INTERACTIVE CONTROL ENGINEERING
Principles and Applications

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

CHAPTER 4

Batch Crystallizer Simulator for Trajectory Control Inside a Metastable Zone

Pages 83–110

Edited by John D. Hedengren

How to cite this chapter. "Batch Crystallizer Simulator for Trajectory Control Inside a Metastable Zone," 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. 4, pp. 83–110.

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

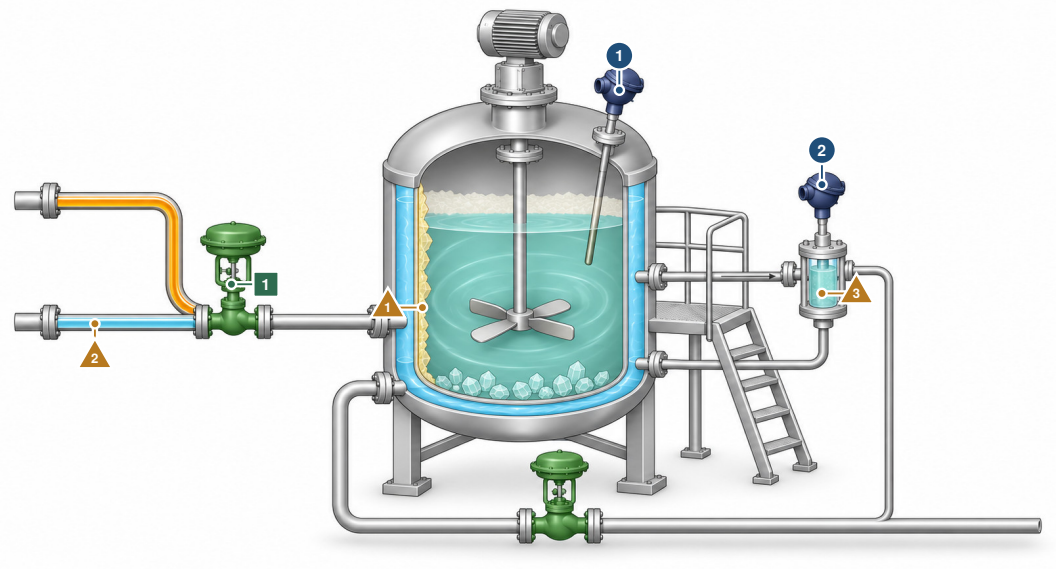
The app for this chapter
https://apmonitor.com/apps/batch-crystallization/

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

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

CHAPTER 4

Batch Crystallizer Simulator for Trajectory Control Inside a Metastable Zone



The controller moves 1 three-way valve blending hot and chilled utility for the jacket

The controller reads 1 batch temperature probe 2 liquor concentration analyzer

What disturbs it 1 scale growing on the cold wall 2 chilled utility warming up 3 analyzer bias

The valve cools the batch along a planned curve, but the crystals are decided by how far the liquor sits above saturation, which only the analyzer sees. Illustration, not to scale.

IN PLAIN TERMS **Growing Large Crystals by Cooling Slowly at First**

The main idea. Cooling a hot liquid brings a dissolved solid out as crystals. A few seed crystals are added so that the new solid grows on them, and the crystal size is judged once, after the batch.

The challenge. Cooled fast while the seeds are small, the liquid makes a cloud of tiny new crystals instead. The thermometer can follow its cooling plan closely all the while, so the shape of the plan matters most.

Why it matters. Tiny crystals clog the filter and dry badly, so a batch that looked fine on every chart can still fail later.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 4.2, the two try-its and the two facts in Section 4.6, and the key points
Study	50 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 4.6), Technology in Practice (Section 4.8), Limitations (Section 4.9) and the open challenge

The control question A seeded cooling crystallization is graded once, after the batch, on crystals no instrument in the loop can see. If three cooling curves on one well-tuned temperature loop give mean sizes that differ by a factor of 4.3, what should the control system be designed around: the tracking error, the trajectory, or the constraints the trajectory meets?

Learning objectives

- Compute the solubility, the relative supersaturation and the metastable limit of the batch from the solubility anchors.
- Derive the cooling rate a seed surface can absorb, and explain why the cubic curve suits a seeded batch.
- Explain why peak supersaturation orders the product and time outside the metastable zone does not.
- Compare the single loop and the cascade with a wall subcooling clamp on an encrusting vessel, by batch time and scale.
- Judge what supersaturation control and a predictive jacket supervisor gain and what each costs in analyzer dependence, fines and valve travel.

Prerequisites Sections 1.4.1 (batch against continuous), 1.4.3 (trajectory tracking), 1.4.4 (path constraints), 1.6.2 (the PI law) and 1.6.1 (engine, speed and seed). Section 1.6.3 is needed only for the predictive supervisor.

The app https://apmonitor.com/apps/batch-crystallization/

See also Volume 3, Chapter 2 for a loop that rejects only the disturbances entering its measurement path, Volume 3, Chapter 3 for the inner-loop speed a cascade requires, and Volume 3, Chapter 8 for a particle size that answers to a residence time and is measured late.

4.1 Introduction

Most plants in an introductory process control course are evaluated while they run. A temperature holds a setpoint, a level holds a band, a composition holds a specification, and the integrated error of the loop is a fair measure of how the plant was operated [1]. A seeded cooling crystallization is evaluated once, after the batch ends, by a particle size analyzer applied to the crystals on the filter. No quantity the control system measured appears in that evaluation. The loop regulates batch temperature because batch temperature has a transmitter, and the engineering content of the recipe is the choice of a temperature trajectory whose consequence is the required size distribution.

The simulator of this chapter makes the gap between the controlled variable and the product objective measurable. Three programmed cooling curves span the same 28 °C in the same 180 minutes on the same single temperature controller, and they give mass-weighted mean sizes of 86.3, 159.2 and 373.4 µm at yields of 62.5, 62.1 and 62.2 %. The temperature tracking gives the opposite ordering. The batch that follows its curve most closely, with a peak deviation of 0.142 °C, finishes at 96.3 % fines by number. The batch with the best product accumulates the largest integrated error of the three and the largest valve travel. This comparison is the central result of the chapter, and it is not available from a transfer function between valve and temperature, because the objective does not appear in that transfer function.

Crystallization determines the solid form of a pharmaceutical product. Size distribution and habit set the filtration rate, the drying time, the flow and compaction behavior in a tablet press and the dissolution rate that governs bioavailability, so the attribute is registered and the crystallizer is the unit that fixes it [2, 3]. A batch that makes 96.3 % fines by number does not fail at the crystallizer. It fails two unit operations later, at the filter and the dryer, after three hours of temperature trends that met the tracking specification.

The class of problem is broader than crystallization. It is batch trajectory tracking against a path constraint, where the objective functional is evaluated only at the final time and the controlled variable is a proxy for a quantity no online instrument reports. Section 1.4.1 sets out these properties of a batch on the kettle of Chapter 1. Batch reactors with a selectivity objective, fed-batch fermentations, polymerization with a molecular weight distribution target and curing operations share this structure.

The form of the optimal input for such problems is a documented general result [4], and the crystallization instance carries a review literature of its own [5]. Two features recur. The optimal input is usually determined by constraints rather than by a sensitivity, so most of the design work is deciding which constraint is active and when. The online measurement is chosen for the transmitter that exists, so an added measurement changes the achievable objective rather than the achievable tracking error.

The model rests on three conservation statements, a solute mass balance, energy balances on the contents and the jacket, and a population balance carried by its first five moments [6], with relative supersaturation as the single driving force for growth and nucleation [7]. The experiments built on the model exercise trajectory tracking, cascade structure justified by time-scale separation, anti-windup and output constraints, measurement lag and bias, and inferential control from a secondary measurement, all standard topics in process dynamics and control texts [1, 8].

This chapter describes a browser-hosted seeded batch crystallizer in the detail needed to rebuild it, and establishes on that plant that the shape of the trajectory decides the product while the tuning of the loop following that trajectory does not. The analysis rests on two measured results, that the ordering of the products reverses the ordering of the tracking errors, and that peak supersaturation orders the outcomes while time spent outside the metastable zone does not. The crystallization kinetics, the method of moments and the cooling profiles are all standard.

Section 4.2 traces the metastable zone, programmed cooling and in-process measurement. Section 4.3 sets down the nine states, a population balance, a wall film and an instrument chain, with their constants, in enough detail for the plant to be reproduced. Section 4.4 compares the three programmed cooling profiles with the cooling rate the crystal surface can absorb. Section 4.5 describes three classical controller architectures on one valve and a predictive supervisor of the jacket. Section 4.6 reports the fourteen scripted scenarios: the cooling curves and a tuning sweep, the ordering of the outcomes, encrustation, supersaturation control, the utility limit, the graded challenge and the predictive comparison. Section 4.8 describes the probe, the valve and the control system of a real crystallizer. Sections 4.7 and 4.9 place the results in plant practice and state the limitations, Section 4.10 summarizes, and the exercises follow.

4.2 Background

The central concept in this domain predates automatic control of it. Its first measurement, told in the history segment below, separated saturation from a further condition at which crystals appear spontaneously. The band between the two is the metastable zone, and it is the operating corridor for a seeded batch. Inside it the crystals already present grow. Outside it the solution nucleates new crystals, a few micrometers across and in numbers that dominate the population count.

> ⧖ HISTORY **1906: A Shower of Crystals Marks the Limit**
>
> In 1906 Henry Miers and Florence Isaac followed the refractive index of cooling sodium nitrate solutions, read by total reflection in a glass prism dipped into the liquid. A stirred solution first grew a few crystals slowly. At a lower temperature a shower of crystals appeared throughout the liquid, and the concentration fell suddenly. They took the two stages for Ostwald's metastable and labile conditions, and they drew a supersolubility curve, nearly parallel to the solubility curve, where a solution passes from one to the other. A crystal put into a metastable solution grows, and if the cooling is slow enough the labile condition is never reached and no cloud of crystals forms [9].

Measurement of the zone width is method dependent. The polythermal method cools a solution at a known rate and records the temperature at which crystals first appear, and the width recovered depends on the cooling rate, on the volume, on the stirring and on the thermal history of the sample [10]. Later analysis established that the width measured this way is set by how long the first nuclei take to grow to a detectable size, so it is a kinetic quantity rather than a thermodynamic boundary [11]. The upper

edge of the corridor is therefore taken from an experiment that does not resemble the batch, and no plant instrument reports its current position.

Cooling policy became a control question once the zone was understood. An unattended jacketed vessel cools on a natural exponential, fastest at the start, which places the highest demand for solute removal at the moment the vessel holds the least crystal surface. Mullin and Nývlt derived the cooling profile that keeps supersaturation constant through a seeded batch and showed it to be strongly convex, slow at the start and fast at the end, and reported that programmed cooling gives a coarser product than either natural or linear cooling at the same endpoints [12]. Jones formulated the same question as an optimal control problem on a population balance and obtained profiles of the same shape [13]. The result has been confirmed experimentally many times, and a controlled comparison on seeded potassium alum found the product mean size to increase from natural through linear to programmed cooling with the seed charge held fixed [14]. Seeding practice moved on the same evidence, from an operator habit to a seed charge and a seeding temperature chosen so that the supersaturation available at the start matches what the seed surface can absorb [7, 15].

Two elements were still missing. The first was a model that could carry a distribution through a dynamic simulation without solving a partial differential equation, which the population balance and its moment transformation supplied, and which remains the standard representation for crystallizer control studies [6]. The second was a measurement. A programmed cooling curve is designed for one seed loading, one growth rate and one solubility, and a plant that departs from any one of them makes the curve wrong. The arrival of in-situ attenuated total reflectance infrared spectroscopy, calibrated through a chemometric model to report liquor concentration during the batch, made the supersaturation itself available as a controlled variable, and closed-loop supersaturation control of a batch crystallizer was demonstrated on that basis [16]. Fujiwara and coworkers separated the two design philosophies that followed, a first-principles route that fits kinetics and optimizes a profile against them, and a direct design route that abandons the kinetic model and holds a concentration trajectory in the phase diagram [17]. Process analytical technology guidance made in-process measurement of this kind a regulatory expectation rather than a research practice [3]. Direct nucleation control went one step further by counting particles with a focused beam reflectance probe and moving the temperature to hold the count, dissolving fines deliberately when there are too many [18]. A review of the field places supersaturation control and direct nucleation control as the two approaches now in industrial use, both motivated by the absence of the product objective from the measured loop [19].

4.3 Crystallizer Model

Crystallizer CR-101 is a 100 liter jacketed vessel charged with 100 kg of water holding 14.0 kg of a dissolved organic solute, which is exactly saturation at 50 °C. The charge is cooled to 48 °C, seeded with 25 g of 50 µm crystal, and taken to 20 °C over 180 minutes. Cooling comes from a jacket fed by a three-way valve that blends a 60 °C hot utility with a 5 °C chilled one. Figure 4.1 shows the equipment, the tag numbers and every symbol the equations below use. Temperature transmitter TT-101 reports the controlled variable and analyzer AT-101 reports the liquor concentration that the third architecture needs, while valve position is the one quantity any of the three controllers can move.

Table 4.1 sorts the signals into the roles a control narrative uses. Units are those of a batch sheet. Time is in minutes, temperature in degrees Celsius, concentration in kilograms of solute per kilogram of solvent, crystal size in meters and displayed in micrometers, moments per kilogram of solvent, heat in kilojoules and valve position in percent open to the chilled utility. Nine states are integrated.

4.3.1 Solubility, Supersaturation and the Metastable Limit

Solubility is a quadratic in temperature, anchored on two measured points and a curvature rather than on three fitted coefficients, so the parameters correspond to what a solubility measurement produces.

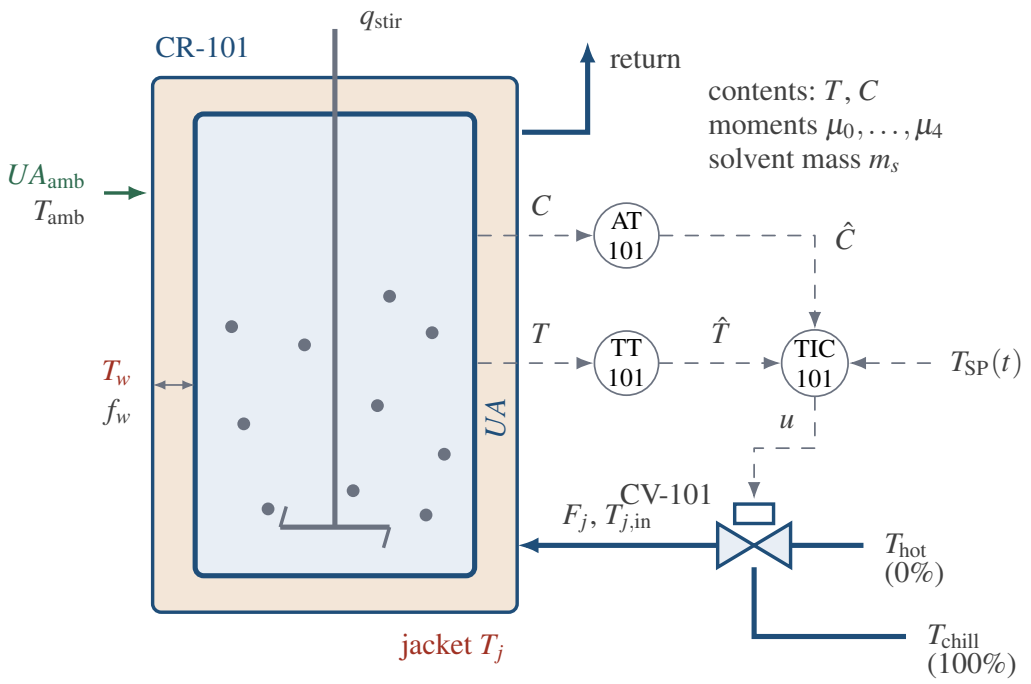


Figure 4.1 Jacketed batch crystallizer CR-101 with its three-way coolant valve. The contents hold solvent mass m_s at temperature T and dissolved solute C, with the crystal population carried as the five moments μ_0 through μ_4. Heat leaves through UA into the jacket at T_j, which is fed at F_j from valve CV-101 blending hot utility T_{hot} at 0 % with chilled utility T_{chill} at 100 %. The crystallizer-side wall sits a fraction f_w of the driving force below the bulk at T_w, which is where encrustation starts. TT-101 reports the batch temperature and AT-101 the liquor concentration, and TIC-101 follows the programmed trajectory $T_{SP}(t)$.

Table 4.1 Classification of Variables. Nominal Values Are Those of the Base Recipe at the Instant the Seed Is Charged.

Symbol	Role	Nominal	Range
T (TT-101)	CV, measured	48.0 °C	20–48 °C
σ (AT-101)	CV (SSC-101), measured	0.062	0.02–0.32 SP
u (CV-101)	MV	20.6 %	0–100 %
T_{chill}	DV, measured	5 °C	2–20 °C
probe bias	DV, unmeasured	0 °C	±2 °C
analyzer bias	DV, unmeasured	0 g/kg	±5 g/kg
T	state	48.0 °C	—
T_j	state	48.4 °C	$\geq T_{chill}$
C	state	0.1400 kg/kg	≥ 0
μ_0	state	2.57×10^6/kg	≥ 0
$\mu_1 \ldots \mu_4$	state	—	≥ 0
ϕ (fouling)	state	0	0–1

The driving force for everything that follows is the relative supersaturation of Equation (4.1).

$$C^*(T) = a_0 + a_1 T + a_2 T^2, \qquad \sigma = \frac{C - C^*(T)}{C^*(T)} \tag{4.1}$$

where C is the dissolved solute in kg/kg of solvent, C^* is the solubility at the same units and σ is dimensionless. Anchor points of 52 g/kg at 20 °C and 140 g/kg at 50 °C with $a_2 = 4.0 \times 10^{-5}\,\mathrm{kg/(kg\,^\circ C^2)}$ give $a_0 = 3.3333 \times 10^{-2}\,\mathrm{kg/kg}$ and $a_1 = 1.3333 \times 10^{-4}\,\mathrm{kg/(kg\,^\circ C)}$. Every rate law, every limit and every setpoint in the application is written in σ and not in the absolute difference $C - C^*$, which is reported for reference only.

The metastable zone width is quoted as degrees of subcooling, so the limit curve is the solubility curve shifted by ΔT_{msz} and the limit in relative supersaturation follows from Equation (4.2).

$$\sigma_{\max}(T) = \frac{C^*(T + \Delta T_{\mathrm{msz}})}{C^*(T)} - 1 \tag{4.2}$$

Because the solubility curve is quadratic rather than exponential, the limit moves with temperature. At 48 °C it is 0.175, at 40 °C 0.190 and at 20 °C 0.207, so the warm end of the batch is the tight end and the margin matters rather than the supersaturation alone.

4.3.2 Growth and Nucleation

Growth is size independent and nucleation has two mechanisms, as shown in Equations (4.3) and (4.4).

$$G = k_g\,\sigma^{g} \tag{4.3}$$

$$B = \underbrace{k_b\,\sigma^{b}\,\mu_3^{j}}_{\text{secondary}} + \underbrace{k_p\,\sigma^{p}\left[1 + \mathrm{e}^{-(\sigma - \sigma_{\max})/w}\right]^{-1}}_{\text{primary}} \tag{4.4}$$

where G is the linear growth rate in m/min and B is the nucleation rate in nuclei per kilogram of solvent per minute. Secondary nucleation is catalyzed by the crystal surface already in the vessel, so it carries the magma term μ_3^{j} and runs everywhere inside the zone as well as outside it. Primary nucleation is spontaneous and is gated off below the metastable limit. With $p = 5$ and a gate width of $w = 0.003$ the primary rate at 40 °C climbs 25 orders of magnitude between a supersaturation 0.15 below the limit and one 0.05 above it, and it multiplies by seven for every hundredth of supersaturation crossed at the limit itself. The gate is a logistic rather than a step because the fixed-step Runge-Kutta integrator evaluates the right-hand side four times per step, and a genuine discontinuity makes those four evaluations disagree about which side of the limit the batch is on. When the liquor is undersaturated, Equation (4.3) is replaced by a dissolution rate $-k_d|\sigma|$ faded linearly to zero over the last five nucleus diameters as the number mean size falls toward L_0, because the moments have no mechanism for removing a crystal from a population.

4.3.3 Population Balance by the Method of Moments

The full description of a crystal population is a number density $n(L,t)$ governed by a partial differential equation in size and time. The moments are its integrals, $\mu_k = \int_0^{\infty} L^k n(L)\,\mathrm{d}L$, and for size-independent growth the partial differential equation collapses exactly onto the ordinary differential equations of Equation (4.5) with nuclei born at a single size L_0 [6].

$$\frac{\mathrm{d}\mu_0}{\mathrm{d}t} = B, \qquad \frac{\mathrm{d}\mu_k}{\mathrm{d}t} = k\,G\,\mu_{k-1} + B\,L_0^{k}, \quad k = 1 \ldots 4 \tag{4.5}$$

Each moment has a meaning. μ_0 is the crystal count, μ_2 is the quantity proportional to total crystal surface, μ_3 is total crystal volume and therefore mass, and μ_4 is what makes the mass-weighted mean available. Equation (4.5) is not an approximation of the population balance under these kinetics. It is the same equation integrated, and the application verifies the identity $\mathrm{d}\mu_3/\mathrm{d}t = 3G\mu_2$ directly with nucleation switched off.

🕮 **Pause.** Try Exercise 4.2 before reading on.

Solute leaves the liquor exactly as crystal volume appears, which fixes the solute balance as Equation (4.6).

$$\frac{\mathrm{d}C}{\mathrm{d}t} = -\rho_c k_v \frac{\mathrm{d}\mu_3}{\mathrm{d}t} = -\rho_c k_v \left(3G\mu_2 + BL_0^3\right) \tag{4.6}$$

where ρ_c is the crystal density and k_v the volume shape factor. Writing the balance this way conserves total solute by construction rather than approximately, so the yield reported after a three hour run carries no accumulation error. Every quantity the batch is evaluated on is then arithmetic on the same five numbers, as shown in Equation (4.7).

$$L_{10} = \frac{\mu_1}{\mu_0}, \quad L_{43} = \frac{\mu_4}{\mu_3}, \quad CV = \sqrt{\frac{\mu_0 \mu_2}{\mu_1^2} - 1} \tag{4.7}$$

where L_{43} is the mass-weighted mean size a particle size analyzer reports and a formulation specification is written against, L_{10} is the number mean, which falls sharply when a nucleation event adds fines, and CV is the coefficient of variation of the number distribution. The magma density is $M_T = \rho_c k_v \mu_3$ in kilograms of crystal per kilogram of solvent. The yield is $(C_0 - C)/C_0$ on the solute originally dissolved, and the fines fraction is the nucleated crystal count $(\mu_0 - \mu_{0,\mathrm{seed}})/\mu_0$, reported by number because particle count governs filter blinding.

The seed itself is charged as a log-normal cut of declared mean size and coefficient of variation, for which $M_k = L^k (1 + v^2)^{k(k-1)/2}$, so the two numbers reported by a sieve analysis fix all five seed moments. At the base recipe the 25 g charge is 2.57×10^6 crystals per kilogram of solvent, or 2.6×10^8 in the vessel.

4.3.4 Energy Balances, the Jacket and the Wall

The contents and the jacket are each well mixed, and the three-way valve sets the jacket inlet temperature algebraically, as shown in Equation (4.8).

$$\begin{aligned} M c_p \frac{\mathrm{d}T}{\mathrm{d}t} &= -UA_{\mathrm{eff}}(T - T_j) + \Delta H_c \, m_s \rho_c k_v \frac{\mathrm{d}\mu_3}{\mathrm{d}t} \\ &\quad + UA_{\mathrm{amb}}(T_{\mathrm{amb}} - T) + q_{\mathrm{stir}} \\ m_j c_j \frac{\mathrm{d}T_j}{\mathrm{d}t} &= F_j c_j (T_{j,\mathrm{in}} - T_j) + UA_{\mathrm{eff}}(T - T_j) \\ T_{j,\mathrm{in}} &= T_{\mathrm{hot}} + \frac{u}{100}\left(T_{\mathrm{chill}} - T_{\mathrm{hot}}\right) \end{aligned} \tag{4.8}$$

where $M = m_s(1 + C^*(T_{\mathrm{sat},0})) = 114.0\,\mathrm{kg}$ is the total charge, c_p is the slurry heat capacity, ΔH_c is the heat of crystallization and u is the valve position in percent. Both sides of the first line carry kJ/min. The contents have a thermal time constant of $Mc_p/UA = 433.2/90 = 4.81\,\mathrm{min}$ and the jacket $m_j c_j/(F_j c_j + UA) = 54/234 = 0.231\,\mathrm{min}$, a separation of 21 to 1, which is the time-scale condition that justifies a cascade [8]. The fermenter of Volume 3, Chapter 3 measures how fast an inner loop must be when its lags are not separated this well. The process gain from valve position to batch temperature is negative, because 0 % is full hot utility. A percent of valve initially changes the cooling rate by 0.0703 °C/min rather than the temperature by a fixed amount, so the valve-to-temperature response is near-integrating, with self-regulation through the jacket blend and the ambient leak appearing only beyond an open-loop time constant of 7.6 min.

The crystallizer-side wall sits a fixed fraction of the driving force below the bulk, and the liquor touching it is the same liquor at a lower temperature, as shown in Equation (4.9).

$$T_w = T - f_w(T - T_j), \qquad \sigma_w = \frac{C - C^*(T_w)}{C^*(T_w)} \tag{4.9}$$

The wall supersaturation is therefore always higher than the bulk value, and a batch can encrust while the bulk marker is comfortably inside the corridor. Once the wall crosses its own limit, crystal grows

on the wall and takes the heat transfer with it, as shown in Equation (4.10).

$$\begin{aligned} \frac{\mathrm{d}\phi}{\mathrm{d}t} &= k_f \max(0,\ \sigma_w - \sigma_{\max}(T_w))\,(1-\phi) \\ UA_{\mathrm{eff}} &= UA\,(1-\lambda\phi) \end{aligned} \tag{4.10}$$

where ϕ is the scaled fraction of the surface, k_f the encrustation rate and λ the heat transfer lost at full encrustation. Encrustation is the one fault in the application that feeds itself, because scale lowers UA_{eff}, which forces a colder jacket for the same duty, which lowers T_w and grows more scale. The polymerization reactor of Volume 12, Chapter 11 also insulates itself with the product it makes, and its jacket walks colder until it reaches a floor it cannot pass.

4.3.5 Instruments

TT-101 reports the batch temperature after three modifications. A constant bias is added, colored noise whose stationary standard deviation equals the configured value is added on top of it, and the sum is taken through a probe lag τ_T. AT-101 receives the same treatment with a lag of τ_C, which stands for an in-situ probe reporting through a chemometric model about once a minute. The two lags are stepped with $\hat{x}_k = \alpha\hat{x}_{k-1} + (1-\alpha)x_k$ and $\alpha = \exp(-\Delta t/\tau)$, so neither can go unstable at any sample period. The jacket resistance temperature detector sits in a fast-moving stream and is treated as ideal, because the point of the inner cascade loop is that its measurement is the good one. The third architecture computes $\hat{\sigma} = \hat{C}/C^*(\hat{T}) - 1$ from both instruments, so an error in either propagates into it.

4.3.6 Parameters, Initialization and Numerics

The constants are listed in Table 4.2. Only one of them is fixed, the softness of the nucleation gate, which exists for the integrator rather than for the physics. The rest are adjustable inputs, because the experiments concern choosing a trajectory rather than identifying a plant.

Every batch opens in exact thermal balance. The concentration is set to $C^*(T_{\mathrm{sat},0})$, the seed moments are charged from the log-normal expressions, and the jacket temperature is solved from Equation (4.8) with $\mathrm{d}T/\mathrm{d}t = 0$ so that the vessel is not drifting before the batch starts. At the base recipe that gives $T_j(0) = 48.40\,^{\circ}\mathrm{C}$ and a valve at 20.6 %. The jacket starts warmer than the batch, because a room at 22 °C takes 39.0 kJ/min out of a vessel at 48 °C while the agitator and the growing seed together return only 3.1 kJ/min. A batch is declared finished when T reaches $T_{\mathrm{end}} + 0.3\,^{\circ}\mathrm{C}$ with $\sigma < \sigma_{\mathrm{done}}$, or when the clock reaches $t_{\max}$, and the simulation pauses at that moment so the product measures are read at their final values. Temperature alone does not end a batch, because a vessel at the final temperature that still holds supersaturation continues to grow crystal.

The engine of Section 1.6.1 carries the nine states forward by classical fourth-order Runge–Kutta on a fixed 0.002 min (0.12 s) step. Controllers run every 0.02 min (1.2 s) and the historian keeps five samples for every simulated minute. Both noise sources are drawn from one generator whose seed is 20260829 (20260901 in scenario 11) and which is reseeded at every restart, so a noisy batch repeats exactly.

4.3.7 Assumptions

The model is detailed in the population balance and the wall film and simplified elsewhere. Each simplification has a domain where it fails.

One well-mixed vessel, with one exception. The contents have a single temperature and a single concentration, where a real 100 liter crystallizer has a supersaturation gradient near the cooling surface and a slower zone under the impeller. Exactly one piece of that is restored, the wall film of Equation (4.9), because it carries the encrustation constraint. Everything else about mixing is absent, including the effect of agitation rate on secondary nucleation, which in practice is one of the more influential operating variables.

Table 4.2 Model Parameters at the Base Recipe. All Are Adjustable Except the Nucleation Gate Width.

Symbol	Meaning	Value	Unit
m_s	solvent charge	100	kg
$T_{\mathrm{sat},0}$	saturation temperature of charge	50.0	°C
T_0	temperature at seeding	48.0	°C
m_{seed}	seed charge	25	g
L_{seed}	seed mean size	50	µm
CV_{seed}	seed coefficient of variation	0.25	—
C^*_{20}	solubility at 20 °C	52	g/kg
C^*_{50}	solubility at 50 °C	140	g/kg
a_2	solubility curvature	0.040	g/(kg °C^2)
ΔT_{msz}	metastable zone width	5.5	°C
k_g	growth coefficient	22	µm/min
g	growth order	1.10	—
k_d	dissolution coefficient	200	µm/min
k_b	secondary nucleation coefficient	5.0×10^9	—
b	secondary nucleation order	2.0	—
j	magma exponent	1.0	—
k_p	primary nucleation coefficient	1.0×10^{10}	—
p	primary nucleation order	5.0	—
L_0	nucleus birth size	2.0	µm
ρ_c	crystal density	1300	kg/m^3
k_v	volume shape factor	0.50	—
UA	jacket heat transfer	90	kJ/(min °C)
c_p	slurry heat capacity	3.80	kJ/(kg °C)
m_j	jacket holdup	15	kg
c_j	coolant heat capacity	3.60	kJ/(kg °C)
F_j	coolant circulation	40	kg/min
T_{hot}	hot utility temperature	60	°C
ΔH_c	heat of crystallization	100	kJ/kg
UA_{amb}	ambient heat leak	1.5	kJ/(min °C)
T_{amb}	room temperature	22	°C
q_{stir}	agitator heat input	3.0	kJ/min
f_w	process-side share of ΔT	0.30	—
k_f	encrustation rate	0.30	1/min
λ	heat transfer lost when fully scaled	0.70	—
τ_T	temperature probe lag	0.30	min
τ_C	analyzer lag	1.00	min
$t_{\max}$	maximum batch time	300	min
σ_{done}	end-of-batch supersaturation	0.020	—
w	nucleation gate width (hidden)	0.003	—

Moments with size-independent growth. Equation (4.5) is exact for the assumed kinetics, and five numbers give every graded quantity without reconstructing a distribution. What five numbers cannot give is a shape. A batch holding coarse seed and a fresh crop of 2 µm fines is bimodal, and its first five moments are identical to those of some unimodal distribution that is not in the vessel. The application therefore reports the fines separately by number, and the gap between L_{10} and L_{43} is the other signal that the population has split. No plot in the application is a particle size distribution.

No breakage, no agglomeration, no growth rate dispersion. All three widen a real distribution, and agglomeration in particular can double an apparent mean size while degrading product quality. Their absence is why a clean batch here reaches coefficients of variation a plant would rarely see.

One polymorph and no impurities. Solubility is a quadratic and the zone width is a fixed number of degrees, neither depending on impurity level, solvent composition or crystal form. Polymorphic

transformation during a cooling crystallization is a costly failure mode the application cannot represent, and neither can it represent the impurity rejection that is often the actual purpose of the step.

The metastable limit is a smooth gate at a fixed subcooling. Real zone widths are stochastic, depend on how fast the solution was cooled to reach the condition, and narrow as the batch accumulates surface, and what a polythermal measurement returns is a nucleation and growth time rather than a boundary [10, 11]. The application treats the limit as 5.5 °C of subcooling and draws it on the phase diagram so the simplification is visible rather than hidden.

Dissolution leaves the count wrong. The fade near L_0 keeps the solute balance closed through a partial dissolution and keeps the mass right, but the crystal count afterward is not physical and neither is the coefficient of variation computed from it.

Constant physical properties and one lumped fouling number. The heat of crystallization, the ambient leak, the agitator input and the slurry heat capacity do not change as solid appears, which over a 28 °C cooling of a 114 kg batch is worth a fraction of a degree. Real scale grows unevenly and changes the mixing as well as the heat transfer, where Equation (4.10) carries one fraction.

The numerics are more repeatable than an industrial control system. With the integration step and the execution period both frozen, one trajectory appears on every machine, and no distributed control system behaves that way.

4.4 Cooling Profiles and the Admissible Cooling Rate

Three programmed profiles are available, and all three run from T_0 to T_{end} in the same batch time t_b, so the comparison between them is about shape alone, as shown in Equation (4.11).

$$T_{\mathrm{SP}}(t) = \begin{cases} T_0 - \Delta T\,\tau & \text{linear} \\ T_{\mathrm{end}} + \Delta T\,\dfrac{\mathrm{e}^{-4.605\,\tau} - 0.01}{0.99} & \text{natural} \\ T_0 - \Delta T\,\tau^3 & \text{cubic} \end{cases} \tag{4.11}$$

where $\tau = t/t_b$ and $\Delta T = T_0 - T_{\mathrm{end}}$. The natural profile is the exponential an uncontrolled vessel with a fixed jacket temperature follows, rescaled to land exactly on the end point instead of one percent short of it. Over the base recipe the three start at 0.156, 0.724 and 0.000 °C/min and finish at 0.156, 0.007 and 0.467 °C/min, as shown in Figure 4.2.

The arithmetic that separates them is available in closed form. Holding supersaturation constant means $\mathrm{d}\sigma/\mathrm{d}t = 0$, and differentiating $C = C^*(T)(1+\sigma)$ and substituting Equation (4.6) gives the cooling rate the crystals can absorb, as shown in Equation (4.12).

$$\left|\frac{\mathrm{d}T}{\mathrm{d}t}\right|_{\sigma\,\mathrm{const}} = \frac{3\,\rho_c k_v\,G\,\mu_2}{(1+\sigma)\,\mathrm{d}C^*/\mathrm{d}T} \tag{4.12}$$

where the nucleation contribution has been dropped because inside the zone it is negligible against growth. The admissible cooling rate is proportional to μ_2, which is the crystal surface, and the crystal surface is whatever the cooling has produced so far. That proportionality is the central result of this section. At the instant of seeding, $\mu_2 = 6.82 \times 10^{-3}\,\mathrm{m^2/kg}$, $\sigma = 0.062$ and $G = 1.02\,\mu\mathrm{m/min}$, so Equation (4.12) returns 0.0032 °C/min. The linear profile demands 48 times that value and the natural profile 224 times it, while the cubic profile demands essentially nothing at the start. By the end of the batch the crystal surface has grown from 0.68 to 36.2 $\mathrm{m^2}$, a factor of 53, and Equation (4.12) grows with it.

A batch crystallization therefore matches a cooling demand to a solute-removal capacity that grows during the run. The general solution is to ramp the demand as the capacity grows rather than linearly in time, which is also the structure of the solution for an autocatalytic reaction or a fed-batch fermentation, and it is the result obtained in closed form for programmed cooling [12, 13].

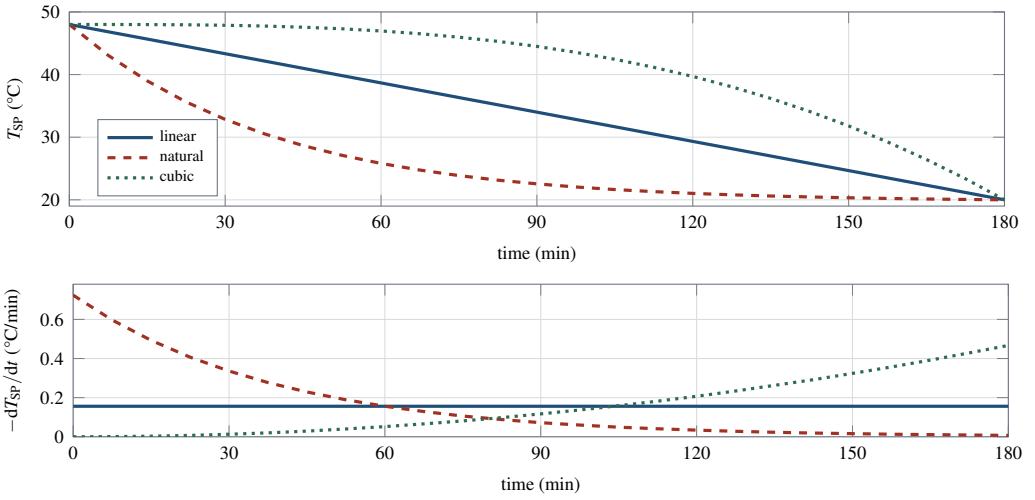


Figure 4.2 The three programmed cooling profiles of Equation (4.11) (linear solid, natural dashed, cubic dotted) and the cooling rates they demand, for the base recipe of 48 °C to 20 °C in 180 minutes. All three cover the same span in the same time. The seed can absorb 0.0032 °C/min at $t = 0$ without any rise in supersaturation, which the natural profile exceeds by a factor of 224 and the linear profile by a factor of 48.

4.5 Controller Design

4.5.1 Classical Architectures on One Valve

The three classical architectures drive the same three-way valve. The single loop is a temperature controller straight onto the valve. The cascade adds an inner jacket temperature loop, so the outer controller's output is a jacket setpoint rather than a valve position. Supersaturation control replaces the programmed trajectory entirely, so the outer controller reads $\hat{\sigma}$ and produces a temperature setpoint that the same temperature controller then follows. Every controller is the velocity-form algorithm of Section 1.6.2, written for this chapter as Equations (4.14) and (4.15).

$$e_k = s\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad s = \pm 1 \tag{4.13}$$

$$\begin{aligned} \Delta u_k^P &= K_c\,(e_k - e_{k-1}) \\ \Delta u_k^I &= \frac{K_c}{\tau_I}\, e_k\, \Delta t_c \\ \Delta u_k^D &= -\,K_c \tau_D\, s\, \frac{\mathrm{PV}_k^f - 2\mathrm{PV}_{k-1}^f + \mathrm{PV}_{k-2}^f}{\Delta t_c} \end{aligned} \tag{4.14}$$

$$\begin{aligned} w_k &= w_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= w_k + I_{k-1} \\ u_k &= \mathrm{slew}(\mathrm{clamp}(c_k, u_{\min}, u_{\max})) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{4.15}$$

where s is the action, PV^f is the measurement passed through a first-order derivative filter of time constant τ_f stepped the same way as the instruments, and T_t is the tracking time. In the notation of Section 1.6.2, s is the action a, the term w_k is the running sum S_k of proportional and derivative increments, and c_k is the output before limiting v_k, with no output bias. The integral used at execution

k was accumulated through execution $k-1$. The demand is clipped to the output limits and then slew limited to $r_{\max}\Delta t_c$ per execution, and only the supersaturation controller uses the slew limit. Anti-windup is back-calculation on the integral alone, and w_k is never truncated. The application sets $T_t = \tau_I$ on all three controllers, so at $\tau_I = 5\,\text{min}$ each execution of the single loop removes 1/250 of the clipped excess, and a loop held against a limit settles with I at the applied output rather than growing with the error. The tracking time affects the reported results where a loop runs against a limit. A recursion that absorbed the whole excess at every execution would give 366 µm at 35 % fines rather than 348 µm at 59 % on the over-tuned row of Table 4.5, and would finish the batch clamped at 1.0 °C in Table 4.8 in 256 rather than 259 minutes. An integral time at the top of its range, 60 min, is read as no reset at all. The integral gain is then zero and the back-calculation term is not applied, because there is no integral to unwind and a proportional-only controller would acquire a permanent bias.

Controller action is the setting that a change of architecture alters. The process gain from valve position to batch temperature is negative, so the single loop and the inner jacket loop both carry $s = -1$. The cascade outer loop is the exception, because its output is a jacket setpoint whose process gain onto batch temperature is positive, and a batch above setpoint needs a lower jacket setpoint, so that loop carries $s = +1$. Changing architecture at run time therefore changes what the outer controller's output means as well as the sign of its error, and carrying the previous integral across would step the valve. An architecture change is therefore treated as a transfer. For one execution every controller runs a tracking step that sets I in Equation (4.15) so that the demanded output equals the valve position already on the plant, the same tracking step that runs at every execution in Manual (Section 1.6.2). Because the integral used at the first automatic execution is the tracked one, the transfer itself moves nothing, and that execution changes the valve by the proportional increment alone. Over the three control executions the bumpless check samples, the valve moves 3.6 %, of which 2.9 % is proportional response and 0.7 % comes from two integral increments, because the measured batch temperature is falling at 4.2 °C/min during the transfer.

4.5.2 Wall Constraint as an Output Limit

The cascade serves a constraint as well as disturbance rejection. By Equation (4.9) the wall is colder than the bulk by $f_w(T - T_j)$, so the only variable that constrains the wall is the jacket temperature. A single loop commands a valve position, which offers no variable on which to state that limit. The cascade outer loop commands a jacket temperature, and its lower output limit is set every execution from the measurement, as shown in Equation (4.16).

$$T_{j,\min} = \hat{T} - \frac{\Delta T_{w,\max}}{f_w} \tag{4.16}$$

The limit therefore moves down the batch with the measured temperature, and anti-windup handles the saturation, because a constraint enforced through the controller's own output limit is seen by the back-calculation term of Equation (4.15), which holds the integral near the floor while the floor is active. Choosing an architecture partly determines which limits can be stated.

4.5.3 Supersaturation Controller

The third architecture's output is a temperature setpoint bounded below by T_{end} and above by T_0, and slew limited to a maximum cooling rate $r_{\max}$. Integral action is required. Holding a constant supersaturation means cooling continuously, and a proportional term alone cannot produce a falling output against a constant error, so the integral term generates the cooling ramp. The slew limit is equally necessary, because it prevents a noisy analyzer from demanding a step change in temperature setpoint. The default settings are $K_{c,\sigma} = 80\,^{\circ}\text{C}$ per unit of supersaturation, $\tau_{I,\sigma} = 15\,\text{min}$ and $r_{\max} = 0.80\,^{\circ}\text{C/min}$. The controller starts with the transfer of an architecture change, because the seeded liquor reads $\hat{\sigma} = 0.062$ against a setpoint of 0.13 and an ordinary first execution would add 5.5 °C of proportional demand. The setpoint starts at T_0 and the integral term opens the ramp at 0.37 °C/min.

Table 4.3 Controller Settings as Shipped. The Single Loop Is the Default and Is Used for Every Cooling Curve Comparison in Section 4.6.

Loop	Action	K_c	τ_I (min)	Output limits
TIC-101, single	−1	12	5.0	0–100 %
TIC-101, cascade	+1	5	4.0	$T_{j,\min}$–55 °C
TJC-101, inner	−1	5	0.3	0–100 %
SSC-101	−1	80	15	T_{end}–T_0

Table 4.3 lists the shipped settings. Derivative time is zero on every loop by default, because a lagged temperature probe on a plant whose setpoint is a smooth ramp offers no benefit from anticipation while amplifying measurement noise.

4.5.4 Predictive Jacket Supervision

The additional controller optimizes a jacket-temperature request rather than replacing the fast jacket loop. Section 1.6.3 describes the predictive option of the apps in general. Receding-horizon constraints are an established industrial control architecture [20]. This implementation uses three measured states $x = (T, T_j, C)$ and the declared clean-vessel constants, with the reduced dynamics

$$\begin{aligned} v &= a[\max(C/C^*(T) - 1, 0)]^g, \\ \dot{T} &= [-UA(T - T_j) + Hv + q_a(T)]/(Mc_p) + b_T, \\ \dot{T}_j &= (u - T_j)/0.3, \qquad \dot{C} = -v. \end{aligned} \tag{4.17}$$

Here u is the jacket request in degrees Celsius, a is the estimated solute-uptake coefficient, $g = 1.1$, b_T is a temperature-rate correction, $H = \Delta H_c m_s$, and $q_a(T) = UA_{\text{amb}}(T_{\text{amb}} - T) + q_{\text{stir}}$. The remaining symbols retain their energy-balance definitions. The initial $a = 3\rho_c k_v \mu_{2,0} k_g$ uses the declared seed surface. Measured concentration change updates a with an 8 minute filter when $\sigma > 0.03$ and uptake exceeds 10^{-8}, within $10^{-7} \le a \le 0.3$. The preceding temperature prediction updates b_T with gain 0.2 and a ± 2 degree-per-minute cap. Neither population moments nor actual fouling, kinetics or utility disturbances are supplied to this estimator. Instrument lags are not inverted. Bounded bias adaptation here does not establish the steady-state conditions needed for offset-free control [21].

Central differences linearize Equation (4.17) at each plan, followed by exact zero-order-hold discretization. Independent analytic Jacobians at 22, 35 and 48 degrees Celsius agree within 4.9×10^{-11} in absolute rate-derivative units. This is a derivative check rather than a plant-prediction residual. Isolated one-degree jacket steps fit a 0.17 minute lag with 0.0155 degree RMS error, while the retained nominal 0.3 minute model gives 0.0949 degree RMS. That model was frozen before the final comparisons.

A plan is computed every 0.5 minutes, or 25 controller scans. Its length in steps of 0.5 minutes is

$$N = \max\big(1,\ \min(120,\ \lceil (t_b - t)/0.5 \rceil)\big),$$

initially at most 60 minutes and shrinking near the recipe endpoint. Twelve graduated input blocks at $N = 120$ reduce the decision count. The objective is

$$\begin{aligned} \min \sum_{k=1}^{N} &\{0.03(T_k - T_e)^2 + 10(\sigma_k - \sigma_r)^2\} \\ &+ 0.05 \sum_j (\Delta u_j)^2 \\ &+ 0.3(T_N - T_e)^2 + 10(\sigma_N - \sigma_r)^2 \\ &+ 10^7 (s_b^2 + s_w^2) + 10^4 (s_b + s_w) \\ &+ 10^4 s_d^2 + 100 s_d. \end{aligned} \tag{4.18}$$

Table 4.4 Three Cooling Curves on One Controller, 48 °C to 20 °C With 25 g of Seed. Numbers Are the Loop and Product Measures at the Moment the Batch Stops.

	Linear	Natural	Cubic
How the loop performed			
Peak deviation (°C)	**0.142**	0.637	0.237
IAE (°C min)	13.43	**13.06**	15.07
Valve travel (% of span)	54	52	66
How the batch performed			
Mean size L_{43} (µm)	159.2	86.3	**373.4**
Coefficient of variation	0.293	0.310	0.308
Fines by number (%)	96.3	99.3	**13.8**
Yield (%)	62.1	62.5	62.2
Peak σ	0.213	0.353	0.150
Time outside the zone (min)	19.1	14.9	**0.0**
Batch time (min)	178.6	152.1	186.8

The final recipe temperature is T_e, and σ_r is the lesser of the operator's request and the bulk limit. Jacket requests remain within the cascade's wall-derived lower bound and hot-side upper bound, with at most three degrees per block-boundary move. Bulk and wall supersaturation limits use the nominal metastable boundary less 0.025, and the third soft limit is the requested wall subcooling. Each $s \geq 0$ is shared across its band's horizon. Existing excess receives an allowance decaying with a 6 minute constant. A fixed 10^{-8} Hessian ridge avoids an unintended absolute-input penalty from scaling the large slack weights. The wall clamp remains active above the optimizer. An invalid status, nonfinite prediction or cost, malformed plan or out-of-box move rejects the plan and restores the tracked cascade.

4.6 Results

The fourteen scenarios are ordered by the concept each one introduces: the batch cooled by hand with the valve wide open (scenario 1), the three cooling curves (2 to 4), cascade and supersaturation control (5 and 6), the recipe faults (7 to 9), the plant and instrument faults (10 to 12), the predictive supervisor (13) and the graded challenge (14). Loading a scenario sets the operating mode, the architecture, the cooling curve, the tunings, the parameter and disturbance overrides and the timed events it requires, and Section 1.6.4 describes how a case is credited. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. Reset restarts a batch from $t = 0$ with the learner's settings kept, which makes most comparisons one change and a second run, while reloading the scenario restores its own settings.

4.6.1 Cooling Curve Comparison

Table 4.4 reports both the loop and the product measures for the three cooling curves, all on the single temperature loop at $K_c = 12$ and $\tau_I = 5$ min, between the same endpoints in the same nominal 180 minutes. Integrated error accumulates over the whole batch in each case, so the observation windows are 178.6, 152.1 and 186.8 minutes and the three IAE figures are not strictly comparable. Peak deviation and valve travel do not have that problem. Peak deviations in this section are taken at the 0.002 min integration step, where the five-per-minute log reads up to 4 % lower. Integrated error is the application's own figure from that log, and valve travel is accumulated at every integration step, as the exported record carries it.

The two halves of Table 4.4 give opposite rankings. The linear batch has the smallest peak deviation in the table, 0.142 °C, and it finishes at 159.2 µm with 96.3 % fines. The natural batch has the lowest integrated error of the three, 13.06 °C min, and the worst product at 86.3 µm and 99.3 % fines. The cubic batch tracks worse than the linear one on both measures and travels the valve 22 %

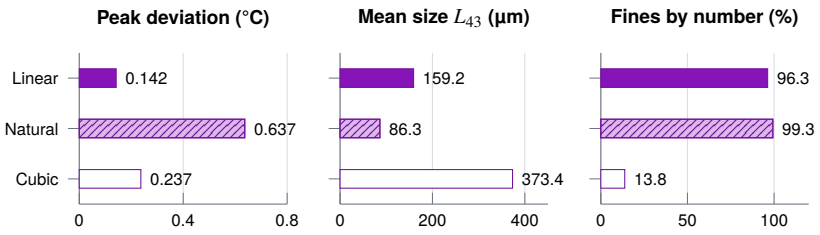


Figure 4.3 The three cooling curves of Table 4.4 on the same loop, between the same endpoints in the same nominal 180 minutes: the loop's peak deviation, the mass-weighted mean size and the fines by number. Notice that the rankings reverse: the linear curve is followed most closely, and the cubic curve, which the loop follows less closely, gives crystals two and a third times as large with a seventh of the fines.

Table 4.5 Tuning Sweep on the Cubic Curve, Scenario 4. Integrated Error Is Accumulated Over Each Batch, Whose Length Is the Last Column.

K_c	τ_I (min)	Peak dev. (°C)	IAE (°C min)	Travel (%)	L_{43} (µm)	Batch (min)
2	20	6.23	503.96	52	373.9	276.5
4	10	1.90	119.43	59	372.6	191.7
12	5	0.237	15.07	66	373.4	186.8
30	2	0.275	**5.23**	80	373.3	186.1
50	1	0.258	8.24	138	373.2	186.0
60	0.5	1.61	105.12	**9114**	348.3	185.3

further, and its crystals are two and a third times larger with a seventh of the fines. Yield differs across the three by four tenths of a point, because yield is set by the final temperature and by almost nothing else.

Figure 4.3 draws the two rankings side by side.

> **TRY IT IN THE APP** **4 · Cubic Cooling, Matched to the Growing Surface**
>
> Load the scenario and press Run. Watch the operating point on the phase diagram travel down between the solubility curve and the metastable limit without touching either line. The batch ends at 186.8 min with 373.4 µm crystals and 13.8 % fines, two and a third times the size of the linear batch, which tracked its curve better.

The cubic curve tracks worse for a structural reason. Its setpoint accelerates, from zero to 0.467 °C/min, and a controller with reset always lags a ramp whose slope is rising, while a linear setpoint has a constant slope that a proportional-plus-integral controller matches with a constant integral term and no residual error. The curve that is easiest to follow is the one that damages the product.

Pause. Try Exercise 4.1 before reading on.

Tuning does not change the result. Table 4.5 sweeps the temperature loop over a fifteen to one range in gain on the cubic curve.

Between a gain of 2 and a gain of 30 the integrated error falls by a factor of 96 and the mean size moves by 0.7 µm. Loose tuning costs batch time, because at a gain of 2 the loop never reaches its own curve and the batch runs 276.5 minutes instead of 186.8. Vessel occupancy, not product quality, is the argument for tuning this loop well.

Table 4.6 Peak Supersaturation Orders the Product and Time Outside the Zone Does Not. All Five Batches Use the Same Kinetics and the Same Vessel.

Batch	Peak σ	Outside (min)	L_{43} (µm)	Fines (%)
Cubic, 25 g seed	0.150	0.0	373.4	13.8
Linear	0.213	19.1	159.2	96.3
Cubic, 6 g seed	0.222	21.5	163.4	99.3
Natural	0.353	14.9	86.3	99.3
Manual, valve wide open	0.706	5.5	42.5	99.9

Table 4.7 Cooling Time on the Cubic Curve, Single Loop. The 150 Minute Batch Never Leaves the Metastable Zone.

t_b (min)	L_{43} (µm)	CV	Fines (%)	Peak σ	Outside (min)
75	121.4	0.376	99.0	0.303	18.8
120	190.3	0.411	97.1	0.226	19.4
150	342.1	**0.550**	**68.2**	0.188	**0.0**
180	373.4	0.308	13.8	0.150	0.0
240	373.8	0.269	10.6	0.104	0.0

The last row reverses the argument, and the damage arrives through an indirect path. At $K_c = 60$ and $\tau_I = 0.5$ min the loop hunts, the valve travels 9114 % of its span, the wall reaches 6.66 °C of subcooling, 77 % of the jacket surface scales over and the product falls to 348.3 µm at 59.1 % fines. The bulk supersaturation leaves the metastable zone for 0.9 min. Over-tuning damages this batch through the coldest surface in the vessel and through the actuator, and not through the variable the controller is regulating.

> **+ DID YOU KNOW? Ninety-Six Times the Error, the Same Crystals**
>
> From the simulator: on the cubic curve, a loop at a gain of 2 accumulates 503.96 °C min of integrated error and one at a gain of 30 only 5.23 °C min, yet their crystals measure 373.9 and 373.3 µm (Table 4.5). The loose loop costs vessel time instead, 276.5 minutes against 186.1.

4.6.2 Ordering of the Outcomes

A candidate explanation for Table 4.4 is time spent above the metastable limit, and the data reject it. Table 4.6 places five batches in order of their peak supersaturation.

Reading the product columns against the time outside the zone gives no ordering. The under-seeded batch spends 21.5 min outside the zone and finishes at 163.4 µm, while the batch cooled manually with the valve wide open spends 5.5 min outside it and finishes at 42.5 µm, four times smaller on a quarter of the exposure. Reading them against the peak supersaturation orders the fines across all five batches and the mean size in all but one pair, the linear and the under-seeded batches, whose peaks differ by 0.009 and whose mean sizes differ by 4 µm on different seed charges. The mechanism is in Equation (4.4). Primary nucleation goes as the fifth power of supersaturation, so a peak of 0.706 against a peak of 0.213 is a rate ratio of $(0.706/0.213)^5 = 400$, and a short excursion to a high peak makes far more crystals than a long excursion to a low one. Cumulative time above the limit is therefore the wrong statistic. This is the general form of any constraint violation whose penalty is a high-order power of the overshoot.

🕮 **Pause.** Try Exercise 4.3 before reading on.

A second correction runs in the opposite direction. Staying inside the zone is necessary but not sufficient. Table 4.7 shortens the batch time on the cubic curve.

Table 4.8 The Same Cubic Recipe in a Vessel With $UA = 18\,\mathrm{kJ/(min\,°C)}$, Scenario 12. The Cascade Rows Keep the Scenario's Outer-Loop Tuning, $K_c = 12$ and $\tau_I = 5\,\mathrm{min}$. The Last Column Is the Largest Departure From the Programmed Curve.

Architecture	Peak ΔT_w (°C)	Scaled (%)	L_{43} (µm)	Batch (min)	Worst error (°C)
Single loop	4.35	**31.2**	375.5	189.0	1.01
Cascade, 1.5 °C clamp	1.48	2.7	373.8	220.5	7.02
Cascade, 1.0 °C clamp	0.99	**0.0**	373.5	**259.1**	10.37

The 150 minute row carries the result. The bulk supersaturation never crosses the limit, the mean size falls only 8 % from its best value, and the batch still finishes at 68.2 % fines with a coefficient of variation of 0.550, the worst in the application. Those fines are entirely secondary nucleation, catalyzed by the seed surface and scaling as σ^2 by Equation (4.4), which runs everywhere inside the zone as well as outside it. Between 180 and 150 minutes the fines rise fivefold, and a further 30 minute reduction removes most of what remains of the product. The last 17 % of the batch time secures almost all of the product quality, a trade that is invisible wherever batch time is scheduled without product data of this kind.

> ✦ **DID YOU KNOW? Ten Minutes on the Clock**
>
> The 2004 guidance on process analytical technology observed that most pharmaceutical processes run to time-defined end points, such as blending for 10 minutes, which ignore physical differences in the raw materials [22]. A batch cooled on a fixed clock is one: the 150 minute row of Table 4.7 never leaves the metastable zone and still finishes at 68.2 % fines.

Scenario 9 separates two effects that are easy to confuse and that the seeding literature keeps apart for the same reason [14, 15]. The recipe seeds at 48 °C, two degrees below saturation, where σ is 0.062 and the seed grows from the first minute. At 50 °C the liquor is exactly saturated, the seed survives intact, and the batch still collapses to 181.3 µm at 96.4 % fines, because the curve now has 30 °C to cover instead of 28 and nothing is banked at the start. At 52 °C the liquor is undersaturated at $\sigma = -0.057$, the seed dissolves at 11 µm/min and is gone inside five minutes, and the batch crosses the limit at 111.6 min, an hour and a half after the seeding error, finishing at 113.7 µm with 98.3 % fines. Neither case produces a visible symptom at seeding time, because the vessel is clear and no instrument on the plant reports that the seed has dissolved.

4.6.3 Encrustation and the Wall Subcooling Clamp

Scenario 5 warms the chilled utility from 5 to 12 °C at $t = 130\,\mathrm{min}$. The single loop takes 0.552 °C of peak deviation and 14.2 °C min of integrated error over the remainder of the batch, and the cascade takes 0.245 °C and 9.3 °C min, because the inner loop sees the disturbance in its own measurement within seconds instead of waiting for a 114 kg charge with a 4.81 min time constant to change temperature. The product is identical either way at 373 µm and 13.8 % fines, which is the honest report. On that batch the disturbance never reached the crystals. The tank of Volume 3, Chapter 2 shows the general rule behind the faster rejection: a loop rejects the disturbances that enter its own measurement path.

Scenario 12 is the case the cascade clamp addresses. The recipe that works in the stainless development vessel is moved into a glass-lined train where $UA = 18\,\mathrm{kJ/(min\,°C)}$ instead of 90, which is five times as much temperature difference for the same duty, and nothing else is changed. Every product criterion passes while the vessel encrusts, as shown in Table 4.8.

On the single loop the bulk never comes closer than 0.04 to its limit and the product is 375.5 µm at 13.9 % fines, while the wall crosses its own limit at 116.5 min, stays over for 65 min, and 31.2 %

of the heat transfer surface is scaled over at the end of the batch. A technology transfer protocol that verifies the trajectory rather than the heat transfer coefficient verifies the part that was never at risk. The last column of Table 4.8 distinguishes an active constraint from a tuning fault. At a 1.0 °C clamp the loop falls 10.37 °C behind its own curve because the constraint is active, and the clean jacket costs 70 minutes of additional vessel time.

> **TRY IT IN THE APP** **12 · A Glass-Lined Vessel, and the Wall That Scales**
>
> Load the scenario and press Run on the single loop. Watch the wall subcooling strip, not the bulk. The bulk never comes closer than 0.04 to its limit and the crystals reach 375.5 µm, yet the wall crosses its own limit at 116.5 min and 31.2 % of the surface scales over. Then run the cascade with a 1.0 °C clamp: no scale, and 259.1 min.

4.6.4 Supersaturation Control

Supersaturation control gives 373.6 µm at 10.9 % fines in 193.3 minutes against 373.4 and 13.8 % in 186.8 for the cubic curve, so on the nominal batch the improvement is small but real and costs 6.5 minutes. The distinguishing feature is the setpoint trace rather than the product numbers, because the cooling curve is generated from a measurement instead of from an assumed kinetic model [16, 17]. The supersaturation setpoint sets the trade between batch time and fines, and the trade is monotone until the batch runs against the metastable limit. At 0.09 the batch takes 279.6 minutes for 8.1 % fines, at 0.13 it takes 193.3 for 10.9, and at 0.17 it takes 149.2 min for 13.5. A setpoint of 0.20 lies above the limit of 0.175 at the warm end, but the reset term cools too slowly to lift σ above 0.173, and the batch finishes in 126.7 minutes at 28.5 % fines. At 0.21 the bulk comes within 0.004 of the limit without crossing it and finishes at 69.9 % fines, because the nucleation gate of Equation (4.4) is 0.003 wide. The throughput optimum sits just below a limit whose position moves with temperature, which is why industrial practice holds a margin [19].

The argument for the third architecture appears when the plant differs from what the curve assumed. With the seed charge cut from 25 to 6 g, the fixed cubic curve gives 163.4 µm at 99.3 % fines and supersaturation control gives 587 µm at 47 % fines, taking 293 minutes to do it, from scenario 8 with the architecture changed or from scenario 6 with the seed changed. The controller has no measurement of the seed charge. It responds only to the supersaturation rising faster than the setpoint allows.

The cost of the third architecture is a dependency on the analyzer, and scenario 11 fails that instrument four ways against a cascade running the cubic curve, which never reads the analyzer and finishes in 187 minutes at 74 % valve travel. Each failure has its own signature. Noise of 0.5 g/kg raises valve travel from 66 to 935 % with the product unchanged at 373.6 µm, which is actuator wear attributable to the control structure. A bias of +3 g/kg makes the supersaturation read high, so the controller acts on driving force it does not have, cools too slowly, and overruns from 193.3 to 262.5 min at slightly lower valve travel, 60 against 66 %. Both together give 263.6 minutes and 1300 % of travel. An analyzer frozen at 100 min is worse than a biased one, because the reading holds at a value the controller has already reached, the error goes to zero, the cooling stops, and at the 300 minute cutoff the yield is 18.2 % instead of 62.2 with the temperature exactly on the setpoint the failed loop requested. The column of Volume 5, Chapter 3 measures what the lag of a composition analyzer alone takes from a loop.

4.6.5 Utility Temperature Limit

Scenario 10 takes the batch to 15 °C rather than 20 °C, which is worth 5.5 points of yield, 67.7 against 62.2 %, and then warms the chilled utility to 18 °C at $t = 110$ min. By Equation (4.8) the jacket inlet is bounded below by the utility temperature at any valve position, so a 15 °C batch is unreachable once the utility reaches 18 °C. The valve is pinned wide open for 133 min, the batch falls 3.9 °C behind its curve, and at the 300 minute cutoff the yield is 65.0 %. The crystals are unharmed at 378.8 µm and 14.6 % fines, because falling *behind* a cooling curve is always safe for the product. The

Table 4.9 Measured Complete-Batch Comparisons

Case	Law	Time min	L_{43} μm	Fines %	Wall excess min
Nominal	Cascade	187.0	373.3	13.8	0.0
	Supersat.	193.3	373.6	10.9	0.0
	MPC	126.9	372.5	15.6	0.0
Encrusted	Cascade	222.0	374.2	11.6	16.7
	Supersat.	195.0	375.0	11.0	0.0
	MPC	181.1	368.9	16.7	0.0
Analyzer fault	Cascade	187.0	373.3	13.8	0.0
	Supersat.	263.1	373.9	8.4	0.0
	MPC	152.6	372.5	13.2	0.0
Warm utility	Cascade	300[a]	378.8	14.6	0.0
	MPC	300[a]	378.4	15.7	0.0
Challenge	Cascade	172.4	373.5	15.1	0.0
	MPC	134.6	372.4	15.3	0.0

[a]Maximum clock, endpoint not reached. Wall excess is time above the plant's wall metastable limit, evaluated after the run.

same event on a batch taken only to 20 °C saturates the valve for 6.5 min, falls 0.56 °C behind and finishes on time at 62.2 % yield. The additional yield requires operating the jacket near its capacity, which leaves no margin for a warmer utility. A controller resting on a limit is no longer controlling, and the trend shows nothing abnormal because the batch is only running late.

4.6.6 The Graded Challenge

The graded challenge combines the full sequence of design decisions. It ships a 150 minute linear recipe on a slow single loop ($K_c = 5$, $\tau_I = 12$ min) in a vessel with half the heat transfer against a 10 °C utility, and it passes when the batch finishes within 180 minutes at a mean size of at least 320 µm, under 20 % fines and no more than 1.5 °C of wall subcooling. As shipped it finishes in 154 minutes at 148 µm, 172 µm short of the size target, and 97.1 % fines while holding its straight line to within 1.03 °C the whole way. The scenario's suggested settings, the cubic curve on a cascade with the wall clamped at 1.5 °C, stretched to a 165 minute curve, finish in 172 min at 374 µm, 15.1 % fines, a coefficient of variation of 0.32, 62.3 % yield, no time outside the zone and a peak wall subcooling of 1.44 °C. The margins on the four targets are 8 min, 54 µm, 4.9 points of fines and 0.06 °C. A single loop passes only in a narrow window, the cubic curve over 165 minutes at the shipped slow tuning, where the lag of the loop holds the wall at 1.48 °C and the batch finishes at 178 min. The same curve at $K_c = 12$ and $\tau_I = 5$ min takes the wall to 1.57 °C and fails. The order of importance is the trajectory, which decides the product, then the architecture, which decides which constraints can be stated, and the tuning last.

4.6.7 Predictive Control Comparison

Table 4.9 uses the same seed, instrument faults and endpoint rule for each structure. Cascade uses $K_c = 5$, $\tau_I = 4$ minutes and jacket $K_c = 5$, $\tau_I = 0.3$ minutes. Supersaturation control uses its shipped $\sigma_{SP} = 0.13$, $K_c^S = 80$, $\tau_I^S = 15$ minutes and 0.8 degree-per-minute rate limit. The MPC settings remain those of Section 4.5.4 in every case. Suggested challenge settings are retained before each structure change. The warm-utility rows end at the 300 minute maximum clock without reaching the required 15 degree endpoint. They are not completed batches.

The nominal gain costs valve movement and fines. Travel rises from 66.4 to 277.6% while fines rise from 13.8 to 15.6%, despite zero bulk or wall metastable-limit time under either law. Figure 4.4 shows the distinct cooling paths and cumulative movement, the cascade as a solid line and the MPC as a dashed one. On the encrusted vessel MPC eliminates the cascade's 16.7 minutes above the wall

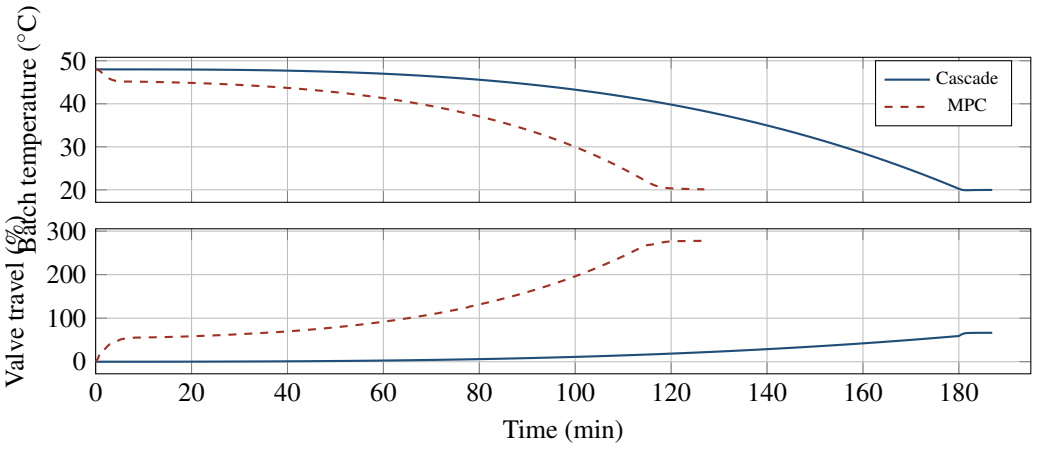


Figure 4.4 Nominal runs continued to their respective endpoints, cascade solid and MPC dashed. The 126.9 minute predictive batch moves the valve 277.6%, compared with 66.4% over the cascade's 187.0 minutes. Product quality is evaluated at those endpoints.

boundary but raises fines from 11.6 to 16.7%. The reduced model carries neither the population objective nor actual fouling, so its band constraints do not certify quality.

An eight-setting hindsight cascade grid varies recipe time over 120, 140, 160 and 180 minutes and wall subcooling over 1 and 1.5 degrees. Its fastest row finishes in 122.4 minutes with 97.1% fines. At a 160 minute recipe it finishes in 167.5 minutes with 16.7% fines. These rows delimit this grid, not the best achievable classical policy. Halving the integration step preserves the nominal ordering. Four independently stepped nominal and encrusted rows reproduce time, size and fines exactly. Across the 33-row comparison, no MPC solve is rejected. Chrome averages 0.685 milliseconds per plan over 360 solves, 0.137% of a core at unit playback and 8.22% at the maximum 60-fold playback.

4.7 Discussion

The phase diagram of Figure 4.5 is the display on which a crystallizer is read. A temperature trend shows temperature falling and a concentration trend shows concentration falling, and neither indicates whether the batch is near its limit, while the position of the operating point between the solubility curve and the metastable limit reports the margin directly. A second operating point at the same concentration and the wall temperature shows the wall crossing the limit on encrusting batches while the bulk remains inside the zone.

Crystal size distribution is not an internal number. It sets the cake resistance on the filter, the solvent retained through the dryer, the flow and compaction of the powder in a tablet press and the dissolution rate that governs bioavailability, which is why it is a registered attribute and why the crystallizer is the unit that fixes the solid form [2]. A batch making 96.3 % fines by number becomes a filtration problem, a drying problem and a formulation problem at once, hours after the crystallizer released it.

A compliant batch record does not carry that attribute. A record captures the setpoint, the measurement, the operator actions, the alarms and the deviations, every one of which is a signal the control system already had. Against a half-degree deviation limit the linear batch of Table 4.4 passes cleanly on 0.142 °C of peak deviation, is released by automated review, and produces a powder with 96.3 % fines. Nothing in the record is false and nothing in it is useful. An in-process attribute that responds to the failure has to come from somewhere else, and without a particle instrument the only candidate is concentration, which is the argument for measuring supersaturation even where it is not controlled [3, 19]. The application makes the same point about the temperature probe, where a +1 °C

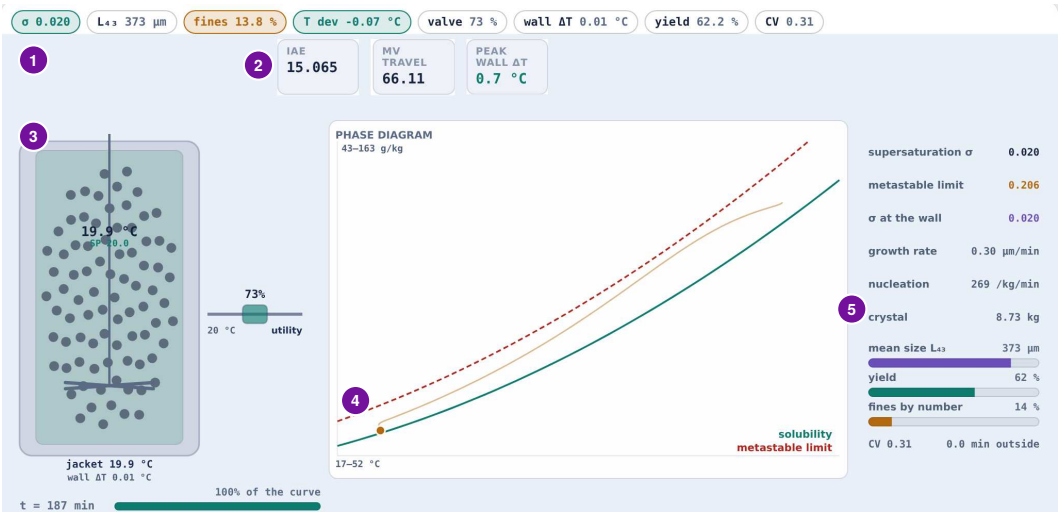


Figure 4.5 Scenario 4 at the end of batch, 186.8 min after the seed was charged on the cubic cooling curve (Table 4.4), cropped to the plant view. The picture is a composite: three tiles of the app's performance panel are pasted into the empty band above the phase diagram. (1) The badges: a mean size of 373 µm, 13.8 % fines, a yield of 62.2 % and a variation coefficient of 0.31. (2) The pasted tiles, the loop and the batch side by side: 15.065 °C min of integrated error, 66.11 % of valve travel and a peak wall subcooling of 0.7 °C. (3) The vessel, its jacket and the utility valve. (4) The phase diagram shows the solubility curve, the metastable limit 5.5 °C to its left, and the path the batch took between them from 140 g/kg at 48 °C down to the endpoint at the lower left. The path never touches either line. (5) The product readout, with zero minutes outside the metastable zone.

bias makes the loop regulate the reading rather than the batch, so the batch ends at 18.9 °C instead of 19.9 and the yield rises 1.2 points with every number in the record on target.

4.8 Technology in Practice

What Measures It

Three readings close the loops of this chapter. TT-101 reads the batch temperature through a 0.3 min probe lag, the jacket detector of the cascade sits in the fast coolant stream and is treated as ideal, and AT-101 reads the liquor concentration, the one reading that sees the driving force. In situ sensors such as infrared spectrometers and laser backscattering probes became commercially available in the 1990s, and the review of the field names them first among the reasons crystallization control grew from then on [19]. Process analytical technology, in the words of the regulator, designs, analyzes and controls manufacturing through timely measurements of critical quality and performance attributes, with the goal of ensuring final product quality [22].

> **TECHNOLOGY** **Sensor: An Infrared Probe That Reads the Liquor**
>
> Ordinary infrared spectroscopy sends light through the medium, and in a slurry the crystals block the beam. An attenuated total reflection probe reflects the beam inside its own crystal, and only an evanescent wave reaches about 1 to 10 micrometers into the liquid, where few crystals intrude. A chemometric calibration relates the absorbances at many wavenumbers, together with the temperature, to the concentration. A background spectrum taken in humid laboratory air drifts, because water absorbs widely in the mid-infrared. Accuracies as high as 0.000 14 g of solute per g of solvent were reported by early 2002 [19]. AT-101 stands for such a probe as a 1 min lag with a settable bias and noise.

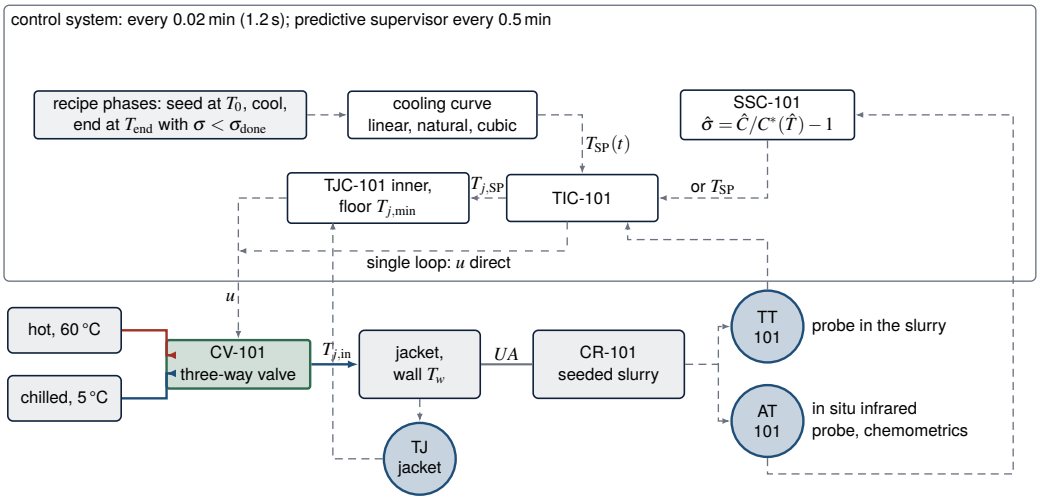


Figure 4.6 The chapter's loops on plant hardware. In the cascade, TIC-101 reads the batch probe TT-101 and sets the jacket setpoint of TJC-101, whose lower limit $T_{j,min}$ keeps the wall subcooling within its clamp, and TJC-101 moves the three-way valve CV-101. The single loop skips the inner loop and sends u to the valve directly (dashed). Supersaturation control replaces the programmed curve with a temperature setpoint computed from the infrared probe AT-101. Notice that the probe that sees the driving force enters only through SSC-101.

What Moves It

One valve moves everything. Where two streams mix, one three-way valve can replace a pair of complementary split-ranged valves [23], and CV-101 blends the hot and chilled utilities into the jacket inlet. The chilled utility's temperature is a floor that no valve position passes, which is the utility limit of scenario 10, and the valve's travel is a cost the chapter counts: analyzer noise of 0.5 g/kg raised it from 66 to 935 % of span with the product unchanged.

> **TECHNOLOGY** **Final Element: A Three-Way Mixing Valve**
>
> A three-way globe valve has three ports, and its plug sets how far two of them connect with the third, so it can split one stream or mix two into one. Mixing hot and cold water to control a temperature is a textbook use of it [23]. A pneumatic valve with excess stem friction moves in jerky steps and ignores small reversals of its signal, which is called stiction, and a loop that strokes a valve back and forth wears out its stem packing early [23]. CV-101 has neither friction nor travel time: 0 % sends hot utility and 100 % chilled. Its travel is counted, and the over-tuned row of Table 4.5 strokes it 9114 % of its span.

Where the Controller Runs

Batch operations follow recipes developed in the laboratory, a sequence of operations such as heating, cooling, reaction and crystallization performed in a pre-specified order [24], and they are controlled from stand-alone controllers, programmable logic controllers, distributed control systems and personal computers [24]. The batch control committee of ISA describes a functional model that relates recipe management, scheduling, sequential control, regulatory control and safety interlock systems [25]. Read in those terms, seeding, cooling and the end-of-batch rule are the sequential control of one recipe, the cooling curve and its time are recipe parameters, and TIC-101, TJC-101 and SSC-101 are the regulatory control. Supervisory and lower-level loops run with cycle times from milliseconds to minutes [26]. The controllers here execute every 0.02 min, which is 1.2 s, and the predictive supervisor plans every 0.5 min, as Figure 4.6 draws it on plant hardware.

IN PRACTICE **Supersaturation Control on Plant Hardware**

In concentration feedback control, the controller computes the supersaturation from the measured concentration, the temperature and the solubility curve, and sends a temperature setpoint to a lower-level temperature controller, typically a cascade of proportional-integral loops. The usual name, supersaturation control, is a misnomer, because a concentration is measured. The recipe becomes a path in the phase diagram, and the temperature profile follows from it on standard tracking control [19]. Regulatory guidance treats the end of a batch the same way: a process end point is not a fixed time but the achievement of the desired material attribute, within an evaluated window of process times [22]. SSC-101 over TIC-101 is that structure.

What Is Modernizing It

Three developments bear on this chapter. Direct nucleation control holds the number of crystals, counted in situ by focused beam reflectance, at a target through alternating phases of supersaturation and dissolution, and it adapts to variations in seed quality and to accidental seeding from crust with almost no prior model [19]. The supersaturation setpoint itself can be chosen by batch-to-batch iterative learning [19], which suits a process where remedial corrections are often impossible once off-specification crystals exist [24], and which Limitations notes this application lacks. Smart valve positioners record total stem travel and can predict when the stem packing will wear out [23], the hardware counterpart of the valve travel this chapter reports beside every batch.

4.9 Limitations

No number in this application is evidence about a real crystallization. The kinetic constants in Table 4.2 were chosen to make the mechanisms observable rather than fitted to a substance, so nothing reported here is a prediction about any material. There is no distribution shape, no agglomeration, no breakage, no growth rate dispersion, one polymorph and no impurity, which between them remove most of what makes an industrial crystallization difficult to develop. The vessel is perfectly mixed apart from a single wall film, so local supersaturation near the coolant surface, the effect of agitation rate on secondary nucleation and the suspension of the largest crystals are all absent. Encrustation is one fraction scaling one coefficient where real scale grows unevenly. AT-101 is one lag, one bias and band-limited noise, where a real in-situ probe reports through a chemometric model whose error depends on temperature, solid loading and the very supersaturation being measured, and it has no drift, no probe fouling and no recalibration. There is no particle instrument at all, so direct nucleation control, which is an industrial architecture beyond those implemented here [18], cannot be demonstrated. The predictive supervisor does not optimize the population moments, and no batch-to-batch learning of the kind in Sections 1.4.5 and 1.4.6 is implemented. The features that survived the simplification are the ones a controller has to be designed around, a setpoint that is a trajectory, a path constraint no instrument reports, a capacity that grows during the run, a constraint that binds first on an unmeasured surface, and an objective evaluated only at the end. Section 4.3.7 records where each assumption fails.

The predictive controller assumes the clean-vessel heat balance and a 0.3 minute closed-jacket response. Its three-state estimate does not reconstruct the instrument lags, crystal population or encrustation. Neither its soft constraints nor its bounded bias imply nonlinear invariance, terminal stability or product-quality guarantees.

OPEN CHALLENGE **Seeing the Crystals at Plant Solids Loadings**

Image analysis classifies crystals by size and shape, through a window or with an in situ video microscope. Most in situ imaging gives quantitative results only at low solids, typically 5 to 7 % at most, and future work aims at hardware and software that stay reliable in denser suspensions [19]. Guidance on process analytical technology names particle size and shape as attributes that are hard to sample representatively [22]. Progress would be a size measurement a loop could close on through the whole batch.

4.10 Conclusion

This chapter has used a seeded batch cooling crystallization as the example of a process that follows a trajectory rather than holding a setpoint and whose product is evaluated only after the run ends. The simulator integrates nine states, the batch and jacket temperatures, the dissolved solute, the first five moments of the crystal population under size-independent growth, and a wall encrustation fraction, with a metastable limit drawn 5.5 °C of subcooling above the solubility curve and primary nucleation gated on it at a fifth-power dependence. Three cooling curves between 48 °C and 20 °C in 180 minutes on the same single loop give 86.3, 159.2 and 373.4 µm of mass-weighted mean size at 99.3, 96.3 and 13.8 % fines, while the same three batches give peak tracking deviations of 0.637, 0.142 and 0.237 °C, so the best-tracked batch is the second-worst product. Sweeping the proportional gain from 2 to 30 moves the integrated error by a factor of 96 and the mean size by 0.7 µm, and pushing it to 60 damages the product through 9114 % of valve travel and 6.66 °C of wall subcooling rather than through the bulk supersaturation. Two measured results contradict the simplest explanation, that peak supersaturation orders the outcomes while cumulative time outside the zone does not, and that a 150 minute batch which never leaves the zone still finishes at 68.2 % fines from secondary nucleation alone. The kinetic constants are chosen to make the mechanisms observable rather than fitted to a substance, so no number reported here is a prediction about any material. The predictive jacket supervisor reduces the nominal batch from 187.0 to 126.9 minutes while increasing fines from 13.8 to 15.6% and valve travel from 66.4 to 277.6%. The app does not offer a particle count measurement, so direct nucleation control cannot be compared with supersaturation control, and it has no nonlinear model predictive control against the metastable zone path constraint, no moving horizon estimation of the growth and nucleation parameters, no solubility that depends on impurity level and no agglomeration. What the crystallizer teaches is the distinction between a controlled variable and a product objective.

KEY POINTS

- A seeded cooling crystallization is graded once, after the batch, on crystals that no instrument in the loop can see.
- Three cooling curves on the same loop give 86.3, 159.2 and 373.4 µm of mean size, and the best-tracked batch is the second-worst product.
- Peak supersaturation orders the outcomes and cumulative time outside the zone does not, and a 150 minute batch that never leaves the zone still finishes at 68.2 % fines.
- The trajectory decides the product, the architecture decides which constraints can be stated, and the tuning comes last.

Where else this applies. The autoclave of Chapter 6 also runs a recipe of temperature ramps whose product is decided inside the load: the recipe is run on the part that lags and limited by the part that leads, as this batch is limited by its coldest surface, the wall.

References

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

[2] N. Variankaval, A. S. Cote, and M. F. Doherty, “From form to function: Crystallization of active pharmaceutical ingredients,” *AIChE Journal*, vol. 54, no. 7, pp. 1682–1688, 2008.

[3] L. X. Yu, R. A. Lionberger, A. S. Raw, R. D’Costa, H. Wu, and A. S. Hussain, “Applications of process analytical technology to crystallization processes,” *Advanced Drug Delivery Reviews*, vol. 56, no. 3, pp. 349–369, 2004.

[4] D. Bonvin, “Optimal operation of batch reactors: a personal view,” *Journal of Process Control*, vol. 8, no. 5–6, pp. 355–368, 1998.

[5] J. B. Rawlings, S. M. Miller, and W. R. Witkowski, “Model identification and control of solution crystallization processes: a review,” *Industrial & Engineering Chemistry Research*, vol. 32, no. 7, pp. 1275–1296, 1993.

[6] A. D. Randolph and M. A. Larson, *Theory of Particulate Processes: Analysis and Techniques of Continuous Crystallization*, 2nd ed. San Diego, CA: Academic Press, 1988.

[7] J. W. Mullin, *Crystallization*, 4th ed. Oxford, UK: Butterworth-Heinemann, 2001.

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

[9] H. A. Miers and F. Isaac, "The refractive indices of crystallising solutions, with especial reference to the passage from the metastable to the labile condition," *Journal of the Chemical Society, Transactions*, vol. 89, pp. 413–454, 1906.

[10] J. Nývlt, R. Rychlý, J. Gottfried, and J. Wurzelová, "Metastable zone-width of some aqueous solutions," *Journal of Crystal Growth*, vol. 6, no. 2, pp. 151–162, 1970.

[11] N. Kubota, "A new interpretation of metastable zone widths measured for unseeded solutions," *Journal of Crystal Growth*, vol. 310, no. 3, pp. 629–634, 2008.

[12] J. W. Mullin and J. Nývlt, "Programmed cooling of batch crystallizers," *Chemical Engineering Science*, vol. 26, no. 3, pp. 369–377, 1971.

[13] A. G. Jones, "Optimal operation of a batch cooling crystallizer," *Chemical Engineering Science*, vol. 29, no. 5, pp. 1075–1087, 1974.

[14] N. Doki, N. Kubota, A. Sato, and M. Yokota, "Effect of cooling mode on product crystal size in seeded batch crystallization of potassium alum," *Chemical Engineering Journal*, vol. 81, no. 1–3, pp. 313–316, 2001.

[15] N. Kubota, N. Doki, M. Yokota, and A. Sato, "Seeding policy in batch cooling crystallization," *Powder Technology*, vol. 121, no. 1, pp. 31–38, 2001.

[16] H. Grön, A. Borissova, and K. J. Roberts, "In-process ATR-FTIR spectroscopy for closed-loop supersaturation control of a batch crystallizer producing monosodium glutamate crystals of defined size," *Industrial & Engineering Chemistry Research*, vol. 42, no. 1, pp. 198–206, 2003.

[17] M. Fujiwara, Z. K. Nagy, J. W. Chew, and R. D. Braatz, "First-principles and direct design approaches for the control of pharmaceutical crystallization," *Journal of Process Control*, vol. 15, no. 5, pp. 493–504, 2005.

[18] A. N. Saleemi, C. D. Rielly, and Z. K. Nagy, "Comparative investigation of supersaturation and automated direct nucleation control of crystal size distributions using ATR-UV/vis spectroscopy and FBRM," *Crystal Growth & Design*, vol. 12, no. 4, pp. 1792–1807, 2012.

[19] Z. K. Nagy and R. D. Braatz, "Advances and new directions in crystallization control," *Annual Review of Chemical and Biomolecular Engineering*, vol. 3, no. 1, pp. 55–75, 2012.

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

[21] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[22] U.S. Food and Drug Administration, "Guidance for industry. PAT: A framework for innovative pharmaceutical development, manufacturing, and quality assurance," U.S. Department of Health and Human Services, Food and Drug Administration, Rockville, MD, Tech. Rep., Sep. 2004.

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

[24] D. Bonvin and G. François, "Control and optimization of batch chemical processes," in *Coulson and Richardson's Chemical Engineering*. Elsevier, 2017, pp. 441–503.

[25] International Society of Automation, "ISA88: Batch control, committee scope and purpose," www.isa.org, standards committees, 2026, accessed 22 November 2026.

[26] 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.

4.11 Exercises

4.1 Concept check · 1 min

Scenarios 2, 3 and 4 run the linear, natural and cubic cooling curves between the same endpoints in the same 180 minutes, on the same single temperature loop with the same tuning. Which statement matches the results?

(a) The linear batch tracks its curve most closely and therefore gives the coarsest crystals

(b) The cubic batch gives the coarsest crystals although it tracks its curve worse than the linear one

(c) The three batches give nearly the same crystals, because they share the endpoints and the yield

(d) The natural batch gives the coarsest crystals, because it cools fastest while the seed is fresh

4.2 Concept check · 1 min

Why can the app carry the whole crystal population with the five ordinary differential equations of Equation (4.5) instead of a partial differential equation in size and time?

(a) Nucleation is negligible inside the zone, so only growth has to be carried

(b) The size distribution is assumed log-normal for the whole batch, like the seed

(c) Growth does not depend on size and nuclei are born at one size, so the moments close

(d) The fines are counted separately by number, so the distribution needs no shape

4.3 Concept check · 2 min

Table 4.6 lists five batches with the same kinetics in the same vessel, from the cubic batch on 25 g of seed to the batch cooled by hand with the valve wide open. Which single statistic orders their fines across all five?

(a) The time the bulk supersaturation spends outside the metastable zone

(b) The integrated absolute error of the batch temperature loop

(c) The total travel of the coolant valve over the batch

(d) The peak value of the bulk relative supersaturation

4.4 Theory and calculation · 6 min

The solubility of Equation (4.1) is anchored at $C^* = 52\,\mathrm{g/kg}$ at 20 °C and 140 g/kg at 50 °C, with the curvature $a_2 = 4.0 \times 10^{-5}\,\mathrm{kg/(kg\,°C^2)}$. The charge holds $C = 0.140\,\mathrm{kg/kg}$ and is seeded at 48 °C. The metastable zone is $\Delta T_{\mathrm{msz}} = 5.5\,°\mathrm{C}$ wide.

(a) Find the coefficients a_1 and a_0. *Checkpoint:* $a_1 = 1.333 \times 10^{-4}\,\mathrm{kg/(kg\,°C)}$ and $a_0 = 3.333 \times 10^{-2}\,\mathrm{kg/kg}$.

(b) Find the solubility and the relative supersaturation at the moment of seeding. *Checkpoint:* $C^* = 0.1319\,\mathrm{kg/kg}$ and $\sigma = 0.0615$.

(c) Find the metastable limit $\sigma_{\max}$ of Equation (4.2) at 48 °C and at 20 °C, and say which end of the batch is the tight one. *Checkpoint:* 0.175 at 48 °C and 0.207 at 20 °C. The warm end is tight.

4.5 Theory and calculation · 7 min

At the instant of seeding the batch is at 48 °C with $\sigma = 0.0615$ and a seed surface $\mu_2 = 6.82 \times 10^{-3}\,\mathrm{m^2/kg}$. Take $k_g = 22\,\mu\mathrm{m/min}$, $g = 1.10$, $\rho_c = 1300\,\mathrm{kg/m^3}$, $k_v = 0.50$ and the solubility coefficients of Table 4.2 and the previous exercise.

(a) Find the growth rate G of Equation (4.3). *Checkpoint:* 1.02 µm/min.

(b) Find $\mathrm{d}C^*/\mathrm{d}T$ at 48 °C and the cooling rate the seed can absorb at constant supersaturation, Equation (4.12). *Checkpoint:* 3.973×10^{-3} kg/(kg °C) and 0.0032 °C/min.

(c) For the base recipe, 48 °C to 20 °C in 180 min, find the initial cooling rate of the linear and the natural profiles of Equation (4.11) and their ratios to the result of (b). Give the initial and final

rates of the cubic profile. *Checkpoint:* 0.156 and 0.724 °C/min, ratios 48 and 224. Cubic 0 and 0.467 °C/min.

4.6 Theory and calculation · 5 min

The charge is $M = 114.0\,\mathrm{kg}$ of slurry with $c_p = 3.80\,\mathrm{kJ/(kg\,°C)}$ behind a jacket of $m_j = 15\,\mathrm{kg}$, $c_j = 3.60\,\mathrm{kJ/(kg\,°C)}$ and $F_j = 40\,\mathrm{kg/min}$. The process-side share of the wall temperature difference is $f_w = 0.30$.

(a) With $UA = 90\,\mathrm{kJ/(min\,°C)}$, find the time constant of the contents and of the jacket, and their ratio. *Checkpoint:* 4.81 min and 0.231 min, about 21 to 1.

(b) Find how far below the measured batch temperature the cascade's jacket floor of Equation (4.16) sits for wall limits of 1.5 and 1.0 °C. *Checkpoint:* 5.0 °C and 3.33 °C.

(c) In the glass-lined vessel of scenario 12, $UA = 18\,\mathrm{kJ/(min\,°C)}$. With the jacket held on the 1.0 °C floor and a clean wall, estimate the heat removed and the cooling rate it buys, ignoring the heat of crystallization and the ambient leak. Compare it with the final rate of the cubic curve. *Checkpoint:* 60 kJ/min and about 0.14 °C/min, against 0.467 °C/min.

4.7 App challenge · 8 min

Load scenario 2, Linear cooling, the curve everybody writes first. Keep Auto, the shipped single loop and the default seed, and set the speed to 20, which runs a 180 min batch in about 9 s.

(a) Run until the batch pauses itself. Read IAE on the scorecard and the Mean size L_{43}, Fines and Batch time tiles. *Checkpoint:* IAE 13.434 °C min, 159 µm, 96.3 %, 179 min.

(b) Load scenario 4, Cubic cooling, matched to the growing surface, keep its settings and the speed of 20, and run until the batch pauses itself. Read the same four values. *Checkpoint:* IAE 15.065 °C min, 373 µm, 13.8 %, 187 min.

(c) Explain in two sentences why the batch with the larger integrated error gives the better product. *Checkpoint:* The product follows the supersaturation the trajectory imposes, not the tracking error.

Give IAE to within 0.01 °C min, sizes to within 1 µm and fines to within 0.1 point.

4.8 App challenge · 9 min

Load scenario 12, A glass-lined vessel, and the wall that scales, which ships $UA = 18\,\mathrm{kJ/(min\,°C)}$ on the single loop. Keep Auto and the default seed, and set the speed to 20.

(a) Run until the batch pauses itself. Read the Peak wall ΔT, Mean size L_{43}, Fines and Batch time tiles, the IAE on the scorecard and the last value of the column `foul_frac` in the exported CSV. *Checkpoint:* 4.3 °C, 375 µm, 13.9 %, 189 min, IAE 30.42 °C min, 0.312.

(b) Press Reset, set Control architecture to Cascade and Wall subcooling limit to 1.0 °C before pressing Run, and run until the batch pauses itself. Read the same values. *Checkpoint:* 1.0 °C, 374 µm, 9.9 %, 259 min, IAE about 625 °C min, 0.

(c) The IAE of the cascade batch is about twenty times that of the single loop. Is the cascade mistuned? *Checkpoint:* No. The clamp is active and the jacket cannot remove heat any faster.

Give temperatures to within 0.1 °C, sizes to within 1 µm, times to within 1 min and the scaled fraction to within 0.005.

4.9 App challenge · 8 min

Load scenario 13, A plan inside the metastable band, which opens on the MPC architecture. Keep Auto and the default seed, and set the speed to 20.

(a) Run until the batch pauses itself. Read the Batch time, Mean size L_{43} and Fines tiles and the last value of the column `valve_travel_pct` in the exported CSV. *Checkpoint:* 127 min, 372 µm, 15.6 %, 277.6 %.

(b) Press Reset, set Control architecture to Cascade before pressing Run, and run until the batch pauses itself. Read the same values. *Checkpoint:* 187 min, 373 µm, 13.8 %, 66.4 %.

(c) The MV travel on the scorecard of run (a) reads about 156, well below the CSV value. Explain the difference, and say which figure to quote for valve wear. *Checkpoint:* The scorecard sums the five logged samples a minute, and the CSV column accumulates every integration step. Quote the CSV.

Give times to within 1 min, fines to within 0.1 point and travel to within 1 %.

4.10 Application area · 3 min

A plant's automated review releases any batch whose temperature stayed within 0.5 °C of its programmed curve. Use the linear batch of Table 4.4 to explain why such a release can be wrong, and name the measurement on CR-101 that could have flagged the batch while it was running.

4.11 Application area · 4 min

A production planner proposes to finish every batch at 15 °C instead of 20 °C for the extra yield. The chilled utility is normally at 5 °C and can reach 18 °C on a summer afternoon.

(a) Estimate the yield at each endpoint from the solubility curve alone, with $C_0 = 0.140\,\mathrm{kg/kg}$ and the coefficients of the first theory exercise. *Checkpoint:* 62.9 % at 20 °C and 68.3 % at 15 °C.

(b) Using scenario 10 as reported in the chapter, say what a warm afternoon does to the 15 °C batch and what the operator sees. *Checkpoint:* The batch cannot finish. The valve sits wide open and the trend only shows the batch running late.

4.12 Review the agent · 5 min

An AI assistant reviews scenario 4 and writes: "The cubic batch has an IAE of 15.07 °C min against 13.43 °C min for the linear batch, so the loop is too slow for the accelerating setpoint. Raise the proportional gain from 12 to 60 and cut the integral time from 5 to 0.5 min. Tighter tracking of the cubic curve will push the mean size above 373 µm." Find the flaw and say how you would check it.