INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 2

Exothermic Reactor Simulator for Temperature Control Across an Ignition Boundary

Pages 22–52

Edited by John D. Hedengren

How to cite this chapter. "Exothermic Reactor Simulator for Temperature Control Across an Ignition Boundary," in *Nonlinear Dynamics, Oscillations, and Stabilization*, John D. Hedengren, Ed., vol. 7 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 22–52.

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

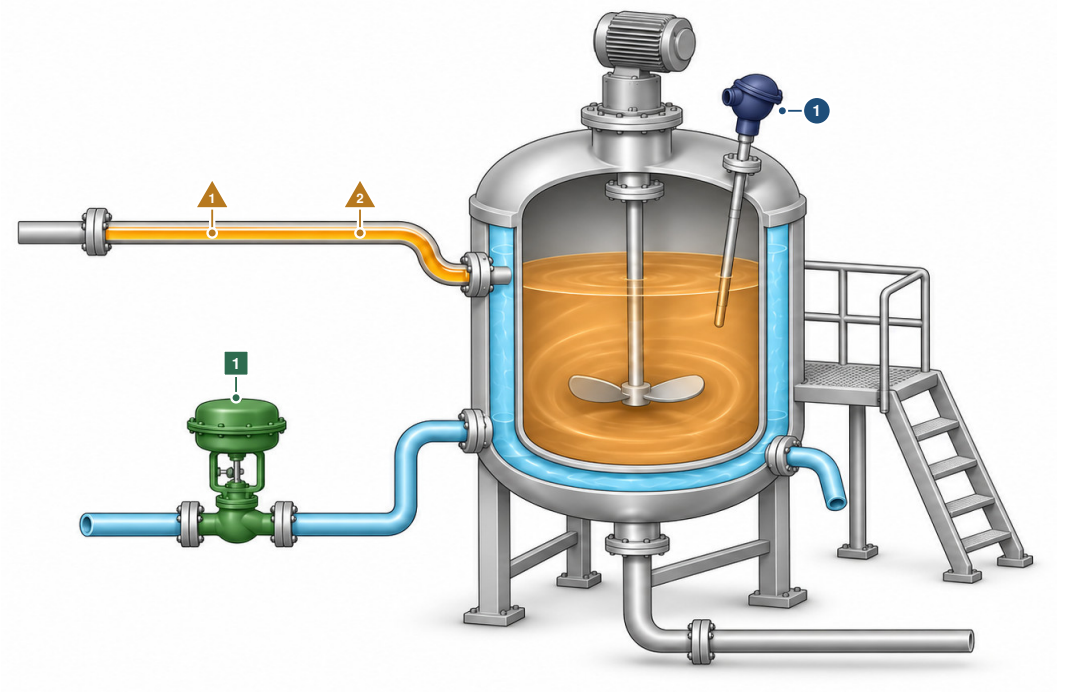
The app for this chapter
https://apmonitor.com/apps/reactor/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v07-nonlinear-stabilization/#reactor

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

CHAPTER 2

Exothermic Reactor Simulator for Temperature Control Across an Ignition Boundary



The controller moves 1 valve that sets the jacket temperature

The controller reads 1 temperature of the stirred liquid

What disturbs it 1 warmer feed 2 richer feed, more reactant in it

Notice that the jacket around the liquid is the controller's only way to remove the heat the reaction releases, and the probe in the liquid is all it reads. Illustration, not to scale.

IN PLAIN TERMS A Reactor That Cannot Rest Where It Works

The main idea. A reactor turns liquids into a chemical, and the reaction gives off heat that a cooling jacket around the tank carries away. The hotter the liquid gets, the faster it reacts and the more heat it gives off.

The challenge. At its usual working temperature the tank cannot stay put on its own, so a controller must keep moving the jacket. A target raised one step too far sends it past its alarm before the controller brings it back.

Why it matters. Real runaway accidents come from the same race between heat made and heat removed.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on reactive incidents (Section 2.1), the history in Section 2.2, the fact on the pendulum (Section 2.4.2), the two try-its in Section 2.7 and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 2.7), Technology in Practice (Section 2.10), Limitations (Section 2.11) and the open challenge

The control question Heat generation in an exothermic reactor grows exponentially with temperature while jacket cooling grows only linearly, so the default setpoint is open-loop unstable and a jacket change of a few kelvin can ignite the vessel. How large a setpoint move can one fixed PI tuning make before the reactor runs past its alarm, and what does a predictive controller built on the linearized balances add?

Learning objectives

- Locate the steady states, the two folds and the Hopf point from the van Heerden construction and the steady-state locus.
- Evaluate the Jacobian at an operating point and classify each steady state from its eigenvalues.
- Explain why a reverse-acting loop holds a branch whose steady-state gain is negative.
- Identify the cold branch from a jacket step test and tell its chord gain from the tangent gain.
- Measure the largest safe setpoint step, and judge detuning, the predictive controller and the moving horizon estimator against it.

Prerequisites Sections 1.4.2 (folds), 1.4.3 (sign from the immediate response), 1.4.4 (eigenvalues), 1.4.9 (Hopf point) and 1.6.2 (velocity form). Section 1.6.3 is needed only for the predictive controller and its estimators.

The app https://apmonitor.com/apps/reactor/

See also Volume 9, Chapter 4 for an offset-free predictive controller whose Kalman filter carries an input bias, and Volume 9, Chapter 3 for a step test that returns a different plant at each operating point. Volume 13, Chapter 4 and Volume 13, Chapter 9 for ignition and branch multiplicity in chemical plants, and Volume 8, Chapter 2 for a runaway held off by dosing.

2.1 Introduction

Most of the plants a first control course meets have a process gain that can be written down once. A tank fills faster when the valve opens further, a heater raises a temperature in proportion to the power it draws, and the constant of proportionality changes slowly enough that a controller commissioned at one operating point still works at another. An exothermic reactor is not that kind of plant. The heat it releases carries the Arrhenius factor $e^{-E/RT}$, so generation climbs exponentially with the very variable that the cooling jacket is trying to hold, while removal through the jacket climbs only in proportion to the temperature difference across the wall [1]. Two curves of that shape can cross once, twice or three times. Where they cross three times the middle crossing is a steady state that no operator can hold without feedback [2, 3], and the region between the outer crossings is a set of operating temperatures that exist only while a controller holds them.

+ DID YOU KNOW? 167 Reactive Incidents, 108 Deaths

From January 1980 to June 2001 the United States Chemical Safety Board counted 167 serious incidents of uncontrolled chemical reactivity, 48 of which killed 108 people. It attributed 35 percent of them to runaway reactions, in which heat is generated faster than cooling media and surroundings can remove it [4], the crossing of the two curves in Figure 2.2(a).

The consequence for control is sharper than the usual warning about nonlinearity. The steady-state gain of the reactor of this chapter passes through $+1.72$ K/K at its nominal operating point, rises without bound at 335.65 K, reappears as -5.56 K/K a few kelvin higher, and returns to $+2.38$ K/K on the hot branch. A proportional gain that gives a well-damped response at low conversion is the same number that produces an excursion of more than 110 K at high conversion, and the excursion happens on the way to a setpoint that the loop then holds perfectly. Runaway in this plant is a property of the transient rather than of the destination.

Exothermic reactors have generated a literature of their own because of that behavior. Steady-state multiplicity and its bifurcation structure were mapped in detail for the single first-order reaction, together with the limit cycles that appear once the ignited state loses stability [5]. Control work has

treated the reactor as the standard hard case for input-output linearization [6], for robust stabilization by classical proportional-integral feedback [7], for gain scheduling across an operating envelope [8], and for constrained predictive control written specifically to keep a batch away from its runaway boundary [9], the problem Volume 8, Chapter 2 takes up for a semi-batch reactor whose dosing is the safety layer. Industrial practice adds a layer the theory usually leaves out, because the jacket temperature that these formulations manipulate is itself the setpoint of a slave loop on coolant flow or coolant mixing [10]. However, a student meeting the reactor for the first time in a course has no way to watch the ignition happen, to see the same controller succeed and then fail on a larger setpoint change, or to find the jacket temperature at which the cold operating point stops existing.

The model rests on two conservation statements over one well-mixed volume, a mole balance on the reactant and an energy balance whose exotherm term is nonlinear in the state that the balance itself determines. The dilution rate, the adiabatic temperature rise and the jacket rate constant are the only three lumped groups in the model, and each is a transport rate divided by a capacity, which is the way a first course builds a dynamic model [11, 12]. What this plant adds to that familiar structure is steady-state multiplicity, an open-loop-unstable branch, a process gain that changes sign, and an actuator that reaches both of its limits inside one transient. The behavior matters industrially because surveys of thermal-runaway incidents keep finding inadequate temperature control and inadequate knowledge of the heat release among the recurring causes [13, 14].

This chapter specifies an interactive exothermic stirred-reactor simulation in enough detail to rebuild it, and reports what its shipped constants actually do rather than what the textbook case is usually drawn as. Section 2.2 sets out the safety record and the cascade that frame the problem, and Section 2.3 gives the balances, the constants and the adaptive explicit integrator. The characterization in Sections 2.4 and 2.5 locates the two fold points, the Hopf point on the hot branch and the open-loop-unstable temperature interval, shows that the ignited branch at the published jacket temperature is not an attractor, and identifies the cold branch from a step test. Section 2.6 describes the PI loop, a linear model predictive controller, its two estimators and a relay experiment. Section 2.7 measures one controller at eight setpoints, where it changes character between two of them, and against two upsets that ignite the reactor in manual. The predictive controller is compared with the shipped loop in Section 2.8, and Section 2.9 hands it a reactor that has already run away, under two estimators of the concentration it cannot measure, while Section 2.10 describes the plant's instruments, valves and control system. The application's eleven guided cases each carry a completion requirement (Section 1.6.4). Sections 2.11 to 2.13 state the limitations and place the results, and Section 2.14 holds the exercises.

2.2 Background

The instrument that matters in an exothermic reactor is the one that reads the liquid temperature, and the history of reactor safety is largely the history of learning what that reading does and does not say. Early continuous processes were run with the operator judging the reaction from the jacket steam or cooling water and from the product, because a thermowell reading inside a stirred vessel lags the hottest part of the contents and gives no warning at all of reactant that has accumulated without reacting. Loss prevention emerged as a discipline in the chemical industry only after a sequence of events made the point that hazard is a property of the process rather than of the care taken by the people running it, and the systematic study of what had actually gone wrong began in the 1960s and 1970s [15].

Runaway reactions have kept their place in that record. A survey of thermal-runaway incidents investigated in the United Kingdom found that the causes clustered in a small number of repeated failures, including inadequate knowledge of the reaction chemistry and its heat release, inadequate temperature control, and inadequate agitation, rather than in exotic mechanisms [13]. A later review across a wider set of incidents reached the same conclusion and added that the majority arose during normal production rather than during commissioning or maintenance [14]. What both surveys

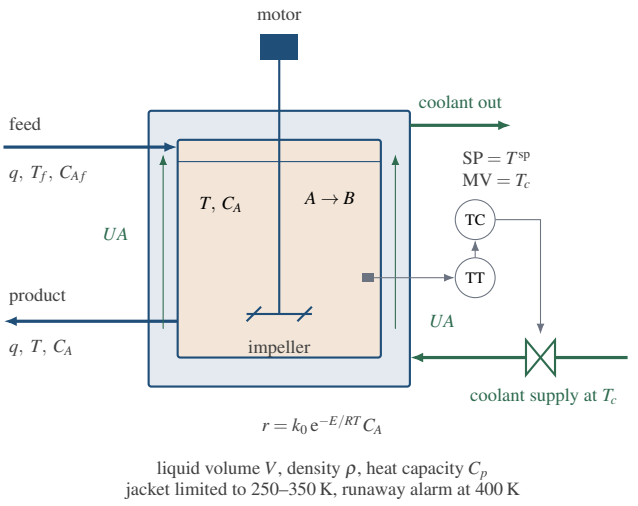


Figure 2.1 Jacketed continuous stirred-tank reactor as simulated. The temperature transmitter reads the well-mixed liquid and the controller sets the jacket temperature directly, so the coolant valve and its slave loop are treated as ideal.

describe is a plant whose heat generation exceeded its heat removal capacity for long enough that no control action could recover it, which is exactly the crossing of the two curves that the theory of autothermic operation had already described [1].

> ⧖ HISTORY **2007: Lost Cooling and a New Course Requirement**
>
> In December 2007 a 2450-gallon batch reactor making a gasoline additive in Florida burst and its contents exploded, killing four people and injuring 32. The United States Chemical Safety Board found that a runaway exothermic reaction had occurred and that a loss of sufficient cooling had most likely caused it. The cooling system lacked the redundancy to survive a single failure, and the root cause was that the company had not recognized the runaway hazard of the chemical it was making [16]. The board asked for reactive hazard awareness to be added to the requirements of chemical engineering degrees. In October 2011 the accreditation board for engineering programs approved changes that require proficiency in all chemical process hazards [17].

The control response to that reading developed alongside it. Reactor stability was placed on an analytical footing by linearizing the coupled mass and energy balances and examining the eigenvalues of the resulting matrix [2], and the extension to a controlled reactor followed immediately [3]. Modern reactor temperature control is structured as a cascade in which the master loop compares the reactor temperature against setpoint and asks for a jacket temperature, while a slave loop delivers that jacket temperature by mixing hot and cold utility or by throttling coolant [10]. The application described here implements the master loop and treats the jacket temperature as though the slave loop were perfect, which is the standard first approximation and the one that makes the thermal behavior visible without a second set of tuning constants.

2.3 Reactor Model

The simulated unit is a perfectly mixed, constant-volume, liquid-phase vessel with a cooling jacket, an agitator, one feed stream and one product stream drawn at the same volumetric rate. A single irreversible first-order reaction $A \rightarrow B$ runs in the liquid and releases heat. The arrangement, together with each symbol used in the balances that follow, is shown in Figure 2.1.

Table 2.1 Variable Classification, Units and Ranges

Symbol	Role	Unit	Range or default
T	CV, state	K	200 to 900 (clamped)
C_A	state, shown	mol/m^3	0 to 20 (clamped)
$\hat{T}$	measurement	K	T plus noise
T_c	MV	K	250 to 350
T^{sp}	setpoint	K	280 to 450, default 340
T_f	DV	K	300 to 380, default 350
C_{Af}	DV	mol/m^3	0.5 to 2, default 1
q	parameter	vol/min	50 to 200, default 100
r	derived	mol/m^3-min	Equation (2.1)
X	derived	–	$(C_{Af}-C_A)/C_{Af}$

2.3.1 Variables

The reactor temperature is the only controlled variable and the jacket temperature is the only manipulated variable, which makes this a single-input single-output plant with two states. The classification is given in Table 2.1.

2.3.2 Governing Equations

The reaction rate per unit volume follows the Arrhenius law with a first-order dependence on the reactant concentration, as shown in Equation (2.1).

$$r = k_0\, \mathrm{e}^{-E/(RT)}\, C_A \tag{2.1}$$

where r is the rate of consumption of A in mol/m^3-min, k_0 is the pre-exponential factor in 1/min, E/R is the activation energy divided by the gas constant in K, T is the liquid temperature in K and C_A is the reactant concentration in mol/m^3. The exponential is evaluated at $\max(T, 1\,\mathrm{K})$ so that the argument can never divide by zero.

A mole balance on A over the constant liquid volume gives Equation (2.2), and an energy balance on the same volume gives Equation (2.3).

$$\frac{\mathrm{d}C_A}{\mathrm{d}t} = \frac{q}{V}\left(C_{Af} - C_A\right) - r \tag{2.2}$$

$$\frac{\mathrm{d}T}{\mathrm{d}t} = \frac{q}{V}\left(T_f - T\right) + \frac{m\Delta H}{\rho C_p}\, r + \frac{UA}{V\rho C_p}\left(T_c - T\right) \tag{2.3}$$

where q is the volumetric feed and product flow, V is the liquid volume, C_{Af} and T_f are the feed concentration and feed temperature, $m\Delta H$ is the heat released per mole reacted, ρ and C_p are the liquid density and heat capacity, UA is the product of the overall heat-transfer coefficient and the jacket area and T_c is the jacket temperature. The three terms of Equation (2.3) are the sensible load carried in by the feed, the exotherm and the jacket duty, and all three are reported in K/min.

Only three lumped groups appear in the balances. Their values at the default constants are the dilution rate $q/V = 1.0/\mathrm{min}$, the adiabatic temperature rise per unit of reactant $m\Delta H/(\rho C_p) = 209.2\,\mathrm{K}$ per mol/m^3, and the jacket rate constant $UA/(V\rho C_p) = 2.0921/\mathrm{min}$. Because the individual constants enter only through these groups, the published values are carried literally in the units in which they were published, and consistency of Equation (2.3) requires UA to be read as 5×10^4 J/(min K) rather than in watts.

Table 2.2 Model Parameters, Values and Visibility

Symbol	Meaning	Value and unit	Adjustable
q	feed and product flow	100 vol/min (50–200)	yes
V	liquid volume	100 vol (50–200)	yes
ρ	liquid density	1000 (800–1200)	yes
C_p	liquid heat capacity	0.239 (0.15–0.4)	yes
$m\Delta H$	heat of reaction	5×10^4 J/mol (2×10^4–8×10^4)	yes
E/R	activation temperature	8750 K (7000–10000)	yes
k_0	pre-exponential factor	7.2×10^{10} 1/min (10^{10}–10^{11})	yes
UA	jacket UA	5×10^4 J/min-K (2×10^4–10^5)	yes
σ	transmitter noise	0.15 K (0–1)	yes
T_f	feed temperature	350 K (300–380)	upset
C_{Af}	feed concentration	1 mol/m^3 (0.5–2)	upset
$T_{c,\min}$	jacket lower limit	250 K	no
$T_{c,\max}$	jacket upper limit	350 K	no
T_{alarm}	runaway alarm	400 K	no
h	integration substep	0.005 min	no
T_s	control period	0.05 min	no

The jacket duty and the single-pass conversion follow algebraically from the state, as shown in Equations (2.4) and (2.5).

$$Q = UA\,(T_c - T) \tag{2.4}$$

$$X = \frac{C_{Af} - C_A}{C_{Af}} \tag{2.5}$$

where Q is negative whenever heat leaves through the jacket and X is clipped to the interval 0 to 1. Neither is displayed by the application, which shows and exports the reactant concentration C_A instead.

The temperature transmitter adds a zero-mean Gaussian error to the state, as shown in Equation (2.6).

$$\hat{T}_k = T_k + \sigma\,\varepsilon_k, \qquad \varepsilon_k \sim \mathcal{N}(0,1) \tag{2.6}$$

where $\hat{T}_k$ is the reading the controller acts on and σ is the transmitter standard deviation. One deviate is formed at each controller execution by the Box–Muller transform of two draws from a seeded generator, which restarts from the same seed whenever a case is loaded or reset, so a run repeats exactly at any value of σ. The controller, the recorded trend and the exported file all act on $\hat{T}$, while the noise-free state T is retained separately so that a measured excursion can be separated from its noise. The runaway alarm compares the noise-free T against 400 K and is shown only while the temperature is above it, with no latch, no acknowledgement and no action on the plant.

2.3.3 Parameters and Initial Conditions

Every constant is listed in Table 2.2. Nine of them are adjustable model parameters, the feed temperature and feed concentration are operator-set upsets, and the rest are fixed. A changed parameter or upset reaches the model at the next controller execution, and the feed flow is a parameter rather than an upset.

Every case starts at steady operation. The four manual cases start from the steady state that the default constants produce with the jacket held at 300 K, namely $C_A = 0.87725\,\text{mol/m}^3$ and $T = 324.475\,\text{K}$, with the reaction running at 0.1227 mol/(m^3 min) and a single-pass conversion of 12.27 %. The seven automatic cases start at the steady state of their 340 K setpoint, with $C_A = 0.67598\,\text{mol/m}^3$ and the jacket at the 302.818 K that holds it. Loading a case or pressing Reset solves that start again for the constants, upsets, setpoint and manual jacket temperature then in force, the learner's changes included: in automatic it is the steady state of the setpoint, and in manual the stable steady state that the manual jacket temperature holds. A manual jacket between the ignition fold and the Hopf point of

Section 2.4 holds no stable state, so the reactor then starts from the cold state of the 300 K jacket and ignites into the relaxation cycle that the same jacket gives when it is set after Run, between 332.07 K and 466.7 K at 304 K. A setting chosen before Run is therefore part of the starting state, and it gives the same run as the same setting carried through Reset. Reset keeps every setting the learner has changed, while reloading the case restores the case's own settings.

2.3.4 Numerics

The right-hand side is integrated with explicit Euler on an outer substep of $h = 0.005$ min. The rate constant $k_0 \mathrm{e}^{-E/RT}$ grows by three orders of magnitude across the temperature range the model permits, so a fixed step that is comfortable on the cold branch is unstable on the hot one. The integrator therefore subdivides its own step against the fastest decay rate in the system, as shown in Equation (2.7).

$$\begin{aligned} \lambda &= \frac{q}{V} + \frac{UA}{V\rho C_p} + k_0 \mathrm{e}^{-E/(RT)} \\ n &= \operatorname{clamp}\left(\left\lceil \frac{h\lambda}{0.2} \right\rceil ,\ 1,\ 8000 \right) \end{aligned} \tag{2.7}$$

where n is the number of inner Euler updates taken over h and 0.2 is the stability margin on $h\lambda$. At the nominal operating point $\lambda = 3.23/\mathrm{min}$ and $n = 1$, so the update is plain Euler. At 450 K the rate constant alone is 258.7/min, so $n = 7$ and the integrator subdivides. The state is clamped to $C_A \in [0,\ 20]\,\mathrm{mol/m^3}$ and $T \in [200,\ 900]$ K after every inner update, which is a numerical guard rather than a physical constraint, because the adiabatic ceiling $T_f + (m\Delta H/\rho C_p)C_{Af}$ is 559 K at the default feed and 768 K at the largest feed concentration the model allows.

Each browser frame advances the simulated clock by the elapsed wall time multiplied by the speed setting, from 0.5 to 10 min of simulated time per second, in whole substeps of h. The controller executes once at $t = 0$ and then at the start of every tenth substep, so the control period is $T_s = 0.05$ min of simulated time, or 3 s. Because the controller counts substeps rather than wall-clock milliseconds, the closed-loop trajectory does not depend on the speed setting or on browser timer jitter. The historian records one row every 0.1 min, every second controller execution, and the export is a comma-separated file of thirty-three columns, eleven holding time, mode, case, setpoint, measured and noise-free reactor temperature, jacket temperature, a saturation flag, the two feed upsets and the reactant concentration, eight more for the predictive controller of Section 2.6.3, and fourteen for the estimator choice and the relay experiment of Sections 2.6.4 and 2.6.5. A row stamped t carries the noise-free temperature at t and the measured temperature that the controller execution at t acts on, while its jacket temperature is the one held over the control period that ended at t, so a jacket move made by an execution first appears on the next row.

2.3.5 Assumptions

Perfect mixing. The liquid is one lumped temperature and one lumped concentration. The assumption fails wherever the agitator is undersized or the viscosity rises with conversion, which is where real runaways start, because a hot pocket reacts faster than the bulk and the thermowell never sees it. The simulator cannot produce that failure mode at all.

Constant volume and constant physical properties. Density and heat capacity do not change with temperature or composition, and the product flow equals the feed flow at all times. Over the 200 K span the model permits, a real liquid changes ρC_p by tens of percent and would change the thermal time constant with it.

Single irreversible first-order reaction. There is no reverse reaction, no side reaction and no decomposition path. Decomposition at elevated temperature is the mechanism that turns many industrial runaways from an excursion into an explosion, and its absence is why the reactor in this simulation always comes back.

Ideal jacket. The jacket temperature is whatever the controller asks for, instantly and uniformly, within its limits. A real jacket has a coolant residence time, a fouling resistance that drifts, and a slave loop with its own dynamics and its own saturation [10]. The consequence is that the manipulated variable here carries no lag of its own, so the only dynamics in the loop are the reactor dynamics.

Ideal measurement. The transmitter is instantaneous and carries white noise only. A thermowell has a time constant of tens of seconds, which is the same order as the reactor time constant identified in Section 2.5, and adding it would change the achievable tuning.

Liquid phase only. There is no vapor space, no pressure state and no relief path. Nothing in the model represents what happens above the boiling point, so temperatures reported above roughly 400 K should be read as a measure of how far the energy balance has run away rather than as a prediction.

2.4 Steady-State Multiplicity and Stability

Setting Equation (2.2) to zero solves for the concentration at any temperature, because the mass balance is linear in C_A, as shown in Equation (2.8).

$$C_A^{ss}(T) = \frac{(q/V)\,C_{Af}}{q/V + k_0 \mathrm{e}^{-E/(RT)}} \tag{2.8}$$

Substituting Equation (2.8) into the energy balance splits Equation (2.3) into a generation term and a removal term, as shown in Equation (2.9).

$$\begin{aligned} Q_{\mathrm{gen}}(T) &= \frac{m\Delta H}{\rho C_p}\, k_0 \mathrm{e}^{-E/(RT)}\, C_A^{ss}(T) \\ Q_{\mathrm{rem}}(T, T_c) &= \frac{q}{V}\left(T - T_f\right) + \frac{UA}{V\rho C_p}\left(T - T_c\right) \end{aligned} \tag{2.9}$$

where both are expressed in K/min so that a steady state is a point at which the two are equal. This is the construction introduced for autothermic operation [1] and it is drawn for the shipped constants in Figure 2.2(a).

2.4.1 Steady States and the Fold Points

The generation curve is sigmoidal rather than purely exponential, because the reactant runs out. Below 330 K the rate constant is small enough that C_A^{ss} stays near C_{Af} and Q_{gen} doubles for roughly every 8 K, while above 380 K conversion exceeds 87 % and the curve flattens against the adiabatic ceiling of 209.2 K/min. The removal line has slope $q/V + UA/(V\rho C_p) = 3.0921/\mathrm{min}$ and shifts left as the jacket warms. Three intersections occur wherever the removal line cuts the steep middle part of the generation curve, as shown in Figure 2.2(a).

Solving Equation (2.3) for T_c instead of for T removes the multiplicity, because the jacket temperature that holds a given reactor temperature is unique, as shown in Equation (2.10).

$$\begin{aligned} T_c^{ss}(T) = T - \frac{V\rho C_p}{UA}\Bigg[&\frac{q}{V}\left(T_f - T\right) \\ &+ \frac{m\Delta H}{\rho C_p}\, k_0 \mathrm{e}^{-E/(RT)} C_A^{ss}(T) \Bigg] \end{aligned} \tag{2.10}$$

This expression gives the jacket temperature that the controller must find at each setpoint, 302.818 K at 340 K, and its graph is the S-shaped locus in Figure 2.2(b). The locus turns over twice, the picture of Section 1.4.2 with two folds instead of one. The ignition fold sits at $T_c = 303.229\,\mathrm{K}$ and $T = 335.654\,\mathrm{K}$, above which the cold branch ceases to exist, and the extinction fold sits at

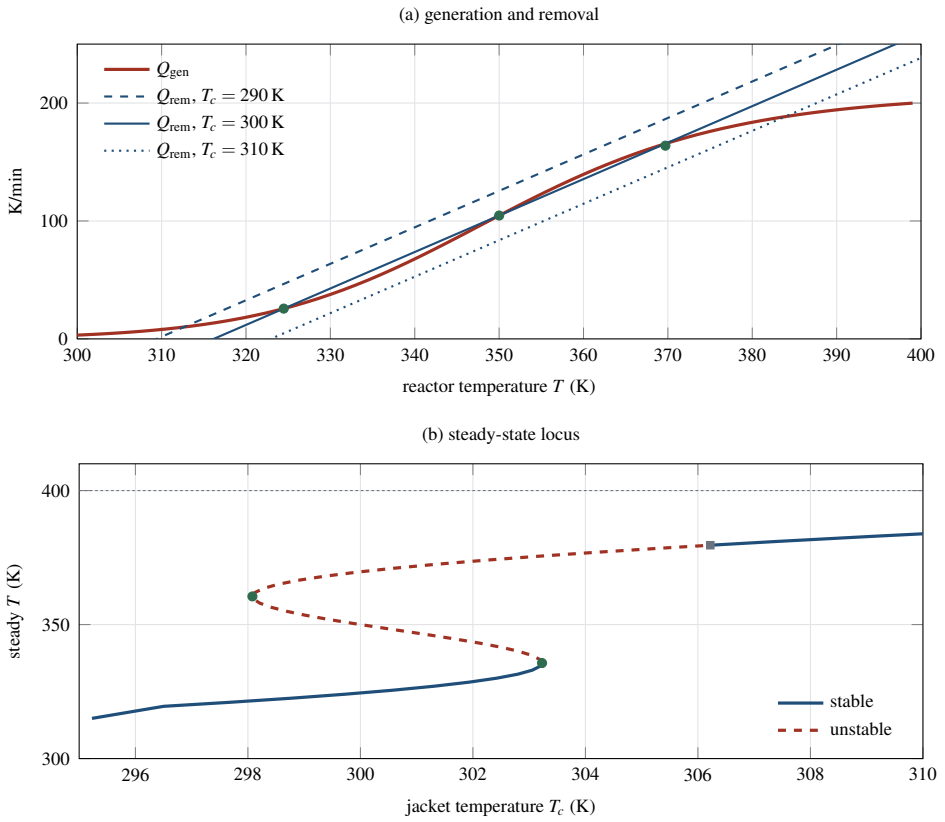


Figure 2.2 Thermal structure of the shipped reactor. Panel (a) is the van Heerden construction at three jacket temperatures, with the three steady states at $T_c = 300\,\mathrm{K}$ marked. Panel (b) is the locus of steady reactor temperature against jacket temperature, with the ignition and extinction folds as circles, the Hopf point at 379.61 K as a square and the 400 K alarm as a horizontal line.

$T_c = 298.080\,\mathrm{K}$ and $T = 360.511\,\mathrm{K}$, below which the hot branch ceases to exist. Three steady states therefore exist only for jacket temperatures inside a window 5.15 K wide.

At the published jacket temperature of 300 K the reactor has steady states at 324.475 K, 350.006 K and 369.705 K, which is the textbook arrangement. The stability of the third one is not. Steady-state multiplicity of the same kind sets the ignition floor of the catalyst beds in Volume 13, Chapter 4 and the two conversion branches of the reactive column in Volume 13, Chapter 9.

2.4.2 Stability of the Three Steady States

Linearizing Equations (2.2) and (2.3) about a steady state gives the Jacobian in Equation (2.11).

$$A = \begin{bmatrix} -\dfrac{q}{V} - k & -\dfrac{kE}{RT^2} C_A \\ \dfrac{m\Delta H}{\rho C_p} k & -\dfrac{q}{V} - \dfrac{UA}{V\rho C_p} + \dfrac{m\Delta H}{\rho C_p} \dfrac{kE}{RT^2} C_A \end{bmatrix} \tag{2.11}$$

where $k = k_0 \mathrm{e}^{-E/(RT)}$ is the rate constant at the operating temperature. The only positive entry is the lower right one, and it is the derivative of the exotherm with respect to temperature. Instability is therefore always the statement that this single term has outgrown the sum of the dilution rate and the jacket rate constant.

Evaluating Equation (2.11) at the three steady states available at $T_c = 300\,\mathrm{K}$ gives eigenvalues of $-1.049 \pm 0.539\mathrm{j}$ at 324.475 K, -0.454 and $+2.834$ at 350.006 K, and $+1.357 \pm 1.540\mathrm{j}$ at 369.705 K. The middle state is a saddle, as expected. The upper state is an unstable focus, which is not what the textbook drawing of three crossings usually shows, and the application's instructions page says so. The hot branch stays unstable until a Hopf bifurcation at $T_c = 306.220\,\mathrm{K}$ and $T = 379.611\,\mathrm{K}$, where the eigenvalues are $\pm 3.702\mathrm{j}$ and the natural period is 1.70 min. Open-loop, therefore, the reactor is unstable at every temperature between 335.654 K and 379.611 K, an interval 43.96 K wide that contains the default setpoint of 340 K.

> ✦ **DID YOU KNOW? A Reactor and a Pendulum, Both Runaway**
>
> An open textbook of industrial instrumentation files exothermic reactions with the inverted pendulum as runaway processes. Both carry internal positive feedback: gravity pulls an off-center pendulum farther off center, and a hotter reaction runs faster and releases more heat [18]. Chapter 7 holds its pendulum upright by feedback, as this chapter holds 340 K.

Integrating the model confirms the consequence. With the jacket held at 300 K, a reactor released 0.1 K either side of the ignited steady state, or from any temperature between 350.1 and 380 K at the steady-state composition, returns to 324.475 K every time. There is no second stable operating point at the published jacket setting. Holding the jacket instead at 304 K, just above the ignition fold, produces a sustained relaxation oscillation rather than a hot steady state, measured in the application as a cycle between 332.07 K and 466.3 to 466.7 K with a period of 8.39 min. A stable ignited state appears only once the jacket is warmer than 306.220 K, and it is then hotter than 379.6 K, which is inside the alarm margin.

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

2.4.3 Disturbances That Remove the Cold Branch

The same fold analysis run against the two operator-set upsets and the feed flow gives the point at which the cold operating point disappears while the jacket sits at 300 K. Raising the feed temperature removes it at $T_f = 356.8\,\mathrm{K}$, raising the feed concentration removes it at $C_{Af} = 1.141\,\mathrm{mol/m^3}$, and raising the feed flow removes it at $q = 126.3$ vol/min. A warmer feed only shifts the locus of Equation (2.10) by $(q/V)\,\Delta T_f/[UA/(V\rho C_p)]$, so the 15 K feed step scripted in the application is equivalent to warming the jacket by 7.17 K at the design feed. Each of these is a small move from the default. A feed concentration 14 % above design is not an unusual upset, and at the design jacket temperature it removes the cold operating point entirely.

2.5 Dynamic Analysis

2.5.1 Steady-State Gain Along the Locus

The steady-state gain of the temperature loop is the reciprocal of the slope of Equation (2.10), and Figure 2.2(b) shows immediately that this slope passes through zero twice. Table 2.3 lists the gain and the open-loop eigenvalues along the locus. The gain is 1.72 K/K at the nominal point, 3.43 K/K only 5.5 K higher, and unbounded at the ignition fold. Past the fold it is negative, because a warmer jacket on the middle branch pushes the reactor to a lower steady temperature. On the hot branch it recovers a positive value near 1 K/K.

A negative steady-state gain on an open-loop-unstable plant would ordinarily argue for a direct-acting controller, and the application uses a reverse-acting one, with $e_k = T^{\mathrm{sp}} - \hat{T}_k$ and a positive K_c, at every operating point. The reason, as for the arm of Section 1.4.3, is that the sign of the immediate response and the sign of the eventual response are different. The jacket enters only the temperature equation, so the instantaneous sensitivity of the measurement to a jacket move is $UA/(V\rho C_p) = 2.0921/\mathrm{min}$ and is positive everywhere. Warming the jacket always warms the reactor first. Only after the concentration has had time to fall does the loss of reactant pull the steady

Table 2.3 Steady-State Gain and Open-Loop Eigenvalues Along the Locus

T (K)	T_c^{ss} (K)	dT/dT_c (K/K)	eigenvalues (1/min)
324.475	300.000	+1.72	-1.049 ± 0.539j
330.000	302.423	+3.43	-0.642 ± 0.576j
335.654	303.229	$\pm\infty$	−0.281, 0
340.000	302.818	−5.56	−0.517, +1.077
350.000	300.002	−3.25	−0.454, +2.834
360.511	298.080	$\mp\infty$	0, +3.336
369.705	300.000	+2.38	$+1.357 \pm 1.540$j
379.611	306.220	+1.23	± 3.702j
390.000	316.269	+0.91	-2.953 ± 4.833j

temperature back down. The controller acts on the first of these because its 0.05 min period is far shorter than the residence time of 1 min, which is why a reverse-acting loop stabilizes an operating point whose steady-state gain is negative.

Pause. Try Exercise 2.1 before reading on.

2.5.2 Identification From a Step Test

Manual mode exists so that the plant can be identified before it is controlled. The application's first case holds the jacket at 300 K with the reactor at its steady state, and stepping the jacket to 303 K at $t = 1$ min moves the reactor from 324.475 K to 332.602 K, so the chord gain over the step is $K_p = 8.127/3 = 2.71$ K/K, the time to 63.2 % of the change is $\tau_p = 1.29$ min, and the noise-free temperature moves on the first 0.005 min substep after the jacket does, so $\theta_p = 0$. The reactor has no transport delay because the jacket acts on the liquid directly.

The chord gain over that 3 K step is 58 % larger than the tangent gain of 1.72 K/K at the starting point, which is a direct measure of how much the plant curves inside a step small enough to look linear on the trend. The application's case 2 makes the smaller step from 300 to 302 K at $t = 1$ min, and the reactor settles at 328.702 K, a chord gain of 2.11 K/K reached 63.2 % of the way 0.94 min after the step. A step of 4 K does not produce a usable fit at all, because 304 K is past the ignition fold and the reactor leaves the cold branch. The test therefore identifies the cold branch only, where Table 2.3 gives a positive steady-state gain, so the 2.71 K/K chord is a property of that branch and not a measurement of the negative gain that the same plant carries a few kelvin higher. A step test that returns a different plant at each operating point is also the lesson of the friction stir weld in Volume 9, Chapter 3.

2.5.3 Tuning From the Fitted Model

Internal model control tuning for a first-order model with negligible dead time gives $K_c = \tau_p/[K_p(\tau_c + \theta_p)]$ and $\tau_I = \tau_p$, where τ_c is the requested closed-loop time constant [11]. Asking for $\tau_c = 0.1\,\tau_p$ cancels the time constant and leaves $K_c = 10/K_p = 3.7$ whatever τ_p turns out to be, and $\tau_I = \tau_p = 1.29$ min then gives $K_I = K_c/\tau_I = 2.9$. The application's default tuning, which is also the suggested tuning of its challenge, is the rounded pair $K_c = 4$ K/K and $\tau_I = 1.333$ min, or $K_I = 3$/min, and it is used for every closed-loop result in Section 2.7 unless a case states otherwise. The identification is performed at 324.5 K and the tuning is then applied across a range in which Table 2.3 shows the gain changing by a factor of three and then changing sign, which is exactly the situation that gain scheduling addresses [8] and that the application makes measurable without providing the schedule.

2.6 Controller Design

2.6.1 Velocity-Form Algorithm

Feedback is applied by the velocity-form PID algorithm of Section 1.6.2, with the proportional term on the error and the derivative on the measurement, executed every $T_s = 0.05\,\text{min}$. In this chapter's symbols each execution forms the increments of Equations (2.12) to (2.14).

$$e_k = T_k^{\text{sp}} - \hat{T}_k \tag{2.12}$$

$$\Delta s_k = K_c\,(e_k - e_{k-1}) - K_c\tau_D\,\frac{\Delta^2 \hat{T}_k}{T_s} \tag{2.13}$$

$$\Delta I_k = \frac{K_c}{\tau_I}\,e_k\,T_s \tag{2.14}$$

where $\Delta^2 x_k = x_k - 2x_{k-1} + x_{k-2}$ is the second backward difference, e_k is the error at execution k, K_c is in kelvin of jacket per kelvin of error, τ_I and τ_D are in minutes, $K_I = K_c/\tau_I$ and $K_D = K_c\tau_D$. The derivative time defaults to zero, so the shipped controller is proportional-integral. Because the derivative acts on the measurement, a setpoint step reaches the jacket only through the proportional term, as a step of $K_c\,\Delta T^{\text{sp}}$ at the next execution. The control period T_s is the Δt of Chapter 1, the running sum of the increments of Equation (2.13) is its S_k, and a gain change acts on the whole error at once.

The controller is reverse acting ($a = +1$ in Equation (1.12)) at every operating point in the ISA sense, because the error of Equation (2.12) is $T^{\text{sp}} - \hat{T}_k$ and K_c is positive, so the jacket demand falls as the measurement rises. Section 2.5 gives the justification for using that action on a plant whose steady-state gain is negative over part of its range.

2.6.2 Limits, Windup and Initialization

The increments are accumulated and the demand is clipped to the jacket range, as shown in Equation (2.15).

$$\begin{aligned} s_k &= s_{k-1} + \Delta s_k \\ c_k &= u_{\text{bias}} + s_k + I_{k-1} \\ u_k &= \min(T_{c,\max},\,\max(T_{c,\min},\,c_k)) \\ I_k &= I_{k-1} + \Delta I_k + \frac{T_s}{\tau_t}\,(u_k - c_k) \end{aligned} \tag{2.15}$$

where s_k accumulates the proportional and derivative increments without limit, I_k is the integral state, c_k is the jacket temperature the algorithm demands, u_{bias} is the bias, fixed when a case is loaded or reset at the manual jacket temperature then in force, and $\tau_t = \tau_I$ is the tracking time of the back-calculation. The integral used at execution k was accumulated through execution $k-1$. The correction is applied only while $K_I \neq 0$. Setting τ_I to the top of its range, 30 min, switches integral action off, and the loop is then proportional only, with neither integral nor correction. Away from the jacket limits the recursion is the positional law, as Section 1.6.2 explains. While the jacket is clipped, substituting the second line into the fourth gives $I_k = I_{k-1} + (T_s/\tau_I)(u_k - u_{\text{bias}} - I_{k-1} - D_k)$, where D_k is the derivative term, so the error cancels and the integral relaxes toward the value that would hold the jacket at its limit with no proportional action, by 3.75 % of the gap per execution at the default $\tau_I = 1.333\,\text{min}$ [19]. After a long saturation the jacket therefore leaves the limit when the error changes sign. There is no rate limit on the jacket temperature and no filter on the derivative term, so derivative action operates directly on the noisy reading and the transmitter noise has to be set to zero before τ_D is useful.

Loading a case or pressing Reset initializes the controller at the case's manual jacket demand u_{man}, as shown in Equation (2.16).

$$u_{\text{bias}} = u_{\text{man}}, \quad I_{-1} = 0, \quad s_{-1} = 0, \quad c_0 = u_{\text{man}} + K_c\,e_0 \tag{2.16}$$

The previous proportional term is zero at that point, so the first execution applies the whole proportional term, which is the positional law about the manual jacket temperature. A case that opens in automatic starts at the steady state of its setpoint with u_{man} at the jacket temperature that holds it, so e_0 is only the transmitter error and the first execution moves the jacket by $K_c \sigma \varepsilon_0$, 0.64 K in the shipped cases at $K_c = 4\,\mathrm{K/K}$. In manual the demand tracks the operator's jacket temperature, and in automatic the manual station follows the jacket, so transfers are bumpless in both directions, as Section 1.6.2 describes. The exception is a setpoint changed together with the transfer, which passes $K_c \Delta T^{\mathrm{sp}}$ to the jacket at the first automatic execution.

The jacket temperature that holds a setpoint matters more on this plant than on a wide-range loop. Only 3.23 K of the 100 K jacket span separates the published setting of 300 K from the ignition fold, so every jacket temperature that holds a cold-branch setpoint above the nominal one lies inside three percent of the actuator span. The learner has no bias to set. When a case starts in automatic the application solves Equation (2.10) itself and seeds the bias with the result, 302.818 K at the 340 K default, so the integral starts at zero, and after a setpoint move or an upset the integral term alone finds the new value. The setpoints the jacket can hold at the design feed run from 282.513 K at its lower limit to 416.427 K at its upper limit.

2.6.3 Model Predictive Control

A linear model predictive controller (MPC) can set the jacket temperature in automatic mode in place of the PI loop, which stays the default everywhere but cases 9 and 10. It works as Section 1.6.3 describes, planning the jacket temperature at every execution, $T_s = 0.05\,\mathrm{min}$ apart, and applying the first move [20, 21]. No step test can supply its model, since the test of Section 2.5 reaches only the cold branch and the reactor leaves its 340 K design point with no input change at all. The model is instead the Jacobian of Equation (2.11) at 340 K with the nominal constants and the design feed, the matrix behind that row of Table 2.3, completed by the jacket term of Equation (2.3), as shown in Equation (2.17).

$$\begin{aligned} \frac{\mathrm{d}\xi}{\mathrm{d}t} &= A_m\,\xi + b_m\,(T_c - \bar{T}_c + \delta), \qquad \tilde{T} = \bar{T} + \xi_2 \\ A_m &= \begin{bmatrix} -1.4793 & -0.024526 \\ 100.279 & 2.0388 \end{bmatrix}, \qquad b_m = \begin{bmatrix} 0 \\ 2.0921 \end{bmatrix} \end{aligned} \tag{2.17}$$

Here $\xi = [\xi_1 \;\; \xi_2]^{\mathsf{T}}$ holds the deviations of C_A and T from $\bar{C}_A = 0.67598\,\mathrm{mol/m^3}$ and $\bar{T} = 340\,\mathrm{K}$, $\bar{T}_c = 302.818\,\mathrm{K}$ holds that state, $\tilde{T}$ is the predicted temperature, δ is an unmeasured disturbance in kelvin of jacket temperature, and the entries keep the units of Equation (2.11), 1/min on the diagonal and in b_m. The eigenvalues of A_m are $+1.0766$ and $-0.5171/\mathrm{min}$, a time constant of 0.929 min for the unstable mode, and its steady-state gain is $G_m = -5.559\,\mathrm{K/K}$. The model thus holds both signs that Section 2.5 separates, the positive first response of 2.0921/min to a jacket move and the negative steady response, so a plan that raises the reactor temperature has to warm the jacket first and then hold it colder than before.

The estimate $\hat{\delta}$ removes offset. A steady-state Kalman filter on Equation (2.17), sampled with a zero-order hold at T_s and extended with δ as a random walk, the input-disturbance model of offset-free predictive control [22], which the boiler of Volume 9, Chapter 4 uses against a fuel bias, runs at every execution in either mode and under either controller, as shown in Equation (2.18).

$$\begin{aligned} \bar{\xi}_k &= \Phi\,\hat{\xi}_{k-1} + \gamma\left(u_{k-1} - \bar{T}_c + \hat{\delta}_{k-1}\right) \\ \eta_k &= \hat{T}_k - \bar{T} - \bar{\xi}_{2,k} \\ \hat{\xi}_k &= \bar{\xi}_k + \ell_\xi\,\eta_k, \qquad \hat{\delta}_k = \hat{\delta}_{k-1} + \ell_\delta\,\eta_k \end{aligned} \tag{2.18}$$

Here $\Phi = \mathrm{e}^{A_m T_s}$ and γ are the sampled model, u_{k-1} is the jacket temperature applied over the period just ended, $\hat{T}_k$ is the reading of Equation (2.6), η_k is the innovation, and $\ell_\xi = [-0.00097 \;\; 0.3595]^{\mathsf{T}}$

and $\ell_\delta = 0.4553$ are the limits of the filter's Riccati recursion when the measurement variance is $r = \sigma^2 = 0.0225\,\mathrm{K}^2$, the state noise variance $10^{-6}r$ and the disturbance noise variance $q_\delta = 10\,r/G_m^2 = 0.00728\,\mathrm{K}^2$. Because δ enters where the jacket does, the 15 K feed step is booked as the 7.17 K of Section 2.4, and a steady load leaves no offset, while $\hat{\xi}_1$ supplies an estimate of C_A that is valid near one operating point.

A quadratic program is posed at every decision on Equation (2.17) sampled at T_s, over a horizon $T_p = 4\,\mathrm{min}$ of $N = 80$ steps, about four time constants of the unstable mode. Twelve jacket temperatures v_j, held over blocks of $(1,1,1,1,2,2,3,4,6,9,15,35)$ steps, minimize Equation (2.19) subject to the hard limits $250\,\mathrm{K} \le v_j \le 350\,\mathrm{K}$.

$$J = \frac{T_s}{\tau_n}\left[\sum_{k=1}^{N}\left(\frac{\tilde{T}_k - T^{\mathrm{sp}}}{y_s}\right)^2 + N\left(\frac{\tilde{T}_N - T^{\mathrm{sp}}}{y_s}\right)^2\right] + \kappa\sum_{j=1}^{12}\left(\frac{v_j - v_{j-1}}{u_s}\right)^2 \tag{2.19}$$

Here $\tilde{T}_k$ is predicted k steps ahead with $\hat{\delta}$ held, v_0 is the jacket temperature applied now, $\tau_n = 1\,\mathrm{min}$, $y_s = 2\,\mathrm{K}$ is the scorecard band, $u_s = 100\,\mathrm{K}$ is the jacket span and $\kappa = 0.3$ is the move suppression, called λ in the application but renamed because λ is the decay rate of Equation (2.7). The weights are 0.0125 per predicted error, 1.0 on the last and 3×10^{-5} per move. No temperature constraint, soft band, rate limit or feed preview enters, so the 400 K alarm and the ignition fold reach the plan only through the model. The decisions are departures from $\bar{T}_c$, a ridge of 10^{-9} relative to the weights makes the Hessian positive definite, and the dual active-set method of [23] finds the optimum. A rule shared by the seven single-loop studios fixed κ and q_δ before any comparison, taking the largest κ within 10 % of the best true-value IAE among settings with no more jacket travel than the PI loop in cases 3 and 4.

The horizon hardly matters here. Case 9 scores 2.694 K min at horizons of 1, 2, 4 and 8 min, with or without the terminal term, and the jacket trajectories differ by at most 0.015 K, because the plan made at the setpoint step already predicts the temperature within 0.04 K of the setpoint 0.25 min ahead. The formulation matters more. With absolute jacket temperatures as decisions and the solver's default ridge, 10^{-9} times one plus the largest diagonal entry of the Hessian, which grows with the square of predictions that the unstable mode multiplies by $\mathrm{e}^{1.0766\times 8} \approx 5500$ over 8 min, the ridge pulls every decision toward 0 K. Rebuilt that way outside the application, case 9 scores 2.712, 13.24 and 469.87 K min at horizons of 4, 6 and 8 min, the last with the reactor falling to 293.58 K on its cold branch while every plan is reported optimal, a defect removed before release.

The PI loop takes over whenever the solver fails to return an optimum or the move it would apply is not finite or lies beyond a jacket limit, and because it tracks the applied jacket at every execution it holds the jacket without a step change until a later plan is accepted. With a rejection forced 0.2 min after the setpoint step of case 9 and the transmitter noise removed, the handover moves the jacket by 0.26 K, within the 0.08 to 0.34 K of the next three PI moves, and the returning plan by 1.22 K. The first plan is computed while a case loads, so the MPC acts from $t = 0$. It reads only the measurement, its own jacket move, the setpoint and the jacket limits, and its model is built once from the nominal constants, so a constant the learner changes reaches the plant of Equations (2.2) and (2.3) and never Equation (2.17).

2.6.4 Moving Horizon Estimation

The Kalman filter of Equation (2.18) is linear in the readings, so nothing keeps its estimate of C_A inside the physical range once the reactor is far from 340 K. The application therefore offers a second source for the state each plan starts from, a moving horizon estimator (MHE) that fits the same sampled model to a window of readings under bounds on the states [24]. The Kalman filter remains the default and keeps running whichever estimator is chosen. With the augmented state $x_j = [\xi_{1,j} \;\; \xi_{2,j} \;\; \delta_j]^{\mathsf{T}}$

and the window $j = s, \dots, k$ with $s = k - N$, every execution solves Equation (2.20) over x_s and $w_s, \dots, w_{k-1}$.

$$\begin{aligned} \min \; & \|x_s - \bar{x}_s\|^2_{\bar{P}^{-1}} + \sum_{j=s}^{k-1} \|w_j\|^2_{Q^{-1}} \\ & + \frac{1}{r} \sum_{j=s}^{k} \left(\hat{T}_j - \bar{T} - x_{2,j}\right)^2 \\ \text{s.t. } & x_{j+1} = \Phi_a x_j + \gamma_a \left(u_j - \bar{T}_c\right) + w_j \\ & 0 \le \bar{C}_A + x_{1,j} \le C_{Af}^{\max} \\ & T_{c,\min} \le \bar{T} + x_{2,j} \le T_f^{\max} + \frac{m\Delta H}{\rho C_p} C_{Af}^{\max} \end{aligned} \tag{2.20}$$

Here $\Phi_a = \left[\begin{smallmatrix} \Phi & \gamma \\ 0 & 1 \end{smallmatrix}\right]$ and $\gamma_a = [\gamma^{\mathsf{T}} \;\; 0]^{\mathsf{T}}$ carry δ as the random walk of Equation (2.18), u_j is the jacket temperature applied from execution j to $j+1$, w_j is the process noise with $Q = \mathrm{diag}(10^{-6} r, 10^{-6} r, q_\delta)$, $\bar{x}_s$ is the estimate the MHE made at $s-1$ carried one step through the model, and $\bar{P}$ is the filter's steady a-priori covariance, held fixed as the arrival cost. The window spans $N = 15$ executions, 0.75 min and 16 readings, so the problem has 48 decisions and 64 bound rows, and the dual active-set method of [23] solves it. The bounds come from declared values only. $C_{Af}^{\max} = 2\,\mathrm{mol/m^3}$ is the richest feed the model allows, $T_{c,\min} = 250\,\mathrm{K}$ is the coldest jacket, and the upper temperature bound of 798.41 K is the adiabatic limit of that feed at the hottest feed temperature, $T_f^{\max} = 380\,\mathrm{K}$. The disturbance δ is not bounded, and a constant the learner changes reaches the MHE only through the reading, as it reaches the filter.

Without an active bound, Equation (2.20) is the filter's own least-squares problem and the two estimates coincide. Over the 599 executions of a 30 min run of case 9 they differ by at most 7.9×10^{-11} of the transmitter's standard deviation, a residue of the filter's Riccati recursion, which stops at a relative change of 10^{-11}, because with the covariance iterated to its fixed point the difference falls to 3.2×10^{-13}. Run under the MPC with each estimator, the ten earlier cases agree to within 6.1×10^{-8} K of jacket temperature. A bound is active in only one of them, case 5, which stays in manual while the filter's concentration estimate falls as far as 0.21 mol/m^3 below zero at 222 executions. The MHE thus adds its bounds and nothing else, and on this reactor only a runaway brings the filter to them (Section 2.9).

A window whose solution is not optimal, not finite or outside the bounds is rejected, and the estimate is then the Kalman update of the MHE's one-step prediction, clipped to the bounds and marked in the log. No case produced a rejection. Forced at the execution at 3.5 min of case 10 with the noise removed and a bound active, a rejection moves the jacket by 23.4 K at that execution, the largest difference from an unfaulted run over the next 32, and because the rejected estimate later serves as the prior of a window, a concentration difference of 2.8×10^{-3} and 1.4×10^{-3} mol/m^3 returns 16 and 32 executions later. An estimate took a median of 142 µs and a 99th percentile of 512 µs under Node.js over the 5470 estimates of those ten cases. In Chrome, on the application's own frame loop during case 9, the estimates took 0.76 and the plans 0.31 percent of one processor core at one minute of process time per second, and 5.23 and 1.71 percent at ten, measured while other work loaded the same computer.

2.6.5 Relay Autotuning

The PI settings can also be measured rather than derived, by the relay experiment of [25], which the application runs as an action from the present state in automatic and never as a mode a case is in. The relay sets the jacket $d = 5\,\mathrm{K}$, five percent of its span, above or below the mean jacket temperature applied over the last 1 min, a tenth of the trend window, and switches each time the reading crosses the setpoint by more than $\varepsilon = 3\sigma = 0.45\,\mathrm{K}$. A cycle counts when each of its halves lasts at least four executions, and once two counted cycles of amplitude a agree within five percent, the describing

function and the Tyreus–Luyben rule give the settings of Equation (2.21).

$$K_u = \frac{4d\sin(\pi D)}{\pi\sqrt{a^2-\varepsilon^2}}, \qquad K_c = \frac{K_u}{3.2}, \qquad \tau_I = 2.2\,P_u \tag{2.21}$$

Here K_u is the ultimate gain in kelvin of jacket per kelvin, D the fraction of each period that the jacket spends at its upper level, which is one half for a symmetric cycle, and P_u the measured period. The experiment ends after 30 cycles or 20 min, and the PI loop then resumes from the applied jacket without a step. Where the band is wide against the swing, the band rather than the loop sets the cycle, which the application flags once $\psi = \arcsin(\varepsilon/a)$, the phase by which the cycle falls short of -180°, reaches 10°.

On this reactor the relay cannot measure the loop. Started from the steady state of case 3 without its setpoint step, its halves last a median of three executions with the transmitter noise and two without it, so 29 of the 30 cycles are rejected and the experiment ends after 4.35 min with the tuning unchanged. Over ten further noise seeds it converged in four, with K_u from 6.13 to 6.90, P_u from 0.389 to 0.413 min and ψ from 23.5° to 26.1°, cycles set by the band. The same relay on the sampled model of Equation (2.17) without noise cycles at $P_u = 0.399$ min, so the fast cycle belongs to the loop and not to the noise. Holding each relay level for at least four executions, 0.2 min, makes the experiment converge for the default seed and all ten others, but the hold then sets the period, $P_u = 0.45$ min or 2.25 holds. The resulting $K_c = 1.65\,\mathrm{K/K}$ and $\tau_I = 0.99$ min score 14.6 and 9.53 K min of measured IAE over the first 30 min of cases 3 and 4, against 5.51 and 6.02 K min for the shipped tuning, and holds of 5, 6 and 8 executions score 24.3, 30.4 and 444 K min on case 3. The application therefore applies no hold and states the reason with the action.

2.7 Results

The application keeps its own scorecard, but its integral of absolute error runs from $t = 0$, so the numbers in this section are computed over fixed windows from the exported columns and from the noise-free state at the 0.005 min substep. Peaks, final values and settling times are taken from the noise-free temperature at the substep, because the export samples every 0.1 min and misses as much as 21 K of an ignition peak in Table 2.5, and the integral of absolute error is taken from the exported measured column that the controller actually sees. Settling time is the last instant, on the 0.005 min grid of the substep, at which the temperature is more than 0.5 K from setpoint, measured from the step, and the integral of absolute error is accumulated over a fixed 12 min window after the step. The alarm entry records whether the noise-free temperature passes 400 K. Each of the eight experiments starts from one of the application's eleven cases with its shipped noise seed, as listed in Table 2.4.

2.7.1 Open-Loop Step Tests

Experiment 1 produces the identification reported in Section 2.5. The reactor moves from 324.475 K to 332.602 K on a 3 K jacket step, rising to 333.219 K 5.14 min after the step, 0.62 K or 7.6 % of the change past its final value, and staying within 2 % of the change from 8.27 min after the step. That single overshoot is the behavior of the underdamped cold branch, whose eigenvalues at 332.6 K are -0.420 ± 0.479j, and the first-order fit with $\tau_p = 1.29$ min captures its dominant part.

Experiment 2 changes the jacket by one kelvin more. The reactor leaves the cold branch, peaks at 469.27 K 3.70 min after the step, and then settles into a sustained oscillation between 332.07 K and 466.3 to 466.7 K with a period of 8.39 min, spending about 0.5 min of each cycle above the 400 K alarm. One kelvin of jacket temperature is the difference between a step response and a relaxation oscillation, which is the whole content of the fold at 303.229 K presented as an experiment. A larger step from case 1, from 300 to 320 K at $t = 1$ min, goes past the Hopf point as well as the fold. The reactor ignites to 505.03 K 0.68 min after the step, stays above 400 K for 0.86 min, and is within 0.5 K of the stable hot steady state at 393.31 K from 1.9 min after the step, with only 0.060 mol/m^3 of reactant left.

Table 2.4 Experiments and the Cases They Start From

#	Case	Mode	Action
1	1	manual	jacket 300 to 303 K at $t = 1$ min, step test
2	1	manual	jacket 300 to 304 K at $t = 1$ min, past the ignition fold
3	1	manual, PI	setpoint 340 K and automatic together at $t = 1$ min
4	1	PI	setpoint 340 K and automatic chosen before Run, then setpoint steps at $t = 1$ min to 345, 350, 351, 352, 355, 360, 365 and 371 K
5	1	PI	as 4 with $K_c = 1.6$, 4 and 8 K/K chosen before Run and a step to 354 K
6	4, 5	PI, manual	feed temperature 350 to 365 K at $t = 1$ min
7	1	PI, manual	feed concentration 1.0 to 1.2 mol/m^3 at $t = 1$ min, in automatic as 4
8	11	PI	as 6 with $K_c = 0.8$ and 4 K/K

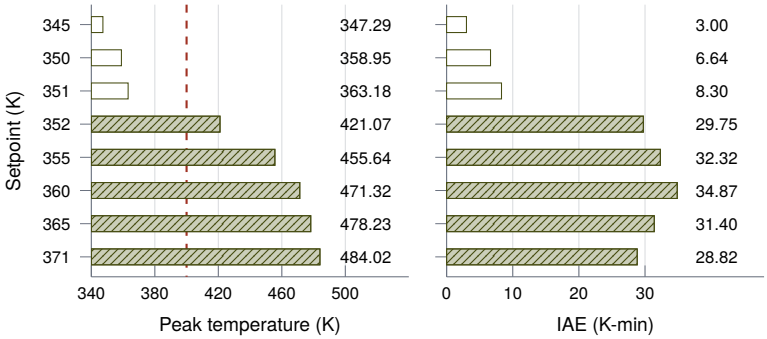


Figure 2.3 One PI tuning on eight setpoint steps from 340 K, the peak temperature (left, dashed line at the 400 K alarm) and the integral of absolute error (right) of Table 2.5. White bars stay below the alarm and hatched bars pass it. Notice the cliff between the 351 and 352 K setpoints, where one kelvin more of setpoint lifts the peak from 363.18 to 421.07 K and the error from 8.30 to 29.75 K min.

2.7.2 One Tuning Across Eight Setpoints

Experiment 3 makes the move from the cold branch that the automatic cases no longer make at their start. Setting the setpoint to 340 K and switching case 1 to automatic together at $t = 1$ min passes $4 \times 15.5 \approx 62$ K to the jacket at the first automatic execution, so the jacket goes to its 350 K limit for one control period, the reactor crosses the ignition fold at 335.654 K under feedback 0.13 min after the transfer, peaks at 342.67 K 0.50 min after it, and stays within 0.5 K of setpoint from 0.97 min after it.

Experiment 4 is the central measurement. Choosing the 340 K setpoint and automatic in case 1 before Run starts the reactor at the steady state of 340 K with the jacket at 302.818 K, the way every automatic case opens. The setpoint is then stepped at $t = 1$ min with $K_c = 4\,\mathrm{K/K}$, $\tau_I = 1.333$ min and $\tau_D = 0$ unchanged, so the 345 K step is case 3 of the application, run for run. The application credits case 3 at $t = 2.0$ min, after the peak, once the temperature has stayed within ± 2 K of 345 K for 0.5 min. Table 2.5 reports the result, and Figure 2.3 draws its peaks and errors against the alarm.

Table 2.5 One PI Tuning, Eight Setpoint Steps From 340 K

T^{sp} (K)	peak (K)	overshoot (K)	IAE (K-min)	settling (min)	alarm
345	347.29	+2.29	3.00	0.845	no
350	358.95	+8.95	6.64	1.505	no
351	363.18	+12.18	8.30	1.545	no
352	421.07	+69.07	29.75	2.355	yes
355	455.64	+100.64	32.32	2.330	yes
360	471.32	+111.32	34.87	2.000	yes
365	478.23	+113.23	31.40	2.105	yes
371	484.02	+113.02	28.82	2.255	yes

The transition is not gradual. An 11 K step overshoots by 12.18 K and never trips the alarm, and a 12 K step overshoots by 69.07 K and trips it, with the integral of absolute error rising from 8.30 to 29.75 K min over the same one-kelvin change of setpoint. In every case the loop reaches the setpoint and holds it, so a metric computed only at the end of the run would report all eight tests as successful. The mechanism is visible in the jacket column. The step moves the jacket at once by $K_c\,\Delta T^{sp}$, to 346.4 K for the 11 K step and to its 350 K limit for the 12 K step, and the reactor, already above the ignition fold, climbs at close to 100 K/min while the controller drives the jacket back down. With the jacket at its 250 K limit, the exotherm exceeds the largest removal the jacket can supply above 362.5 K when 0.50 mol/m^3 of reactant is left and above 364.5 K when 0.45 mol/m^3 is left. The 11 K step turns at 363.18 K with the jacket at 251.1 K. The 12 K step reaches the jacket limit at the execution 0.35 min after the step, with the reactor already at 366.2 K and 0.499 mol/m^3 of reactant, and the reactor then runs to 421.07 K, burning the reactant down to 0.018 mol/m^3 while the jacket sits on its lower limit for 0.6 min.

Where that boundary falls does not depend on the transmitter noise. Repeating Experiment 4 with $\sigma = 0$ and with eight other seeds of the generator puts the transition between the 11 K and the 12 K step every time. The 11 K step peaks between 362.61 and 363.19 K over those seeds, while the 12 K step peaks between 406.17 and 420.34 K, and at 416.67 K without noise, so the noise changes how far the ignited reactor runs but not whether it ignites. What does not move is the abruptness, which is one kelvin of setpoint between a well-damped response and an alarm.

Cases 6 and 7 test the remedy that first suggests itself on the 5 K step of case 3. Reducing the proportional gain from 4 to 1.6 K/K (case 6) raises the peak from 347.29 K to 351.67 K, more than twice the step, and the reactor stays within 0.5 K of setpoint only from 6.095 min after the step, against 0.845 min at the reference gain. Doubling the gain to 8 K/K (case 7) lowers the peak to 346.36 K and settles in 0.51 min. Experiment 5 repeats the comparison on a 14 K step from 340 to 354 K at $t = 1$ min, which ignites the reactor at the reference gain. Reducing the gain to 1.6 K/K raises the peak from 450.78 K to 460.13 K, and the loop never settles. It falls into a relaxation cycle between about 338 and 416 K with a period of 2.10 min, which passes the alarm 15 times in the first 30 min. Doubling the gain to 8 K/K keeps the reactor off the hot branch altogether, with a peak of 363.04 K, or between 361.90 and 363.59 K over five other noise seeds. Linearized about 354 K with $\tau_I = 1.333$ min, the loop is stable only for $K_c > 1.69\,\mathrm{K/K}$, against 0.57 K/K at 340 K and 1.04 K/K at 345 K, because the open-loop instability grows along the middle branch, so the conservative gain cannot hold the new setpoint at all. On the larger 31 K step to 371 K, detuning to $K_c = 1\,\mathrm{K/K}$ and $\tau_I = 1.25$ min, or $K_I = 0.8/\mathrm{min}$, lowers the first peak by 5.95 K, to 478.07 K against 484.02 K, but that gain is again below the stability bound of 1.45 K/K there, and the reactor passes 400 K 21 times in the 30 min after the step. The excursion is governed by how fast the controller can remove heat after ignition, and the operating point itself needs a minimum gain to exist, so detuning buys no safer approach on this plant. Raising the gain does not remove the largest excursions either. Steps to 360 K and above pass the alarm at every K_c from 4 to the 20 K/K top of its range.

Table 2.6 The Same Upset With and Without Feedback

Upset	Mode	peak (K)	final (K)	outcome
T_f to 365	PI	342.15	340.02	returns to setpoint
T_f to 365	manual	484.29	380.75	ignites, alarm
C_{Af} to 1.2	PI	341.06	340.02	returns to setpoint
C_{Af} to 1.2	manual	515.38	392.20	ignites, alarm

Peak after the upset and mean noise-free temperature over the last 5 min of the 30 min after it. PI rows start at the steady state of 340 K, and manual rows hold the jacket at 300 K from the steady state at 324.475 K.

> **TRY IT IN THE APP** **6 · Conservative Tuning**
>
> Load the scenario and press Run. With the gain cut from 4 to 1.6 K/K, the 5 K setpoint step peaks at 351.67 K, more than twice the step, and stays within 0.5 K of setpoint only from 6.095 min after it. Then load 7 · Aggressive tuning. At 8 K/K the peak is 346.36 K and the loop settles in 0.51 min.

2.7.3 Feedback Against Igniting Upsets

Cases 4 and 5 run the same upset twice, a feed temperature step from 350 to 365 K at $t = 1$ min, once with the loop closed at 340 K and once with the jacket held by hand at 300 K and the reactor on its cold branch at 324.475 K. Experiment 7 does the same from case 1 with a feed concentration step from 1.0 to 1.2 mol/m^3 at $t = 1$ min. Table 2.6 gives the outcome.

Both upsets are past the fold values computed in Section 2.4, so in manual the cold branch no longer exists and the reactor has to leave it. Under feedback the same upsets are absorbed with a peak deviation of 2.15 K and 1.06 K respectively, or 1.88 to 2.14 K and 1.06 to 1.17 K over five other noise seeds, because the controller lowers the jacket by 7.2 and 6.5 K, to the 295.65 K and 296.34 K that Equation (2.10) requires, before the exotherm can build. The feed concentration case ends at the hot steady state of 392.20 K with 93.6 % conversion. The feed temperature case ends at 380.75 K after an oscillation that decays with a period of 1.62 min. A feed 15 K warmer acts like a jacket 7.17 K warmer, which places this hot state just past the Hopf point of Section 2.4, where its eigenvalues are -0.238 ± 3.895j and the predicted period is 1.61 min. By hand the same upset is a race. Case 5 asks the learner to keep the cold branch, and a jacket lowered to 293 K, where the cold steady state with the warmer feed is 324.77 K, does so if it is set no later than $t = 2.05$ min, and the reactor then peaks at 339.37 K. Set at 2.1 min, the same jacket is too late and the reactor ignites to 450.86 K before it returns.

> **TRY IT IN THE APP** **5 · Recover by Hand**
>
> Load the scenario and press Run without touching the jacket. At 1 min the feed warms from 350 to 365 K, the cold operating point disappears, and the reactor temperature climbs past the 400 K alarm to a peak of 484.29 K (Table 2.6). Then load 4 · Disturbance rejection, the same upset under the PI loop, and watch it peak at 342.15 K.

Feedback has to be strong enough to act before the exotherm does, and the application's challenge, case 11, makes that its target. It applies the feed temperature step of case 4 with the proportional gain cut to 0.8 K/K, which still stabilizes 340 K in the linear sense, and asks for the measured temperature to stay within ± 2 K of 340 K over the last third of a 4 min run, from $t = 2.64$ min. With the shipped gain the reactor ignites to 445.08 K 0.96 min after the upset, falls back to 325.23 K once the reactant is spent, and is still 15.07 K from setpoint in the judged window, 7.5 times the band. With the suggested $K_c = 4\,\mathrm{K/K}$ the worst measured error in the same window is 0.573 K, inside the band with 1.43 K to spare, and the worst noise-free error is 0.438 K. The suggested tuning gives the same run whether it is set before Run or carried through Reset.



Figure 2.4 Stirred Reactor Control Studio during a setpoint move that ignites the reactor, a composite of the plant view's vessel and readouts (left, scaled down) and the trends (right), with the controls and the Learn panel cut. Case 1 opens in manual at the published steady state of 324.5 K, the setpoint is set to 340 K with the loop transferred to automatic at $t = 4$ min (2), and the setpoint is raised to 355 K at $t = 13$ min. The trends show the reactor igniting above the 400 K alarm (3) while the jacket sits on its 250 K limit (4, shaded) and the reactant concentration collapses (5). The run peaked at 456.0 K (446.9 K in the 0.1 min trend) with the jacket on its limit for 0.5 min, values the screen does not print. At $t = 30$ min the readouts (1) give 354.8 K (355.1 K measured, at the top right of the trends) with the jacket at 299.9 K and 0.414 mol/m^3 of reactant left.

Table 2.7 Five Cases Run Twice, Under the PI Loop and Under the Predictive Controller, With Identical Transmitter Noise

Case	Controller	IAE (K min)	IAE$_T$	OS (K)	e_{max}	$t_{0.5}$ (min)	Travel (K)	Limit (%)
3, setpoint	PI	2.68	1.90	2.29	—	0.90	119.2	0
	MPC	1.86	1.10	0.75	—	0.50	89.2	0
4, feed	PI	3.26	2.75	—	2.15	1.70	93.5	0
	MPC	2.32	1.57	—	2.05	0.70	72.0	0
8, limit	PI	41.89	41.79	—	—	—	71.6	80
	MPC	42.80	42.70	—	—	—	95.7	87
9, unstable	PI	4.38	3.86	2.29	2.33	2.00	127.7	0
	MPC	2.69	2.05	0.75	2.27	0.60	103.5	0
11, challenge	PI, $K_c = 0.8$	57.94	57.80	—	93.44	—	105.8	7.5
	PI, $K_c = 4^a$	2.26	2.30	—	2.15	1.70	39.9	0
	MPC	1.42	1.15	—	2.05	0.70	36.6	0

IAE uses the measured temperature, as the scorecard does, and IAE$_T$ the noise-free one. OS is the largest true overshoot of the new setpoint, e_{max} the largest true deviation after the feed step, $t_{0.5}$ the time from the last event until the true temperature stays within ± 0.5 K of setpoint (a dash where it never does), Travel the summed absolute jacket change and Limit the percentage of time with the jacket on a bound. [a]Suggested, $\tau_I = 1.333$ min.

One 30 min session carries the whole argument of this chapter, as shown in Figure 2.4. The measured reactor temperature is plotted against its setpoint, with the jacket temperature between its two limits and the reactant concentration on the same time axis below it, so the interval the jacket spends on its lower limit can be read against the ignition peak and the collapse of the reactant concentration.

2.8 Predictive Control at the Unstable Operating Point

Table 2.7 compares the two controllers on cases 3, 4, 8 and 9 over 10 min each and on the 4 min challenge. The controller is selected before the run starts, which gives the two runs of a case identical events and identical transmitter noise. The PI loop keeps $K_c = 4$ K/K, $\tau_I = 1.333$ min and $\tau_D = 0$, except that the challenge loads $K_c = 0.8$ K/K and suggests 4, while every MPC run uses the settings of Section 2.6.3. Unlike Section 2.7, values are read on the 0.1 min log, as the scorecard reads them, with the IAE over the whole case, so case 3 under the PI loop scores 2.68 K min here and 3.00 K min in Table 2.5. An instant finer than the log is named as an execution.

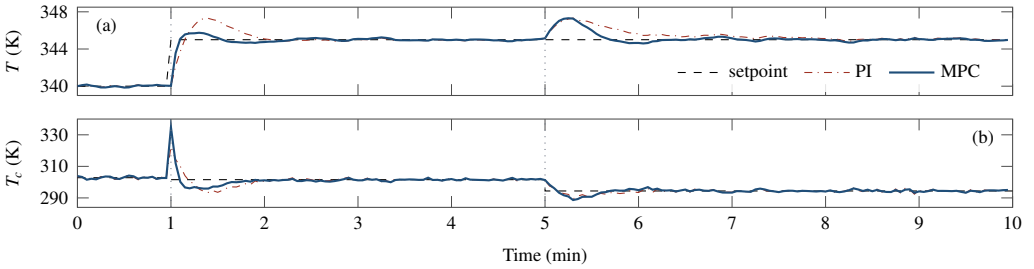


Figure 2.5 Case 9 under the predictive controller and the PI loop with identical transmitter noise, at every controller execution (0.05 min, twice the rate of the trend), the PI loop dash-dotted and the MPC solid. Dotted lines mark the setpoint step to 345 K and the 15 K feed step. (a) Noise-free temperature and setpoint (dashed). (b) Jacket temperature and the steady value that holds the setpoint (dashed).

Case 3 steps the setpoint from 340 to 345 K at 1 min, and the plan made at that execution already holds the maneuver that the negative steady gain requires. Its first block raises the jacket to 334.76 K, its second holds 312.28 K, every later block lies below the 302.82 K that held 340 K, and the last holds 301.84 K, near the 301.92 K at which the model holds 345 K (the balances need 301.57 K). The applied jacket follows, 334.76 K for one execution (311.96 K on the log), below its start from 1.10 min, lowest at 295.90 K. The PI loop makes the same reversal by feedback, from a kick to 322.41 K to a low of 293.46 K, but over the two executions after the kick it lowers the jacket by 8.05 and 4.38 K where the MPC lowers it by 22.80 and 11.77 K. The MPC peaks at 345.75 K against 347.29 K, stays within ± 0.5 K from 0.50 min after the step against 0.90 min, and moves the jacket 25 % less.

In case 4 the estimate $\hat{\delta}$ passes half of the 7.17 K that the 15 K feed step is worth 0.4 min after it and 90 % after 0.6 min, and it ends at 7.08 K, the same under the PI loop, so the difference lies in how the plan uses it. The peaks nearly agree, 342.05 against 342.15 K, but the MPC is back within ± 0.5 K 0.70 min after the upset against 1.70 min, with 23 % less jacket travel.

Case 8 is the one the MPC loses, by 2.2 %. When the setpoint drops to the unreachable 280 K, both controllers put the jacket on its 250 K floor and leave the reactor at 282.51 K, but the MPC lifts the jacket to 255.02 K at 1.55 min and to 261.38 K at 1.95 min. Its model, linearized 43 to 53 K above the reactor, keeps the unstable mode of the middle branch, so the plan made at 1.50 min raises the jacket to 313.69 K within two executions and still predicts a dip to 279.73 K, while the reactor, on its stable cold branch, never falls below 282.51 K. The MPC gathers 0.95 K min more IAE than the PI loop between 1 and 3.3 min, more than it loses by, and its estimate books the error of the model as a load of 63.07 K.

Case 9 repeats case 3 to 5 min, execution for execution, and then warms the feed by 15 K, as shown in Figure 2.5. Both controllers then peak about 2.3 K above the setpoint, 2.27 against 2.33 K, but the MPC is back within ± 0.5 K 0.60 min after the upset against 2.00 min, and the scorecard IAE falls by 39 % with 19 % less jacket travel. Both jackets end near the 294.40 K that holds 345 K against the warmer feed, at 295.04 and 294.92 K. Because the MPC never reads K_c, the challenge under the MPC is the first 4 min of case 4 again, with a worst measured error of 0.41 K in the judged window against 15.07 K for the shipped $K_c = 0.8\,\mathrm{K/K}$ and 0.573 K for the suggested tuning.

Repeating the sweep of Table 2.5 under the MPC measures what the model buys beyond 345 K. On a 0.5 K grid, the largest setpoint step that keeps the noise-free temperature below the alarm is 11.5 K under the PI loop and 13.5 K under the MPC. At 352 K the PI loop peaks at 421.07 K and the MPC at 359.25 K, because from the same first move to the 350 K limit its plan cuts the jacket to 306.27 K at 1.10 min, where the PI loop is at 316.95 K. From a 14 K step up both ignite the reactor, with peaks at 360, 365 and 371 K within 2.1 K of each other. Held at the new setpoint from 10 to 31 min, the MPC moves the jacket a third less but more loosely, a root-mean-square true error of

Table 2.8 Case 10 Under the Two Estimators, Handed Over at the Scripted 3 min and Later

t_h (min)	Estimator	IAE (K min)	IAE_h (K min)	t_2 (min)	Travel (K)	$\min \hat{C}_A$ (mol/m^3)	e_C	e_T (K)
3	Kalman	123.1	54.7	1.60	397	−0.241	0.119	5.37
	MHE	118.0	49.6	1.54	364	0	0.040	5.48
3.5	Kalman	117.5	19.7	1.00	270	−0.212	0.096	2.10
	MHE	112.2	14.7	0.95	170	0	0.041	1.97
4	Kalman	123.1	6.28	0.575	162	−0.212	0.067	0.655
	MHE	123.4	6.47	0.665	165	0	0.042	0.736

IAE is the scorecard's over the 10 min run on the measured temperature, and IAE_h the noise-free one from the hand-over t_h to 10 min. t_2 is the time from the hand-over after which the true temperature stays within ±2 K of 340 K, Travel the summed jacket change, $\min \hat{C}_A$ the lowest concentration estimate of the chosen estimator, and e_C and e_T the root-mean-square errors of its concentration and temperature estimates after the hand-over, on the 0.005 min substep.

0.115 against 0.101 K at 345 K and 0.187 against 0.137 K at 360 K, where the plant is more unstable than the model.

Of 880 plans none was rejected, and a bound was active only in 176 plans of case 8. A plan took 0.05 to 0.22 ms on average under Node.js and 199 and 212 µs in two Chrome runs of case 9, wall-clock values from one desktop computer, at most 0.42 % of a processor core at one minute of process time per second.

2.9 Estimation After a Runaway

Case 10 gives the predictive controller a reactor that has already run away. It opens in manual on the cold branch at 324.475 K with the jacket at 300 K, warms the feed from 350 to 365 K at 1 min as case 5 does, and switches to automatic under the MPC at 3 min. By then the noise-free temperature has peaked at about 484 K and fallen to 418.1 K, with 0.015 mol/m^3 of reactant left. The case loads the MHE and asks for a second run with the Kalman filter, so the two runs share their events and transmitter noise. Table 2.8 reports both, together with runs in which the hand-over is moved to 3.5 and 4 min.

With the hand-over at 3 min, the filter's concentration estimate is below zero from the execution at 2.95 to the one at 4.05 min, 23 executions, down to −0.241 mol/m^3, while the true concentration after the hand-over lies between 0.015 and 0.677 mol/m^3. The MHE holds its estimate on the zero bound instead, with a bound active at 30 executions from 2.95 to 4.40 min, and its concentration error after the hand-over is 0.040 against 0.119 mol/m^3. The plan built on that estimate brings the reactor within ±2 K of setpoint 1.54 min after the hand-over against 1.60 min, lowers the IAE after the hand-over from 54.7 to 49.6 K min, or 9 %, and moves the jacket 8 % less. At the execution at 3.5 min, marked in Figure 2.6, the MHE's plan starts from zero concentration where the filter's starts from −0.237 mol/m^3, and the jacket is set to 273.2 K against 278.2 K, while the true concentration is 0.026 and 0.030 mol/m^3.

The gain is small and depends on when the reactor is handed over. Both estimators feed the same model, linearized 78 K below the temperature at the hand-over, and the MHE's temperature estimate is no better, with an error after the hand-over of 5.48 against 5.37 K. Handed over at 3.5 min, the MHE lowers the IAE after the hand-over from 19.7 to 14.7 K min, but at 4 min it raises it from 6.28 to 6.47 K min. Over ten other noise seeds the MHE gives the lower IAE after the hand-over in all ten runs at 3 and 3.5 min and in none at 4 min, with the 3 min values spread over 54.7 to 55.0 against 49.6 to 49.9 K min. Halving the integration substep gives 53.4 against 48.7 K min at 3 min and 6.34 against 6.52 K min at 4 min, the same ranking. A concentration of the right sign does not repair the model either. In the rows logged at 3.2, 3.3 and 3.4 min the end of the MHE's plan lies below absolute zero, down to −2861 K at 3.3 min, and the filter's plan reaches −2477 K at the same row.

Changing the estimator in the middle of the case moves the plan only where a bound is active. Chosen at 3.5 min of a Kalman run with the noise removed, when the filter places the concentration at −0.238 mol/m^3, the MHE returns zero, moves the temperature estimate from 398.8 to 385.1 K, and

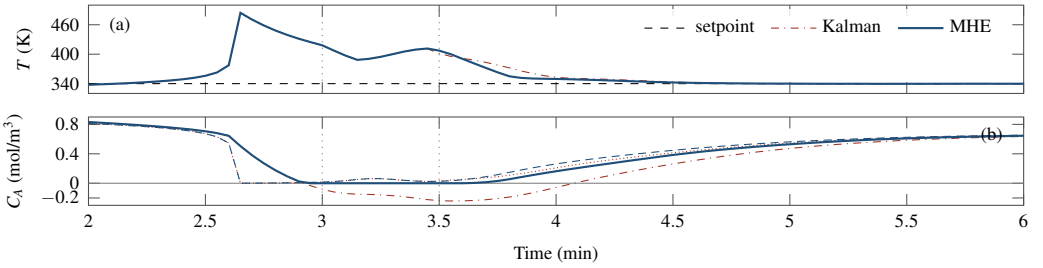


Figure 2.6 Case 10 with the MHE and with the Kalman filter, identical transmitter noise, at every controller execution. Dotted vertical lines mark the hand-over to the predictive controller at 3 min and the execution at 3.5 min discussed in the text. (a) Noise-free temperature with the Kalman filter (dash-dotted) and the MHE (solid), and the setpoint (dashed). (b) The estimate the plan starts from, Kalman dash-dotted and MHE solid, and the true concentration of each run, Kalman dotted and MHE dashed.

sets the jacket 52.3 K away from that of a run that stays on the filter. The difference is the bound acting on the estimate, not a bump in the transfer. Switched back one minute later, the plan again starts from the filter's estimate, which differs from the other run's by 0.086 mol/m^3 only because the plant moved differently in between.

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

2.10 Technology in Practice

What Measures It

The one controlled variable is the liquid temperature, and in a plant it is read by an element in a thermowell. Resistance thermometers are preferred at lower temperatures, where they are most accurate, and become inaccurate above about 700 °C as their sheath degrades. Thermocouples work from about −200 °C to 2320 °C without a power source, but their insulation wears with use and needs periodic checks to keep their accuracy [26]. The chapter's transmitter is instantaneous and adds white noise of standard deviation 0.15 K, while Section 2.3.5 notes that a real thermowell has a time constant of tens of seconds, the same order as the reactor's own.

> **TECHNOLOGY Sensor: A Temperature Element in a Thermowell**
>
> A reactor's temperature is read by a resistance thermometer or a thermocouple inside a thermowell, a pressure-tight sheath that lets the element be replaced without depressurizing the vessel. A platinum resistance thermometer covers about −200 °C to 800 °C and is chosen for its accuracy and precision [26]. The well's mass adds a first-order lag, an air gap at its tip adds a second one, and a lag severe enough can compromise the stability of feedback. A well in a vessel sees less turbulence than one in a pipe, which slows the heat transfer and leaves more fouling on it. A thermocouple most often fails open, and without burnout detection the trend then shows a wildly varying false temperature, while a burnout mode drives the reading high or low on purpose [18].

What Moves It

The chapter manipulates the jacket temperature itself. In a plant that temperature is the setpoint of a slave loop that mixes hot and cold utility or throttles coolant (Section 2.2), and its final elements are the control valves on the utility lines. A cascade works only if the slave variable responds faster than the master, and the slave controller is configured and tuned first [18]. The ideal jacket of the model, which reaches any temperature from 250 to 350 K at once, is the limit of an infinitely fast slave loop, so the jacket limits are the only property of the final element that the chapter's results feel.

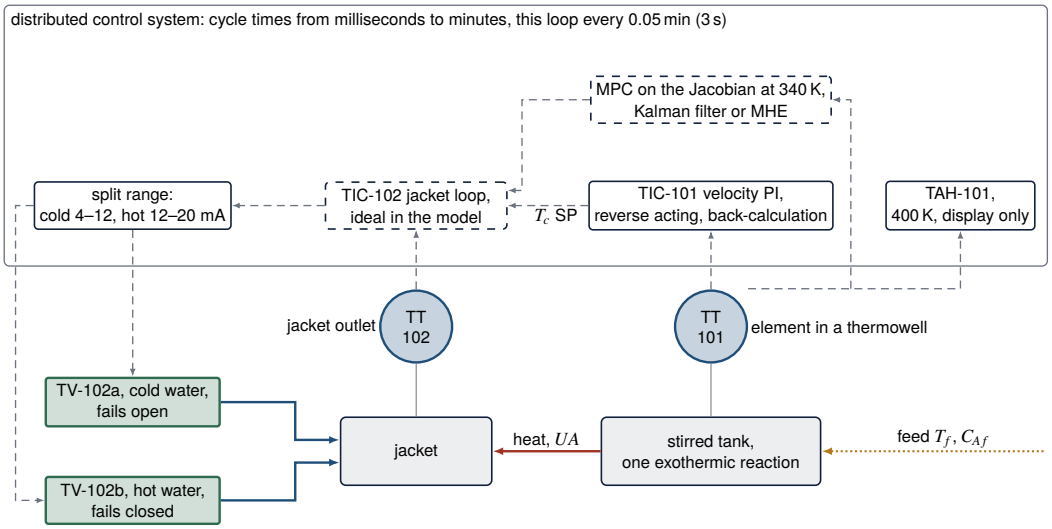


Figure 2.7 The reactor temperature loop on plant hardware. A temperature element in a thermowell (TT-101) feeds the velocity-form PI loop TIC-101 in a distributed control system, or the dashed predictive controller and its estimator when they are selected. Either one sets the jacket temperature, which a plant holds with a slave loop (TIC-102, dashed) driving a cold-water valve that fails open and a hot-water valve that fails closed through a split range. The runaway alarm only displays. Notice that the chapter's ideal jacket removes everything below the jacket setpoint, so the slave loop, its sensor and both valves are the hardware the model does not carry.

TECHNOLOGY **Final Element: Split-Range Jacket Valves**

A jacket that must both heat and cool can be fed by a hot-water valve and a cold-water valve split-ranged on one controller output, the cold valve on the lower half of the 4 to 20 mA signal and the hot valve on the upper half, so the two are never open together. Each fail position is chosen for safety. The hot valve fails closed and the cold valve fails open, so a lost controller signal or a lost air supply gives the reactor full cooling [18]. In the 2007 accident water was injected into the jacket and allowed to boil, and a supply valve failing closed or a drain valve failing open was among the single points of failure of the cooling [16].

Where the Controller Runs

A reactor temperature loop runs in a distributed control system, where supervisory and lower-level loops execute continuously with cycle times from milliseconds to minutes, and where loops can be nested so that the setpoint of one comes from the variable of another [27]. The chapter's PI loop executes every 0.05 min, or 3 s, in the velocity form with back-calculation, and the predictive controller takes the same slot when it is selected. Figure 2.7 draws the loop on plant hardware.

> **IN PRACTICE Controlling the Hottest Reading**
>
> An exothermic reactor can carry several temperature transmitters, with a high-select function passing only the hottest reading to the controller that throttles the coolant. Three transmitters on one variable can instead feed a median select, which keeps a valid measurement when one of them fails [18]. The reading is itself a point of failure. The 2007 investigation listed a faulty temperature indication among the single-point failures of the cooling system, and the procedures then in use gave no emergency instruction for a loss of cooling [16]. The chapter's one perfectly mixed temperature has no hot spot for a selector to find.

What Is Modernizing It

The first development is the predictive controller inside the control system itself. A survey of industrial practice counted the integration of predictive control into distributed control platforms among the major recent changes, and noted that an unstable process cannot in general be represented by the impulse-response models of early products, while state-space or autoregressive models can [20]. That is why the chapter's controller is built on the state-space Jacobian of Equation (2.17) rather than on a step test. The second is the estimator. The same survey found the bias update of industrial controllers the best assumption only for stable plants without disturbance information, and named the Kalman filter as the richer framework for feedback, with better performance shown for a state-space model and an optimal observer [20]. The chapter's filter carries the input disturbance δ of Equation (2.18), and its moving horizon estimator adds bounds on the states.

2.11 Limitations

Nothing in this reactor is intended for design work. The model carries two states and no vapor space, so it cannot represent boiling, pressure rise, relief or the secondary decomposition that turns an industrial excursion into an explosion, and the temperatures it reports above the boiling point of a real liquid are an index of how far the balance has diverged rather than a prediction. The lumped-parameter assumption removes the hot spot, which is the mechanism behind a large share of real runaways. The jacket is ideal, so there is no coolant residence time, no fouling and no slave loop, and the cascade structure that industrial reactors actually use cannot be built or tuned [10]. The transmitter is instantaneous, so the thermowell lag that limits achievable tuning on a real vessel is absent. The noise generator restarts from one fixed seed in every case, so a noisy run replays exactly but a learner cannot sample a second noise realization. The setpoint range runs from 280 to 450 K although the jacket limits hold the reactor only between 282.5 K and 416.4 K at the design feed, and nothing in the application marks a setpoint outside that interval as unreachable. The actuator-limit case asks for 280 K, and the reactor stops at 282.51 K, the coldest the jacket can hold. The reaction is a single first-order step with fixed stoichiometry, so nothing in the model changes when the feed composition changes except the amount of heat available. The four features that were kept, the Arrhenius rate, the two-state balance, the hard jacket limits and the short control period, are there because the loop has to be tuned around every one of them.

The predictive controller adds limits of its own. Its model is Equation (2.11) frozen at 340 K, right about the first response to a jacket move and wrong about what the Arrhenius term changes away from that point, which costs it case 8 and leaves it unable to prevent ignition on steps of 14 K and more. It has no temperature constraint, sees a feed upset only through $\hat{\delta}$, keeps its nominal model when UA or k_0 is changed, and was compared on one noise seed. Its moving horizon estimator bounds the states of that same linear model, so it corrects the sign of the concentration estimate but not the model, and after the 3 min hand-over of case 10 its temperature estimate is no closer to the truth than the filter's, 5.48 against 5.37 K root-mean-square. The relay experiment cannot measure this loop at the 0.05 min control period.

OPEN CHALLENGE When a Nonlinear Predictive Controller Pays

A 2003 survey of industrial predictive control found nonlinear applications mostly in chemicals, polymers and air and gas processing, and much smaller in size and scope than linear ones, likely because of the computing the nonlinear algorithms need. Since nonlinear processes are hard to model from plant data, it expected first-principles models to become necessary, and it named among the open questions of practice when the added expense of a nonlinear controller can be justified [20]. The linear controller of this chapter loses case 8 and cannot prevent ignition on large steps. Progress would be a test, made before commissioning, of how far a linearized model can be trusted on a plant whose gain changes sign.

2.12 Discussion

The gap between a controller that reaches setpoint and a controller that is safe on the way there is not an artifact of a small simulation. Surveys of thermal-runaway incidents report inadequate temperature control and inadequate understanding of the heat release as recurring causes, and report that most events happened during routine production rather than during an unusual operation [13, 14]. Loss prevention as a discipline grew out of exactly this observation, that the hazard is carried by the process and shows up when the operating point moves rather than when someone makes a mistake [15]. The simulation reproduces the structure of that hazard in its smallest form. The plant has an operating region that exists only under feedback, the boundary of that region is crossed by a disturbance of a few percent, and the excursion that follows is over before an operator could act.

The control literature has answered the problem in several ways that the application makes it natural to ask about. Input-output linearization cancels the Arrhenius term explicitly and delivers a loop whose gain no longer moves with the operating point [6]. Robust proportional-integral designs give conditions under which classical feedback stabilizes the reactor over a stated region despite kinetic uncertainty [7]. Gain scheduling interpolates a family of linear controllers across the envelope and is the industrial compromise between the two [8], and constrained predictive control writes the thermal limit into the optimization as a constraint [9]. None of these is implemented here as cited, since the predictive controller of Section 2.6.3 is linear and writes no thermal limit into its plan. What is implemented is the plant that motivates all of them, with the measurement needed to show that a single fixed tuning cannot cover its operating range.

Both balances are differentiable, and the ignition boundary is a property of the model rather than a constraint added to it. A nonlinear predictive controller solving Equations (2.2) and (2.3) over a receding horizon could therefore treat the alarm and the jacket range as explicit constraints instead of meeting them as a saturated actuator [28]. A moving horizon estimator on the same two equations could reconstruct C_A from the temperature history, which is the soft-sensing problem of an industrial reactor with only a thermowell. The bounded linear estimator of Section 2.6.4 marks the limit of the one-point approach, since after a runaway it corrects the sign of C_A while the plan built on it still predicts temperatures below absolute zero.

2.13 Conclusion

The reactor of this chapter is an exothermic jacketed stirred tank simulated in a browser with two states, an Arrhenius first-order reaction, a jacket temperature limited to 250 to 350 K as the manipulated variable, two feed upsets and a 400 K runaway alarm. Analysis of the shipped constants places the ignition fold at $T_c = 303.229\,\mathrm{K}$, the extinction fold at $T_c = 298.080\,\mathrm{K}$ and a Hopf bifurcation on the hot branch at $T_c = 306.220\,\mathrm{K}$, so three steady states exist only inside a jacket window 5.15 K wide and the ignited branch at the published jacket temperature of 300 K is an unstable focus rather than a second place to operate. Every reactor temperature between 335.654 K and 379.611 K is open-loop unstable, including the application's own default setpoint of 340 K. A step test in manual identifies $K_p = 2.71\,\mathrm{K/K}$ and $\tau_p = 1.29\,\mathrm{min}$ with no measurable dead time, and internal model control turns that fit into the shipped tuning of $K_c = 4\,\mathrm{K/K}$ and $\tau_I = 1.333\,\mathrm{min}$. With that tuning held fixed, an 11 K setpoint step overshoots by 12.18 K while a 12 K step overshoots by 69.07 K and trips the alarm

at every noise seed tried, and on a 14 K step reducing the gain to 1.6 K/K raises the peak from 450.78 to 460.13 K and leaves the reactor in a relaxation cycle through the alarm. Feedback absorbs a feed concentration upset that ignites the same reactor in manual, holding the peak deviation to 1.06 K against an excursion to 515.38 K. A linear predictive controller on Equation (2.11) at 340 K plans the warm-then-cold jacket move and raises the largest step kept below the alarm from 11.5 to 13.5 K, but loses at the jacket floor by 2.2 %. Its moving horizon estimator keeps the concentration nonnegative after a runaway and lowers the error after a 3 min hand-over from 54.7 to 49.6 K min, but not after a 4 min one. The natural extensions are a coolant slave loop, a thermowell lag, a choice of noise seed and a gain-scheduled or nonlinear predictive controller for comparison against the fixed tuning. The application is an accessible platform for measuring how far a single linear controller can be trusted on a plant whose gain changes sign.

KEY POINTS

- The reaction's heat climbs faster with temperature than the jacket can remove it, so the default 340 K exists only while a controller holds it.
- The steady-state gain is 1.72 K/K at the nominal point and −5.56 K/K at 340 K, yet a reverse-acting loop holds it, because a warmer jacket always warms the reactor first.
- With one fixed tuning an 11 K setpoint step overshoots by 12.18 K, while a 12 K step overshoots by 69.07 K and trips the alarm.
- A linear predictive controller raises the largest step kept below the alarm from 11.5 to 13.5 K but loses at the jacket floor.

Where else this applies. Chapter 4 controls a wheel at a friction peak where the process gain changes sign, and Volume 8, Chapter 2 holds a runaway off with the feed instead of the jacket.

References

[1] C. van Heerden, "Autothermic processes," *Industrial & Engineering Chemistry*, vol. 45, no. 6, pp. 1242–1247, 1953.

[2] O. Bilous and N. R. Amundson, "Chemical reactor stability and sensitivity," *AIChE Journal*, vol. 1, no. 4, pp. 513–521, 1955.

[3] R. Aris and N. R. Amundson, "An analysis of chemical reactor stability and control—I," *Chemical Engineering Science*, vol. 7, no. 3, pp. 121–131, 1958.

[4] U.S. Chemical Safety and Hazard Investigation Board, "Hazard investigation: Improving Reactive Hazard Management," U.S. Chemical Safety and Hazard Investigation Board, Washington, DC, Tech. Rep. 2001-01-H, 2002. [Online]. Available: https://www.csb.gov/assets/1/20/reactivehazardinvestigationreport.pdf

[5] A. Uppal, W. H. Ray, and A. B. Poore, "On the dynamic behavior of continuous stirred tank reactors," *Chemical Engineering Science*, vol. 29, no. 4, pp. 967–985, 1974.

[6] C. Kravaris and C.-B. Chung, "Nonlinear state feedback synthesis by global input/output linearization," *AIChE Journal*, vol. 33, no. 4, pp. 592–603, 1987.

[7] J. Alvarez-Ramirez and R. Femat, "Robust PI stabilization of a class of chemical reactors," *Systems & Control Letters*, vol. 38, no. 4-5, pp. 219–225, 1999.

[8] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

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

[10] M. G. Balaton, L. Nagy, and F. Szeifert, "Jacket temperature control of a batch reactor using three different temperature levels," *Industrial & Engineering Chemistry Research*, vol. 52, no. 5, pp. 1939–1946, 2013.

[11] 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.

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

[13] P. F. Nolan and J. A. Barton, "Some lessons from thermal-runaway incidents," *Journal of Hazardous Materials*, vol. 14, no. 2, pp. 233–239, 1987.

[14] R. Saada, D. Patel, and B. Saha, "Causes and consequences of thermal runaway incidents—will they ever be avoided?" *Process Safety and Environmental Protection*, vol. 97, pp. 109–115, 2015.

[15] T. A. Kletz, "The origins and history of loss prevention," *Process Safety and Environmental Protection*, vol. 77, no. 3, pp. 109–116, 1999.

[16] U.S. Chemical Safety and Hazard Investigation Board, "Investigation report: T2 Laboratories, Inc. runaway reaction, Jacksonville, Florida, December 19, 2007," U.S. Chemical Safety and Hazard Investigation Board, Washington, DC, Tech. Rep. 2008-3-I-FL, 2009. [Online]. Available: https://www.csb.gov/assets/1/20/t2_final_copy_9_17_09.pdf

[17] ——, "CSB applauds AIChE's response to including reactive hazard awareness in college chemical engineering curriculum," News release, 2012. [Online]. Available: https://www.csb.gov/csb-applauds-aiches-response-to-including-reactive-hazard-awareness-in-college-chemical-engineering-curriculum/

[18] 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.

[19] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

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

[21] 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.

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

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

[25] K. J. Åström and T. Hägglund, "Automatic tuning of simple regulators with specifications on phase and amplitude margins," *Automatica*, vol. 20, no. 5, pp. 645–651, 1984.

[26] A. Boghossian, J. Brown, and S. Zak, "3.2: Temperature sensors," in *Chemical Process Dynamics and Controls*, P. Woolf, Ed. University of Michigan open textbook, LibreTexts, 2023.

[27] 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.

[28] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

2.14 Exercises

2.1 ○ Concept check · 2 min

At the default setpoint of 340 K the steady-state gain from the jacket to the reactor temperature is −5.56 K/K. The PI loop is still reverse acting there, with the error $T^{\mathrm{sp}} - \hat{T}$ and a positive K_c, as it is on the cold branch where the gain is +1.72 K/K. Why does that loop hold 340 K?

(a) A warmer jacket always warms the reactor first, and the controller acts on that response long before the reactant runs down

(b) The integral term reverses the sign of the loop on the middle branch, so the proportional term only has to damp the response

(c) The negative gain belongs to the noise-free temperature, while the controller acts on a measurement whose gain stays positive

(d) The jacket limits keep the reactor on the cold branch, where the steady-state gain has the sign the controller assumes

2.2 Concept check · 1 min

With the jacket held at 300 K, the reactor has steady states at 324.475, 350.006 and 369.705 K. At which of them can it rest with the jacket held there and no controller?

(a) At all three, because each of them balances heat generation against heat removal

(b) Only at 324.475 K, because the middle state is a saddle and the upper one an unstable focus

(c) At 324.475 and 369.705 K, the outer two crossings of the van Heerden construction

(d) Only at 369.705 K, because the transmitter noise ignites the cold state sooner or later

2.3 Concept check · 2 min

In case 10 the predictive controller takes over a reactor that has already run away, once with the Kalman filter and once with the moving horizon estimator (MHE). Which statement agrees with Section 2.9?

(a) The MHE estimates the temperature far better, so its plan lowers the IAE after every hand-over time tried

(b) The MHE plans with the nonlinear balances, so its plan no longer predicts temperatures below absolute zero

(c) The MHE holds the concentration estimate at zero where the filter's goes negative, and its gain depends on the hand-over time

(d) The two estimators disagree in every case, because the MHE fits 15 executions while the filter uses only the last

2.4 Theory and calculation · 7 min

Take the default constants: $q/V = 1/\text{min}$, $m\Delta H/(\rho C_p) = 209.205\,\text{K}$ per mol/m^3, $UA/(V\rho C_p) = 2.09205/\text{min}$, $k_0 = 7.2\times10^{10}/\text{min}$, $E/R = 8750\,\text{K}$ and $C_{Af} = 1\,\text{mol/m}^3$.

(a) Find the rate constant k and the steady concentration at $T = 340\,\text{K}$. *Checkpoint:* $k = 0.4793/\text{min}$ and $C_A = 0.6760\,\text{mol/m}^3$.

(b) Evaluate the four entries of the Jacobian of Equation (2.11) at this state. *Checkpoint:* -1.4793, -0.02453, 100.28 and 2.0388.

(c) Find its eigenvalues and the steady-state gain from the jacket to the reactor temperature, $G = -cA^{-1}b$ with $b = [0\ \ 2.0921]^{\mathsf{T}}$ and $c = [0\ \ 1]$. *Checkpoint:* $+1.0766$ and –0.5171/min, and $G = -5.56\,\text{K/K}$.

2.5 Theory and calculation · 8 min

Use the model A_m, b_m of Equation (2.17) at 340 K, with the reactor temperature as the output and the loop closed with the reverse-acting law $u = K_c e$ or the continuous PI law $u = K_c(e + \int e\,\mathrm{d}t/\tau_I)$. Ignore the control period.

(a) Write the transfer function from the jacket to the reactor temperature and find the smallest proportional gain that stabilizes the loop. *Checkpoint:* $2.0921(s+1.4793)/(s^2 - 0.5595s - 0.5567)$, and $K_c > 0.267\,\text{K/K}$.

(b) With $\tau_I = 1.333\,\text{min}$, find the characteristic polynomial and the smallest stabilizing K_c from the Routh test. *Checkpoint:* $K_c > 0.57\,\text{K/K}$.

(c) The challenge loads $K_c = 0.8\,\text{K/K}$. Is that gain stable at 340 K, and why can it still fail after a feed upset? *Checkpoint:* Stable, yet the bound is 1.04 K/K at 345 K.

2.6 Theory and calculation · 5 min

The step test of Section 2.5 gives $K_p = 2.71\,\text{K/K}$, $\tau_p = 1.29\,\text{min}$ and $\theta_p = 0$.

(a) Find the internal model control tuning with $\tau_c = 0.1\,\tau_p$. *Checkpoint:* $K_c = 3.7\,\text{K/K}$, $\tau_I = 1.29\,\text{min}$, $K_I = 2.9/\text{min}$.

(b) The loop runs at steady state at 340 K with the jacket at 302.818 K and $K_c = 4\,\mathrm{K/K}$. Find the smallest setpoint step whose proportional kick puts the jacket on its 350 K limit at the first execution. *Checkpoint:* 11.8 K.

(c) While the jacket is clipped, by what fraction of the gap does the integral relax at each execution with $\tau_I = 1.333\,\mathrm{min}$, and how long does it take to close half the gap? *Checkpoint:* 3.75 %, about 18 executions or 0.9 min.

2.7 App challenge · 8 min

Load scenario 2, Open-loop actuator step, and change no setting. The reactor starts in Manual at 324.5 K with the jacket at 300 K, and the jacket steps to 302 K at 1 min. Run to 20 min at the default seed and export the CSV file.

(a) Read the true temperature, column `true_pv`, at 6 min and at 20 min, and the largest value in the column. *Checkpoint:* 328.79 and 328.70 K, largest 328.85 K near 4.7 min, within 0.01 K.

(b) Compute the chord gain of the step, and find the time after the step at which `true_pv` passes 63.2 % of its change. *Checkpoint:* 2.11 K/K, and between the rows at 1.9 and 2.0 min, about 0.94 min after the step.

(c) Compare the chord gain with the tangent gain of 1.72 K/K at the start and with the 2.71 K/K of a 3 K step, and say what a 4 K step would give. *Checkpoint:* 23 % and 58 % above the tangent; no fit at all.

2.8 App challenge · 8 min

Load scenario 1, Reference operating point. Before Run, set the setpoint to 340 K and switch to Auto, so that the reactor starts at the steady state of 340 K with the jacket at 302.8 K. Use the default seed and the default tuning. Run, pause at any time between 1 and 3 min, set the setpoint to 351 K, run for 4 min more and export the CSV file. Then press the reload button, make the same settings before Run and repeat with a step to 352 K.

(a) For the 351 K run, read the largest `true_pv` after the step and the lowest `mv`. *Checkpoint:* 363 K, within 0.6 K; lowest jacket 250.0 to 254.5 K, on the limit for at most one row.

(b) For the 352 K run, read the same two values and count the rows at which `mv` is 250 K. *Checkpoint:* 405 to 423 K; 5 or 6 rows, 0.5 to 0.6 min on the limit.

(c) Explain why one kelvin of setpoint separates the two outcomes.

2.9 App challenge · 7 min

Load the challenge, Challenge · return to target, the eleventh case. The proportional gain is loaded at 0.8 K/K, the feed temperature steps from 350 to 365 K at 1 min, and the run is judged on the measured temperature from 2.64 to 4 min. Use the default seed, run each case to 4 min, and export the CSV file each time. For each run find the largest $|\mathtt{pv} - \mathtt{sp}|$ in the judged window and read the IAE tile of the scorecard.

(a) Run the case as loaded. *Checkpoint:* 15.07 K and 57.94 K min.

(b) Press Reveal suggested tuning, press Reset and run again. *Checkpoint:* 0.573 K and 2.26 K min.

(c) Press the reload button, set the Controller to MPC and run again. *Checkpoint:* 0.41 K and 1.42 K min.

Accept an error within 0.01 K and an IAE within 0.02 K min.

2.10 Application area · 4 min

The simulated transmitter is instantaneous, while a real thermowell has a time constant of tens of seconds. As a rough guide, treat the thermowell as a delay τ_d and use the limit $p\,\tau_d < 1$ of Section 1.4.7 for a loop around one unstable pole p.

(a) Find the largest such delay at 340 K, where $p = 1.0766/\text{min}$, and at 350 K, where Table 2.3 gives $p = 2.834/\text{min}$. *Checkpoint:* 55.7 s and 21.2 s.

(b) Compare both with the 3 s control period and with a thermowell of 30 s, and state what that means for operating higher on the middle branch. *Checkpoint:* The control period uses about 5 % of the allowance at 340 K.

2.11 Application area · 3 min

The adiabatic ceiling $T_f + (m\Delta H/\rho C_p)\,C_{Af}$ is the temperature the reactor would reach if all of its feed reacted with no cooling.

(a) Find the ceiling at the design feed and after the feed concentration rises to $1.2\,\text{mol/m}^3$, with $T_f = 350\,\text{K}$. *Checkpoint:* 559.2 K and 601.0 K.

(b) With the jacket at 300 K the cold branch disappears at $C_{Af} = 1.141\,\text{mol/m}^3$. How far is the design feed from that fold, and what does the app's 400 K alarm lack that a plant's protection would need? *Checkpoint:* 14 %.

2.12 Review the agent · 5 min

An AI assistant reviews Table 2.5 and writes: "The 12 K step overshoots by 69 K because $K_c = 4$ is too aggressive near the fold. Before moving the setpoint from 340 to 354 K, lower K_c to 1.6 K/K. The linearized loop at 340 K is stable for any K_c above 0.57 K/K, so 1.6 keeps a margin of almost three, and the approach will be slower and safer." Find the flaw and say how you would check it.