Jose, A. T. Aguayo, J. M. Arandes, and J. Bilbao, "Stable operation conditions for gas-solid contact regimes in conical spouted beds," *Industrial & Engineering Chemistry Research*, vol. 31, no. 7, pp. 1784–1792, 1992.

[18] P. A. Demkowicz, D. W. Marshall, and J. D. Hunn, "Advanced gas reactor fuel specification technical bases," Idaho National Laboratory, Idaho Falls, ID, Report INL/RPT-23-71992, Revision 0, 2023, oSTI 1983878.

[19] 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.

[20] 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.

[21] B. Conry, K. I. Montoya, R. S. Heldt, E. G. Herbert, and E. López-Honorato, "Linking spout fluidization hydrodynamics to pyrolytic carbon deposition characteristics in a fluidized bed chemical vapor deposition reactor," *Chemical Engineering Science*, vol. 334, p. 124319, 2026.

[22] P. A. Demkowicz and I. Glagolenko, "Advanced fuels campaign FY25 coated particle fuel workshop summary," Idaho National Laboratory, Idaho Falls, ID, Report INL/RPT-25-87367, 2025, oSTI 17121540.

8.11 Exercises

8.1 Concept check · 1 min

Coater CV-301 carries a thermocouple, three mass flow controllers, a differential pressure cell and a load cell. Which quantity does one of these instruments read during a batch?

(a) The mean silicon carbide thickness, through the integrated estimate

(b) The spouting ratio U/U_{ms}, through the fluidization loop FIC-301

(c) The mass of the bed, through load cell WT-301 under the whole coater

(d) The coating efficiency, through the gap between estimate and truth

8.2 Concept check · 1 min

Which statement about the balance below is right?

$$C_b \frac{\mathrm{d}T_b}{\mathrm{d}t} = UA\,(T_w - T_b) - C_g\,(T_b - T_g) + \Delta H_L\,\dot{m}_{dep}$$

(a) It is the susceptor balance, and its last term is the radiative loss to the room

(b) It is the growth law of one layer, and its last term is the deposit on one size class

(c) It is the thermocouple lag, and its last term is the bias added to the reading

(d) It is the bed balance, and its last term is the heat released as the deposit forms

8.3 Concept check · 2 min

In scenario 10 the coater holds 20 % more kernels than the batch ticket records, and the endpoint is Inferential + mass. Why does the bed-mass correction fail to rescue the batch?

(a) The deposit is spread over more particles, so the measured gain matches the model and the factor stays near one

(b) The load cell is filtered with a time constant of one minute, so the extra mass is reported too late to act on

(c) The factor is clamped between 0.25 and 3, and the factor this error needs lies outside that range

(d) The gate holds the factor at one until 6 g has been deposited, which this batch never reaches

8.4 Theory and calculation · 6 min

The silicon carbide layer has $E_a = 230\,\mathrm{kJ/mol}$ and $T_{ref} = 1500\,^{\circ}\mathrm{C}$, with $\mathscr{R} = 8.314\,\mathrm{J/(mol\,K)}$. In mid-layer of a nominal batch the two branches of Equation (8.9) are $J_{kin} = 3.09 \times 10^{-3}\,\mathrm{kg/(m^2\,min)}$ and $J_{sup} = 4.90 \times 10^{-3}\,\mathrm{kg/(m^2\,min)}$.

(a) Find the factor by which the kinetic branch changes when the bed runs at 1450 °C and at 1550 °C. *Checkpoint:* 0.636 and 1.534.

(b) Find the combined flux J at 1500 °C and at 1450 °C, with the supply branch unchanged, and their ratio. *Checkpoint:* 1.89×10^{-3} and 1.40×10^{-3} kg/(m^2 min), ratio 0.740.

(c) A 59.5 min hold lays 34.9 µm of silicon carbide at 1500 °C. Predict the thickness the same hold lays at 1450 °C, and the time the inferential endpoint needs there to reach 35.0 µm. Compare with the 26.2 µm and 81.1 min of scenario 7. *Checkpoint:* 25.8 µm and 80.6 min.

8.5 Theory and calculation · 7 min

A fresh charge is $m_0 = 200\,\mathrm{g}$ of uranium oxycarbide kernels of $d_k = 500\,\mu\mathrm{m}$ and $\rho_k = 10\,900\,\mathrm{kg/m^3}$ in a column of $D_c = 75\,\mathrm{mm}$ with an 8 mm nozzle and voidage $\varepsilon = 0.42$. At the buffer condition, 1250 °C, the gas density is $\rho_g = 0.294\,\mathrm{kg/m^3}$ and 5.5 slpm of acetylene flows. Take $g = 9.81\,\mathrm{m/s^2}$.

(a) Treating every kernel as a 500 µm sphere, find the particle count N and the surface area A_b. *Checkpoint:* 280 300 particles and 0.220 m^2.

(b) Find the bed depth H and the minimum spouting velocity U_{ms} of Equation (8.6). *Checkpoint:* 7.16 mm and 0.228 m/s.

(c) Find the carrier flow that holds a spouting ratio of 2.20, using Equation (8.5). *Checkpoint:* 18.4 slpm.

8.6 Theory and calculation · 8 min

The estimator's mean particle has a kernel radius $\hat{r}_0 = 250\,\mu\mathrm{m}$, the count is $\hat{n} = 279\,865$ and the buffer density is $\rho_1 = 1050\,\mathrm{kg/m^3}$.

(a) Find the bed mass gain the estimator expects for a 95 µm buffer. *Checkpoint:* 31.3 g.

(b) The load cell reports a gain of 22.5 g over the buffer. Invert it through the geometry, as Equation (8.21) does, and compare with the shortcut that scales 95 µm by the ratio of the gains. *Checkpoint:* 73.7 µm against 68.3 µm.

(c) In scenario 8 the true bed mass rises by 24.86 g over the buffer layer. Invert it and compare with the Buffer tile, 79.67 µm. *Checkpoint:* 79.6 µm.

8.7 App challenge · 9 min

Load scenario 7, The SiC step runs fifty degrees cold, which opens on the Timer endpoint with SiC temperature at 1450 °C. Use the default seed and the default speed of 20, which plays 20 min of the batch each second. The run pauses itself when the batch completes, and the notes record when each layer starts and ends.

(a) Run as loaded. Read SiC mean, and the SiC start and end times in the notes. *Checkpoint:* 26.22 µm, 4.0 min to 63.5 min, rejected.

(b) Press Reset, set Endpoint to Inferential and run again. Read the same quantities and the batch time. *Checkpoint:* 35.03 µm, 4.0 min to 85.1 min, accepted in 139 min.

(c) Find the ratio of the two layer durations, compare it with the ratio of rates in the theory exercise on the silicon carbide branches, and say which term of Equation (8.19) extends the layer. *Checkpoint:* 0.734.

Accept the thicknesses within 0.05 µm and the times within 0.1 min.

8.8 App challenge · 9 min

Load scenario 8, The estimate lies, and nothing on the screen looks wrong, in which the coating efficiency is 0.72 and the endpoint is Inferential. Use the default seed and speed, and run each batch until it stops by itself.

(a) Run as loaded. Read Buffer, IPyC / OPyC, SiC mean, Inferred − actual and the batch time in the last note. *Checkpoint:* 79.67, 34.9 / 35.4 and 32.02 µm, a gap of 28.35 µm, 191 min, rejected.

(b) Set Endpoint to Inferential + mass, press Reset and run again. Read the same tiles, then export the CSV file and read `k_corr` in the last row. *Checkpoint:* 97.17, 41.2 / 39.9 and 34.84 µm, a gap of −0.79 µm, accepted in 219 min, factor 0.890.

(c) Explain why the second batch takes longer, and why the first gave no warning on its temperature, spouting and coating trends. *Checkpoint:* Each layer runs until the corrected estimate reaches target, 28 min in all.

Accept thicknesses within 0.05 µm, the factor within 0.001.

8.9 App challenge · 8 min

Load scenario 12, The bed goes slack during the SiC layer, which starts at the purge before the silicon carbide layer with Spouting ratio setpoint at 1.30. Use the default seed and speed, and run each batch until it stops by itself.

(a) Run as loaded. Read SiC mean, SiC 1σ and the verdict. *Checkpoint:* 35.11 µm, 2.89 µm, rejected on SiC variability.

(b) Press Reset, set the setpoint to 1.50 and run. Repeat at 1.55. Make each setting before the layer starts at 4 min. *Checkpoint:* 2.06 µm rejected; 1.86 to 1.87 µm accepted.

(c) Explain why the mean stays on target while the spread changes, and why raising the setpoint above 2.0 no longer helps. *Checkpoint:* The spread is 0.61 µm at 2.0 and above.

Accept SiC 1σ within 0.01 µm of the stated values or range.

8.10 Application area · 3 min

A plant manager notes that every mean in scenario 12 at a ratio of 1.30 was in tolerance, and proposes dropping the silicon carbide spread from the acceptance criteria. Give two reasons from the chapter for keeping it.

8.11 Application area · 3 min

For each fault below, name the measurement on the coater that reveals it and say what that measurement compares: a coating efficiency of 0.72, a charge 20 % heavier than its ticket, and a carrier leak that delivers 62 % of the indicated flow.

8.12 Review the agent · 5 min

An AI assistant is asked how to pass the challenge, A batch with the efficiency unknown, and writes: “The coating efficiency is the only fault that costs quality, and only the load cell bears on it. Set Endpoint to Inferential + mass and leave the other settings as loaded. The correction will absorb the efficiency error, so the batch will be accepted, if slowly.” Find the flaw and say how you would check it.