INTERACTIVE CONTROL ENGINEERING
Principles and Applications

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

CHAPTER 2

Semi-Batch Exothermic Reactor Simulator With Dosing as the Safety Layer

Pages 23–54

Edited by John D. Hedengren

How to cite this chapter. "Semi-Batch Exothermic Reactor Simulator With Dosing as the Safety Layer," 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. 2, pp. 23–54.

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

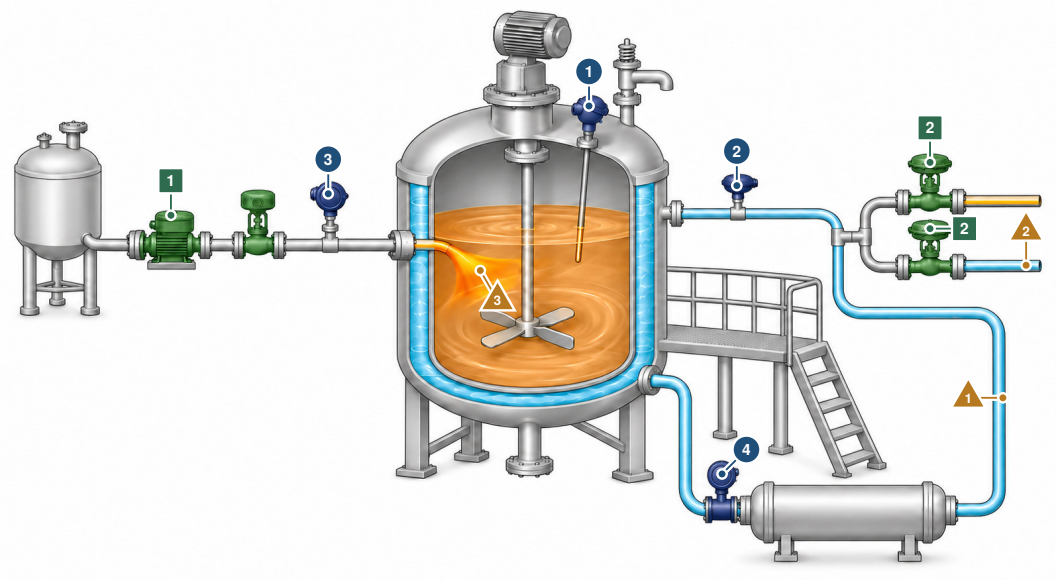
The app for this chapter
https://apmonitor.com/apps/semi-batch-reactor/

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

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

CHAPTER 2

Semi-Batch Exothermic Reactor Simulator With Dosing as the Safety Layer



The controller moves 1 pump that doses the second reagent 2 cooling-water and hot-water valves on the jacket loop

The controller reads 1 batch temperature 2 jacket temperature 3 feed flow meter 4 low-flow switch on the jacket loop

What disturbs it 1 jacket circulation lost 2 warmer cooling water 3 weaker catalyst, so fed reagent reacts later

Notice that the feed enters the stirred liquid while every utility stays outside the wall, so only the dosing pump decides how much unreacted reagent the batch holds. Illustration, not to scale.

IN PLAIN TERMS The Danger Is the Feed Not Yet Used

The main idea. This reactor is filled with one liquid and then fed a second one slowly, because the two give off heat as they react. Feeding slowly lets a cooling jacket around the tank keep up with that heat.

The challenge. If the tank is kept too cool, the fed liquid piles up without reacting while the thermometer looks perfect. Should the cooling then fail, the stored liquid reacts all at once and the tank can overheat until a safety disc bursts.

Why it matters. The feed controller, not the temperature controller, decides how much hidden heat a batch carries at any moment.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on a stirrer left off (Section 2.1), the history at the end of Section 2.2, the two try-its and the fact in Section 2.6, and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the measured cases (Section 2.6, through the challenge), Technology in Practice (Section 2.7), Limitations (Section 2.8) and the open challenge

The control question A batch dosed cold can hold its setpoint within a fraction of a kelvin while the reagent fed into it piles up unreacted. What protects a semi-batch reactor against a cooling failure, the temperature loop or the dosing controller, and how far can that protection trust an estimate?

Learning objectives

- Compute the adiabatic temperature rise of the unreacted feed, the MTSR and the criticality class from the balances.
- Explain why the highest MTSR of a dose has a minimum in the process temperature, and why a colder batch can be the more dangerous one.
- Compare fixed, temperature-limited and accumulation-limited dosing through a cooling failure, a weak catalyst and warm cooling water.
- Predict the drift of a heat-flow calorimeter from the error in its assumed heat-transfer coefficient.
- Judge what a predictive dosing controller gains over a one-step optimizer, and what relinearizing along its plan does not fix.

Prerequisites Sections 1.4.2 (recipes and guards) and 1.4.4 (path constraints and back-off), 1.6.2 (the PI law) and 1.6.1 (engine, speed and seed). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/semi-batch-reactor/

See also Volume 4, Chapter 2 for a protective override that rests at its clamp until a limit nears, Volume 4, Chapter 6 for a safety limit that sets the line speed, Volume 7, Chapter 2 for the ignition boundary of a cooled exothermic reactor, Volume 13, Chapter 4 for catalyst beds run close to an ignition floor, and Volume 13, Chapter 2 for an inventory that accumulates unseen until a purge removes it.

2.1 Introduction

A semi-batch reactor is charged with one reagent and fed the other, and when the reaction is strongly exothermic the feed rate is fixed before anything else, because the dose caps the rate at which heat can be released [1–3]. The recipe of this chapter doses 1200 kg of a reagent B, 8000 mol of it, into 3000 kg of a reagent A in solvent held at 70 °C. At 200 kJ/mol the dose carries 1.6 GJ, enough to heat the finished batch by 211.6 K if it were released at once. Spread over four hours it releases about 111 kW, against a jacket that could remove 194 kW with its cooling water at 8 °C.

The misconception this reactor corrects is that a batch on its temperature setpoint is a safe batch. What endangers it is not the heat being made but the heat not yet made, the B that has been fed and has not reacted [4]. If the cooling fails, that B reacts without heat removal and the batch climbs toward the maximum temperature of the synthesis reaction (MTSR), the quantity on which the thermal assessment of a cooling failure rests [5, 6]. Measured on the model, the recipe at 70 °C never holds more than 234 mol of unreacted B, and its MTSR never passes 76.2 °C. The same four-hour dose into a batch held at 30 °C keeps the true temperature between 30.0 and 30.2 °C. Yet at 2 h 30 min it holds 4398 mol of B unreacted and an MTSR of 160.4 °C, 30 K above the temperature at which the product's decomposition would reach its maximum rate within a day. The colder batch is the more dangerous one, and its temperature trend does not show it.

DID YOU KNOW? A Stirrer Left Off

In a semi-batch plant described by the British safety regulator, an operator filled the reactor with the first reagent and began adding the second without starting the stirrer. When the agitator was started, the reaction was so violent that the reactor burst from overpressure, and 11 people were killed [7].

The problem is one of a general class. A controller that has to limit an inventory it cannot measure, whose release after a failure is the hazard, appears wherever a reactor is fed rather than charged. Semi-batch emulsion polymerization has the same shape, where calorimetry estimates the conversion and feedback on it has kept monomer from accumulating after a deliberate inhibition [8], and so do

fed-batch systems in which online optimization is joined to runaway prediction and prevention [9]. In each case the estimate comes from a balance, and calorimetric state estimators for reactor diagnosis and control form a family of methods of their own [10]. Feed profiles for a second-order reaction under safety constraints have been optimized [11] and computed online from a heat balance [12]. A transfer function from feed to temperature carries none of this, because in it the feed is a load and the accumulation is not a state.

The thermal safety literature presents MTSR and the criticality classes as quantities of a design assessment [6], and the optimization literature presents the feed profile as the solution of a problem [12]. This chapter works instead with a reactor that holds its setpoint while its criticality class moves from 2 to 5 and then loses its cooling. The app computes the accumulation, MTSR and Stoessel's classes live from the true state, logs the calorimeter's estimate beside the truth, and ends a cooling failure at the rupture disc when the accumulation says it will. The kinetics, the criticality classes, heat-flow calorimetry, cascade and split-range control and accumulation-limited dosing are all standard methods. The model has 18 states, and its energy and mole ledgers close to round-off. Its decomposition is anchored on a declared threshold through the adiabatic time to maximum rate. Its calorimeter sees the accumulation only through the heat-transfer coefficient it assumes, and its four shipped dosing strategies act on one constraint. An optional predictive controller on the dose is measured against those strategies on the same runs, and a second setting of it, which relinearizes its model along its own plan, is measured against the single linearization. A sweep of the process temperature finds a minimum of the highest MTSR near 66 °C. The ten cases and the challenge separate the heat being made from the heat stored, the temperature controller from the safety layer, and the estimate from the truth.

The chapter is laid out as follows. Section 2.2 gives the history of accumulation, calorimetry and the cooling-failure scenario. Section 2.3 sets out the model, from the kinetics and the balances to the criticality classes and the events. Section 2.4 measures the highest MTSR against the process temperature. Section 2.5 describes the cascade, the calorimeter, the dosing strategies and the interlocks, with the predictive option in Sections 2.5.8 and 2.5.9. Section 2.6 runs the cases and the challenge and compares the predictive controller with the shipped strategies (Sections 2.6.9 and 2.6.10). Section 2.7 describes the plant's instruments, dosing pump and control system, Section 2.8 states the limitations, Section 2.9 concludes, and the exercises of Section 2.10 run the cases.

2.2 Accumulation, Calorimetry and the Cooling-Failure Scenario

Feeding one reagent into the other became the standard way to run a strongly exothermic liquid-phase reaction because it hands the release of heat to a pump. In 1986 the limits of safe operation of a semi-batch reactor were compared with those of a continuous stirred tank [1], and the conditions for thermally safe operation of a semibatch reactor were mapped for slow liquid-liquid reactions in 1990 [2]. Safety and runaway prevention in batch and semibatch reactors were reviewed together in 2006 [3].

The assessment rests on measurements of the decomposition that a runaway can set off. An accelerating rate calorimeter for thermal hazard evaluation was described in 1980 [13], and the adiabatic time to maximum rate that such an instrument measures defines T_{D24}, the temperature at which that time is 24 h. The same time can be estimated from dynamic differential scanning calorimetry (DSC) [14]. Chemical reaction engineering was turned toward safety in 1988 [5], and the hazards linked to accumulation in semi-batch reactors were assessed in 1996 [4]. Stoessel's textbook sets out the cooling-failure scenario, which orders the process temperature, MTSR, T_{D24} and the maximum temperature for technical reasons (MTT) into five criticality classes, together with its use in process design [6].

The estimate that a dosing controller needs came from calorimetry. Heat-flow calorimetry was developed as a tool for process optimization and process safety at once [15], and calorimetric state estimators, which infer the reaction rate from a reactor's heat balance, were reviewed for diagnosis

and control in 1992 [10]. In 1997 on-line calorimetry estimated the conversion of a semibatch emulsion copolymerization, and feedback on that estimate avoided monomer accumulation, and with it a potential runaway, when an inhibitor was added on purpose [8]. Optimal feed profiles for a second-order reaction in a semi-batch reactor under safety constraints followed in 1999 [11]. By 2001 the feed that maximizes conversion was being computed on-line from a heat balance, bounded by the heat release in normal operation and by the temperature the batch could reach after a cooling failure [12]. Online model-based optimization has since been combined with runaway prediction and prevention in fed-batch systems [9].

The fourth case of the app follows an accident. On December 19, 2007, T2 Laboratories in Jacksonville, Florida, was making its 175th batch of a gasoline additive when a runaway reaction burst the reactor, and the explosion and fire killed four employees and injured 32 people [16]. The U.S. Chemical Safety and Hazard Investigation Board found that a loss of sufficient cooling likely caused the runaway, that the cooling system was susceptible to single-point failures, that the relief system could not relieve a runaway, and that the company had not recognized the hazard. Its recommendations went to AIChE and ABET, to add reactive hazard awareness to the requirements for baccalaureate chemical engineering programs [16]. The Board had found that only 11 % of the universities it surveyed required process safety education in the core of the degree [16].

> ⧖ HISTORY **2008: A Feed Started Before the Vessel Was Hot**
>
> In August 2008 a 4500-gallon pressure vessel exploded in West Virginia during the restart of a pesticide unit, after an outage to upgrade the unit's control system, and two employees died. The vessel was meant to be filled with solvent first and held near 135 °C, the temperature at which the incoming material decomposed fast enough not to accumulate, and interlocks kept its feed valve shut until it was hot and its circulation was running. During the restart the fill and the heat-up were skipped and the interlocks were bypassed. Hours after the feed began the vessel was still at 60 to 65 °C, the material piled up, and the runaway overwhelmed the relief system [17].

2.3 The Reactor as Modeled

Reactor R-101 holds 6 m^3. It is charged with 3000 kg of A in solvent, 12 000 mol of A, and dosed with 1200 kg of feed carrying 6.667 mol/kg of B, so that A stands in 50 % excess over the 8000 mol of B. The feed pump P-103 delivers up to 600 kg/h. The jacket loop holds 150 kg of water, circulated at 2 kg/s by P-102, and a split-range pair of valves replaces part of its flow with cooling water at 8 °C or steam-heated water at 140 °C. The arrangement, with every variable that the equations use, appears in Figure 2.1, and Table 2.1 sorts the variables by role.

The state vector has 18 entries. Six carry the process, the moles of A, of unreacted B and of product, the batch and jacket temperatures, and the split-range signal behind its actuator. Twelve are ledgers integrated by the same step, and two of them, the feed mass m_{fed} and the product decomposed n_X, also enter the volume, the heat capacity and the heat-transfer law. The other ten book every joule released by the synthesis and the decomposition, removed into the jacket, lost to the room, brought in by the feed and added by the utilities (with the cooling-water and hot-water shares apart), and every mole of B fed and reacted. Units are SI inside every rate, and the panels show kg/h, kW, kW/K and MJ/K.

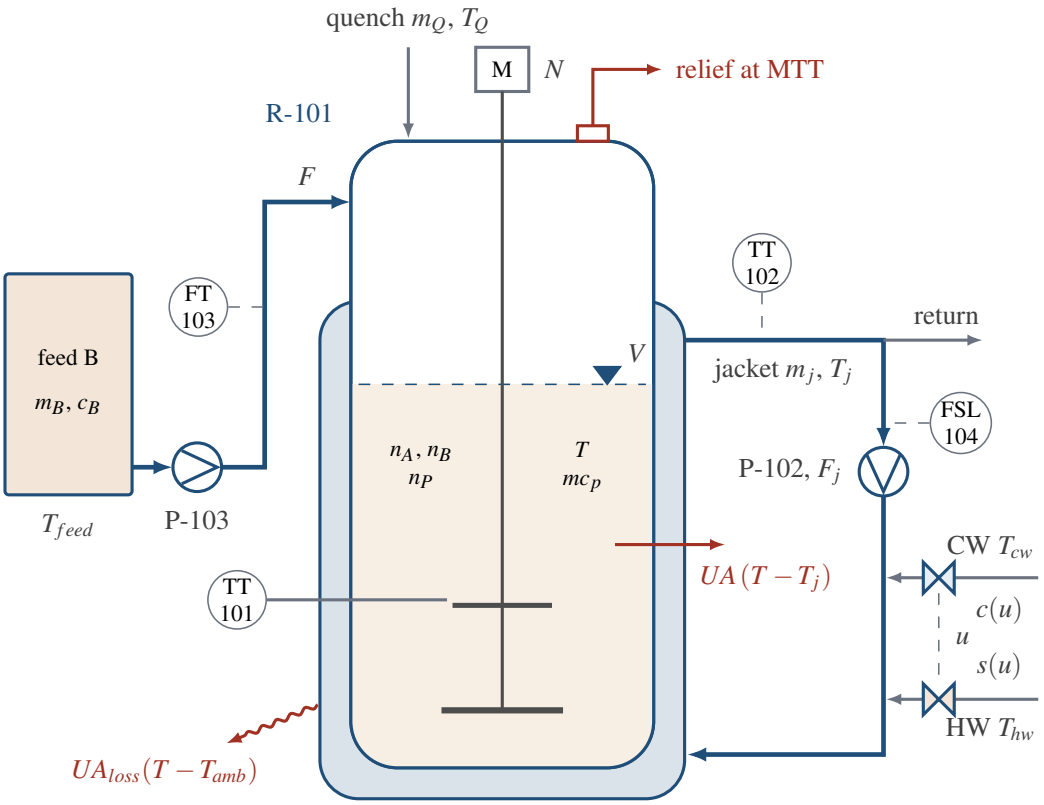


Figure 2.1 Semi-batch reactor R-101 as modeled. Feed of strength c_B at T_{feed} is dosed at F by P-103 through FT-103 into a batch of volume V holding n_A, n_B and n_P at temperature T, stirred at speed N. The jacket, holding m_j of water at T_j, takes $UA\,(T - T_j)$ from the batch and circulates at F_j through P-102, with fractions $c(u)$ and $s(u)$ of its flow replaced by cooling water at T_{cw} and hot water at T_{hw} on the split-range signal u. TT-101 and TT-102 are the batch and jacket thermometers and FSL-104 the low-flow switch of the loop. The rupture disc relieves the batch at MTT, the quench adds m_Q of solvent at T_Q, and the batch loses $UA_{loss}(T - T_{amb})$ to the room.

Table 2.1 Variable Roles, Reset Values and Ranges in Case 2

Variable	Role	Reset value	Range
T	CV, TT-101	70.0 °C SP	30–90
T_j	CV, TT-102	70.2 °C	SP 5–135
MTSR_{est}	CV, AIC-104	SP 80 °C	0–90 K below
$T_{j,SP}$	MV of TIC-101	70.2 °C	5–135
u	MV, split range	50.09 %	0–100
F	MV, feed	0, then 300 kg/h	0–600
T_{cw}	DV, cooling water	8 °C	5–40
a, activity	DV, unmeasured	100 %	25–150
f, fouling	DV, unmeasured	0 %	0–60
feed strength	DV, unmeasured	100 %	80–120
N/N_0	DV, agitator	100 %	20–100
F_j	DV, circulation	2 kg/s	0–100 %
n_A, n_B, n_P	states	12000, 0, 0 mol	—
T, T_j, u	states	70.0, 70.2 °C, 50.09 %	—
m_{fed}, n_X	states, ledgers	0 kg, 0 mol	—
ten ledgers	states	0	—

2.3.1 Kinetics

The synthesis A + B → P is second order, and the product decomposes in a first-order reaction, as shown in Equation (2.1).

$$\begin{aligned} r &= k(T)\,\frac{n_A\,n_B}{V}, \qquad r_d = k_d(T)\,n_P \\ k(T) &= k_{70}\,\frac{a}{100}\,\exp\Big[-\frac{E_r}{R}\Big(\frac{1}{T_K}-\frac{1}{343.15}\Big)\Big] \\ k_d(T) &= k_{d,24}\,\exp\Big[-\frac{E_d}{R}\Big(\frac{1}{T_K}-\frac{1}{T_{D24,K}}\Big)\Big] \end{aligned} \tag{2.1}$$

where r and r_d are in mol/s, T_K is the absolute temperature, $k_{70} = 2.4 \times 10^{-6}\,\mathrm{m^3/(mol\,s)}$, $E_r = 120\,\mathrm{kJ/mol}$, a is the catalyst activity in percent of nominal, $E_d = 300\,\mathrm{kJ/mol}$ and $R = 8.314\,\mathrm{J/(mol\,K)}$. The activation energy of the synthesis makes it 13.5 times slower at 50 °C than at 70 °C and 257 times slower at 30 °C, and 2840 times faster at 150 °C. The decomposition constant is anchored on the declared $T_{D24} = 130\,^\circ\mathrm{C}$ through the zero-order estimate of the adiabatic time to maximum rate, as shown in Equation (2.2).

$$\mathrm{TMR}_{ad}(T) = \frac{mc_p\,R\,T_K^2}{(-\Delta H_d)\,k_d(T)\,n_P\,E_d} \tag{2.2}$$

The constant $k_{d,24}$ is chosen so that $\mathrm{TMR}_{ad}(T_{D24}) = 24\,\mathrm{h}$ for the finished batch, with $mc_p = 7.56\,\mathrm{MJ/K}$ and n_P equal to the 8000 mol of the nominal dose, which gives $k_{d,24} = 1.97 \times 10^{-7}\,\mathrm{/s}$ with $(-\Delta H_d) = 250\,\mathrm{kJ/mol}$. Near T_{D24} the decomposition doubles its rate about every 3.1 K. The finished batch would reach its maximum rate in 222 h at 120 °C, 2.9 h at 140 °C, 23 min at 150 °C and 3.4 min at 160 °C, and at 70 °C the decomposition runs 6.3 million times slower than at T_{D24}. A heat release that grows exponentially with temperature against a removal that grows only linearly is also what sets the ignition boundary of the continuous reactor in Volume 7, Chapter 2.

2.3.2 Mole Balances and the Accumulation

With the feed mass flow F in kg/s and c_B the moles of B it carries per kilogram, the mole balances and the batch's volume and heat capacity follow Equation (2.3).

$$\begin{aligned} \frac{\mathrm{d}n_A}{\mathrm{d}t} &= -r, \qquad \frac{\mathrm{d}n_B}{\mathrm{d}t} = F\,c_B - r, \qquad \frac{\mathrm{d}n_P}{\mathrm{d}t} = r - r_d \\ \frac{\mathrm{d}n_X}{\mathrm{d}t} &= r_d, \qquad \frac{\mathrm{d}m_{fed}}{\mathrm{d}t} = F \\ V &= (m_{A0} + m_Q)/\rho_A + m_{fed}/\rho_B \\ mc_p &= (m_{A0} + m_Q)\,c_p + m_{fed}\,c_{pB} \end{aligned} \tag{2.3}$$

Here c_B is the label value of 6.667 mol/kg, a reagent of 150 g/mol, times the feed strength in percent over 100, $\rho_A = 900\,\mathrm{kg/m^3}$, $\rho_B = 1000\,\mathrm{kg/m^3}$, $c_p = c_{pB} = 1800\,\mathrm{J/(kg\,K)}$, and m_Q is the quench mass, zero until the quench fires. The unreacted B, n_B, is the accumulation, what has been fed less what has reacted. The heat capacity grows from 5.40 MJ/K for the charge to 7.56 MJ/K for the finished batch, and the volume from 3.33 to 4.53 $\mathrm{m^3}$, 56 to 76 % of the vessel.

2.3.3 Energy Balances and the Split Range

The batch and the jacket water exchange heat across the wall, as shown in Equation (2.4).

$$\begin{aligned} mc_p \frac{\mathrm{d}T}{\mathrm{d}t} &= (-\Delta H_r)\, r + (-\Delta H_d)\, r_d - UA\,(T - T_j) \\ &\quad + F\, c_{pB}\,(T_{feed} - T) - UA_{loss}\,(T - T_{amb}) \\ m_j c_{pj} \frac{\mathrm{d}T_j}{\mathrm{d}t} &= F_j c_{pj}\left[c(u)\,(T_{cw} - T_j) + s(u)\,(T_{hw} - T_j)\right] \\ &\quad + UA\,(T - T_j) \end{aligned} \tag{2.4}$$

with $(-\Delta H_r) = 200\,\mathrm{kJ/mol}$, $T_{feed} = 25\,^{\circ}\mathrm{C}$, $UA_{loss} = 20\,\mathrm{W/K}$ to a room at $T_{amb} = 20\,^{\circ}\mathrm{C}$, $m_j = 150\,\mathrm{kg}$, $c_{pj} = 4180\,\mathrm{J/(kg\,K)}$ and F_j equal to 2 kg/s times the circulation disturbance, and zero while the circulation has failed. The jacket inlet is the circulating water with a fraction $c(u)$ of it replaced by cooling water or $s(u)$ by hot water, so that at the crossover, with both valves shut, the loop only recirculates. The openings follow the signal as shown in Equation (2.5).

$$\begin{aligned} c(u) &= \mathrm{clamp}\Big(\frac{u_{lo} - u}{u_{lo}}, 0, 1\Big), \qquad u < u_{lo} \\ s(u) &= \mathrm{clamp}\Big(\frac{u - u_{hi}}{100 - u_{hi}}, 0, 1\Big), \qquad u > u_{hi} \\ u_{lo} &= X - \tfrac{1}{2}\Delta u, \qquad u_{hi} = X + \tfrac{1}{2}\Delta u \\ \tau_v \frac{\mathrm{d}u}{\mathrm{d}t} &= u_{cmd} - u \end{aligned} \tag{2.5}$$

where c and s are zero outside the stated ranges, the crossover is $X = 50\,\%$ and the dead band $\Delta u = 0$ as shipped (both adjustable), and $\tau_v = 10\,\mathrm{s}$ is the lag of the valve actuator behind the command u_{cmd}. The jacket inlet temperature $T_j + c\,(T_{cw} - T_j) + s\,(T_{hw} - T_j)$ rises continuously and monotonically with u, but its slope differs on the two sides. With the jacket at 50 °C it is 0.84 K per percent on the cooling side and 1.8 K per percent on the hot side, which one slave tuning has to live with. The pressurizer of Volume 18, Chapter 3 shares one controller output between two unequal actuators in the same way.

The batch enthalpy mc_pT, referenced to 0 °C with one heat capacity per kilogram and heats of reaction independent of temperature, closes against the ledgers exactly. Over case 2 to its completion the synthesis releases 1592 MJ, the jacket takes 1500 MJ and the room 16 MJ, the cooling water carries away 1505 MJ while the hot water brings 1.45 MJ, and the energy balance closes to 3.2×10^{-16} relative.

2.3.4 Heat Transfer

The conductance of the wall follows the wetted area and a viscosity factor, as shown in Equation (2.6).

$$UA = U_0\,(1 - \varphi_v x_P)\,(a_0 + a_1 V)\,(N/N_0)^{0.6}\,(1 - f) \tag{2.6}$$

where $U_0 = 500\,\mathrm{W/(m^2\,K)}$, $\varphi_v = 0.30$, $x_P = (n_P + n_X)/n_{B,nom}$ is the product made as a fraction of the 8000 mol dose, $a_0 = 2.6\,\mathrm{m^2}$ is the bottom head, $a_1 = 2.22\,\mathrm{m^2}$ per $\mathrm{m^3}$ of batch, N/N_0 is the agitator speed as a fraction of design and f the fouling fraction. In case 2, UA is 5.00 kW/K at the start of the dose, 4.84 half-way and 4.49 at its end, the area gaining 27 % while U loses 29 %. The law, the agitator exponent and the fouling factor are illustrative. In series with its circulation the jacket's effective conductance, $UA\,F_j c_{pj}/(UA + F_j c_{pj})$, is 3.13 kW/K at the start of the dose, so full cooling could remove 194 kW from a 70 °C batch with water at 8 °C and 125 kW with water at 30 °C.

2.3.5 The Hazard and the Criticality Classes

The accumulation turns into temperature through Equation (2.7).

$$\Delta T_{ad,acc} = \frac{(-\Delta H_r)\, n_B}{m c_p}, \qquad \mathrm{MTSR} = T + \Delta T_{ad,acc} \tag{2.7}$$

MTSR is where the batch would go if the cooling stopped now and all the unreacted B reacted with no heat removed. It rises 26.5 K for every 1000 mol of accumulation in the finished batch, and 37.0 K in the charge, where mc_p is smallest. The class follows from the order of the four temperatures of the cooling-failure scenario [6], as the application computes it. Class 1 has $T_p < \mathrm{MTSR} < \mathrm{MTT} < T_{D24}$, class 2 $T_p < \mathrm{MTSR} < T_{D24} < \mathrm{MTT}$, class 3 $T_p < \mathrm{MTT} < \mathrm{MTSR} < T_{D24}$, class 4 $T_p < \mathrm{MTT} < T_{D24} < \mathrm{MTSR}$, and class 5 has T_{D24} below both MTSR and MTT, where T_p is the process temperature. Here MTT is 170 °C, the boiling point of the solvent at the set pressure of the rupture disc. Because it lies above T_{D24} only classes 2 and 5 can occur, and the batch moves from 2 to 5 the moment MTSR passes 130 °C. T_{D24}, MTT, the MTSR margins and the class thresholds are teaching values. Real ones come from adiabatic calorimetry and DSC [13, 14], assessed by the method of Stoessel [6].

2.3.6 Events, Initialization and Numerics

Discrete events act after each integration step. A cooling failure sets $F_j = 0$ for 10 min or until the circulation is restored, and the water standing in the jacket then warms toward the batch. The quench adds $m_Q = 1000$ kg of solvent at $T_Q = 20$ °C at once, as shown in Equation (2.8).

$$T^{+} = \frac{m c_p T + m_Q c_p T_Q}{m c_p + m_Q c_p} \tag{2.8}$$

It cools by mixing and dilutes, and carries no inhibitor. When T reaches MTT the rupture disc bursts, the batch is lost and the state is frozen, because the model does not describe the venting. The decomposition is flagged as having taken over once $T > T_{D24}$ and its heat $q_d = (-\Delta H_d)\, r_d$ exceeds both the synthesis heat $q_r = (-\Delta H_r)\, r$ and the heat the batch is losing, $q_j + q_{loss}$. The dose is complete when the measured feed total reaches 1200 kg, and the batch is complete once the dose is done, the circulation is running and n_B has fallen below 0.5 % of the nominal dose.

The reset state is the charge at its starting temperature T_0, the setpoint or 30 °C in the challenge, with nothing fed and the jacket holding the charge against its heat loss. That gives $T_j = T_0 + UA_{loss}(T_0 - T_{amb})/UA$, 70.2 °C for a 70 °C charge, with the hot-water valve just cracked at $u = 50.09$ %. A setpoint chosen before Run, directly or carried through Reset, also moves the starting charge, while one changed during a run leaves the charge where it was and the batch is heated or cooled to the new value. The engine of Section 1.6.1 keeps time in minutes for this app and multiplies every rate by 60. The plant advances by fixed fourth-order Runge-Kutta steps of 1 s. In a runaway the reaction time scale falls to seconds, so each step is divided into $n = \lceil \max(\lambda h, |\dot{T}|h/1\,\mathrm{K}, 1) \rceil$ equal substeps, at most 400, with $\lambda = k(n_A + n_B)/V + k_d$ and $h = 1$ s, and n is chosen from the state at the start of the step so that a run replays exactly. The controllers act every 5 s, the historian logs every 2 s, and all of the noise is seeded at 20260926. The speed selector runs 1, 5, 20, 60 or 300 simulated minutes per second, 5 by default, at which the four-hour dose plays out in 48 s. Halving the step moves the cold batch's MTSR at 2 h 30 min by less than 1×10^{-8} K and the relief of case 4 by 0.5 s, and tightening the substep bound from 1 to 0.2 leaves the synthesis peak of case 4 at 157.1398 °C. The energy, jacket and mole ledgers close below 1×10^{-9} relative through a normal batch, a runaway to the relief and a quench, one of the 61 checks the application passes.

Table 2.2 collects the parameters. The disturbances of Table 2.1 are sliders and events of their own, and the scenarios set them at reset or fire them at stated times.

2.3.7 Assumptions

The batch is one well-mixed liquid with one heat capacity per kilogram for every component, heats of reaction that do not change with temperature and a reaction that does not change the volume. Real

Table 2.2 Model Parameters. The Show Model Panel Offers the Visible Ones as Sliders, While the Hidden Ones Stay Out of It

Symbol	Meaning	Value	Class
m_{A0}, n_{A0}	charge of A in solvent	3000 kg, 12 000 mol	visible
m_B	feed to dose	1200 kg	visible
$c_{B,lab}$	B per kilogram of feed	6.667 mol/kg	hidden
c_p	heat capacity of the batch	1800 J/(kg K)	visible
c_{pB}	heat capacity of the feed	1800 J/(kg K)	hidden
ρ_A, ρ_B	densities of charge and feed	900, 1000 kg/m^3	hidden
T_{feed}	feed temperature	25 °C	visible
k_{70}	synthesis rate constant, 70 °C	2.4×10^{-6} m^3/(mol s)	visible
E_r	synthesis activation energy	120 kJ/mol	visible
$-\Delta H_r$	heat of the synthesis	200 kJ/mol	visible
T_{D24}	decomposition threshold	130 °C, teaching value	visible
E_d, $-\Delta H_d$	decomposition	300, 250 kJ/mol	visible
MTT	boiling point at the disc	170 °C, teaching value	visible
V_r	reactor volume	6 m^3	hidden
U_0	clean U at the start	500 W/(m^2 K)	visible
φ_v	fall in U at full conversion	0.30	visible
a_0, a_1	wetted area law	2.6 m^2, 2.22 m^2/m^3	hidden
m_j, F_j	jacket water, circulation	150 kg, 2.0 kg/s	visible
c_{pj}	jacket water heat capacity	4180 J/(kg K)	hidden
T_{hw}	steam-heated water	140 °C	visible
τ_v	split-range actuator lag	10 s	hidden
UA_{loss}, T_{amb}	heat loss, room	20 W/K, 20 °C	hidden
F_{max}	feed pump capacity	600 kg/h	visible
m_Q	quench mass	1000 kg	visible
T_Q	quench temperature	20 °C	hidden

batches change heat capacity and density by a few percent as they convert, and a poorly mixed feed point can hold unreacted B that no well-mixed model shows.

The synthesis is an elementary second-order reaction with an Arrhenius constant chosen for teaching rather than fitted to a chemistry, and the catalyst activity simply scales it. A reaction with an induction period or an autocatalytic step accumulates in ways this model does not reproduce.

The decomposition is first order in the product and anchored through the zero-order estimate for the finished batch. A batch holding less product takes longer to reach its maximum rate at the same temperature, 30 min rather than 21 min at 150.6 °C for the batch of case 4, and the estimate says nothing about autocatalytic decompositions.

The rupture disc ends the model at MTT. Two-phase venting, disc sizing, the pressure history and boiling below MTT are not modeled, because the reactor is treated as closed and rated up to the disc.

The jacket is one well-mixed loop with a constant circulation, and its standing water keeps absorbing heat during a lost circulation without being allowed to boil. The heat-transfer law, the viscosity factor, the agitator exponent and the fouling factor are illustrative, and the sizes, rates and limits were picked to teach one recipe, not to design a plant.

2.4 Where the MTSR Is Lowest

A cold batch holds more B and a hot batch starts higher, and between the two the highest MTSR of a dose has a minimum in the process temperature. Figure 2.2 measures it. Each point is a run of the cold-start configuration with the charge initialized at the setpoint, the fixed 300 kg/h profile from $t = 0$ and the high-temperature feed stop switched off, both set before Run, so that only the temperature changes. The highest MTSR falls from 206.7 °C at 30 °C to a minimum of 75.6 °C at 66 and 67 °C, and rises again to 93.5 °C at 90 °C.

Below the minimum the accumulation dominates, 6679 mol at the end of the dose at 30 °C against 234 mol at 70 °C. Above it the batch temperature dominates, and from 75 °C upward the highest MTSR comes in about the first 8 min of the dose, during the overshoot as the feed starts, rather than

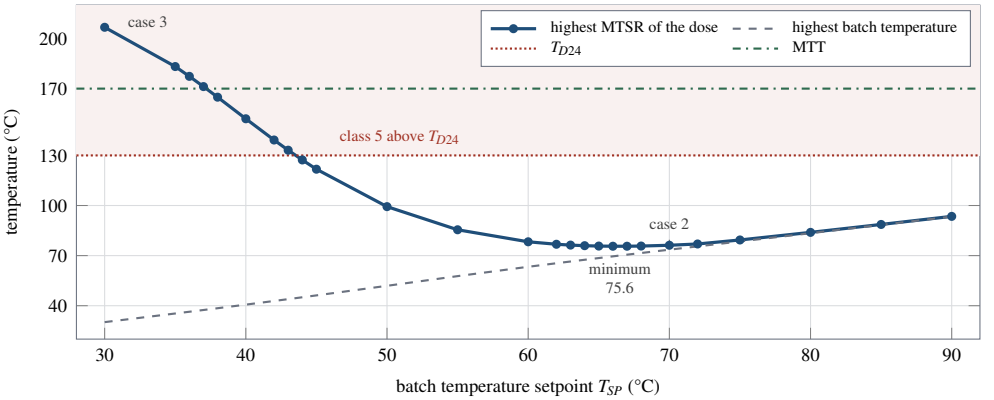


Figure 2.2 Highest MTSR of the fixed four-hour dose against the batch temperature setpoint, with the highest batch temperature of each run. Every point is an engine run of case 3 with the setpoint changed and the high-temperature feed stop off before Run, so that the charge starts at the setpoint. The highest MTSR is the solid line with dots, the highest batch temperature the dashed line, T_{D24} the dotted line and MTT the dash-dotted line. The shaded band lies above T_{D24}, where the batch is in criticality class 5.

at its end. The curve crosses MTT near 37 °C and T_{D24} near 43.5 °C, so every setpoint of 43 °C or less ends its dose in class 5, while the shipped 70 °C sits 3 to 4 K above the minimum with a highest MTSR 0.6 K above it.

A quasi-steady argument places the minimum. When the reaction keeps up with the feed, the unreacted B is roughly the molar feed divided by the pseudo-first-order rate constant, as shown in Equation (2.9).

$$n_B \approx \frac{F c_B V}{k(T)\, n_A}, \qquad \mathrm{MTSR} \approx T + \frac{(-\Delta H_r)\, F c_B V}{m c_p\, k(T)\, n_A} \tag{2.9}$$

With the Arrhenius law the derivative of this MTSR with temperature vanishes where $\Delta T_{ad,acc} = RT_K^2/E_r$, which is 8.0 K near 67 °C. At the end of the dose, where the highest MTSR occurs, the estimate of n_B is within 6 to 13 % of the engine's between 62 and 70 °C, 247 mol against 234 at 70 °C. The engine's own $\Delta T_{ad,acc}$ passes 8.0 K between 67 °C, where it is 8.6 K, and 68 °C, where it is 7.7 K. The measured curve is flat within 0.11 K from 65 to 68 °C, which is why the minimum is best read as a band. Running below it trades a few kelvin of process temperature for a large accumulation.

2.5 Implemented Control System

Three PI loops, a heat-flow calorimeter, a dosing controller with four shipped strategies and an optional fifth, and a layer of interlocks run every $\Delta t_c = 5$ s, as shown in Figure 2.3. Within a physics step they act in the order of Figure 1.10, and the plant's own events (the failure, the quench, the relief) follow the integration.

2.5.1 Transmitters

No controller reads the plant's amounts. TT-101 lags the true batch temperature by a 15 s first-order lag and adds first-order colored noise with a standard deviation of 0.05 K and a 30 s correlation time. TT-102 lags the jacket by 6 s with 0.05 K of noise correlated over 18 s. The Coriolis meter FT-103 lags the feed by 3 s with 0.3 % of multiplicative noise correlated over 12 s, and its integral is the measured total that ends the dose. FSL-104 trips when the true circulation falls below 20 % of design.

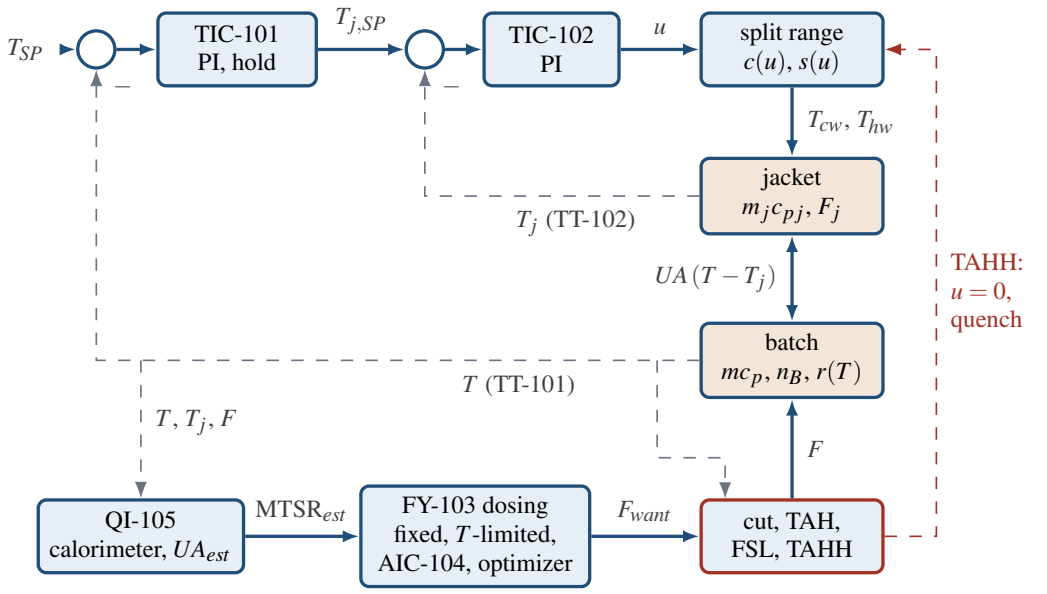


Figure 2.3 Control structure as implemented. TIC-101 sets the jacket setpoint and TIC-102 the split-range signal. QI-105 estimates MTSR from TT-101, TT-102, FT-103 and an assumed UA_{est}, FY-103 turns it into the wanted feed under the chosen strategy, and the temperature cut and the interlocks limit or stop the feed, the high-high trip also commanding full cooling and the quench.

Table 2.3 Loop Pairings, Tunings and Limits as Shipped

Tag	CV	MV	γ	K_c	τ_I, output
TIC-101	T (TT-101)	$T_{j,SP}$	+1	4	20 min, 5–135 °C
TIC-102	T_j (TT-102)	u	+1	3	2 min, 0–100 %
AIC-104	MTSR_{est}	F	+1	40	15 min, 0–F_{ceil}

Limit $L = 80\,°\mathrm{C}$, $F_{ceil} = 600\,\mathrm{kg/h}$, step $h = 5\,\mathrm{min}$
Profile 300 kg/h, cut 2 K with 0.3 K of hysteresis
TAH 80 °C (clears at 78), FSL 20 %, TAHH 115 °C

Gains in K per K, % per K and kg/h per K, in the order of the rows.

2.5.2 The Shared PI Law

All three loops run the framework's velocity-form PI law, Equations (1.14) to (1.17) of Section 1.6.2, with $\Delta t = \Delta t_c$. The action, written γ in Table 2.3, is +1 for all three loops and the bias is zero, with the integral seeded at reset to the actuator's reset value. The derivative term is unused in all three loops, which leaves the running sum equal to the proportional term. Windup is handled on the integral alone, by back-calculation whose tracking time is τ_I, the clamp limits only what is emitted, and none of the loops is rate limited. A loop in MANUAL tracks the operator's value and returns to automatic without a bump. The pairings and shipped settings are listed in Table 2.3.

2.5.3 The Temperature Cascade

TIC-101 turns the batch temperature error into the jacket setpoint, limited to 5 to 135 °C, and TIC-102 moves the split-range signal between 0 and 100 %. Both act directly, so a colder batch raises the jacket setpoint and a colder jacket raises the signal toward the hot water. With the integral hold on, TIC-101 runs with $\tau_I = \infty$ while the signal sits at 0 or 100 % or FSL-104 reports lost flow, so its integral neither

grows nor unwinds while the jacket cannot answer, a form of conditional integration [18]. In Local the jacket loop holds its own setpoint and TIC-101 tracks.

The application identifies no control model, and a step from the case 2 reset with the jacket loop closed in Local and the dose start moved to 120 min measures the two plants the cascade sees. Stepped from its reset value of 70.2 to 65 °C, the jacket reaches 63 % of the change 24 s after the new setpoint reaches TIC-102, with a 0.83 K undershoot. The batch follows the jacket with a gain of 0.995 and a 63 % time of 17.9 min, which is $mc_p/UA = 18.0$ min at the start of the dose and grows to 28.1 min at its end. The inner loop is about 45 times faster than the plant the outer loop sees, which is what makes the cascade worth having [18]. Volume 3, Chapter 3 measures how fast an inner loop must be on a fermenter whose lags lie within a factor of two of one another.

2.5.4 The Heat-Flow Calorimeter

QI-105 infers the reaction heat from the batch's heat balance using only TT-101, TT-102, FT-103 and a conductance it has to assume, as shown in Equation (2.10).

$$\begin{aligned} q_{est} &= UA_{est}(T_m - T_{j,m}) + (mc_p)_{est}\,\dot{T}_f \\ &\quad - F_m c_{pB}(T_{feed} - T_m) + UA_{loss}(T_m - T_{amb}) \\ n_{B,est} &= c_{B,lab}\, m_{fed,m} - \textstyle\sum q_{est}\,\Delta t_c/(-\Delta H_r) \\ \mathrm{MTSR}_{est} &= T_m + (-\Delta H_r)\max(0, n_{B,est})/(mc_p)_{est} \end{aligned} \tag{2.10}$$

where the subscript m marks a measurement, $\dot{T}_f$ is the slope of TT-101 through a first-order filter of 1 min, $m_{fed,m}$ is the measured feed total and $(mc_p)_{est}$ is built from it. The assumed conductance UA_{est} is the clean-jacket law of Equation (2.6) without fouling or agitator terms, evaluated at the measured volume and the calorimeter's own conversion and multiplied by a correction that is 100 % as shipped. A True UA option, which no real calorimeter has, hands it the plant's conductance for comparison. Every term but UA_{est} is measured, every constant is the declared value of Table 2.2, and a wrong conductance drifts the estimate by $n_{B,est} - n_B \approx -\int (UA_{est} - UA)(T - T_j)\,dt/(-\Delta H_r)$. A jacket that has fouled since the calibration removes less heat than the calorimeter believes, so it credits the reaction with heat it did not make and reports an accumulation that is too low, the failure in the safe-looking direction. In case 2, with nothing fouled, the estimate stays within 21.1 mol of the truth and its MTSR within 0.63 K.

2.5.5 Dosing Strategies

FY-103 sets the feed. The fixed profile spreads the dose evenly, $F_{prof} = m_B/t_{dose}$, which is 300 kg/h for four hours, from the dose start until the measured total reaches 1200 kg. Temperature-limited dosing runs the same profile but cuts it while TT-101 is more than $\delta_T = 2$ K above the setpoint, resuming 0.3 K below that. Accumulation-limited dosing gives the feed to AIC-104, a PI loop on MTSR_{est} with its setpoint the limit $L = T_{D24} - \Delta_{MTSR}$, 80 °C for the shipped margin of 50 K, and it tracks the applied feed whenever it is not in charge. The limit on the estimate is a path constraint in the sense of Section 1.4.4, held below T_{D24} with a back-off, and like the explosive-limit override of Volume 4, Chapter 6 it lets a safety limit rather than the recipe set the throughput. Here the limited quantity is not measured at all and comes from an estimate. The one-step optimizer takes the largest feed whose predicted MTSR one step h ahead meets the same limit, with the reaction rate held at its calorimetric value $r_{est} = \max(0, q_{est})/(-\Delta H_r)$, as shown in Equation (2.11).

$$\begin{aligned} F^* = \max F \quad &\text{subject to} \quad 0 \le F \le F_{ceil} \text{ and} \\ T_m + \frac{(-\Delta H_r)\left[n_{B,est} + (F c_{B,lab} - r_{est})\,h\right]}{(mc_p)_{est} + F c_{pB} h} &\le L \end{aligned} \tag{2.11}$$

The constraint is linear in F, so the optimum is closed form. With $\ell = L - T_m$, $\alpha = h\,[(-\Delta H_r)c_{B,lab} - \ell\, c_{pB}]$ and $\beta = \ell\,(mc_p)_{est} - (-\Delta H_r)(\max(0, n_{B,est}) - r_{est}h)$, the feed is $F^* = \mathrm{clamp}(\beta/\alpha, 0, F_{ceil})$

when $\alpha > 0$, and otherwise F_{ceil} if $\beta \geq 0$ and zero if not. Holding the rate fixed is conservative for a second-order reaction, which speeds up as B accumulates. The optimizer is a one-step constrained maximization and not a model predictive controller. Both limited strategies keep the temperature cut behind them, every strategy delivers the last of the dose exactly, and the application's own test confirms that the optimizer's feed meets the limit with equality and that 1 % more breaks it.

2.5.6 Interlocks

Above the loops, and in MANUAL as well as in automatic, TAH stops the feed while TT-101 reads above 80 °C and releases it 2 K below. FSL-104 stops the feed while the circulation is lost. TAHH at 115 °C is latched until the interlocks are reset, stopping the feed for good, driving the signal to full cooling and the jacket setpoint to its 5 °C floor, and firing the quench when it is armed. The interlocks stop the hazard from growing. None of them removes the B already in the batch.

2.5.7 Simulation Model, Loop Plants and Controller Models

The simulation model is the one of Section 2.3. The plants the loops see are the measured lags above, the 24 s inner loop and the 18 min to 28 min batch, and no shipped controller uses an identified model. The shipped controllers' internal models are the clean-jacket conductance and the heat balance of Equation (2.10), the adiabatic MTSR of Equation (2.7), which ignores the jacket water and the room and is conservative for that reason, and the one-step prediction of Equation (2.11). None of them contains the kinetics, which the predictive option of Section 2.5.8 carries in a model of its own.

2.5.8 Model Predictive Control

A fifth dosing strategy, model predictive control (MPC), is an option and never the default of any case, and only scenarios 9 and 10 load it. Every $T_s = 30\,\mathrm{s}$, six scans of Δt_c, it plans the feed over $N = 10$ steps (5 min) in seven blocks whose lengths are 1, 1, 1, 1, 2, 2 and 2 steps. As in Section 1.6.3, only the first block's feed is applied, here held for 30 s, and the initial plan is computed on the priming call at $t = 0$ [19, 20]. Its request replaces the wanted feed of FY-103, so the temperature cut, TAH, FSL, TAHH and the quench act on it as on the other strategies. Nonlinear predictive control has been applied to a semi-batch reactor to avoid thermal runaway [21], while this controller solves a quadratic program on a model relinearized at every decision.

The model is written from the balances of Section 2.3 with the declared values of Table 2.2 and the design disturbances (water at 8 °C, full activity and circulation), never the run's values, so a weak catalyst, a fouled jacket or warm water is mismatch it is not told about. Its six states are the unreacted B and A, the batch and jacket temperatures, the integral I of TIC-101 and the feed total, as shown in Equation (2.12).

$$\begin{aligned}
\frac{\mathrm{d}n_B}{\mathrm{d}t} &= F c_{B,lab} - \hat{r}, \quad \frac{\mathrm{d}n_A}{\mathrm{d}t} = -\hat{r}, \quad \frac{\mathrm{d}m_{fed}}{\mathrm{d}t} = F \\
mc_p \frac{\mathrm{d}T}{\mathrm{d}t} &= (-\Delta H_r)\,\hat{r} - UA_{est}(T - T_j) + F c_{pB}(T_{feed} - T) \\
&\quad - UA_{loss}(T - T_{amb}) + mc_p\,\hat{d}_T \\
\frac{\mathrm{d}I}{\mathrm{d}t} &= \frac{K_{c,T}}{\tau_{I,T}}(T_{SP} - T) \\
\frac{\mathrm{d}T_j}{\mathrm{d}t} &= \mathrm{clamp}\big((T_{j,SP} - T_j)/\tau_s,\; g_c,\; g_h\big)
\end{aligned} \tag{2.12}$$

Here $\hat{r} = \hat{\theta}\,k(T)\,n_A n_B / V$ at full activity, UA_{est} is the calorimeter's clean-jacket law, $T_{j,SP} = K_{c,T}(T_{SP} - T) + I$ clamped to 5 to 135 °C with the tuning of TIC-101 ($K_{c,T} = 4$, $\tau_{I,T} = 20\,\mathrm{min}$, I held at a jacket limit when the hold is on), and g_c and g_h are the rates of T_j at full cooling and full hot water, $[UA_{est}(T - T_j) + F_j c_{pj}(T_x - T_j)]/(m_j c_{pj})$ with T_x equal to $\hat{T}_{cw}$ or T_{hw}. The lag $\tau_s = 22\,\mathrm{s}$ stands for the closed jacket loop. Steps of 10 K up and down in its setpoint, on case 2 in Local with no feed, fit a first-order-plus-dead-time model with a gain of 1.02, a dead time of 0.16 min and a lag of 0.17 to

0.20 min, with residuals of 2.8 to 3.3 % of the step, and τ_s lumps the two. With water at 30 °C and the jacket near 45 °C, as in scenario 9, the true loop takes 91 s to reach 63 % of a step, because the cooling valve's authority falls with $T_j - T_{cw}$, an error that matters only while the valve is in its range.

With its estimates settled and the feed stepped to a new constant rate, the model's batch temperature stays within 0.54 K of the truth over 60 min in six open-loop checks. On a jacket fouled by 15 % it stays within 0.15 K while the predicted MTSR runs up to 6.94 K low, because the calorimeter's error passes into the estimates. Replaying the feed the plant received in scenario 9 is harder. From 2 min, with $\hat{T}_{cw}$ still at 8 °C, the error within 5 min reaches 0.84 K on the MPC's run and 1.29 K on the optimizer's, from 10 to 180 min it is at most 0.23 and 0.30 K, and over the 30 min from 200 min, across the end of the dose, it reaches 1.47 K.

At each decision the nonlinear model is integrated with no feed over the horizon by fourth-order Runge-Kutta, its Jacobian at the measured point (equal to a hand derivation within 3×10^{-9}) gives the feed's effect, the pair is discretized exactly over T_s, and an affine term per step makes the linear free response equal the nonlinear one. Over the 436 decisions of scenario 9 the planned temperature differs from the nonlinear model under the same plan by 0.42 K on average and 1.45 K at most. In the 210 decisions at full cooling the plan over-predicts the peak of its own feed by 0.50 K on average and 1.22 K at most, the safe direction, while the first plan, with the valve at the crossover, under-predicts it by 1.45 K.

Four estimates make the prediction offset free through an augmented disturbance model [22], all from instruments the calorimeter already reads, and all updated on every scan under every strategy. The multiplier $\hat{\theta}$ is the calorimeter's reaction rate over the model's, each filtered over 10 min, so the weak catalyst of case 7 appears as $\hat{\theta} = 0.25$. The water temperature $\hat{T}_{cw}$ comes from the jacket's heat balance while the commanded cooling valve is more than 15 % open, filtered over 5 min. At each decision $\hat{d}_T$ moves by 0.3 of the 30 s prediction residual of T, within ±7.2 K/h, and the bias b_M, the calorimeter's MTSR less the model's, makes the limit mean what it means to AIC-104. No Kalman filter is used. Each plan solves Equation (2.13) for the block feeds F_j in units of 100 kg/h.

$$
\begin{aligned}
\min \; & \sum_{k=0}^{N-1} -w_k F_k + S \sum_{j=0}^{6} (F_j - F_{j-1})^2 \\
& + \textstyle\sum_{i=M,T} \left(\rho\, \varepsilon_i^2 + \sigma \varepsilon_i \right) \\
\text{s.t. } & 0 \le F_k \le \min(F_{ceil}, F_{max}) \\
& T_s \textstyle\sum_k F_k \le m_B - m_{fed,m} \\
& \widehat{\mathrm{MTSR}}_k \le \max\left(L,\, \widehat{\mathrm{MTSR}}_k^0\right) + \varepsilon_M, \quad \varepsilon_M \ge 0 \\
& \hat{T}_k \le \max\left(T_{SP} + \delta_T - \beta,\, \hat{T}_k^0\right) + \varepsilon_T, \quad \varepsilon_T \ge 0
\end{aligned}
\tag{2.13}
$$

where F_k is the feed of the block holding step k (in kg/s in the dose row), F_{-1} the feed applied now, $w_k = 0.01\,(1 - k/2N)$ a reward that favors the near steps, $S = 0.001$ the move weight, $\hat{T}_k$ and $\widehat{\mathrm{MTSR}}_k$ the predictions over steps 1 to N (the latter linearized at the measured point, plus b_M), and the superscript 0 the free response with no feed. The ceiling sits $\beta = 0.25$ K below the cut, at 71.75 °C, and $L = 80$ °C. The feed bounds and the rest of the dose are hard, and the two bands are soft at $\rho = 10^4$ per K^2 and $\sigma = 10^3$ per K, with one slack per band for the whole horizon. The larger of the limit and the free response makes each band a funnel, so a slack counts only what the feed adds, and both were zero at every decision of every case. A flat reward spread the rest of the dose over the horizon and never finished it, hence the fall of w_k to half. Four plans of scenario 9 re-solved by an independent ADMM method agree to 3.2×10^{-6}, and a plan of nine variables takes 0.10 to 0.15 ms on average in Node.js.

The settings were chosen once on scenario 9, over 60 combinations of horizons from 5 to 60 min, back-offs from 0 to 1 K and move weights from 0.001 to 0.1, by a rule fixed beforehand: no cut, no TAH, true $T \le 72$ °C and true MTSR ≤ 81 °C, then the earliest end of the dose, ties within 1 % going

to the smaller feed travel. No other case is retuned. The end of the batch, 200 to 430 min away, lies beyond the solver's 120 steps, and within that limit the rule chose 5 min because every 30 and 60 min candidate was cut or ended its dose later, so the end of the dose enters through the hard row on the rest of the dose rather than a shrinking horizon.

A plan is accepted only when its solver status is optimal, all of its outputs are finite, its slacks are not negative and its moves lie within bounds. Otherwise the run is marked and AIC-104, whose integral has tracked the applied feed, doses until the next valid plan. Forced at five instants with the noise off, the hand-over moves the feed by at most 3.02 kg/h, the size of AIC-104's own next move (3.16 kg/h). The return is not bumpless (4 to 116 kg/h) but lies within the plan's ordinary changes between decisions, 83 kg/h on average and 412 kg/h at most. The controller reads only TT-101, TT-102, FT-103 and its total, the calorimeter, its own commands and the declared parameters, and the application's tests catch each of 16 doctored copies that read the plant or skip a protection under the MPC alone.

2.5.9 Linearization Along the Plan

The MPC has a second setting, its linearization, single in every case but case 10. Linearizing once per decision at the current state is the successive linearization of nonlinear MPC [23], and it holds the feed's effect at the measured point for the whole 5 min, although the plan's own feed moves the batch away from that point. The synthesis speeds up by 13.0 % per kelvin at 70 °C, and the jacket stops following its setpoint once the cooling valve is fully open. The iterated setting keeps the single plan as iteration 1 and relinearizes the same model along the path that plan predicts. With Equation (2.12) written as $\dot{x} = f(x,F)$, $x = (n_B, T, T_j, I, m_{fed}, n_A)$, and $\bar{F}_k$ the previous iteration's plan, each step is linearized at its own predicted state and planned feed, as shown in Equation (2.14).

$$\begin{aligned} \bar{x}_0 &= x_m, \qquad \bar{x}_{k+1} = \Phi(\bar{x}_k, \bar{F}_k) \\ A_k &= e^{J_k T_s}, \qquad B_k = \textstyle\int_0^{T_s} e^{J_k \tau}\, d\tau\; g_k \\ c_k &= \bar{x}_{k+1} - A_k \bar{x}_k - B_k \bar{F}_k \\ \hat{x}_{k+1} &= A_k \hat{x}_k + B_k F_k + c_k, \qquad \hat{x}_0 = x_m \\ \widehat{\mathrm{MTSR}}_k &= \nabla \mathrm{MTSR}(\bar{x}_k)\,(\hat{x}_k - \bar{x}_k) + \mathrm{MTSR}(\bar{x}_k) + b_M \end{aligned} \tag{2.14}$$

Here x_m is the state the decision starts from, Φ the model over T_s at constant feed (two Runge-Kutta substeps), J_k and g_k the derivatives of f in x and F at $(\bar{x}_k, \bar{F}_k)$ by the central differences used at the measured point, and $\hat{T}_k$ the second entry of $\hat{x}_k$, for $k = 1$ to N. The affine term c_k makes the model reproduce the plan's own path, so a plan that stops moving is a fixed point and a converged plan predicts what the nonlinear model gives for it. The weights, bounds, dose row, soft prices and funnels of Equation (2.13) are iteration 1's.

With A_k and B_k changing along the horizon, the condensed program carries products $A_{k-1} \cdots A_{j+1} B_j$ where the shared solver, written for one model, carries matrix powers. The application states it on the same dense active-set QP solver, and given one model at every step it returns the shared solver's plan bit for bit at all 351 decisions of case 10, while an independent log-barrier solve agrees within 0.0036 kg/h. The iteration stops when no planned feed moves by 1 kg/h, below the noise of FT-103, or after three relinearizations, both values declared before any run. A program that fails the acceptance test, or a plan still moving after the third relinearization, falls back to iteration 1's plan, with a mark in the historian. Each relinearization is a Gauss-Newton step on the plan, the step that real-time iteration takes once per sample [24], repeated here until the plan settles, between linear and fully converged nonlinear MPC [25]. A plan of case 10 perturbed by up to 100 kg/h per step is predicted within 3.2×10^{-4} K of the nonlinear model at the median decision, and within 3.2×10^{-5} K at 10 kg/h, while the single linearization misses by 0.51 K at every size because it does not reproduce the plan's own path. The converged plans of case 10's dose match the nonlinear model within 6.2×10^{-5} K, where the single plans at the same 324 decisions err by up to 1.66 K.

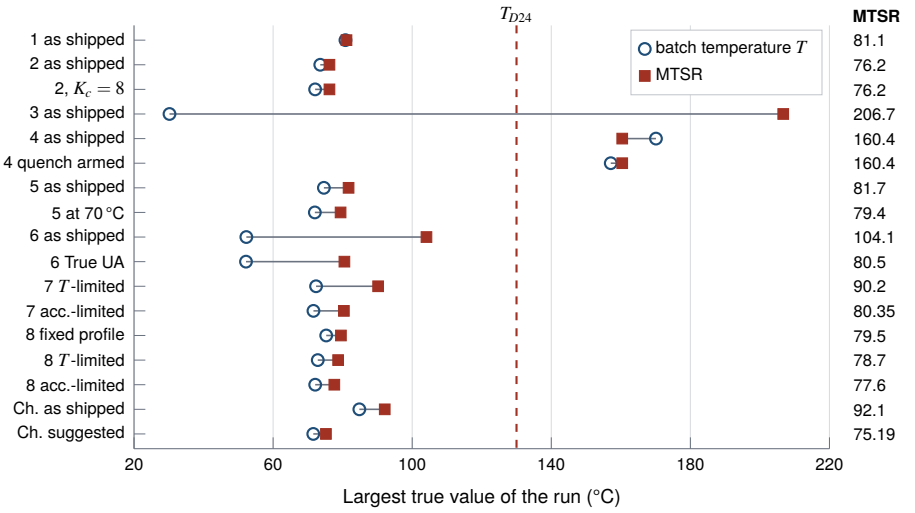


Figure 2.4 The largest true batch temperature (open circles) and the largest true MTSR (filled squares) of every run of Table 2.5, with T_{D24} dashed and the MTSR values at the right. The line between the two marks is the heat the batch stores as unreacted B. Notice case 3, whose batch never passes 30.2 °C while its MTSR reaches 206.7 °C, and case 4, whose temperature passes its MTSR only because the run continues to the relief (footnote a of the table).

While case 10 doses, 324 of its 331 plans converge, in 3.58 programs on average and at most four. Before the pump starts 12 of 19 do not, moving by about 295, then 9 to 17, then 2 to 3.4 kg/h. A decision takes 0.57 ms in Node.js against 0.12 ms for the single plan. In Chrome on the page's frame loop, under the load of other jobs on the host, it takes 1.3 to 1.5 ms at 1× speed, 0.3 % of a core, and at 300 simulated minutes per second, about ten decisions to a frame, the MPC takes 36 to 37 % of a core while the page holds 60 frames a second. A switch between the settings, or a fallback, moves the feed only at a decision and is not ramped. With the noise off, a fallback or a switch to the single plan steps the feed by up to 248 kg/h and the return by up to 255 kg/h, against ordinary moves of up to 279 kg/h, and over the next 10 min the batch is at most 0.10 K warmer than without it, with no cut.

2.6 Measured Cases

The scenario sequence of Table 2.4 runs through ten cases to a graded challenge. Every case starts from the shipped settings of Tables 2.3 and 2.2 and departs from them only as its row states. A departure is made at $t = 0$ before Run, and the application gives the same run when it is carried through Reset. The engine itself, run at its default seed, supplied each value quoted here, and times are clock times from the start of the batch. Table 2.5 gathers the largest true values of each run, and Figure 2.4 sets the largest MTSR of each run beside its largest batch temperature.

2.6.1 A Batch Whose Heat Follows the Pump

Case 1 is run in MANUAL, the pump at 300 kg/h and the split-range signal left at 50 % with both utility valves shut. The reaction makes 122 kW and no utility removes it, so the batch reaches 73.7 °C at 5 min and 79.5 °C at 10 min, and the TAH stops the feed at 10.7 min, when TT-101 reads 80.1 °C. The heat follows the pump because at 70 °C, with A in excess, the reaction takes B almost as fast as it arrives. After 10 min, 331 mol have been fed and 24 mol are unreacted, and stopping the pump then drops the reaction heat from 122 to 4.9 kW by 12 min. The *Hold the temperature* helper, pressed at 5 min, sets the signal to 33.9 %, the cooling valve 32 % open, and the batch peaks at 75.1 °C and is back at 69.2 °C by 60 min.

Table 2.4 Scenario Sequence of the Application

#	Scenario and configuration	Teaching point
1	MANUAL, 300 kg/h, valve at 50 %	the heat follows the pump
2	Cascade, fixed profile from 10 min	a feed-controlled batch
3	Charged and dosed at 30 °C	accumulation, no symptom
4	Case 3, circulation lost 2:30 to 5:00	an honest runaway
5	Case 4, accumulation-limited	safe, and slow
6	50 °C, AIC-104, U fouled 15 %	a fooled calorimeter
7	Temperature-limited, activity 25 %	no temperature symptom
8	Water at 30 °C, fixed 3 h profile	the jacket sets the pace
9	Case 8 dosed by the MPC	the cut planned ahead
10	Case 2 dosed by the MPC, iterated	a plan exact to its model
Ch.	Cold charge, two utility upsets	graded challenge

Table 2.5 Largest True Accumulation, MTSR and Temperature of Each Run, With the End of the Dose and the Outcome. Runs Stop at 8 h, at Completion or at the Relief, and Last 1 h in Case 1 and 7 h in the Challenge

Case and configuration	n_B (mol)	MTSR (°C)	T (°C)	Dose ends	Outcome
1 as shipped	49	81.1	80.8	—	TAH, 10.7 min
2 as shipped	234	76.2	73.56	4:10	done 4:26
2 with $K_c = 8$	234	76.2	72.08	4:10	done 4:25
3 as shipped	6679	206.7	30.2	4:00	class 5 at 1:44
4 as shipped	4398	160.4[a]	170.1	—	lost 3:55
4 quench armed	4398	160.4	157.1	—	quenched 3:10
5 as shipped	1641	81.7	74.6	—	563 kg by 8 h
5 at 70 °C[b]	324	79.4	72.0	6:47	done 7:08
6 as shipped	2056	104.1	52.3	3:42	limit +24.1 K
6 True UA	1145	80.5	52.2	5:54	limit +0.5 K
7 T-limited	765	90.2	72.35	4:09	done 5:47
7 acc.-limited	391	80.35	71.6	6:00	done 7:13
8 fixed profile	181	79.5	75.27	3:00	done 3:08
8 T-limited	275	78.7	72.84	3:34	done 3:50
8 acc.-limited	235	77.6	72.1	4:45	done 5:00
Ch. as shipped	1196	92.1	84.8	4:15	worst 22.76 K
Ch. suggested[c]	198	75.19	71.55	3:43	worst 1.61 K

n_B, MTSR and T are the largest true values of the run, and times are h:min. [a] At the failure. Once no B is left the running maximum is T itself, 170.1. [b] Setpoint raised before Run, so the charge starts at 70 °C. [c] Optimizer, 55 K margin, 450 kg/h ceiling, dose from 20 min, batch loop $K_c = 6$ and $\tau_I = 15$ min, integral hold on.

In case 2 the cascade holds 70 °C and the pump starts at 10 min. Before the feed the jacket sits 0.2 K above the batch and the noise keeps the signal between 49.3 and 51.0 % around the crossover. The signal crosses 50 % for good at 10.3 min and settles near 35 %, between 33 and 37 % with the noise, the cooling valve 26 to 33 % open. The feed is a load step of about 110 kW that the PI cascade can correct only after it has warmed the batch, so the batch peaks at 73.56 °C at 19.2 min and returns within 0.5 K of its setpoint for the rest of the dose at 64 min. With the batch loop gain doubled to 8 the peak is 72.08 °C at 16.4 min and the return comes at 47.6 min. From 70 min to the end of the dose the jacket runs 20.5 to 22.0 K below the batch and removes 96.5 to 105 kW. The unreacted B never exceeds 234 mol, 2.9 % of the dose, reached as the dose ends at 4 h 10 min, MTSR never passes

76.2 °C, and the batch is complete at 4 h 26 min. At 70 °C this batch is feed-controlled.

2.6.2 Accumulation Without a Symptom

Case 3 charges the batch at 30 °C and doses it on the same profile without heating it. During the dose the true temperature stays between 30.0 and 30.2 °C, TT-101 reads between 29.9 and 30.3 °C with its noise, the signal stays between 38.7 and 50.6 %, and the cooling valve never opens more than 22.5 %, because the reaction never makes more than 29.5 kW. At 1 h 44 min MTSR passes T_{D24} with 91 % of the B fed so far unreacted, and the class moves from 2 to 5 at the same moment (the calorimeter's class 13 s later). At 2 h 30 min, 4398 mol are unreacted, 55 % of the dose, with MTSR at 160.4 °C, and at 4 h 6677 mol, 83 %, with MTSR at 206.7 °C, above MTT. Nothing on the temperature trend changes, which is the point of the case and the left end of Figure 2.2.

> **TRY IT IN THE APP** **3 · Cold Start**
>
> Load the scenario and press Run. Watch the trends: the batch temperature stays flat near 30 °C while the true MTSR climbs. At 1 h 44 min it passes the 130 °C decomposition threshold and the criticality class steps from 2 to 5, with no change on the temperature trend. By 4 h the MTSR stands at 206.7 °C.

2.6.3 An Honest Runaway

In case 4 the jacket circulation fails at 2 h 30 min on the cold batch of case 3 and returns at 5 h, with the quench not armed. FSL-104 stops the feed at the scan that coincides with the failure, with 4398 mol unreacted, the most of the run, and MTSR at 160.4 °C. That stops the hazard from growing and does nothing about the B already there. For 35 min the batch creeps from 30.1 to 45.4 °C, warming the water standing in the jacket with it, until 3 h 05 min. The synthesis then accelerates, to 57.8 °C at 3 h 09 min and 75.0 °C at 3 h 10 min, and in the next 13 s it runs to 157.1 °C, gaining 37.6 K in the last one-second step and releasing 250 MW at the start of that step. TT-101, behind its 15 s lag, reads 78.6 °C at the true peak and trips the TAH at 100.9 °C and the TAHH at 116.7 °C, 2 s and 7 s after it, commanding full cooling that cannot flow. The synthesis peak lies 3.2 K below the MTSR at the failure because the standing water and the room took some heat during the creep, which is why MTSR is the conservative number. The jacket water, then at 59.9 °C, absorbs heat from the batch, which falls to 150.56 °C at 3 h 17 min, 20.6 K above T_{D24}, where the decomposition's 18.8 kW exceeds the 18.7 kW the batch loses and its time to maximum rate for this batch is about 30 min. The batch then heats itself, to 154.8 °C at 3 h 40 min and 159.9 °C at 3 h 50 min, and reaches MTT at 3 h 55 min with the decomposition making 699 kW (between 3 h 54 min and 3 h 55 min on eight other seeds). The rupture disc bursts and the batch is lost 65 min before the cooling would have returned. The run at 230 min is shown in Figure 2.5. Armed, the quench fires at the same high-high trip and brings the batch from 156.7 °C to 127.9 °C, below T_{D24}, and the decomposition never takes over. The reactor survives and the batch is scrapped.

The screenshot also shows the calorimeter's limits. During the jump its filtered slope and its integral lag the truth, so its MTSR spikes to 214.4 °C at 190.6 min while no B is left. Afterwards it credits the decomposition's heat to the synthesis, and its estimate of the unreacted B has fallen to –387 mol by 230 min, which the badges show as zero.

Pause. Try Exercise 2.1 before reading on.

2.6.4 The Same Failure, Accumulation-Limited

Case 5 repeats case 4 with AIC-104 holding the calorimeter's MTSR at 80 °C. At 30 °C the limit lets the pump run only as fast as the slow reaction takes the B, so by 2 h 30 min it has fed 276 kg against 750 kg on the fixed profile, with 1476 mol unreacted, MTSR at 80.0 °C and the pump at 28 kg/h. After the failure the stored B reacts slowly and the batch peaks at 74.6 °C at 4 h 25 min, class 2 throughout and 55 K below T_{D24}. When the cooling returns at 5 h the dose resumes, and by 8 h the pump runs at 23 kg/h with 563 kg fed, the true MTSR having peaked at 81.7 °C as AIC-104 brought

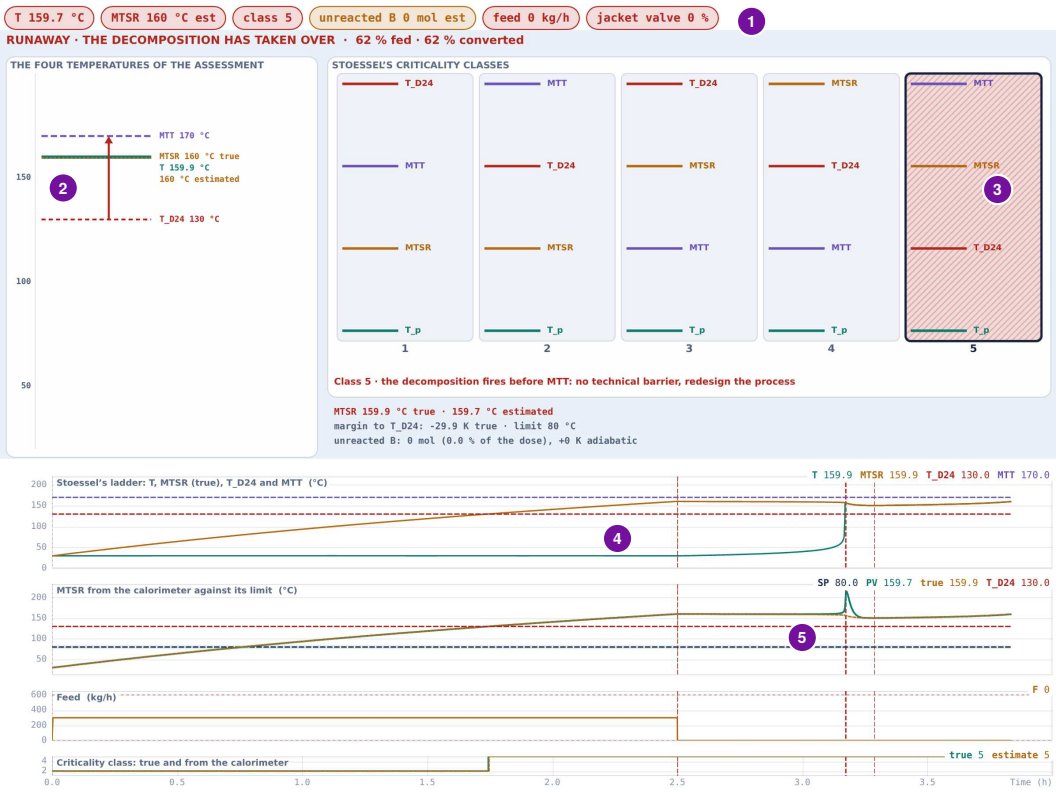


Figure 2.5 The semi-batch reactor in case 4, cooling failure at the wrong moment, at $t = 3.83$ h (230 min) with no control changed, a composite of the Safety view (top) and the trends (bottom), with the header, the controls and the Learn panel cut. The status line reads runaway, the decomposition has taken over, with 62 % of the feed in and 62 % converted, and the badges (1) give TT-101 at 159.7 °C, the calorimeter's MTSR at 160 °C, class 5, no unreacted B estimated, no feed and the split-range valve at 0 %. The ladder of the four temperatures (2) places the true batch at 159.9 °C and the true MTSR at 160 °C (160 °C estimated) between T_{D24} at 130 °C and MTT at 170 °C, with the decomposition's arrow drawn from T_{D24} up to MTT, and class 5 is picked out and hatched among Stoessel's five (3), over a line reading that the decomposition fires before MTT with no technical barrier. The margin line reads −29.9 K to T_{D24} against the 80 °C limit. The trend strips, over a 4.0 h window, show the true MTSR climbing while the batch holds 30 °C (4), to its value at the failure, 160.4 °C in Table 2.5, and the feed of the fixed profile stopping at 2.5 h when the circulation is lost. The batch creeps and then jumps, the calorimeter's MTSR spikes during the jump against its 80 °C limit (5), and the class steps from 2 to 5 long before the failure in the truth and the estimate. Dashed markers stand at the failure, at the crossing of T_{D24} with the TAH and TAHH trips, and at the decomposition taking over.

the feed back. It is safe, and very slow. Raising the setpoint to 70 °C before Run, which starts the charge at 70 °C, shows where the time goes. The same limit leaves 53 mol unreacted at the failure, the batch peaks at 72.0 °C at 57 min, and the dose finishes at 6 h 47 min, two and a half hours of it spent waiting for the cooling.

> **TRY IT IN THE APP** **5 · The Same Failure, Accumulation-Limited**
>
> Load the scenario and press Run, then watch the feed. The accumulation limit lets the pump run only as fast as the cold batch reacts, 28 kg/h at 2 h 30 min against 300 kg/h on the fixed profile. When the circulation fails, the batch peaks at 74.6 °C and stays in class 2. Compare the feed trend with 4 · Cooling failure at the wrong moment.

2.6.5 A Calorimeter Fooled by Fouling

Case 6 runs at 50 °C, where the reaction is 13.5 times slower than at 70 °C, dosed by AIC-104 at the same 80 °C limit, with a jacket that has lost 15 % of its U since the calorimeter was calibrated. The calorimeter overestimates the heat removed, counts too much B as reacted, and AIC-104 feeds into an accumulation it cannot see. The dose finishes at 3 h 42 min instead of the 5 h 54 min it takes with the true conductance. At that moment the calorimeter assumes 4.72 kW/K where the jacket has 4.18 kW/K, reports 1192 mol unreacted where 2054 mol are, and puts MTSR at 81.3 °C while the true value is 104.0 °C. The departure of −861.9 mol matches the −863.8 mol predicted from the integral of the conductance error to 0.2 %. With the True UA option or a correction of 85 % the dose takes 5 h 54 min and the true MTSR never exceeds 80.5 °C. At 70 °C the same fouling hardly matters, because the reaction keeps up with the feed, and the true MTSR peaks at 81.0 °C with the setpoint raised before Run and the charge started at 70 °C.

Pause. Try Exercise 2.2 before reading on.

2.6.6 A Weak Catalyst

Case 7 runs at 70 °C with a catalyst at a quarter of its activity and temperature-limited dosing. The cut acts only twice, at 8.6 and 23.9 min, and the batch never goes above 72.35 °C, less than the 73.56 °C of case 2. Yet the unreacted B reaches 765 mol against 234 mol, MTSR reaches 90.19 °C, and the dose finishes at 4 h 09 min as if nothing were different. A slower reaction releases its heat later rather than hotter. Accumulation-limited dosing reads the slower reaction in the calorimeter and slows the feed, so MTSR peaks at 80.35 °C, the unreacted B at 391 mol, and the dose takes 6 h. On identical seeds the temperature-limited dose overshoots the 80 °C limit by 10.2 K and the accumulation-limited dose by 0.35 K. Anything that slows the reaction makes the two strategies disagree in the same way, an inhibitor in the feed or a missed initiation as well as a weak catalyst.

Pause. Try Exercise 2.3 before reading on.

2.6.7 Warm Cooling Water

Case 8 runs the fixed profile in three hours, 400 kg/h, with cooling water at 30 °C. After the first 30 min the jacket removes 130 to 138 kW, all that full cooling can, while the reaction makes 142 to 146 kW until the last minute of the dose. The split-range signal sits at 0 % for 93 % of the dose, and the batch runs at 73.9 to 74.8 °C with a peak of 75.27 °C. Temperature-limited dosing cuts the feed 14 times, holds the batch below 72.84 °C, and stretches the dose from 3 h 00 min to 3 h 34 min. Accumulation-limited dosing, with the same cut behind it, finishes at 4 h 45 min. The jacket, not the recipe, sets the pace, and a saturated temperature loop has no handle on it but the feed.

2.6.8 The Challenge

The challenge charges the batch at 30 °C, to be heated to 70 °C and dosed, warms the cooling water from 8 to 30 °C at 2 h, and loses the circulation for 10 min at 2 h 30 min. The target is the dose finished by 4 h 00 min, the true MTSR never above 130 °C, the true batch temperature within 2 K of 70 °C from 45 min until the dose is complete, and no relief or quench, judged when the batch is complete. Settings are changed before Run, and a changed model parameter or disturbance voids the attempt (Section 1.6.4). As shipped, the dose starts at once on the four-hour profile with a sluggish batch loop ($K_c = 1.5$, $\tau_I = 60$ min, no hold). At 45 min the batch is still at 47.2 °C, and the B fed

into it cold ignites at about 1 h and takes it to 84.8 °C, tripping the TAH at 1 h 00 min. The worst deviation after 45 min is 22.76 K, 20.76 K outside the band, MTSR reaches 92.1 °C and the dose finishes at 4 h 15 min, 15 min late. The suggested configuration of Table 2.5 starts the dose at 20 min, once the batch is hot, and runs the one-step optimizer against a 75 °C limit under a 450 kg/h ceiling, with the batch loop at 6 and 15 min and its integral held while the jacket cannot answer. The dose then finishes at 3 h 43 min (223.0 min), 17.0 min early, with a worst deviation of 1.61 K, 0.39 K inside the band, never more than 198 mol unreacted and MTSR never above 75.19 °C, 54.8 K under its bound. AIC-104 in place of the optimizer, at the same limit and ceiling, stretches the dose to 5 h 04 min and misses the band by 0.68 K. A fixed profile can pass as well, because the failure comes at a known time. With the suggested batch loop, a fixed 3.5 h profile started at 10 min finishes at 3 h 50 min within 1.85 K, while the same profile started at 20 min misses the deadline by 6 s. A second ten-minute failure, fired by hand outside the graded attempt, tests a configuration further. At 3 h 20 min it leaves the optimizer's MTSR at 75.2 °C and its dose done at 3 h 53 min, although the batch strays 2.08 K, while the fixed profile reaches an MTSR of 77.1 °C, strays 2.70 K and ends its dose after 4 h 00 min. Over second failures every 5 min from 25 min to 3 h 45 min the optimizer's MTSR stays below 75.5 °C, but the band is missed at 26 of the 41 times, by up to 1.39 K, so the accumulation limit protects the MTSR wherever the failure falls but not the band.

> ✦ DID YOU KNOW? **Faster and Safer at Once**
>
> From the simulator. In the challenge the shipped settings end the dose at 4 h 15 min with up to 1196 mol of B unreacted and an MTSR of 92.1 °C, the suggested ones at 3 h 43 min with at most 198 mol and 75.19 °C (Table 2.5). The faster batch stores less heat.

2.6.9 Predictive Dosing Against the One-Step Optimizer

Scenario 9 is the warm-water batch of case 8 dosed by the MPC, with the 80 °C MTSR limit, a 600 kg/h ceiling and the cut at 72 °C. Table 2.6 runs it, the dosing cases and the challenge under the MPC and the shipped strategies with the same seed, instruments, noise and events, to the completion of the batch or at most 600 min (420 min in the challenge). Neither structure sees anything in advance, the shipped strategies keep their shipped tuning (the suggested one in the challenge), and the MPC is never retuned.

On the warm-water batch the MPC is never cut and the batch never passes 71.87 °C, while the one-step optimizer is cut 21 times, the first at 2.9 min, spends 74.4 min of the dose above 72 °C and peaks at 73.20 °C at 4.3 min, as shown in Figure 2.6. From 45 min to the end of the dose the MPC holds the batch 0.73 K from its setpoint on average against 1.48 K. It is slower. Its dose ends at 218.0 min against 207.9 min, its batch is complete at 231.6 min against 225.8 min, 5.8 min later, and it moves the pump more, 36 199 kg/h of summed feed changes against 29 060 kg/h. The optimizer gains its time by letting the batch run above the cut, and a hotter batch sheds more heat into a jacket that is already at full cooling. On five other noise seeds the MPC is never cut and completes the batch at 231.4 to 233.2 min, while the optimizer is cut 21 times on each and completes at 225.0 to 225.8 min, and halving the integration step leaves both counts and both completion times unchanged.

Neither part of the controller produces the result alone. Without the temperature ceiling the MPC is cut 21 times at every look-ahead from one to ten steps, the optimizer's behavior. With the ceiling kept and a single 30 s step of look-ahead it is cut 19 times, because a change of feed reaches the batch temperature over minutes rather than seconds, and two steps cut it twice and four steps once. One move held over the full 5 min is never cut, but its dose ends at 290.6 min, 73 min later, so the seven moves are what keep it fast. Longer horizons do not help on this batch. The 10 to 240 min sweep is scripted, because the interactive horizon control stops at 60 min. Each horizon is cut once, between 2.9 and 3.9 min. The cause is not isolated, but each cut falls while $\hat{T}_{cw}$ is still between 14.5 and 15.9 °C on its way to 30 °C, when the plan counts on a colder jacket than the plant has. Holding $\hat{\theta}$ at 1 or dropping $\hat{d}_T$ or the MTSR band moves the end of the dose by at most 0.5 min without a

Table 2.6 Shipped Dosing Strategies Against the MPC, Seed 20260926. Times in Minutes, Temperatures in Degrees Celsius

Case	Structure	Dose ends	Batch done	Highest T	Highest MTSR	Worst dev. (K)	Cuts	Above 72 °C	Travel (kg/h)
9, water at 30 °C	MPC	218.0	231.6	71.87	79.39	1.87	0	0.0	36199
	one-step optimizer	207.9	225.8	73.20	80.04	3.00	21	74.4	29060
	accumulation-limited	285.1	300.0	72.10	77.58	2.10	9	7.1	8274
	optimizer, 1 min prediction step[a]	207.3	222.9	73.46	79.99	3.09	21	76.8	26403
	temperature-limited, 220 min, cut at 71.25 °C[a]	255.1	272.6	71.99	76.71	1.59	12	0.0	8509
	accumulation-limited, 175 kg/h ceiling[a]	416.9	427.3	71.91	73.77	0.55	0	0.0	381
	MPC, no temperature ceiling	209.2	222.5	73.53	79.69	3.03	21	75.2	26501
	MPC, one 30 s step ahead	209.5	224.3	73.24	79.82	3.02	19	44.1	45037
	MPC, one move held over 5 min	290.6	290.6	71.71	77.80	2.94	0	0.0	13930
2, cascade at 70 °C	one-step optimizer	179.3	203.5	73.16	80.57	2.57	17	38.2	23758
	MPC	178.0	202.7	71.85	80.20	1.85	0	0.0	24842
6, fouled jacket	AIC-104 (shipped)	222.1	—	52.27	104.09	2.22	4	0.0	5573
	one-step optimizer	199.4	588.6	53.11	102.85	2.65	5	0.0	12579
	MPC	194.4	582.4	51.92	102.94	1.92	0	0.0	31858
7, weak catalyst	temperature-limited (shipped)	249.1	346.8	72.35	90.19	1.54	2	5.8	1800
	one-step optimizer	352.4	425.3	72.49	80.14	1.53	2	7.5	25168
	MPC	357.0	429.5	71.85	80.19	1.76	0	0.0	29168
Ch., as shipped[b]	fixed profile (shipped)	254.6	292.1	84.77	92.14	22.76	0	9.4	1800
	one-step optimizer	202.1	232.1	77.82	84.21	22.77	6	31.8	13297
	MPC	206.0	221.2	72.97	80.42	23.01	1	7.2	16608
Ch., suggested[b]	one-step optimizer (suggested)	223.0	236.7	71.55	75.19	1.61	0	0.0	10166
	MPC	221.0	233.9	71.83	75.05	1.83	0	0.0	10355
	optimizer, 525 kg/h ceiling[a]	212.8	226.3	71.87	75.33	1.87	0	0.0	11303

Highest values and the minutes above 72 °C during the dose come from the true state at every physics step. The worst deviation, the largest $|T - T_{SP}|$, runs from 45 min to the end of the dose as the scorecard computes it. Cuts count each start of a temperature cut, including a lowered cut's own, and feed travel sums the absolute change of the feed from scan to scan. A dash marks a batch not complete by 600 min. [a] Setting found by search in hindsight. [b] Every run of the challenge as shipped fails its target, and the three runs with the suggested settings pass. The MPC solves 436, 356, 389, 714, 412 and 442 plans in the six cases in table order, and no plan is rejected in any run.

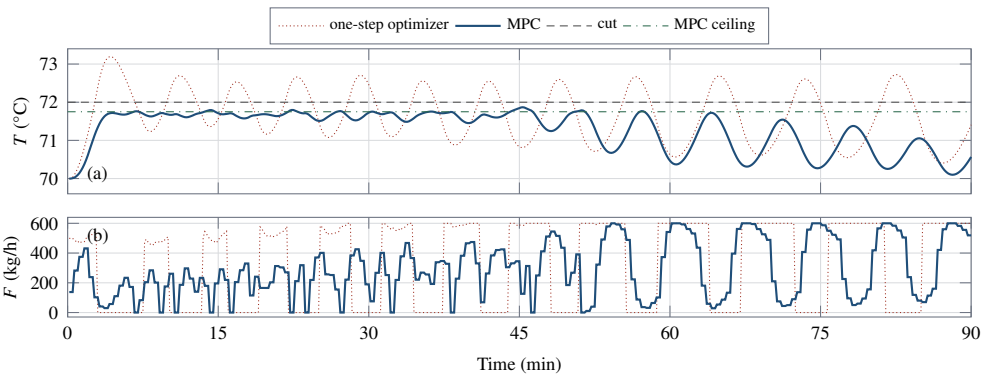


Figure 2.6 Scenario 9, the warm-water batch, over its first 90 min under the MPC (thick solid line) and under the one-step optimizer (thin dotted line) on the same seed. (a) True batch temperature with the 72 °C cut (dashed) and the MPC's 71.75 °C ceiling (dash-dotted). (b) Feed. The optimizer is cut 12 times in this window and 21 times over the dose.

cut, and holding $\hat{T}_{cw}$ at 8 °C ends it 5.1 min sooner but leaves the batch 1.42 K from its setpoint on average and is cut once on two of five other seeds. The back-off has a price of its own. With none the MPC is cut 11 times, and with 1 K it is never cut but its dose ends at 253.0 min, 35.0 min after the shipped 0.25 K.

The shipped strategies keep this batch below 72 °C only with hindsight and more slowly. The fastest such setting found by search, temperature-limited dosing over 220 min with its cut lowered to 71.25 °C, completes the batch at 272.6 min, 41.0 min after the MPC, and is still stopped 12 times by its own lowered cut. Without any cut, accumulation-limited dosing needs a 175 kg/h ceiling and completes at 427.3 min, and the optimizer needs 150 kg/h and 489.2 min.

In case 2 the MPC ends its dose 1.3 min before the optimizer (178.0 against 179.3 min) without a cut against 17. The MPC loses case 7, where it is never cut against twice but its dose ends 4.6 min after the optimizer's (357.0 against 352.4 min), and $\hat{\theta}$ is essential there, since with it held at 1 only 979 kg are fed by 600 min. In the cold cases 3 to 5 the MPC, AIC-104 and the optimizer feed within 3 kg of one another by 600 min. Case 6 fools every structure through the calorimeter. The MPC ends its dose first, at 194.4 min, but the true MTSR reaches 102.94 °C while the calorimeter reports 80.01 °C, against 104.09 °C under AIC-104 and 102.85 °C under the optimizer, and with the True UA option or the 85 % correction the MPC holds 80.18 °C and 80.17 °C. The challenge as shipped defeats every structure, and the MPC ends its dose 3.9 min after the optimizer with a worst deviation of 23.01 K from the cold start. With the suggested settings and the strategy switched to the MPC it passes and ends the dose 2.0 min before the optimizer, but the optimizer under a 525 kg/h ceiling chosen in hindsight ends it at 212.8 min, 8.2 min before the MPC, and passes as well. The MPC moves the pump more than the optimizer in cases 2, 6, 7 and 9 and in the challenge, and less only in the cold cases. A move weight of 0.1 would cut scenario 9's travel to 5189 kg/h and complete the batch 1.4 min later without a cut, but the tuning rule chose 0.001.

Every protection acts on the MPC's request. In every case run to its end no scan feeds through a cut, TAH, FSL or TAHH, and the protections stop requested feed on 1692 scans in each of cases 4 and 5 and on 92 in the challenge. Forced, a TAH at 71 °C stops 7166 scans of requests in scenario 9, a TAHH at 71.5 °C latches at 3.8 min and fires the quench, and case 4 with no MTSR margin reaches the TAHH at 203.0 min, after which no feed enters.

2.6.10 The Plan Linearized Once or Along Its Path

The iterated setting was built on the expectation that the single linearization is most wrong where the accumulation limit binds, and the runs do not bear that out. Every accepted single plan was simulated again on its own nonlinear model over every scenario, each accumulation-limited one also with each disturbance at the limits of its range. In the 46 runs whose MTSR row is active at 20 % or more of the dosing decisions, the plan under-predicts MTSR by at most 0.27 K at the 95th percentile, because MTSR is linear in the unreacted B through Equation (2.7) and the free response is exact. The temperature is where the single plan errs, in case 10 high by 0.48 K at its worst step at the median decision and by up to 1.52 K.

On the case where the accumulation limit binds with the single plan's largest MTSR error, both linearizations hold the true MTSR 0.01 K under its limit, as Table 2.7 shows. Every dose there ends before 4 h 00 min, so the target turns on the temperature band alone. On the default seed the iterated batch peaks at 71.87 °C and passes while the single one reaches 72.11 °C, but over six seeds the iterated batch peaks at 71.84 to 72.30 °C and passes twice, the single batch at 72.00 to 72.19 °C and never. The single linearization with its ceiling 1 K below the cut, the setting the first round's tuning rule picks here, peaks at 71.26 to 71.45 °C and passes on all six seeds for 1.6 min more dose. Where the limit binds the two linearizations tie, and a wider back-off wins.

Where the temperature ceiling binds the iteration helps. In case 10 the dose ends at 175.4 min against 178.0 and the batch is complete at 198.2 min against 202.7, 4.5 min sooner, both batches peaking at 71.86 and 71.85 °C without a cut. The iterated batch runs at full cooling at 98 decisions against 33, spends 88.0 min of the dose above 71.5 °C against 76.5 and moves the pump more (28 422 kg/h of travel against 24 842 kg/h), and on five other seeds it completes 3.1 to 5.6 min sooner. The decision at 120 min shows why, as shown in Figure 2.7. From the same state at full cooling, the single plan asks for 363 kg/h and predicts the batch at the ceiling by the end of the horizon, while the controller's own model under that plan never rises above 71.40 °C and ends 1.54 K lower. The

Table 2.7 The MPC Linearized Once and Along Its Plan, Seed 20260926. Times in Minutes, Temperatures in Degrees Celsius

Structure	Dose ends	Batch done	Highest T	Margin (K)	Cuts	QPs	Target
Ch., suggested, MPC, activity 150 %[a]							
one-step optimizer	204.4	213.5	72.19	−0.44	2	—	no
MPC, single	204.9	212.5	72.11	0.01	1	1	no, 0/6
MPC, iterated	204.4	212.0	71.87	0.01	0	2.40	yes, 2/6
single, 1 K back-off	206.0	213.8	71.31	0.02	0	1	yes, 6/6
one relinearization	204.4	212.0	71.87	0.01	0	2	yes, 2/6
Case 10, case 2's batch with water at 8 °C							
fixed profile (case 2)	250.1	266.1	73.56	3.80	0	—	
MPC, single	178.0	202.7	71.85	−0.20	0	1	
MPC, iterated (case 10)	175.4	198.2	71.86	−0.16	0	3.58	
one relinearization	175.5	198.4	71.86	−0.16	0	2	
Scenario 9, water at 30 °C							
MPC, single (case 9)	218.0	231.6	71.87	0.61	0	1	
MPC, iterated	217.2	229.5	71.92	0.95	0	3.81	
one relinearization	212.5	225.5	71.87	1.05	0	2	

[a] The challenge with its suggested settings, the strategy switched to the MPC (the optimizer in the first row) and the catalyst at 150 %, where the MTSR limit is 75 °C. The margin is the limit less the highest true MTSR, QPs counts the programs of a decision while dosing, and the target column gives the default seed and the count over it and five other seeds.

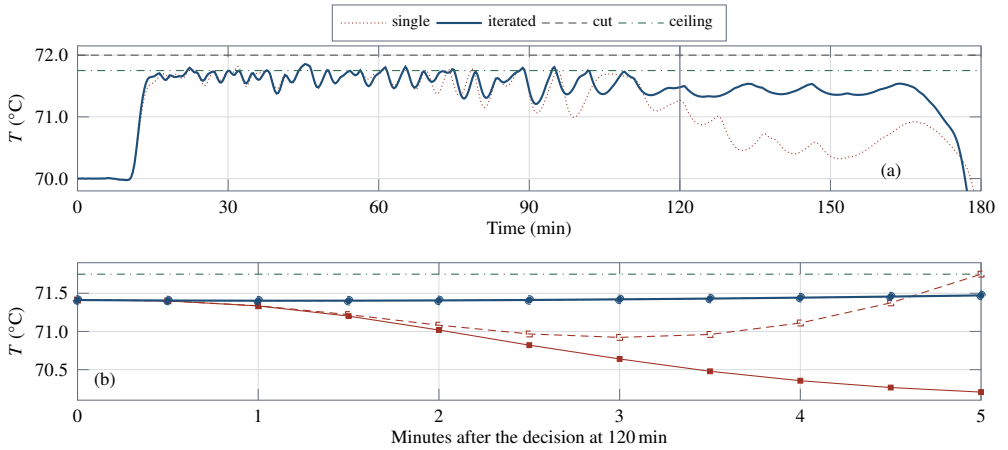


Figure 2.7 Case 10, case 2's batch under the MPC, linearized once and along its plan on the same seed. (a) True batch temperature to 180 min, linearized once (thin dotted line) and along the plan (thick solid line), with the 72 °C cut (dashed) and the 71.75 °C ceiling (dash-dotted). (b) The decision at 120 min, marked in (a), with both plans computed from the same state, squares for the single plan and circles for the iterated one. Dashed lines with open marks are each plan's predicted temperature, and solid lines with filled marks the controller's nonlinear model under that plan.

iterated plan asks for the 600 kg/h pump limit at every step and predicts 71.47 °C, what the model gives. On scenario 9's warm water the batch is complete at 229.5 min against 231.6, and 14 % of the dosing plans keep the single one. One relinearization with no test of convergence does as well on case 10 (198.4 min) and better on scenario 9 (225.5 min) at about half the cost (0.31 ms), but on one of scenario 9's five other seeds it is cut once, where neither linearization is.

Under the worst cases the application can produce, 39 cases run under every structure (every automatic scenario, each disturbance at the limits of its range on case 2's batch and the suggested challenge, the cooling lost for good at 1, 2 and 3 h, a second ten-minute failure every 30 min of the

challenge), the iterated MPC reaches no level the single one does not. Both peak at 74.64 °C, their highest true MTSR is 102.98 and 102.94 °C on case 6's fooled calorimeter, each passes the limit by more than 1 K in three runs, and neither reaches T_{D24}, the TAH or a runaway. Running closer to the ceiling, the iterated batch meets an unplanned event hotter. With the cooling lost for good at 2 h on case 2's batch it peaks at 73.16 °C against 72.74 °C, both cut once.

2.7 Technology in Practice

What Measures It

The dosing loop of this chapter controls a quantity that no instrument reads. QI-105 builds it from two temperatures, a feed flow and a heat-transfer coefficient it assumes, which is what a heat-flow calorimeter does in the laboratory. There the temperature difference between reactor and jacket measures the rate of heat production, and the accumulation of unreacted material in a semi-batch reaction can be measured [7]. Where the heat evolution can be measured or estimated online, feedback can hold it at a setpoint [26]. The two temperature elements are those of the cooled reactor of Volume 7, Chapter 2, here with lags of 15 and 6 s, and the feed meter FT-103, with its 3 s lag and 0.3 % of noise, supplies the total that ends the dosing phase.

> **TECHNOLOGY Sensor: A Coriolis Meter on the Feed**
>
> A Coriolis meter shakes one or more tubes carrying the liquid at their resonant frequency. The inertia of the flowing liquid twists the tubes in proportion to its mass flow, so the meter measures mass flow directly and, from the frequency, the density as well. It needs no straight pipe before or after it and is very accurate and dependable, which is why it is used where the flow is product being bought and sold. Its weaknesses are cost in large sizes, vibration from nearby machines, which disturbs the tubes' own motion, and gas bubbles trapped in bent tubes, which spread the mass unevenly and give errors in flow and density [27]. In this chapter the meter's total ends the dosing phase and enters the calorimeter's count of B, so a meter error passes straight into the estimated accumulation.

What Moves It

Two final elements act on the batch. The split-range pair on the jacket loop is the arrangement drawn for Volume 7, Chapter 2, here with a steeper slope on the hot side than on the cooling side (Section 2.3). The other is the dosing pump. Industrial practice often runs a semi-batch reactor isothermally and controls its temperature by the feed rate of the limiting reactant, which is hard because a larger cold feed first lowers the temperature before the faster reaction raises it [26]. The temperature-limited strategy of this chapter cuts the feed on the temperature in that spirit, while the accumulation-limited one sets it from the estimate instead.

> **TECHNOLOGY Final Element: A Metering Pump**
>
> A small dosing stream is often moved by a metering pump, a pump mechanism, motor and drive electronics in one package, whose 4 to 20 mA control signal sets the flow from zero to a maximum chosen on the pump. Many use a plunger driven at variable intervals, so the flow arrives as discrete pulses rather than a continuous stream. A positive-displacement pump discharges a fixed volume for every revolution, and with its discharge blocked it locks and overloads its motor unless a relief path returns the flow [27]. The chapter's pump P-103 delivers the requested feed up to 600 kg/h with no dynamics of its own, and the interlocks stop it directly.

Where the Controller Runs

Batch control runs on stand-alone single-loop controllers, programmable logic controllers, distributed control systems and personal computers, often in combination [26]. The batch control committee of ISA, whose standards form the ISA-88 series, describes a functional model that relates five types of control: recipe management, scheduling, sequential control, regulatory control and safety interlock

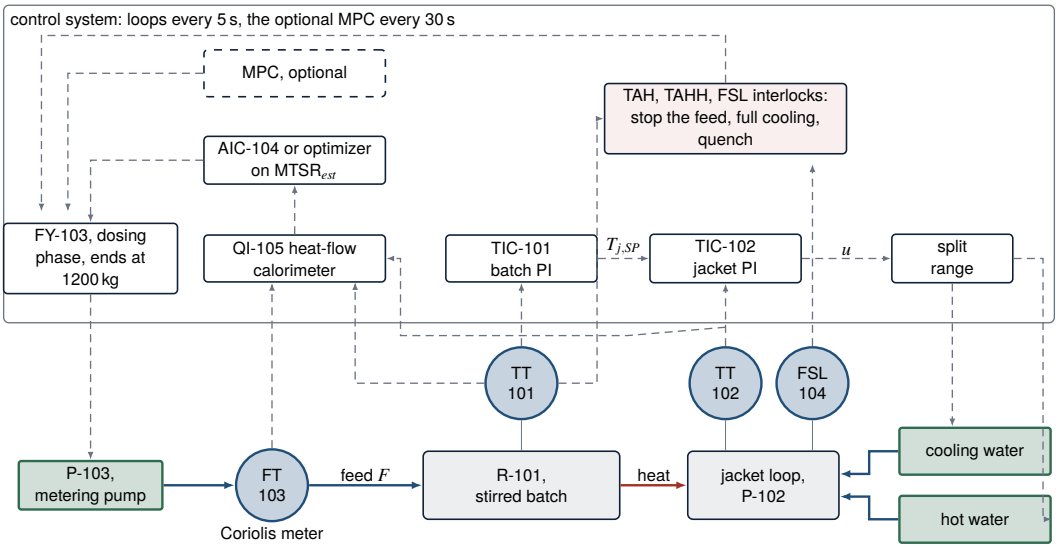


Figure 2.8 The chapter's loops on plant hardware. The cascade TIC-101 and TIC-102 reads the batch and jacket temperatures and moves the cooling-water and hot-water valves through the split range. The calorimeter QI-105 reads both temperatures and the Coriolis meter FT-103, and AIC-104, the one-step optimizer or the dashed optional MPC sets the feed request of the dosing phase, which the metering pump delivers until the measured total reaches 1200 kg. Notice that the interlocks read the same TT-101 as the cascade and act on the same feed, with no second sensor and no separate valve.

systems [28]. In those terms the cascade TIC-101 and TIC-102 is regulatory control, the dosing strategy and the measured total that ends the dose are the sequential control of the dosing phase, and TAH, TAHH and FSL-104 are the safety interlocks. A distributed control system executes such loops with cycle times from milliseconds to minutes [29]. The chapter's loops act every 5 s and the predictive controller every 30 s, as Figure 2.8 draws them on plant hardware.

IN PRACTICE Three Ways to Stop the Feed

A guide of the British safety regulator compares three ways to protect a semi-batch reactor whose temperature is safety critical. In the first an operator watches one thermocouple and shuts off the addition by hand near the upper limit. In the second a temperature controller throttles the reactant valve and an alarm calls for a manual valve to be closed. In the third an independent temperature switch, set below the upper limit, is hardwired to a separate shut-off valve in the reactant line, and sensors of different types avoid a common-mode failure. The guide points to IEC 61508 and IEC 61511 for the integrity such systems need [7].

What Is Modernizing It

Two developments are changing how such a batch is dosed. The first is optimization repeated from batch to batch. On a one-ton industrial reactor, a policy that ran isothermally and then let the reaction finish without jacket cooling, with the switching time adapted from run to run, cut the reaction time by 35 % against the isothermal policy of industrial practice. It did so after only two batches and without off-specification product, keeping a back-off from the final temperature limit for safety [26]. That is the run-to-run idea of Section 1.4.5 applied to a recipe. The second is predictive control on a process model, which the same chapter counts among the advances that have influenced industrial implementation [26], and which the MPC of this chapter brings to the dose on the balances of Equation (2.12).

2.8 Limitations

The kinetics, the heat-transfer law, the fouling and agitator factors and the decomposition parameters are illustrative, and T_{D24}, MTT, the MTSR margins and the class thresholds are teaching values, so the app demonstrates Stoessel's method [6] without assessing any real process. One synthesis and one first-order decomposition leave out induction periods, autocatalysis and side reactions, and the model ends at the rupture disc without pressure, venting or boiling.

The calorimeter keeps the declared law and constants, so its only systematic errors are those a reader introduces: fouling, the agitator, a wrong correction, an unmeasured feed strength or a changed parameter. It attributes all heat to the synthesis and cannot see a decomposition, and it is never recalibrated during a batch.

OPEN CHALLENGE Models Good Enough and Small Enough

A 2017 handbook chapter on batch control found that predictive control and optimization have influenced industrial practice, but that they depend on process models of reasonably high fidelity and modest complexity, which are difficult to achieve. It called for reduced-order models built for control, and for better sensors and analyzers, because estimates can never completely replace measurements [26]. This chapter's calorimeter and predictive controller both keep a declared heat-transfer law that a fouled jacket makes wrong. Progress would be the calorimeter the Conclusion asks for, one that recalibrates its conductance during the batch.

The heat loss and the water in the jacket give a real batch some margin that the adiabatic criterion does not count. With the circulation lost for good at 2 h 30 min (case 3 with the setpoint changed and the quench disarmed before Run, and the open-ended cooling failure fired at that time), a batch dosed at 35 °C survives with MTSR at 148.7 °C, although it passes T_{D24}, while one dosed at 34 °C, with MTSR at 151.7 °C, reaches the disc at 12 h 45 min. The criterion is conservative on purpose. The interlocks act on TT-101 with its lag and have no redundancy or voting, and the optimizer looks one step ahead.

The predictive controller's model leaves out fouling, the agitator, the feed strength, circulation below design, the decomposition and the quench mass, and its 22 s jacket lag is about four times too fast with water at 30 °C. It inherits the calorimeter's errors, and on a fouled jacket its estimates settle at $\hat{\theta} = 1.86$ and $\hat{T}_{cw} = 4.0\,°\mathrm{C}$. Its 5 min horizon does not plan the end of the dose, and its one slack

per band cannot price a violation step by step. It trades 5.8 min of the warm-water batch for the 21 cuts it avoids, and in case 7 and the suggested challenge a classical setting finishes sooner.

The iterated linearization makes the plan exact about the controller's model, not the plant, so it removes none of this mismatch. Three relinearizations leave 2 % of case 10's dosing plans and 14 % of scenario 9's unsettled, where a fourth, outside the declared rule, would leave 0.6 and 6.6 %.

2.9 Conclusion

A semi-batch reactor stores its hazard in the reagent that has been fed and has not yet reacted, and this simulator logs that accumulation, its MTSR and Stoessel's class beside a temperature trend that cannot show them. At 70 °C the fixed four-hour dose never holds more than 234 mol, while at 30 °C the same dose holds its temperature within 0.2 K and accumulates 4398 mol by 2 h 30 min, an MTSR of 160.4 °C. Losing the circulation then runs the batch through a synthesis runaway to 157.1 °C and a decomposition to its rupture disc at 3 h 55 min, while accumulation-limited dosing meets the same failure with MTSR held at 80.0 °C and a peak of 74.6 °C. The highest MTSR of the dose is lowest, 75.6 °C, near 66 °C. A 15 % fouled jacket leaves the calorimeter reporting 1192 mol where 2054 mol are, and the challenge improves from a 22.76 K excursion to a dose finished at 3 h 43 min within 1.61 K. On warm cooling water the optional predictive controller is never cut where the one-step optimizer is cut 21 times, and it completes the batch 5.8 min later. Relinearized along its own plan it ties with one linearization where MTSR binds, because MTSR is linear in the unreacted B, and completes case 10's batch 4.5 min sooner at the same 71.9 °C peak. The kinetics and the thresholds are teaching values, and the decomposition is a single anchored reaction. A fuller model would add an autocatalytic decomposition, a vent model at the disc, and a calorimeter that recalibrates its conductance. The simulator makes the accumulation of a fed reactor visible, and with it the reason that the dosing, not the temperature loop, keeps the batch safe.

KEY POINTS

- A semi-batch reactor stores its hazard in the reagent that has been fed and has not yet reacted, and its temperature trend cannot show it.
- The same four-hour dose holds at most 234 mol unreacted at 70 °C, but 4398 mol by 2 h 30 min at 30 °C, an MTSR of 160.4 °C.
- Accumulation-limited dosing meets the same cooling failure with MTSR held at 80.0 °C and a peak of 74.6 °C.
- A jacket fouled by 15 % leaves the calorimeter reporting 1192 mol where 2054 mol are.

Where else this applies. An artificial pancreas limits the insulin already delivered whose effect has not yet appeared, the insulin on board of Volume 19, Chapter 2, as AIC-104 limits the B that has been fed and has not yet reacted.

References

[1] P. Hugo and J. Steinbach, "A comparison of the limits of safe operation of a SBR and a CSTR," *Chemical Engineering Science*, vol. 41, no. 4, pp. 1081–1087, 1986.

[2] M. Steensma and K. R. Westerterp, "Thermally safe operation of a semibatch reactor for liquid-liquid reactions. Slow reactions," *Industrial & Engineering Chemistry Research*, vol. 29, no. 7, pp. 1259–1270, 1990.

[3] K. R. Westerterp and E. J. Molga, "Safety and runaway prevention in batch and semibatch reactors—A review," *Chemical Engineering Research and Design*, vol. 84, no. 7, pp. 543–552, 2006.

[4] P. Lerena, W. Wehner, H. Weber, and F. Stoessel, "Assessment of hazards linked to accumulation in semi-batch reactors," *Thermochimica Acta*, vol. 289, no. 2, pp. 127–142, 1996.

[5] R. Gygax, "Chemical reaction engineering for safety," *Chemical Engineering Science*, vol. 43, no. 8, pp. 1759–1771, 1988.

[6] F. Stoessel, *Thermal Safety of Chemical Processes: Risk Assessment and Process Design*. Wiley, 2008.

[7] Health and Safety Executive, *Designing and Operating Safe Chemical Reaction Processes*, ser. HSG143. Sudbury, UK: HSE Books, 2000. [Online]. Available: https://www.hse.gov.uk/pubns/priced/hsg143.pdf

[8] I. S. de Buruaga, A. Echevarría, P. D. Armitage, J. C. de la Cal, J. R. Leiza, and J. M. Asua, "On-line control of a semibatch emulsion polymerization reactor based on calorimetry," *AIChE Journal*, vol. 43, no. 4, pp. 1069–1081, 1997.

[9] F. Rossi, S. Copelli, A. Colombo, C. Pirola, and F. Manenti, "Online model-based optimization and control for the combined optimal operation and runaway prediction and prevention in (fed-)batch systems," *Chemical Engineering Science*, vol. 138, pp. 760–771, 2015.

[10] H. Schuler and C.-U. Schmidt, "Calorimetric-state estimators for chemical reactor diagnosis and control: review of methods and applications," *Chemical Engineering Science*, vol. 47, no. 4, pp. 899–913, 1992.

[11] O. Ubrich, B. Srinivasan, P. Lerena, D. Bonvin, and F. Stoessel, "Optimal feed profile for a second order reaction in a semi-batch reactor under safety constraints," *Journal of Loss Prevention in the Process Industries*, vol. 12, no. 6, pp. 485–493, 1999.

[12] ——, "The use of calorimetry for on-line optimisation of isothermal semi-batch reactors," *Chemical Engineering Science*, vol. 56, no. 17, pp. 5147–5156, 2001.

[13] D. I. Townsend and J. C. Tou, "Thermal hazard evaluation by an accelerating rate calorimeter," *Thermochimica Acta*, vol. 37, no. 1, pp. 1–30, 1980.

[14] A. Keller, D. Stark, H. Fierz, E. Heinzle, and K. Hungerbühler, "Estimation of the time to maximum rate using dynamic DSC experiments," *Journal of Loss Prevention in the Process Industries*, vol. 10, no. 1, pp. 31–41, 1997.

[15] W. Regenass, "The development of heat flow calorimetry as a tool for process optimization and process safety," *Journal of Thermal Analysis*, vol. 49, no. 3, pp. 1661–1675, 1997.

[16] U.S. Chemical Safety and Hazard Investigation Board, "T2 Laboratories, Inc. runaway reaction," U.S. Chemical Safety and Hazard Investigation Board, Investigation Report 2008-3-I-FL, Sep. 2009. [Online]. Available: https://www.csb.gov/t2-laboratories-inc-reactive-chemical-explosion/

[17] ——, "Investigation report: Pesticide chemical runaway reaction pressure vessel explosion, Bayer CropScience, LP, Institute, West Virginia, August 28, 2008," U.S. Chemical Safety and Hazard Investigation Board, Washington, DC, Tech. Rep. 2008-08-I-WV, Jan. 2011. [Online]. Available: https://www.csb.gov/assets/recommendation/bayer_report_final.pdf

[18] 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.

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

[20] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[21] A. Kummer, T. Varga, and L. Nagy, "Semi-batch reactor control with NMPC avoiding thermal runaway," *Computers & Chemical Engineering*, vol. 134, p. 106694, 2020.

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

[23] J. H. Lee and N. L. Ricker, "Extended kalman filter based nonlinear model predictive control," *Industrial & Engineering Chemistry Research*, vol. 33, no. 6, pp. 1530–1541, 1994.

[24] M. Diehl, H. G. Bock, J. P. Schlöder, R. Findeisen, Z. Nagy, and F. Allgöwer, "Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations," *Journal of Process Control*, vol. 12, no. 4, pp. 577–585, 2002.

[25] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, "From linear to nonlinear MPC: bridging the gap via the real-time iteration," *International Journal of Control*, vol. 93, no. 1, pp. 62–80, 2020.

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

[27] 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.

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

[29] 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.

2.10 Exercises

2.1 Concept check · 2 min

The same four-hour dose of B is run into one batch held at 70 °C and into another held at 30 °C. Both temperature loops hold their setpoints within a fraction of a kelvin. Which statement about a cooling failure halfway through the dose is right?

(a) The cold batch is safer, because it stays 40 K further from T_{D24} during the dose

(b) The two batches are equally safe, because both hold their setpoints within a fraction of a kelvin

(c) The cold batch is more dangerous, because its reaction lags the feed and the unreacted B raises its MTSR

(d) The cold batch is more dangerous, because its jacket runs near full cooling and has no reserve left

2.2 Concept check · 2 min

Since the calorimeter QI-105 was calibrated, the jacket has fouled and lost 15 % of its heat-transfer coefficient. The calorimeter still uses the clean-jacket law, and AIC-104 doses on its estimate. What happens?

(a) The calorimeter reports too much unreacted B, so AIC-104 slows the feed more than it needs to

(b) The calorimeter reports the right unreacted B, because the error in UA cancels in the heat balance

(c) The calorimeter reports too little reaction heat, so the dose finishes later than it should

(d) The calorimeter reports too little unreacted B, so AIC-104 feeds faster than the true MTSR allows

2.3 Concept check · 1 min

In case 7 the catalyst runs at a quarter of its activity, and the dose is temperature-limited: the feed is cut while TT-101 is more than 2 K above the setpoint. Why does the cut act only twice in the whole dose?

(a) A slower reaction releases its heat later rather than hotter, so the batch seldom rises 2 K above its setpoint

(b) The cut is disabled by temperature-limited dosing while the calorimeter reports a weak catalyst

(c) The batch loop raises its own gain when the catalyst is weak, which holds the batch near the setpoint

(d) The weak catalyst lowers the heat capacity of the batch, so the same heat raises it less

2.4 Theory and calculation · 5 min

The recipe doses 8000 mol of B, with $(-\Delta H_r) = 200\,\mathrm{kJ/mol}$, into a batch whose heat capacity grows from 5.40 MJ/K for the charge to 7.56 MJ/K when the dose is complete.

(a) Find the heat the whole dose carries and the temperature rise of the finished batch if it were all released at once with no heat removed. *Checkpoint:* 1.6 GJ and 211.6 K.

(b) Find the rise of MTSR for every 1000 mol of unreacted B, in the finished batch and in the charge. *Checkpoint:* 26.5 K and 37.0 K.

(c) A batch near the end of its dose sits at 70 °C. How much unreacted B would put its MTSR at $T_{D24} = 130\,^{\circ}\mathrm{C}$? *Checkpoint:* 2268 mol.

2.5 Theory and calculation · 6 min

Use the quasi-steady estimate of Equation (2.9) with $k_{70} = 2.4 \times 10^{-6}\,\mathrm{m^3/(mol\,s)}$, $E_r = 120\,\mathrm{kJ/mol}$, a feed of 300 kg/h carrying 6.667 mol/kg, and the state at the end of the dose: $V = 4.53\,\mathrm{m^3}$, $n_A = 4234\,\mathrm{mol}$ and $mc_p = 7.56\,\mathrm{MJ/K}$.

(a) Estimate the unreacted B and the MTSR at the end of the dose at 70 °C. *Checkpoint:* About 248 mol and 76.6 °C.

(b) Repeat the estimate at 30 °C, where the rate constant is 257 times smaller, and say what the result means. *Checkpoint:* About 64 000 mol, more than the whole dose.

(c) Differentiate the quasi-steady MTSR with respect to T and show that it is lowest where $\Delta T_{ad,acc} = RT_K^2/E_r$. Evaluate this at 67 °C. *Checkpoint:* 8.0 K.

2.6 Theory and calculation · 5 min

At the start of the dose the wall conductance is $UA = 5.00\,\mathrm{kW/K}$, and the jacket circulates 2 kg/s of water with $c_{pj} = 4180\,\mathrm{J/(kg\,K)}$. Each kilogram of feed carries 6.667 mol of B at 200 kJ/mol.

(a) Find the effective conductance of wall and circulation in series, and the most heat full cooling can remove from a batch at 70 °C with cooling water at 8 °C and at 30 °C. *Checkpoint:* 3.13 kW/K, 194 kW and 125 kW.

(b) Case 8 doses the 1200 kg in three hours. Find the heat this feed makes once it reacts as it arrives, and the warmest cooling water that could still remove it at 70 °C. *Checkpoint:* 148 kW and about 22.6 °C.

(c) Find the highest steady feed the warm jacket can carry at 70 °C and the dose time it implies. *Checkpoint:* About 338 kg/h and 3.55 h.

2.7 App challenge · 9 min

Use the default seed and a speed of 60 for both scenarios.

(a) Load scenario 4, Cooling failure at the wrong moment, change no setting and run until the run stops. Read the Batch and Most unreacted B tiles. *Checkpoint:* Lost at 3 h 55 min, 4398 mol.

(b) Load scenario 5, The same failure, accumulation-limited, change no setting and run to 5 h (300 min). Read the Highest T tile, the highest class on the Safety view, and the fed mass, column `fed_kg`, at 2 h 30 min. *Checkpoint:* 74.6 °C, class 2, 276 kg.

Explain why the feed interlock FSL-104, which stops the feed in both runs when the circulation fails, does not decide the outcome.

2.8 App challenge · 8 min

Load scenario 6, Fouled jacket, fooled calorimeter. Use the default seed and a speed of 60, and export the file after each run.

(a) Run until the Dose finished tile shows a time. In the first row with column `dose_done` equal to 1, read the true and estimated unreacted B (`nB_mol`, `nB_est_mol`), and read the Highest MTSR tile. *Checkpoint:* 3 h 42 min, 2054 mol true against 1192 mol estimated, 104.1 °C.

(b) Predict the calorimeter's error at that row from $-\int (UA_{est} - UA)(T - T_j)\,\mathrm{d}t/(-\Delta H_r)$, summing over the rows with columns `UA_est_kWK`, `UA_kWK`, `T_meas_C` and `Tj_meas_C`. *Checkpoint:* About −863 mol, against −862 mol from (a).

(c) Press Reset, set UA used by the calorimeter to True UA before Run, and run until the dose finishes. Read the Dose finished and Highest MTSR tiles. *Checkpoint:* 5 h 54 min and 80.5 °C.

Accept the unreacted B within 5 mol and the predicted error within 2 %.

2.9 App challenge · 8 min

Load scenario 9, MPC: the temperature cut, planned ahead, change no setting, and run at a speed of 60 with the default seed until the run stops at the end of the batch. Read the Dose finished, Batch, and Highest T tiles, and count the FEED CUT marks on the batch temperature trend (or the starts of column `feed_cut` in the exported file). Then press Reset, set Dosing strategy to One-step optimizer before Run, and repeat. Which controller finishes sooner, and what does each pay?

2.10 Application area · 3 min

The disc of R-101 is replaced by one with a lower set pressure, at which the solvent boils at MTT = 120 °C, below $T_{D24} = 130$ °C. Classify, with the definitions of the chapter, the recipe batch at 70 °C whose MTSR never passes 76.2 °C, and the cold batch of case 3 at 2 h 30 min, with MTSR at 160.4 °C. What does the new order of the four temperatures mean for a cooling failure?

2.11 Application area · 4 min

In case 4 with the quench armed, the high-high trip adds 1000 kg of solvent at 20 °C to a batch of 3750 kg at 156.7 °C. The batch and the solvent have the same heat capacity per kilogram, the charge's density is 900 kg/m^3 and the feed's 1000 kg/m^3, and 750 kg of feed is in the batch.

(a) Find the temperature after the quench. *Checkpoint:* 127.9 °C.

(b) Find the quench mass that would bring the batch to 120 °C, and the liquid volume with each quench in the 6 m^3 vessel. *Checkpoint:* About 1376 kg; 5.19 and 5.61 m^3.

2.12 Review the agent · 5 min

An AI assistant reviews case 6 and writes: "The accumulation-limited dose overshoots its MTSR limit by 24 K because AIC-104 is a PI loop that reacts too late. Switch the dosing strategy to MPC and set its linearization to Iterated. A converged plan predicts exactly what the nonlinear model gives for it, so the MPC will hold the true MTSR at the 80 °C limit." Find the flaw and say how you would check it.