INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 2 Practical PID Control

CHAPTER 4

Steam-Heated Exchanger Laboratory with Hidden Parameters for Step Testing and Tuning

Pages 87–118

Edited by John D. Hedengren

How to cite this chapter. "Steam-Heated Exchanger Laboratory with Hidden Parameters for Step Testing and Tuning," in *Practical PID Control*, John D. Hedengren, Ed., vol. 2 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 4, pp. 87–118.

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

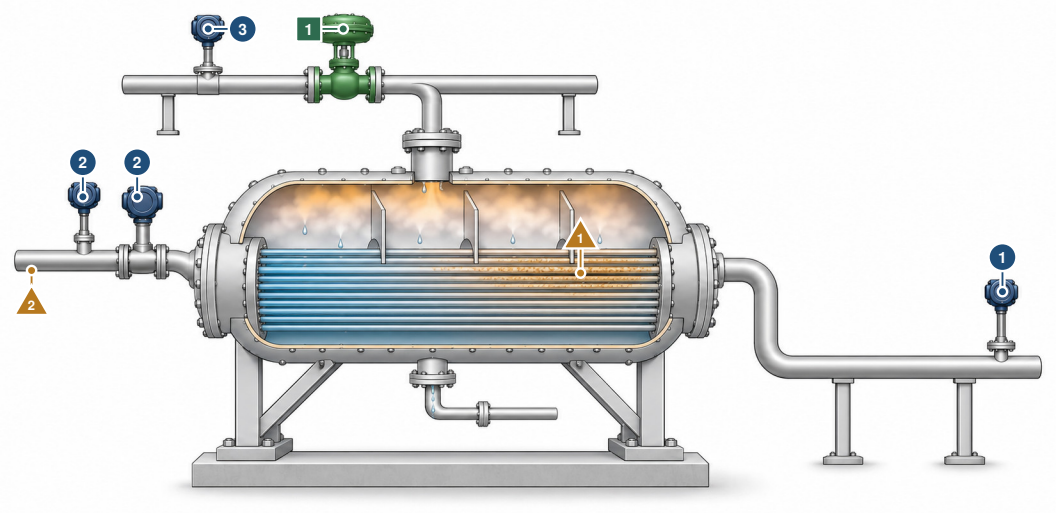
The app for this chapter
https://apmonitor.com/apps/heat-exchanger/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v02-practical-pid/#heat-exchanger

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

CHAPTER 4

Steam-Heated Exchanger Laboratory with Hidden Parameters for Step Testing and Tuning



The controller moves 1 the steam valve

The controller reads 1 the outlet temperature, read far down the pipe 2 the feed's temperature and flow 3 the steam header pressure

What disturbs it 1 deposits building up on the tubes 2 changes in feed rate and feed temperature

The controller moves only the steam valve, and it reads the product temperature well past the bundle, so each move shows up late. Illustration, not to scale.

IN PLAIN TERMS **Tuning a Steam Heater From One Test**

The main idea. Steam flows around a bundle of tubes and warms the liquid inside them. A controller opens or closes the steam valve to hold the liquid at a set temperature.

The challenge. Its settings come from one small test that moves the valve and times the reply, but the reply changes with how much liquid flows. The thermometer sits far down the pipe, and a wait read too short makes the controller too bold, so the valve swings back and forth.

Why it matters. Steam is the most common way a plant warms a stream, and a misread test wears the valve for no gain.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on a real steam heater (Section 4.2), the try-it segment of case 8 and the history of predictive control (Section 4.5), the try-it segment of case 14 (Section 4.6), the fact in the results (Section 4.7) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 4.7), Technology in Practice (Section 4.8), Limitations (Section 4.9) and the two open challenges

The control question The constants of a steam-heated exchanger are hidden, and its outlet temperature loop has a dead time and a gain that both move with the feed rate. What does one valve step test identify on such a plant, and how far can a tuning rule be trusted with the three numbers it returns?

Learning objectives

- Write the energy balances of the exchanger and compute its duty, steam flow and transport delay at an operating point.
- Identify a first-order-plus-dead-time model from a valve step and explain why its dead time exceeds the transport delay.
- Apply the SIMC and reaction-curve rules to the fit and predict what a misread dead time does to each.
- Distinguish a duty limit, fouling and valve stiction from a tuning problem by the controller output.
- Compare a single loop, a cascade, a heat-balance feedforward and a predictive controller built from the same identified model.

Prerequisites Sections 1.4.4 (step test and fit), 1.4.5 (tuning rules), 1.4.9 (windup) and 1.6.2 (the algorithm of the apps). Section 1.6.3 is needed only for the predictive controller. Chapter 2 introduces the step test on a simpler plant.

The app https://apmonitor.com/apps/heat-exchanger/

See also Volume 1, Chapter 2 and Volume 1, Chapter 4 for step tests on a tank and on a heated device. Volume 3, Chapter 2, Volume 3, Chapter 3 and Volume 3, Chapter 6 for the choice of structure, the speed of an inner loop and feedforward with a feedback trim. Volume 13, Chapter 7 for a transport delay that stretches with flow, and Volume 13, Chapter 4 for exchanger bypasses that set reactor inlet temperatures.

4.1 Introduction

Steam condensing on one side of a tube bundle is the most common way a process plant adds heat to a stream. The service appears as feed preheat ahead of a reactor, as reboiler duty on a distillation column, as product conditioning before storage, and as jacket service on a fermenter, and the temperature loop that regulates it is usually the first controller a chemical engineering student designs [1, 2]. That loop is also the one on which a formula-sheet approach fails first, for four reasons that are physical rather than incidental. A transport delay separates the bundle from the transmitter. The steady-state gain is not a constant because the driving force is a log-mean temperature difference that collapses as the outlet approaches the tube wall. The tube-side film coefficient rises with the 0.8 power of flow [3, 4]. The actuator is a valve whose capacity depends on the square root of a pressure drop that the valve itself destroys as it opens.

Those four mechanisms are consequences of three fundamentals that a first course states in general form. Conservation of energy applied to the fluid, to the tube metal and to the shell gives one state equation per inventory, and the number of inventories in series sets the order of the response and therefore the phase lag a controller must survive [1]. Heat crosses each surface in proportion to a driving force, so the input-output gain of the loop is the derivative of a driving force rather than a constant, and it changes with the operating point. Time-scale separation orders the resulting lags, with the tube metal and the process inventory slow, the shell and the valve fast, and the outlet line a pure transport delay, which is what allows a fifth-order plant to be identified as a first-order-plus-dead-time (FOPDT) model at all [5]. The experiments of this chapter therefore cover process order, dead time, empirical identification from a reaction curve, PID tuning and anti-windup, sampled measurements, cascade structure, feedforward from a measured disturbance, gain scheduling, and the distinction between a saturated loop and a detuned one.

A simulator that supplies a FOPDT transfer function and asks for a controller exercises none of that behavior. The application of this chapter inverts the arrangement. Exchanger E-101 is modeled from first principles, and the constants of that model, including the two heat-transfer coefficients, the metal mass, the shell capacitance, the valve capacity and the transport volume, are withheld from the parameter panel, which lists only data-sheet and instrument values. What is available is what a real

unit is instrumented for, an outlet temperature transmitter, an inlet temperature transmitter, a feed flow meter, a steam flow meter, a header pressure gauge and a valve position. A model is obtained the way a commissioning engineer obtains one, by perturbing the plant and reading the answer off a trend, the step test of Section 1.4.4.

Dynamic exchanger models of this kind are well established. Distributed and lumped formulations appear in the simulation literature [6], fouling has been treated as a series resistance since the Kern and Seaton deposition model [7], compound cascade and feedforward structures are the industrial default [8], and nonlinear and predictive controllers have been applied to the same unit [9]. However, the published models are presented with their parameters supplied, which removes the identification step that determines whether the resulting controller works at all. This chapter works with an interactive steam-heated exchanger in which the plant parameters are withheld. The identified model, the simulator equations and the controller's internal model are kept separate, so a misread trend produces a badly tuned loop rather than a hidden correction.

The chapter is organized as follows. Section 4.2 gives the history of the loop and of its models. Section 4.3 gives the five-state model, its instrument chain and every constant needed to rebuild the simulator. Section 4.4 identifies a FOPDT model from a valve step and shows how its three numbers move with the feed rate. Section 4.5 describes the PID algorithm, the single-loop, cascade and feedforward structures, the two tuning rules, and a model predictive controller that plans on the identified model, with an alarm band that can be tightened by the misfit of that model. Section 4.6 derives the three properties that govern the tuning, a dead time and a gain that move with load and a duty limit that looks like a tuning problem. Section 4.7 reports the measured comparisons, which start in manual operation and end in fault diagnosis. Section 4.8 names the devices and the platform of such a loop in a plant. Sections 4.9 and 4.10 state what the model cannot do and what the loop is worth on a plant, and Section 4.12 holds the exercises.

4.2 Background

Automatic regulation of heat-transfer duty preceded the theory that describes it. Nineteenth-century steam users throttled supply valves by hand against a thermometer, and the earliest self-acting devices were direct thermal regulators in which a filled bulb drove the valve with no external power. The development that matters for this chapter is the pneumatic three-term controller of the 1930s, which added reset and pre-act to a proportional response and made it possible to hold a temperature at setpoint rather than near it [10]. Ziegler and Nichols published their reaction-curve and continuous-cycling settings in 1942 for this class of loop [11], and the reaction curve remains the archetype of the step test the application asks the reader to perform. The model-based rules that displaced them derive the same three settings from an assumed model and an explicit closed-loop speed, which makes the quality of the identification the dominant term in the result [12, 13].

The modeling literature followed later. Steady-state exchanger design had long rested on film correlations of the Dittus-Boelter and Sieder-Tate form [3, 4], but the dynamic behavior a controller sees, metal capacitance in series with fluid capacitance plus a transport lag, was analyzed explicitly only as digital simulation became routine [6]. Fouling followed a similar path. The Kern and Seaton deposition-minus-removal picture, reviewed fifty years on by Müller-Steinhagen [7], expresses the deposit as a series resistance and makes the fouling factor a quantity that a controller output can be used to estimate.

The exchanger is a standard control testbed because a single PID loop is not the best available structure for it. Compound feedforward and cascade structures remain the industrial default on outlet temperature services [8], and the nonlinear input-output gain has motivated a long line of predictive and nonlinear controllers [9]. Surveys of control education report a persistent gap between that literature and the textbook treatment, namely the absence of an accessible plant on which the single loop can be measured before the structures that improve it are introduced [14, 15]. Volume 3,

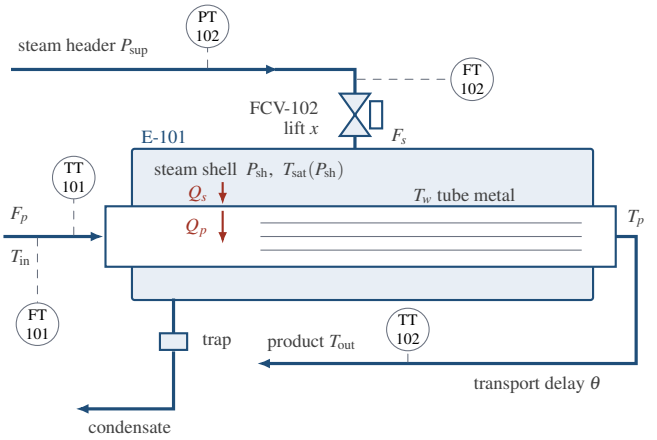


Figure 4.1 Exchanger E-101 with the steam header, control valve FCV-102, the condensing shell, the tube bundle and the outlet transport line. The process stream enters the tubes at T_{in} with flow F_p, receives duty Q_p from the tube metal at T_w, and leaves the bundle at T_p. Steam admitted at lift x against header pressure P_{sup} raises the shell to P_{sh} and delivers duty Q_s to the metal. TT-102 reads the product θ minutes after it leaves the bundle.

Chapter 2 takes up the choice among those structures for a given disturbance, and this chapter measures the single loop first.

> **DID YOU KNOW? A Steam Heater in Acid Service**
>
> In one phosphoric acid plant, studied in 2023, steam at 116 to 125 °C heats the acid from about 70 to about 80 °C at volume flows of up to 3407 m^3/h, and the fouling resistance of the exchanger was computed from plant readings over runs of up to 122 h [16].

4.3 Process Model

Exchanger E-101 heats a process liquid from a nominal 40 °C to a setpoint of 85 °C using saturated steam condensing in the shell. Steam arrives from a header at 4 barg through control valve FCV-102. Condensate leaves through a trap, which is represented only through its effect on the condensing area. The arrangement, the instruments and every symbol that appears in the equations below are shown in Figure 4.1. The controlled variable is the product outlet temperature reported by transmitter TT-102 downstream of the bundle, and the manipulated variable is the position of FCV-102 in every control structure the application offers.

The variables a controller interacts with are classified in Table 4.1. Four alarm limits bound normal operation, TAHH-102 at 95 °C, TAH-102 at 90 °C, TAL-102 at 79 °C, and FAL-102 on steam flow at 250 kg/h. The alarms do not trip the simulation. They latch and annotate the record, which is enough to show that an aggressive test has left the product off specification.

The plant is written in operator units, temperature in °C, pressure in bar absolute, process flow in kg/h, steam flow in kg/min, valve position in percent, with time in minutes. Five states are integrated. Each follows from a conservation law applied to one inventory, the process fluid in the bundle, the tube metal, the steam in the shell, the valve stem and the fouling deposit, which is the standard route from physical fundamentals to a state-space process model [2].

Table 4.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Most Scenarios Start.

Symbol	Role	Nominal	Range
T_{out} (TT-102)	CV, measured	85 °C	60–95 °C SP
u (FCV-102)	MV	44.9 %	0–100 %
F_s (FT-102)	inner CV, measured	1727 kg/h	0–4000 kg/h span
F_p (FT-101)	DV, measured	20 000 kg/h	12 000–30 000 kg/h
T_{in} (TT-101)	DV, measured	40 °C	25–55 °C
P_{sup} (PT-102)	DV, measured	4 barg	1.5–5.5 barg
ϕ fouling	DV, unmeasured	0 %	0–150 %
T_p	state	85.0 °C	—
T_w	state	111.6 °C	—
P_{sh}	state	2.60 bar abs	0.3–P_{sup}
x	state	44.9 %	0–100 %
R_f	state	0	≥ 0

4.3.1 Energy Balances

The process fluid in the bundle is treated as a single well-mixed inventory M_p that receives heat from the tube metal and exchanges sensible heat with the throughput, as shown in Equation (4.1).

$$M_p C_p \frac{\mathrm{d}T_p}{\mathrm{d}t} = \dot{m}_p C_p \left(T_{\text{in}} - T_p\right) + Q_p \tag{4.1}$$

where M_p is the process-side thermal inventory in kg, C_p is the process heat capacity in kJ/kg/K, $\dot{m}_p = F_p/60$ is the process mass flow in kg/min, T_{in} is the inlet temperature and Q_p is the duty crossing from the metal in kJ/min. The tube metal is a second capacitance between the condensing steam and the process fluid, as shown in Equation (4.2).

$$M_w C_w \frac{\mathrm{d}T_w}{\mathrm{d}t} = Q_s - Q_p \tag{4.2}$$

where $M_w C_w$ is the tube-metal heat capacity in kJ/K and Q_s is the duty delivered by condensing steam. Both sides of Equations (4.1) and (4.2) carry units of kJ/min. The two duties are given by Equations (4.3) and (4.4).

$$Q_p = UA_p \cdot \Delta T_{\text{lm}}(T_w - T_{\text{in}}, T_w - T_p) \tag{4.3}$$

$$Q_s = UA_s \left[T_{\text{sat}}(P_{\text{sh}}) - T_w\right] \tag{4.4}$$

where UA_p and UA_s are the process-side and condensing-side conductances in kJ/min/K and T_{sat} is the shell saturation temperature. The process-side driving force is the log mean of the two terminal differences, as shown in Equation (4.5).

$$\Delta T_{\text{lm}}(a,b) = \frac{a-b}{\ln(a/b)} \tag{4.5}$$

The arithmetic mean $\frac{1}{2}(a+b)$ replaces (4.5) whenever either end falls to 1×10^{-9} K or below, or when the two ends agree to within one part in 10^6. That substitution keeps the right-hand side smooth enough for a fixed-step integrator during a transient in which the fluid is briefly hotter than the metal. Retaining the log mean rather than a single lumped ΔT is the mechanism by which the process gain falls as the outlet approaches the wall.

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

4.3.2 Steam Side

The shell is a pressure capacitance charged by the valve and discharged by condensation, as shown in Equation (4.6).

$$C_{\mathrm{sh}}\frac{\mathrm{d}P_{\mathrm{sh}}}{\mathrm{d}t} = \dot{m}_s(x, P_{\mathrm{sh}}) - \frac{Q_s}{\lambda_{\mathrm{eff}}(P_{\mathrm{sh}})} \tag{4.6}$$

where C_{sh} is the shell steam capacitance in kg/bar, $\dot{m}_s$ is the steam admitted in kg/min and λ_{eff} is the latent heat actually delivered in kJ/kg. The shell pressure is bounded by $0.3 \le P_{\mathrm{sh}} \le P_{\mathrm{sup}}$ in two places. The derivative in (4.6) is set to zero when the state is at a bound and the derivative points outward, so the integrator cannot step through a wall and back, and the state is clipped again after each step. Saturated-steam properties are quadratics in $L = \ln P$ with P in bar absolute, fitted over 1–10 bar, as shown in Equations (4.7) and (4.8).

$$T_{\mathrm{sat}} = 99.6 + 26.772\,L + 3.519\,L^2 \quad [^\circ\mathrm{C}] \tag{4.7}$$

$$\lambda = 2258 - 67.33\,L - 16.59\,L^2 \quad [\mathrm{kJ/kg}] \tag{4.8}$$

Both fits sit within 0.4 K and 2 kJ/kg of tabulated saturated-steam properties across that range, which is tighter than any other approximation in the model. Wet steam is handled by $\lambda_{\mathrm{eff}} = \chi\lambda$ with quality χ available as an operator-set input over $[0.8, 1]$ and bounded to $[0.5, 1]$ internally. At the design point the shell settles at 2.60 bar absolute (1.59 barg) and 128.4 °C against a 4 barg header, because nothing forces the shell to sit at header pressure. The shell finds the pressure at which the steam admitted equals the steam condensed.

4.3.3 Valve, Film Coefficients and Transport Delay

Steam flow through FCV-102 obeys an installed characteristic $f(x) = x^a$ with x the fractional lift, a square-root pressure-drop law, and a cap at the header capacity, as shown in Equation (4.9).

$$\dot{m}_s = \min\left(\dot{m}_{s,\max} f(x)\sqrt{\frac{P_{\mathrm{sup}} - P_{\mathrm{sh}}}{\Delta P_{\mathrm{nom}}}},\ \frac{F_{s,\max}}{60}\right) \tag{4.9}$$

where $\dot{m}_{s,\max}$ is the valve capacity at the design pressure drop ΔP_{nom} and $F_{s,\max}$ is the header capacity in kg/h. The flow is zero when $P_{\mathrm{sup}} - P_{\mathrm{sh}}$ is not positive. The exponent a is user-adjustable. Linear trim is $a = 1$, $a = 2$ behaves like equal-percentage trim whose gain rises with lift, and $a = 0.5$ is quick-opening. The stem is a first-order actuator with $\tau_v\,\mathrm{d}x/\mathrm{d}t = x_{\mathrm{cmd}} - x$ and is clipped to 0–100 % after each step.

The process-side conductance carries the turbulent film exponent and the fouling resistance in series, as shown in Equation (4.10).

$$UA_p = \left[\frac{1}{UA_{p0}\,(F_p/F_{p0})^n} + R_f\right]^{-1}, \qquad n = 0.8 \tag{4.10}$$

where UA_{p0} is the clean conductance at the design feed flow F_{p0} and R_f is the fouling resistance in min K/kJ. Equation (4.10) is the Dittus-Boelter and Sieder-Tate scaling written at the equipment level [3, 4]. Fouling resistance is $R_f = (\phi/100)\,R_{f,\max}$ where ϕ is a percentage set directly as an operator input or accumulated by an accelerated ramp $\mathrm{d}\phi/\mathrm{d}t$ that is the fifth integrated state. Condensate flooding reduces the condensing side by $UA_s = UA_{s0}\max(0.05,\, 1 - \varphi/100)$ with φ the flooded fraction in percent.

TT-102 is located downstream of the bundle and therefore measures fluid that left the bundle θ minutes earlier, as shown in Equation (4.11).

$$\theta = \frac{V_{\mathrm{trans}}}{\dot{m}_p/\rho} \tag{4.11}$$

Table 4.2 Model Parameters. Hidden Parameters Are Withheld From the Application's Parameter Panel and Must Be Inferred From Experiments.

Symbol	Meaning	Value	Unit
Hidden, the plant to be identified			
UA_{p0}	process-side UA at design flow	1380	kJ/min/K
UA_{s0}	condensing-side UA	3726	kJ/min/K
M_p	process thermal inventory	1200	kg
M_wC_w	tube-metal heat capacity	500	kJ/K
C_{sh}	shell steam capacitance	3.0	kg/bar
n	tube-side film exponent	0.8	—
V_{trans}	outlet transport volume	0.1	m^3
$R_{f,\max}$	fouling resistance at 100 %	2.6×10^{-4}	min K/kJ
$\dot{m}_{s,\max}$	valve capacity at design ΔP	64.0	kg/min
ΔP_{nom}	design pressure drop	2.4	bar
τ_v	valve actuator lag	0.20	min
F_{p0}	design feed flow	20 000	kg/h
Visible, data sheet and instruments			
C_p	process heat capacity	4.18	kJ/kg/K
ρ	process density	1000	$\mathrm{kg/m}^3$
$F_{s,\max}$	steam header capacity	2600	kg/h
FT-102 span	steam flow transmitter span	4000	kg/h
τ_{TT}	TT-102 lag	0.10	min
τ_{FT}	FT-102 lag	0.05	min
a	installed valve characteristic	1.0	—
Adjustable disturbances			
ϕ	fouling	0	%
φ	condensate flooding	0	%
χ	steam quality	1.0	—
S	valve stiction deadband	0	%
b	TT-102 bias	0	°C
σ	TT-102 noise	0	°C
Numerics			
Δt	integration step (RK4)	0.005	min
Δt_c	controller sample period	0.02	min
—	historian rate	50	samples/min

where V_{trans} is the outlet transport volume in m^3 and ρ is the process density in $\mathrm{kg/m}^3$. The result is clipped to 0–6 min. The delayed temperature receives an optional bias and optional band-limited noise and then passes through a first-order transmitter lag τ_{TT} before the controller reads it. Equation (4.11) carries the principal consequence for control design, because the dead time is not a constant of the plant. It is a fixed volume divided by a variable flow, so it is inversely proportional to production rate.

Every parameter needed to rebuild the simulator is listed in Table 4.2. At the design point the exchanger holds 85.0 °C with the valve at 44.9 %, passing 1727 kg/h of steam for a duty of 1045 kW against 1252 kW available with the valve wide open. The overall conductance is 16.78 kW/K and the transport delay is 0.30 min.

4.3.4 Steady-State Initialization

Every scenario begins at an algebraic steady state for its manual valve position or its setpoint rather than at an arbitrary initial vector. Scenario 5 then sets the fluid in the bundle to 78 °C to start the product cold. Given a valve position, a bisection on P_{sh} with a nested bisection on T_w solves (4.1) through (4.6) with all derivatives set to zero, with T_p following directly from the fluid balance, and a further bisection inverts that map to find the valve position that holds a requested outlet temperature.

Bisection rather than Newton iteration is used because every residual involved is monotone and bisection cannot diverge on an extreme disturbance setting. Each bisection runs 44 iterations, which brackets the shell pressure to better than 10^{-12} bar. The same routine supplies the balancing valve position offered during manual operation and the available-duty figure reported with the performance measures, and it is refreshed no more than four times per simulated minute because it is expensive relative to a single integration step.

4.3.5 Assumptions

The model retains only the mechanisms that change how the loop must be tuned. Each assumption below is stated with the range over which it holds and the condition under which it fails.

1. *Lumped process side.* The fluid in the bundle is one well-mixed inventory with a log-mean driving force to the metal rather than a distributed axial profile. The steady-state duty and the dominant time constant are correct, and the fitted gain at the design point agrees with a steady-state sweep to within 1 %. The first few seconds of a step response are approximate, and the assumption fails outright for any question about local tube-wall temperature, vaporization inside the tubes, or axial fouling distribution.

2. *Saturated shell at a single pressure.* The shell is saturated throughout, with T_{sat} and λ from (4.7) and (4.8). This holds between 1 and 10 bar absolute, which covers the whole adjustable range of the steam header. It fails during a rapid pressure transient, when a real shell is neither uniform nor exactly saturated, and it fails for superheated supply, non-condensables and shell-side maldistribution, none of which are represented.

3. *Lumped shell capacitance.* The steam side is one capacitance in kg/bar, so the pressure state responds as a first-order lag driven by the imbalance between admission and condensation. A real shell has a fixed vapor volume in which density and temperature both change, which makes the true pressure response stiffer than this. The consequence is that the shell contributes slightly more apparent dead time in the simulator than it would on the corresponding unit.

4. *Perfect condensate removal.* The trap is not modeled. Condensate is removed as fast as it forms unless the flooding disturbance is applied, and flooding then removes condensing area in direct proportion. This is adequate for showing loss of duty and the resulting valve drift. It cannot reproduce stall, subcooling, condensate level dynamics or the differential pressure that drives a real trap.

5. *Fouling as one uniform series resistance.* The deposit is a single R_f added in series with the clean film, ramped at a constant rate. The series form is standard and the effect on UA is correct. Real deposits build unevenly, change the tube-side pressure drop as well as the heat transfer, and follow a deposition-minus-removal law with an asymptote rather than a constant rate [7]. The accelerated rate of up to 3 %/min compresses months into minutes and is not a claim about fouling kinetics.

6. *Turbulent film scaling.* The tube-side coefficient follows $F_p^{0.8}$ over the full feed range of 12 000 to 30 000 kg/h, which stays turbulent for a water-like fluid in a bundle of this duty. The correlation is extrapolated and wrong at very low throughput, where the flow would be transitional or laminar and the exponent nearer 0.33.

7. *Transport delay as a moving tap.* The outlet delay is a fixed volume divided by volumetric flow, implemented as a linearly interpolated tap on a history buffer rather than as true plug flow. This reproduces the throughput dependence exactly and is correct in steady state. During a rapid flow change the tap moves faster than fluid physically can, so the delay transitions instantaneously where a real line would blend two residence times.

8. *Constant physical properties.* C_p and ρ are constant except for the deliberate heat-capacity error and the steam correlations. Over a 40 to 95 °C span in liquid water this is accurate to a few percent. It fails for a fluid near a phase change or with a strongly temperature dependent viscosity, where the film coefficient would also move with temperature and not only with flow.

9. *Valve and instrument idealization.* FCV-102 has a first-order actuator lag, an installed characteristic x^a and an optional pure deadband for stiction [17]. Positioner dynamics, hysteresis separate from deadband, and a finite stroking rate are absent. TT-102 and FT-102 are first-order lags with optional bias and noise, so drift, quantization and thermowell fouling are not represented.

10. *Omitted entirely.* Ambient heat loss, wall thermal expansion, the process-side pump characteristic, downstream mixing volumes, tube vibration and mechanical failure other than stiction and fouling are all absent. They are omitted because none of them changes the structure of the loop the application is used to teach.

4.4 Identification from Step-Response Data

No element of the model above is first order, and the only explicit dead time in the path from valve to transmitter is the 0.30 min transport delay of (4.11). A small step in the valve nevertheless produces a response that reads as FOPDT, the model of Equation (1.1), as shown in Equation (4.12).

$$\hat{G}(s) = \frac{\hat{K}_p \, \mathrm{e}^{-\hat{\theta} s}}{\hat{\tau} s + 1} \tag{4.12}$$

where $\hat{K}_p$ is the identified gain in °C/%, $\hat{\tau}$ is the identified time constant and $\hat{\theta}$ is the identified dead time, both in minutes. The identified constants are lumped surrogates. The value of $\hat{\tau}$ combines the fluid and metal capacitances with the shell and valve dynamics, and $\hat{\theta}$ combines the true transport delay (4.11) with the valve stroke, the metal lag, the transmitter lag and the initial curvature of a higher-order response, the effective dead time that Section 1.4.4 describes. At the design point a 5 % step, fitted over the 20 min that follow it, yields $\hat{K}_p = 0.355\,^{\circ}\mathrm{C}/\%$, $\hat{\tau} = 3.02$ min and $\hat{\theta} = 0.94$ min at $R^2 = 0.9997$, against a true transport delay of 0.30 min. The scenarios that need a model preload this one with $\hat{\tau}$ rounded to 3.0 min.

Those three numbers are not properties of the exchanger. The same test repeated from the balancing valve position at 24 000 kg/h returns $\hat{K}_p = 0.22\,^{\circ}\mathrm{C}/\%$, $\hat{\tau} = 2.49$ min and $\hat{\theta} = 0.82$ min, and at 16 000 kg/h it returns $\hat{K}_p = 0.56\,^{\circ}\mathrm{C}/\%$, $\hat{\tau} = 3.82$ min and $\hat{\theta} = 1.10$ min. The gain varies by a factor of 2.6 across a feed range of 16 000 to 24 000 kg/h. Three mechanisms act together. More mass to heat per minute at the same duty is a smaller temperature rise per unit of steam, the log-mean driving force in (4.3) shifts as the outlet moves relative to the wall, and the valve runs further up its travel where the shell has pressured up, so $P_{\mathrm{sup}} - P_{\mathrm{sh}}$ shrinks and the valve contributes less loop gain through (4.9). The shell pressure drop falls from 3.45 bar at 30 % lift to 0.73 bar at full lift. Volume 1, Chapter 2 meets the same lesson on a tank whose fitted gain and time constant move with the operating level.

Both a graphical reading of the reaction curve and a least-squares fit are available for the same record. The fitter locates the largest single-sample jump in the logged valve command, takes a pre-step baseline from the twelve preceding samples, truncates the window at the next command move, and then exploits the fact that (4.12) is linear in $\hat{K}_p$ for any fixed pair $(\hat{\tau}, \hat{\theta})$. A one-parameter least squares gives the optimal gain in closed form for each candidate pair, so only a two-dimensional grid of 4200 candidates plus one local refinement of 2500 candidates is searched. The fit is reported with R^2 and RMSE alongside the manually read constants, so the difference between a graphical reading and a regression is quantified without a numerical library. The fitter returns nothing when the logged step is smaller than 0.5 % or the response window is shorter than 1.5 min.

A separate steady-state regression addresses the nonlinearity itself. Recording a sequence of settled pairs (u, T_{out}) and fitting $T_{\mathrm{out}} = b_0 + b_1 u$ gives R^2 above 0.99 over a narrow span and a poorer

fit over a wide one, because the steady-state map runs from 76.8 °C at 30 % lift to 93.9 °C at full lift while the local slope falls from 0.75 °C/% at 30 % lift to 0.06 °C/% at full lift. Up to twelve points are retained. The regression shows that a linear model identified at one operating point describes the plant only near that point, which is the argument for gain scheduling [18] in Section 1.4.13.

4.5 Controller Design

4.5.1 Velocity Form PID Algorithm

Both loops run the velocity-form PID algorithm [18] of Section 1.6.2, the one implementation that the framework apps share. It is written here as increments per scan in the symbols of this loop, because the results below cite its terms. Error and the three increments are given by Equations (4.13) and (4.14).

$$e_k = \mathrm{SP}_k - \mathrm{PV}_k \tag{4.13}$$

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

where K_c is in %/°C on the temperature loop, τ_I and τ_D are in minutes, and Δt_c is the control period. Derivative comes from a second difference on a filtered measurement, with $\mathrm{PV}_k^f = \alpha \mathrm{PV}_{k-1}^f + (1-\alpha)\mathrm{PV}_k$ and $\alpha = \exp(-\Delta t_c/\tau_f)$ and $\tau_f = 0.05$ min by default, which is stable at any sample period because α is never negative. Section 1.4.8 gives what this filter costs in phase. The increments are combined, the demand is clipped to the output limits and the integral is corrected by back-calculation, as shown in Equation (4.15).

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

Here z_k accumulates the proportional and derivative increments without limit, c_k is the demanded position built on the integral accumulated through the previous scan, and I_k also carries the starting output. These are the sum S_k, the output before limiting v_k and the integral with its bias of Equations (1.24) to (1.26). The last term drives I_k toward the value that reproduces the clipped output with tracking time $T_t = \tau_I$ ($\tau_{I,F}$ on the flow loop), removing $1/150$ of the clipped excess per scan at $\tau_I = 3$ min, so only the integral is corrected. Away from the limits the recursion is algebraically identical to the positional law with the same three gains, but not at a limit, where the integral unwinds over a time of order τ_I, the relaxation of Equation (1.18). Disabling anti-windup sets $T_t = 10^9$ min, so the integral winds up against a saturated output and reset windup can be observed rather than described. In manual operation I_k is reset every scan so that c_k equals the output in force (the valve position, or the measured steam flow for the outer loop of a cascade), and in automatic operation the manual valve setting follows u_k after every controller execution, so transfer is bumpless in both directions.

Output limits are 0–100 % on the valve and on the steam-flow setpoint of the cascade and ±50 % on the feedforward trim, with no rate limit. The anti-windup of TIC-101 knows only its own limits. A valve held wide open by FIC-102 is not reported back to it, and its setpoint can reach 100 % of the 4000 kg/h FT-102 span while the header delivers at most 2600 kg/h (65 % of span), so the outer integral can wind while the inner loop is saturated. The 0–100 % clip on feedforward plus trim has the same effect. Both loops are reverse acting in the ISA sense (Section 1.4.3), because opening the steam

valve raises the temperature, so an outlet temperature above setpoint has to close the valve. With the error written as SP − PV in (4.13), that is what a positive K_c delivers. The controller executes every 0.02 min (1.2 s), independent of the 0.005 min (0.3 s) physics step, which makes the measurement a sampled one and places a half-sample of additional lag in the loop (Section 1.4.11).

4.5.2 Three Architectures, One Actuator

The manipulated variable is the same in all four architectures, the three PID structures described here and the predictive controller of Section 4.5.4. What changes is the calculation that produces it. In the single loop, TIC-101 drives FCV-102 directly. In cascade, TIC-101 outputs a steam-flow setpoint in percent of FT-102 span and FIC-102 drives the valve with $K_{c,F} = 2\,\%/\%$ and $\tau_{I,F} = 0.25\,\text{min}$, which is about the valve stroking time. The inner loop absorbs header-pressure upsets and the valve characteristic before they reach the temperature. The separation of time scales required for cascade control is present, because the flow loop settles in a fraction of a minute against an outer time constant of 3.0 min, a ratio of about ten [1]. Volume 3, Chapter 3 measures what a cascade loses when that separation is absent. In the third architecture the steam demand is computed from the process heat balance, as shown in Equation (4.16).

$$F_{\text{ff}} = K_{\text{FF}} \frac{\dot{m}_p C_p \,(T_{\text{SP}} - T_{\text{in}})}{\lambda_0}, \qquad \lambda_0 = \lambda\,(2.60\,\text{bar}) \tag{4.16}$$

where K_{FF} is a dimensionless trim on the calculation and F_{ff} is converted to percent of FT-102 span. The result passes through a backward-Euler lead-lag with constants τ_{lead} and τ_{lag} and a pure delay θ_{FF}, and it is added to a temperature trim bounded to $\pm 50\,\%$ of span before FIC-102 moves the valve. Equation (4.16) is deliberately built from the nameplate C_p rather than the true one, so a heat-capacity error biases the feedforward calculation and the feedback trim is required rather than optional. Whichever structure is selected, the idle loops keep tracking, so the architecture can be changed mid-run without a bump. This mirrors the compound feedforward-plus-cascade practice reported for industrial exchanger services [8].

Pause. Try Exercise 4.1 before reading on.

4.5.3 Tuning Rules

Two rules are offered, both operating on the identified model and never on the simulator equations. Section 1.4.5 derives both and compares them on one process. The IMC-derived SIMC PI rule [12, 13] is shown in Equation (4.17).

$$K_c = \frac{\hat{\tau}}{\hat{K}_p(\tau_c + \hat{\theta})}, \qquad \tau_I = \min\!\big(\hat{\tau},\, 4(\tau_c + \hat{\theta})\big), \qquad \tau_D = 0 \tag{4.17}$$

where the closed-loop constant τ_c is selected as $2\hat{\theta}$, $\hat{\theta}$ or $0.5\hat{\theta}$ for conservative, moderate and aggressive response. The reaction-curve PID rule [11] is shown in Equation (4.18).

$$K_c = \frac{1.2\,\hat{\tau}}{\hat{K}_p \hat{\theta}}, \qquad \tau_I = 2\hat{\theta}, \qquad \tau_D = 0.5\hat{\theta} \tag{4.18}$$

Both results are rounded to the tuning-entry resolution and clamped to $K_c \le 40\,\%/^\circ\text{C}$, $0.1\,\text{min} \le \tau_I \le 30\,\text{min}$ and $\tau_D \le 3\,\text{min}$ before they are loaded. With the application's τ_c tied to $\hat{\theta}$, both rules' gains scale inversely with the identified dead time: changing 0.94 min to 0.50 min multiplies each by 1.88. Equation (4.18) places $\hat{\theta}$ in a denominator with no additive term, so a dead time read as 0.50 min instead of 0.94 min produces $K_c = 20.3\,\%/^\circ\text{C}$ on the preloaded model ($\hat{K}_p = 0.355\,^\circ\text{C}/\%$, $\hat{\tau} = 3.0\,\text{min}$), against 10.8 %/°C from the same rule on the correctly read dead time and 4.5 %/°C from the SIMC rule on the same model. The application ships that reading error as a scenario that preloads exactly the 20.3 %/°C the rule returns, and every measurement reported below is taken at that value.

> **TRY IT IN THE APP** **8 · Ziegler-Nichols, and an Optimistic Dead Time**
>
> Load the case, which preloads $K_c = 20.3\,\%/^\circ\mathrm{C}$ from a dead time read as 0.50 min, and run it to 24 min. Watch the steam valve trend after the feed turns 8 K colder at 4 min: it rings with a period of 4.3 min, each swing about 0.54 of the one before. Then Reset, enter $\hat{\theta} = 0.94\,\mathrm{min}$, apply the reaction-curve rule before 4 min and run again.

4.5.4 Model Predictive Control on the Identified Model

The fourth architecture replaces TIC-101 with a model predictive controller (MPC) that moves FCV-102 directly. Every existing case keeps its shipped architecture, and only the two cases added with the MPC select it, the comparison case of Section 4.7.2 and the tube case of Section 4.7.3. Every $T_s = 0.1$ min, or five controller executions, the MPC predicts TT-102 over a horizon, chooses a sequence of valve positions, applies the first and holds it until the next decision, the receding-horizon scheme of Section 1.6.3 [19, 20]. Its model joins the identified valve path of Equation (4.12) to a measured feed-flow path in deviation variables, as shown in Equation (4.19).

$$T_{\mathrm{out}} - T_0 = \frac{\hat{K}_p\, \mathrm{e}^{-\hat{\theta} s}}{\hat{\tau} s + 1}\,(u - u_0 + d) + \frac{K_w\, \mathrm{e}^{-\theta_w s}}{\tau_w s + 1}\, w \tag{4.19}$$

where u_0, T_0 and F_0 are the valve command, the TT-102 reading and the FT-101 reading when the estimator starts, $w = F_p - F_0$ is the measured feed deviation in kg/h, d is an unmeasured disturbance in percent of valve travel, and $K_w = -1.275 \times 10^{-3}\,^\circ\mathrm{C\,h/kg}$, $\tau_w = 3.08$ min and $\theta_w = 0.60$ min are the gain, lag and dead time of the feed path.

The valve path is the model entered for the tuning rules, so an identification error reaches the MPC as it reaches the PI, and while no complete model is entered the MPC makes no plan and TIC-101 keeps the valve. In the comparison case it is the preloaded fit of Section 4.4, 0.355 °C/%, 3.0 min and 0.94 min. The feed path was identified once on the simulator at the design point, with the valve held at 44.9 % and no measurement noise, by averaging the responses to feed steps of ± 1000 kg/h (root-mean-square residual 0.32 % of the response), and it is declared as constants, as the nameplate data of Equation (4.16) are. Its gain matches the steady-state slope of the process model at a fixed valve, -1.2735×10^{-3} °C h/kg, which is smaller in magnitude than the -2.25×10^{-3} °C h/kg of a fixed duty because more feed also raises the film coefficient of Equation (4.10). Both dead times are kept as records of what has already been sent, $D_u = 47$ and $D_w = 30$ executions of $\Delta t_c = 0.02$ min, and each plan starts from the state that the commands already in transit will produce, the prediction of a Smith predictor, which is exact for the model.

The disturbance d comes from a steady-state Kalman filter built on the two paths, sampled at Δt_c for a zero-order hold and augmented with d as a random walk, the usual construction for offset-free predictive control [21]. The filter runs at every execution whichever architecture holds the valve, as shown in Equation (4.20).

$$\begin{aligned}
\bar{x}_{u,k} &= \mu_u \hat{x}_{u,k-1} + \hat{K}_p(1-\mu_u)\left(u_{k-1-D_u} - u_0 + \hat{d}_{k-1}\right)\\
\bar{x}_{w,k} &= \mu_w \hat{x}_{w,k-1} + K_w(1-\mu_w)\, w_{k-1-D_w}\\
\nu_k &= T_{\mathrm{out},k} - T_0 - \bar{x}_{u,k} - \bar{x}_{w,k}\\
\hat{x}_{u,k} &= \bar{x}_{u,k} + \ell_u \nu_k, \qquad \hat{x}_{w,k} = \bar{x}_{w,k} + \ell_w \nu_k\\
\hat{d}_k &= \hat{d}_{k-1} + \ell_d \nu_k
\end{aligned} \tag{4.20}$$

where x_u and x_w are the outputs of the two paths in °C, $\mu_u = \mathrm{e}^{-\Delta t_c/\hat{\tau}} = 0.99336$ and $\mu_w = \mathrm{e}^{-\Delta t_c/\tau_w} = 0.99353$, u_{k-1-D_u} and w_{k-1-D_w} are the command and the feed reading D_u and D_w executions before the period just ended, and ν_k is the innovation. The gains $\ell_u = 0.1033$, $\ell_w = 8.4 \times 10^{-7}$ and $\ell_d = 2.667\,\%/^\circ\mathrm{C}$ follow from a measurement variance r, state noise variances of $10^{-6} r$ and a disturbance noise variance $q_d = r/\hat{K}_p^2$, with $r = 1\,^\circ\mathrm{C}^2$ because only the ratios set the gains. Every

load the model lacks, a sagging header, a colder feed, fouling, a failing trap or an error in the model itself, appears in $\hat{d}$ as the valve offset that would produce it, and holding $\hat{d}$ over the horizon removes the offset, so the estimate supplies the integral action. The comparison case has no unmeasured load, so the $\hat{d}$ of −8.28 % of valve reached at its end measures only how far the linear model, identified at 85 °C and 20 000 kg/h, misses the exchanger at 88 °C and 24 000 kg/h.

Each decision solves a quadratic program on the same model resampled at T_s, as shown in Equation (4.21).

$$\begin{aligned} \min_{v,\varepsilon} \; & \frac{T_s}{\hat{\tau}} \sum_{k=1}^{N} \left(\frac{\hat{T}_k - T_{\mathrm{SP}}}{1\,^{\circ}\mathrm{C}} \right)^2 + \kappa \sum_{j=1}^{12} \left(\frac{v_j - v_{j-1}}{5\,\%} \right)^2 + c_\varepsilon \varepsilon^2 \\ \text{s.t. } & 0 \le v_j \le 100\,\%, \qquad \varepsilon \ge 0 \\ & T_{\mathrm{lo}} - \varepsilon \le \hat{T}_k \le T_{\mathrm{hi}} + \varepsilon, \quad k = 1, 3, \ldots, N-1 \end{aligned} \tag{4.21}$$

where $\hat{T}_k$ is the TT-102 reading predicted k steps beyond the dead time with $\hat{d}$ held, $v_1, \ldots, v_{12}$ are valve positions held over twelve blocks whose lengths in steps are $(1,1,1,1,2,2,3,4,6,9,15,75)$, $N = 120$ steps covers 12 min after $\hat{\theta}$, v_0 is the command applied now, $\kappa = 1.5$ is a dimensionless move suppression, ε is one slack in °C priced at $c_\varepsilon = 10^3$ per °C^2, and $T_{\mathrm{lo}} = \min(79\,^{\circ}\mathrm{C}, T_{\mathrm{SP}} - 1\,^{\circ}\mathrm{C})$ and $T_{\mathrm{hi}} = \max(90\,^{\circ}\mathrm{C}, T_{\mathrm{SP}} + 1\,^{\circ}\mathrm{C})$ are TAL-102 and TAH-102 moved, where needed, to 1 °C from the setpoint. The factor $T_s/\hat{\tau}$ keeps the error sum from growing with the horizon, and with these scales $\kappa = 1$ prices one 5 % valve move like a 1 °C error held for one time constant. The valve range is the only hard constraint. The band is soft, so a plan always exists, and moving it off the setpoint keeps a 90 °C setpoint from sitting on TAH-102. There is no rate limit and no terminal weight, because the model is stable and $\hat{\theta} + NT_s = 12.94$ min exceeds four time constants.

> ⧖ HISTORY **Predictive Control Moves Into the Plant, 1973 to 1983**
>
> Control by a sequence of open-loop optimizations was described in the 1960s, but industry put it to work. Engineers at an oil company applied their own predictive method in 1973 and published a version built on a step-response model in 1979. A paper of 1976 described the first applications, among them a power-plant steam generator whose response time varied inversely with load, handled by running the controller with a variable sample time. These controllers penalized valve moves for a less aggressive response and some robustness to model error, and one showed a better feedforward response to inlet temperature than a PID lead-lag. In 1983 the method was posed as a quadratic program with the constraints written explicitly [20]. The predictive controller of this chapter is of that second kind.

With the preview on, the feed enters every prediction as a known input, the FT-101 readings taken θ_w before each period while the record reaches them and the present reading held thereafter, a measured disturbance rather than a forecast. With the preview off the reading is frozen and a feed change reaches the MPC only through $\hat{d}$. The feedforward of Equation (4.16) also reads TT-101, so it receives more measured information than the MPC. The dual active-set solver of the framework [22] solves the program, 13 variables and 145 inequality rows, with a ridge near 10^{-9} on the Hessian diagonal.

The move suppression and the disturbance noise were fixed once, on the comparison case of Section 4.7.2, by the rule applied to the predictive controllers of the apps. Of 55 settings (κ from 0.05 to 5, a factor of 0.01 to 100 on q_d), 25 reject no plan and move the valve no more than the IMC-PI (45.6 %), and among those within 10 % of the best true-temperature IAE (12.71 °C min) the largest κ and then the smallest factor are taken, $\kappa = 1.5$ with a factor of 1 (13.78 °C min). Decision periods of one, two and five executions give true-temperature IAE of 14.29, 13.83 and 13.78 °C min under the same rule, at about 370, 190 and 85 µs per plan in Node.js, and horizons from 6 to 30 min move the scorecard IAE only from 14.12 to 14.00 °C min.

A plan is discarded if the solver cannot certify it as optimal, or if its first move is non-finite or outside 0–100 %. TIC-101 tracks the applied command at every execution and therefore takes over

the valve bumplessly (0.026 % when a rejection is forced 8 min into the comparison case), the record is marked, and control returns to the MPC at the next acceptable plan. Loading a case solves the first plan. The MPC reads only TT-102, the command applied, FT-101, the setpoint, the alarm limits and its own settings, and its model is the learner's entry and the declared feed path, so a change to a plant parameter or a disturbance setting is model mismatch for it, as it would be on a real exchanger.

4.5.5 Alarm Band Tightened by the Learner's Fit

The plan of Equation (4.21) treats the entered model as exact, so near TAH-102 it stakes the band on the step test behind it. A second setting of the MPC first measures how far the entered model misses that step test and then keeps the reading inside the band for any valve-path gain within that misfit. This is constraint tightening, in which one nominal plan is kept and its constraints are moved in by a bound on the model's error [23, 24], rather than a plan made against every model of a set [25].

In manual operation the controller records the command it applied and the TT-102 reading at each execution, finds the step test in that record by the rules of the fitter of Section 4.4 (the largest move of at least 0.5 %, a twelve-execution baseline and at least 1.5 min of response), and replays the entered model through it, as shown in Equation (4.22).

$$\begin{aligned} \delta^2 &= \frac{\sum_i \left(T_{\mathrm{m},i} - \bar{T}_\mathrm{b} - g_i\right)^2}{\sum_i g_i^2} \\ g_i &= \hat{K}_p \Delta u_\mathrm{s} \left(1 - \mathrm{e}^{-(t_i - t_\mathrm{s} - \hat{\theta})/\hat{\tau}}\right) \end{aligned} \tag{4.22}$$

where t_s and Δu_s are the time and size of the step, $\bar{T}_\mathrm{b}$ is the mean baseline reading, $T_{\mathrm{m},i}$ is the reading at execution i of the response, and $g_i = 0$ while $t_i - t_\mathrm{s} \le \hat{\theta}$. The relative misfit δ measures the learner's model against the learner's own data, not against the hidden constants, and equals $|K_p/\hat{K}_p - 1|$ for a gain error alone on a linear plant. The step test of the tube case fits as 0.3596 °C/%, 3.14 min and 0.92 min, and with $\hat{\tau} = 3.0$ min and $\hat{\theta} = 0.94$ min the entries 0.36, 0.54 and 0.252 °C/% give δ of 1.77, 34.30 and 40.82 %, against 33.3 and 42.9 % for a pure gain error of 1.5 and 0.7 times.

Read as the relative error of the valve path, δ places the reading at prediction step k in $\hat{T}_k \pm h_k$ with $h_k = \delta|\tilde{x}_{u,k} - \hat{x}_u|$, where $\tilde{x}_{u,k}$ is the valve-path output that the plan predicts and $\hat{x}_u$ is its estimate from Equation (4.20) at the decision. The half-width is zero where the plan leaves the valve path alone. It grows with the commands still in the dead-time record, sent but not yet seen at TT-102, and with the plan's own moves, to δ times the whole change at the end. Keeping both edges in the band gives linear rows, as shown in Equation (4.23).

$$\begin{aligned} T_\mathrm{lo} - \varepsilon_1 &\le \hat{T}_k + \delta\left(\tilde{x}_{u,k} - \hat{x}_u\right) \le T_\mathrm{hi} + \varepsilon_1 \\ T_\mathrm{lo} - \varepsilon_2 &\le \hat{T}_k - \delta\left(\tilde{x}_{u,k} - \hat{x}_u\right) \le T_\mathrm{hi} + \varepsilon_2 \end{aligned} \tag{4.23}$$

for $k = 1, 3, \ldots, N-1$. They replace the band of Equation (4.21), with $c_\varepsilon(\varepsilon_1^2 + \varepsilon_2^2)$ in place of $c_\varepsilon \varepsilon^2$, each as a further output $(1 \pm \delta)x_u + x_w$ of the same model against the band shifted by $\pm\delta\hat{x}_u$. Everything else is the nominal plan's, and the tube has no weight of its own. One slack per edge for the whole horizon was kept by a rule fixed beforehand, because per-step slack gave the same minutes past the band at 1.5 to 2.1 ms a plan. The program grows from 13 variables and 145 rows to 14 and 266, and a plan takes 104 to 106 µs in Node.js (86 to 88 µs for the nominal band) and 137 to 172 µs on the frame loop of the application in Chrome, or 0.17 % of one core in real time. No step test, or a misfit of 100 % or more, leaves the nominal band, because $1 \pm \delta$ then admits a valve that does nothing or acts in reverse, and from 50 % the controller's note warns that the tube will work the valve hard. The plan display shades the tube, the misfit and the end width are reported and logged, and the misfit is computed from the controller's own record and the entered model alone.

The interval is exact for a gain error on the linear model only once the estimator has settled. Around a linear plant of the model's structure with two thirds of the entered gain ($\delta = 1/3$), each

plan applied open loop to a copy of the plant leaves 4790 of 18 060 judged predictions outside its tube, by up to 0.32 °C and none before the setpoint step, because while $\hat{d}$ still absorbs the gain error the model predicts a drift of x_u that is wrong in full and the tube covers only δ of it.

4.6 Dynamic Analysis

Three properties of the loop follow from the model and govern how it must be tuned.

The first is a dead time that moves with production rate. By (4.11), a higher throughput shortens θ, from 0.375 min at 16 000 kg/h to 0.250 min at 24 000 kg/h. A controller tuned at low rate is therefore conservative at high rate, which preserves stability margin, while a controller tuned at high rate and operated at low rate loses it.

The second is a gain that falls as the load rises. The log-mean driving force and the valve self-throttling act in the same direction, so the steady-state gain falls from 0.561 °C/% at 16 000 kg/h to 0.219 °C/% at 24 000 kg/h and to 0.147 °C/% at 27 000 kg/h, a factor of 3.8 across the range. Combined with the moving dead time, the loop gain of a fixed controller varies by more than a factor of three over the normal operating window, which is the condition that gain scheduling addresses [18]. Chapter 6 builds such a schedule for a loop whose gain moves by orders of magnitude.

The third is a capacity limit that presents as a tuning problem. When required duty approaches available duty the valve saturates at 100 % and the temperature falls short of setpoint. The symptom is a steady offset with a saturated output, which is also the symptom of a wound-up integrator, and the common response is to increase gain. Required and available duty are reported as separate trends together with the percentage of capacity in use, so the two causes can be distinguished from the record. At 27 000 kg/h with the exchanger 70 % fouled and the header at 3.2 barg, the required duty is 1411 kW against 1248 kW available, and the outlet settles at 79.8 °C with the valve wide open. Chapter 3 asks the same question, whether the setpoint is reachable at all, of a composition loop.

Two further behaviors are diagnostic problems rather than tuning problems. Fouling raises R_f slowly, the temperature stays at setpoint, and the controller output drifts upward, so the controller output rather than the controlled variable is the fouling indicator. A ramp to 100 % fouling moves the balancing valve position from 44.9 % to 71.9 % and drops the overall conductance from 16.78 kW/K to 13.30 kW/K while the outlet holds 85 °C. Valve stiction is modeled as a deadband S, where the stem is held while $|u - x| < S$ and, once moving, continues until the command is within $0.2S$, which converts a smooth controller output into the square-wave stem motion characteristic of a sticking valve [17]. The resulting limit cycle is not removable by tuning.

Pause. Try Exercise 4.3 before reading on.

> **TRY IT IN THE APP** **14 · Fouling**
>
> Load the case and press Run. The tubes foul on a ramp while the loop stays in automatic. Watch two trends: the outlet temperature stays at its setpoint, and the steam valve creeps open. The valve, not the temperature, carries the news of the fouling.

4.7 Results

Nineteen scenarios plus a graded challenge are arranged as a laboratory sequence, covering manual operation and finding the operating point, a step test and a regression comparison, a first conservative automatic loop and then a model-based one, a setpoint move and an aggressive-rule comparison, the three load disturbances, cascade against the header upset and feedforward against measured load changes, four diagnostics, a comparison of three controllers designed from one identified model, and the alarm band of that predictive controller set from the learner's own fit. Each scenario sets an initial mode, a tuning, a set of disturbance values and a list of timed events. The disturbances and faults are summarized in Table 4.5, and the three load disturbances are applied at $t = 4$ min from a settled start.

Table 4.3 Scenarios in Selector Order and Completion Requirements, First Half

	Scenario	Required to complete
1	Hand operation	From the steady 30 % start, move the valve after 1 min and hold the outlet within ±1 °C of 85 °C by hand for 3 min, then set the heat-balance position.
2	Operating point	Settle within ±0.5 °C of 85 °C for 2 min in Manual, record five or more steady points spanning at least 15 % of valve travel, then a sixth below 32 % or above 62 %.
3	Step test	Step the valve 5 % from the settled design point after 1 min, hold it 15 min, enter $\hat{K}_p$, $\hat{\tau}$ and $\hat{\theta}$ within 0.28–0.43 °C/%, 2.4–3.8 min and 0.7–1.2 min, then run the least-squares fit.
4	Reading against fit	Fit a 5 % step and then a step of at least 9 %, each over at least 10 min, and load the fit.
5	First try in Auto	Recover the cold start to 84.8 °C with $K_c \leq 1.5\,\%/°\mathrm{C}$, then Reset, select the moderate SIMC rule within 0.5 min and recover again with $K_c \geq 3\,\%/°\mathrm{C}$.
6	SIMC tuning	Switch to Manual and back to Auto after 0.5 min, hold a setpoint of 87 °C for 8 min with $\tau_I \geq 2.9$ min, then reload and repeat the step from 85 °C for 10 min with $\tau_I \leq 0.85$ min.
7	Setpoint test	Run the 85 to 90 °C step at 3 min to $t = 15$ min with $K_c \geq 4\,\%/°\mathrm{C}$, then Reset, select the conservative SIMC rule (K_c near 3.0 %/°C) before the step and repeat to 15 min.
8	ZN, dead time	Run the 8 K inlet drop at 4 min to $t = 24$ min with $K_c \geq 19\,\%/°\mathrm{C}$, then Reset, enter $\hat{\theta} = 0.94$ min and apply the ZN rule before 4 min (K_c of 9 to 13 %/°C) and repeat to 24 min.
9	Feed-flow upset	After the step to 24 000 kg/h at 4 min, hold within ±0.3 °C for 2 min, then fit a Manual step of 4 % or more over 5 min.
10	Inlet upset	After the 8 K inlet drop at 4 min, hold the outlet within ±0.2 °C for 2 min in Auto.

Tables 4.3 and 4.4 list the cases in selector order with the requirement for completing each, judged as Section 1.6.4 describes. Reset keeps the learner's own settings and a reload restores the shipped ones (Section 1.6.1). A setting chosen before Run is part of the starting state and gives the same run as the same setting carried through a Reset, so a step the learner is to watch is made after the run has started. The challenge is judged on the case as given, so changing a model parameter or a disturbance setting, or firing an extra upset, voids the attempt.

Each load disturbance has a computable magnitude. A feed step from 20 000 to 24 000 kg/h raises the required duty by 209 kW and moves the balancing valve from 44.9 % to 59.4 %. An inlet drop of 8 K raises the required duty by 186 kW and the steady steam flow by 329 kg/h, because the latent heat falls as the shell pressure rises. A header pressure drop from 4.0 to 3.2 barg at a fixed valve position cuts the steam flow from 1727 to 1594 kg/h, a fall of 7.7 %, and would settle the outlet at 81.9 °C if the controller did not act.

4.7.1 Measured Comparison of Two Tunings

Table 4.6 reports the application's own performance measures for three controllers subjected to the identical upset, a feed inlet temperature stepped from 40 °C to 32 °C at $t = 4$ min, read at $t = 10$ min. All three were produced by a rule applied to the same identified model. The outer columns are the shipped tunings of the inlet-temperature and tuning-comparison scenarios, and the middle column is the reaction-curve rule reapplied after the fitted dead time is entered, which is the second run the tuning-comparison scenario asks for.

Both reaction-curve controllers beat SIMC on every error measure, with IAE of 3.37 and 3.43 °C min against 8.61 °C min. That advantage follows from the rule, which asks for more gain and adds derivative action, and not from the misread dead time. Reading the dead time 0.44 min short inflates $1.2\hat{\tau}/(\hat{K}_p\hat{\theta})$ by a factor of 1.88, trims the largest deviation by 0.12 °C, leaves IAE slightly worse and multiplies the valve travel by 4.5. The loop rings with a period of 4.3 min, each valve swing about 0.54 of the one before, so by $t = 24$ min it has accumulated an IAE of 6.05 °C min and 187 % of travel against 3.67 °C min and 25.8 % for the correctly read dead time. The remaining margin does not survive a change of operating point. At a feed of 18 000 kg/h, with a higher process gain and a longer dead time, the same upset grows into an oscillation that strokes the valve between 14 % and 100 %, while the PID tuned on the fitted dead time remains stable at 16 000 kg/h. The consequential

Table 4.4 Scenarios in Selector Order and Completion Requirements, Second Half

	Scenario	Required to complete
11	Header sag	After the header falls to 3.2 barg at 4 min, hold the outlet within $\pm 0.2\,^{\circ}\mathrm{C}$ for 2 min with the single loop.
12	Cascade	Run the same sag in cascade and switch to the single loop after 8 min, then reload the case, set $K_{c,F} \leq 0.35\,\%/\%$ before the sag and run in cascade to $t = 12$ min.
13	Feedforward	Run both load changes for 29 min in feedforward with $K_{\mathrm{FF}} = 1$ and no lead, Reset and repeat in cascade, reload and repeat with $K_{\mathrm{FF}} \geq 1.25$, then set a lead of at least 0.35 min.
14	Fouling	Run the 2.5 %/min fouling ramp in Auto to $t = 40$ min, then set fouling to at least 145 % and run until the valve saturates.
15	Stiction	Run 40 min with the 3 % deadband, 10 min with $K_c \leq 2.3\,\%/^{\circ}\mathrm{C}$ and $\tau_I \geq 5.9$ min, 10 min in cascade, then remove it.
16	Instruments	With 0.15 °C noise and a 2 °C bias, run 20 min with $\tau_D = 0.47$ min, Reset and run 20 min with $\tau_D = 0$, reload and run 20 min with $\tau_f \geq 0.25$ min, then set the bias to −1.5 °C or below.
17	Capacity limit	Saturate the valve for 1 min and run past the cleaning at 20 min to 28 min, then Reset, switch anti-windup off at the start and run to 28 min.
18	Three controllers	Run the setpoint step to 88 °C at 3 min and the feed step to 24 000 kg/h at 15 min to $t = 30$ min under the MPC, then under the IMC-PI after a Reset, then under the feedforward with that PI as trim after a Reset.
19	Tube from the fit	Run the step test, the switch to Auto at 11 min, the setpoint step to 89 °C and the feed drop to 17 000 kg/h to $t = 40$ min with the tube, Reset with $\hat{K}_p \geq 0.5\,^{\circ}\mathrm{C}/\%$ and repeat, then Reset with the nominal band and repeat.
20	Challenge	Hold TT-102 within $\pm 1\,^{\circ}\mathrm{C}$ of 85 °C for the full 60 min through five load and supply upsets, with no alarm, no time saturated and less than 250 % valve travel, with the model and disturbances as given.[a]

Every case is an experiment except case 17, a demonstration, and case 20, the graded challenge. [a] With the shipped single-loop tuning ($K_c = 1\,\%/^{\circ}\mathrm{C}$, $\tau_I = 10$ min) the worst error is 4.58 °C, with 18 % valve travel and no saturation. With the suggested SIMC trim on the heat-balance feedforward ($K_c = 4.5\,\%/^{\circ}\mathrm{C}$, $\tau_I = 3$ min, $K_{\mathrm{FF}} = 1$) it is 0.66 °C, 0.34 °C inside the band, with 154 % travel and no saturation. The architecture alone is enough: feedforward with the shipped gains passes at 0.75 °C and 69 % travel.

Table 4.5 Load-Disturbance Scenarios and the Mechanism Each Exercises

Scenario	Change	What it exercises
Feed flow	20 000 → 24 000 kg/h	simultaneous change in gain, dead time and residence time
Inlet temperature	40 → 32 °C	duty step not visible at TT-102 for θ minutes
Steam header	4.0 → 3.2 barg	actuator-side upset detected first by the inner loop
Fouling ramp	R_f increasing	offset-free CV, drifting MV
Stiction	$S > 0$	limit cycle not removable by tuning
Capacity	duty demand > duty available	saturation, not detuning

error was in the identification, and the remaining difference, 24.8 % of valve travel against 17.4 % for SIMC, is a design choice about how much valve motion the service warrants. Figure 4.2 draws the table, and Figure 4.5 shows the misread case in the application at $t = 9$ min.

Performance is reported with the scorecard measures of Section 1.6.4 rather than one score. They are the IAE accumulated from the start of the run, the total valve travel, the fraction of time the output remained at a limit and, after a setpoint step only, percentage overshoot past the step target and settling time against a 1 °C band. The largest deviation is carried in the exported record. The trade is explicit on this service, because a tuning that reduces IAE on the header-pressure upset generally does so by increasing valve travel, which is packing and seat wear.

Table 4.6 Measured Performance After an 8 K Inlet-Temperature Drop, Read at $t = 10$ min. All Three Controllers Come From the Same $\hat{K}_p = 0.355$, $\hat{\tau} = 3.0$ min. Time Outside $\pm 1\,^{\circ}$C Is Taken From the Exported Record.

	SIMC PI $\hat{\theta} = 0.94$	ZN PID $\hat{\theta} = 0.94$	ZN PID $\hat{\theta} = 0.50$
K_c (%/°C)	4.5	10.8	20.3
τ_I (min)	3.0	1.88	1.0
τ_D (min)	0	0.47	0.25
IAE (°C min)	8.61	3.37	3.43
Max. deviation (°C)	2.16	1.36	1.24
Outside $\pm 1\,^{\circ}$C (min)	4.92	1.48	0.84
Valve travel (% of span)	17.4	24.8	112
Time saturated (%)	0	0	0

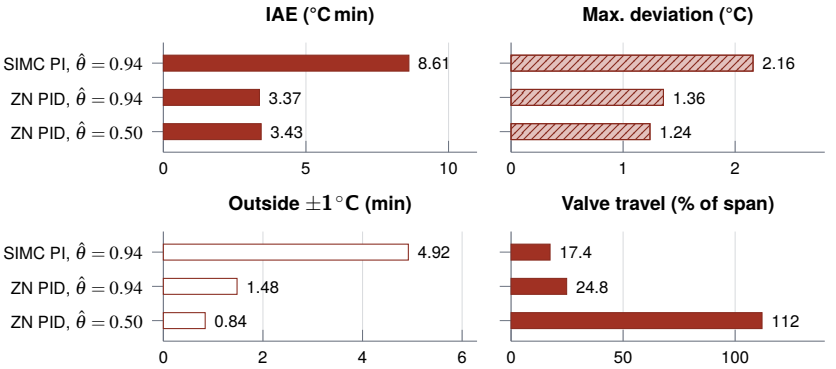


Figure 4.2 The three tunings of Table 4.6, all from the same identified model, after the 8 K inlet drop, read at $t = 10$ min: integrated absolute error and valve travel (solid bars), largest deviation (hatched) and time outside $\pm 1\,^{\circ}$C (open). Notice that the misread dead time leaves the error measures nearly where the correct reading puts them and multiplies the valve travel by 4.5.

> ✦ **DID YOU KNOW? A Bolder Tuning That Only Wears the Valve**
>
> From the simulator: after the 8 K inlet drop, the reaction-curve PID tuned on a dead time read as 0.50 min instead of 0.94 min reaches an IAE of 3.43 against 3.37 °C min at $t = 10$ min, while its valve travel rises from 24.8 to 112 % of span (Table 4.6).

4.7.2 Three Designs From One Identified Model

The three designs that one identification supports, the IMC-PI of Equation (4.17), the same PI as the trim of the heat-balance feedforward, and the MPC of Section 4.5.4, are compared in Table 4.7. Each case is loaded with its own seed and instruments and the structure is chosen before the run, so every design meets the same events. The PI keeps the tuning that the case ships, the feedforward rows use the default lag of 0.1 min and $K_{FF} = 1$ with that PI as the trim, while the MPC runs with the settings of Section 4.5.4 in every case, never retuned. The case added for the comparison preloads the 5 % fit and its SIMC tuning ($K_c = 4.5\,\%/^{\circ}$C, $\tau_I = 3.0$ min), raises the setpoint from 85 to 88 °C at 3 min and the feed from 20 000 to 24 000 kg/h at 15 min, and runs for 30 min.

On the setpoint half of the added case the model is the only information any design holds. The IMC-PI and the MPC stay close, with scorecard IAE of 7.17 and 8.84 °C min at 15 min and overshoots of 0.12 and 0.09 °C, while the feedforward carries the setpoint into its heat balance, opens the valve

Table 4.7 IMC-PI, Heat-Balance Feedforward and MPC From the Same Identified Model, With the Same Seed and Events in Each Case

Case	Design	IAE (°C min)	OS (°C)	e_{max} (°C)	$t_{\pm1}$ (min)	Travel (%)	Sat (%)	Plans
18, SP, feed	IMC-PI	20.64	0.12	1.98	8.34	45.6	0	—
	FF + FIC	17.21	2.34	0.99	4.72	230.6	0	—
	MPC	14.05	0.09	1.08	4.90	40.9	0	300
9, feed	IMC-PI	9.70	—	1.84	3.90	14.6	0	—
	FF + FIC	2.34	—	0.64	0	36.4	0	—
	MPC	2.51	—	0.86	0	19.3	0	300
	MPC[a]	5.53	—	1.69	2.58	17.9	0	300
7, SP	IMC-PI	13.22	0.10	—	3.56	38.1	0	—
	FF + FIC	18.30	2.43	—	6.30	142.5	24.3	—
	MPC	16.21	0.17	—	4.34	30.7	0	150
6, trap	IMC-PI	3.64	—	0.71	0	5.7	0	—
	FF + FIC	1.99	—	0.40	0	22.1	0	—
	MPC	2.22	—	0.66	0	7.9	0	300
8, $\hat{\theta} = 0.5$	ZN PID	6.05	—	1.24	0.86	186.8	0	—
	IMC-PI	6.25	—	1.81	2.60	28.2	0	—
	FF + FIC	12.72	—	1.11	1.98	315.8	0	—
	MPC	6.66	—	2.02	2.84	34.8	0	240
	IMC-PI[b]	11.78	—	2.16	4.90	17.7	0	—
	MPC[b]	7.19	—	1.99	3.30	23.1	0	240
7, SP[c]	IMC-PI	12.92	0.70	—	2.62	71.0	0	—
	MPC	14.51	0.31	—	3.62	43.9	0	150
17, capacity	IMC-PI	118.20	—	5.18	26.46	19.1	80.9	—
	MPC	114.22	—	5.18	21.90	25.9	74.3	280
14, fouling	Cascade	0.89	—	0.07	0	27.8	0	—
	MPC	7.94	—	0.39	0	23.2	0	400
19[d]	FF + FIC	8.40	—	0.66	0	153.7	0	—
	MPC	18.38	—	1.33	2.18	63.1	0	600

IAE and travel from the scorecard over 30, 30, 15, 30, 24, 15, 28, 40 and 60 min in case order. OS is the overshoot of TT-102 past a new setpoint, e_{max} its largest deviation after the load, $t_{\pm1}$ the time with the true product more than 1 °C off, Sat the share of time at a valve limit, Plans those of the MPC, none rejected. Case 6 adds a failing steam trap at 4 min, and in case 8 and row *c* every design comes from the misread fit. [a]Preview off. [b]Fitted $\hat{\theta} = 0.94$ min. [c]Designed from the 24 000 kg/h fit. [d]Challenge, suggested settings, for information.

by 28.0 % at the first execution after the step against 13.5 and 2.4 %, reaches 99.4 % at 4.64 min and overshoots by 2.34 °C. On the feed half the measurement decides. The PI waits for TT-102 and dips 1.98 °C, while the feedforward and the MPC act on FT-101 at once and dip 0.99 and 1.08 °C, the MPC with 22.6 % of valve travel after 15 min against 70.9 %, as shown in Figure 4.3. Over the case the MPC reaches 14.05 °C min against 20.64 and 17.21 °C min with the least travel of the three, and the steam used differs by less than 2 % (1015, 1032 and 1021 kg).

Rerunning the MPC on the added case with one setting changed at a time separates what the model contributes from what the optimization contributes. A move suppression of 0.5 gives an IAE of 11.72 °C min with 57.1 % of travel, and 10 gives 20.05 °C min with 33.5 %. Zero lets the valve chatter through 4214.5 % of travel although every plan stays optimal. The preview switched off raises the IAE from 14.05 to 17.13 °C min. A dead time entered as $\hat{\theta} = 0.5$ min gives the MPC 12.27 °C min with 56.5 % of travel, against 13.87 °C min with 82.0 % for the IMC-PI from the same reading.

The measured feed step of case 9 repeats the second half. The feedforward has both the lowest IAE, 2.34 against 2.51 °C min for the MPC and 9.70 for the PI, and the smaller peak, 0.64 against 0.86 °C, because its inner flow loop moves the steam within a fraction of a minute where the move suppression spreads the correction of the MPC. Without the preview the MPC reaches 5.53 °C min, so the preview accounts for 3.02 of its 7.19 °C min gain over the PI and the estimator for the rest. The failing trap enters on the steam side, where FT-102 at a fixed valve falls 54 kg/h within 0.2 min while the outlet has not yet moved, so the feedforward with its flow loop holds both the smallest peak (0.40 against 0.66 and 0.71 °C) and the lowest IAE (1.99 against 2.22 and 3.64 °C min).

The valve limit of case 17 is where the constraint enters the plan. Both designs hold the valve at

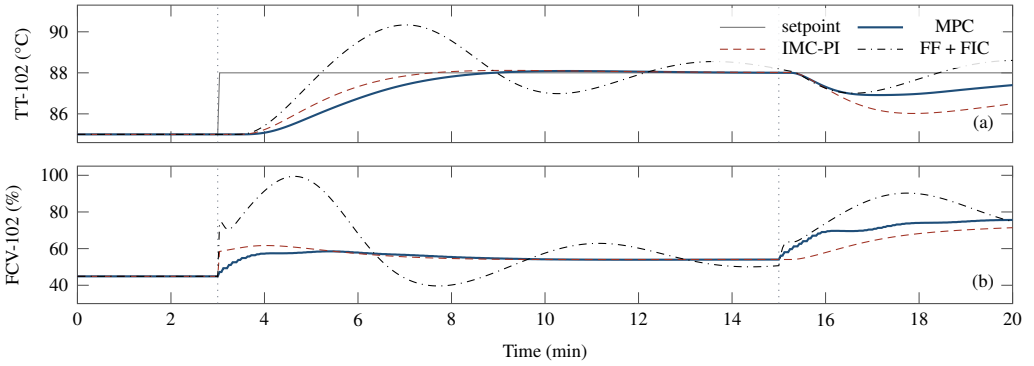


Figure 4.3 The added comparison case under the MPC (thick solid), the IMC-PI (dashed) and the heat-balance feedforward (dash-dotted), all three from the same identified model, to $t = 20$ min. The setpoint (thin gray) rises from 85 to 88 °C at 3 min and the feed from 20 000 to 24 000 kg/h at 15 min (dotted vertical lines). (a) TT-102 and setpoint. (b) Valve command. At $t = 20$ min the scorecard IAE is 12.72, 14.20 and 15.70 °C min and the valve travel 40.2, 41.8 and 215.1 %.

full lift from the start and spend the first 19.98 min outside the band, and 207 of the 280 plans of the MPC have the upper bound active on the first move, with up to 12 active-set iterations. After the cleaning at 20 min the MPC brings the valve off its stop 0.82 min after the event against 2.66 min for the PI, overshoots by 0.97 against 1.85 °C and spends 1.92 against 6.48 min outside the band, which is the whole of the difference in $t_{\pm 1}$. The PI unwinds its integral through the back-calculation of Equation (4.15) on a time scale of τ_I, while the MPC holds no integrator and plans again against the bound from a revised $\hat{d}$ at every decision.

The MPC loses the plain setpoint step of case 7, 16.21 against 13.22 °C min. Its first move is 3.98 % against the proportional kick of 22.50 %, a minute later the valve stands at 65.5 against 73.0 %, and TT-102 comes within 0.1 °C of 90 °C at 9.40 against 8.80 min, because the move suppression chosen to keep the travel of the MPC within that of the PI on the added case also slows its first moves. The model's gain is a smaller part. The steady-state balance gives a secant gain of 0.277 °C/% between 85 and 90 °C, and with only $\hat{K}_p$ changed to 0.28 °C/% the MPC reaches 15.35 and the IMC-PI designed from it 12.30 °C min, so the gap between them stays near 3 °C min. The 24 000 kg/h fit of Section 4.4 used on the same step improves both designs, and the PI stays ahead, 12.92 against 14.51 °C min, at the price of 0.70 against 0.31 °C of overshoot and 71.0 against 43.9 % of travel.

A dead time read as 0.5 min in case 8 helps the PI rules more than the MPC. The rules turn the short $\hat{\theta}$ into larger gains (20.3 and 8.5 %/°C for the reaction-curve PID and the IMC-PI, against 10.8 and 4.5 on the fitted dead time), which on this unmeasured inlet drop give IAE of 6.05 and 6.25 against 6.66 °C min for the MPC, although the reaction-curve PID moves the valve 186.8 % against 34.8 %. The Kalman gains do not depend on $\hat{\theta}$, so the misreading only shortens the record the MPC plans through and moves its IAE from 7.19 to 6.66 °C min. On the fitted dead time the MPC beats the IMC-PI, 7.19 against 11.78 °C min, but not the reaction-curve PID of the preceding comparison (3.67 °C min, 25.8 % of travel).

The fouling ramp of case 14 is lost to the shipped cascade, 7.94 against 0.89 °C min, although neither leaves the 1 °C band. A random-walk disturbance with a fixed filter gain trails a ramp, and the MPC runs on average 0.14 °C below setpoint between 5 and 15 min and 0.29 °C between 25 and 40 min while $\hat{d}$ falls from -3.3 percent of valve at 10 min to -23 percent at 40 min, whereas the inner loop of the cascade answers the steam flow that fouling removes, 159 kg/h over the run at the design valve position, as soon as FT-102 registers it. On the challenge with its suggested settings, reported for information only, the MPC peaks at 1.33 °C, outside the 1 °C band, 2.2 min after the feed inlet temperature falls at 12 min, against 0.66 °C for the suggested feedforward, because the inlet

Table 4.8 Five Designs on the Tube Case at Three Entered Gains, From 11 to 40 min With the Same Seed and Events

Design	t_{TAH} (min)	Peak (°C)	IAE (°C min)	Steam (kg)	Travel (%)	Sat (min)
$\hat{K}_p = 0.252$ (0.7×), $\delta = 40.8\,\%$						
IMC-PI	2.96	90.52	19.30	863.7	74.9	0
FF + FIC	4.60	91.62	28.64	866.0	339.7	2.66
MPC, nominal	0	89.74	16.79	855.0	62.7	0
MPC, tube	0	89.62	16.36	854.5	64.1	0
MPC, band $-h_N$	0	89.57	23.58	850.2	191.2	0.30
$\hat{K}_p = 0.36$ (1.0×), $\delta = 1.8\,\%$						
IMC-PI	4.42	90.70	21.77	864.2	51.0	0
FF + FIC	3.36	91.53	22.07	863.4	198.8	1.96
MPC, nominal	0	89.92	19.33	855.7	53.2	0
MPC, tube	0	89.92	19.33	855.7	53.2	0
MPC, band $-h_N$	0	89.92	19.33	855.7	53.2	0
$\hat{K}_p = 0.54$ (1.5×), $\delta = 34.3\,\%$						
IMC-PI	6.52	90.91	26.41	865.1	35.6	0
FF + FIC	3.74	91.42	22.84	864.0	162.4	0.68
MPC, nominal	1.88	90.16	22.53	856.2	50.1	0
MPC, tube	1.14	90.12	21.76	855.5	61.3	0
MPC, band $-h_N$	0	89.26	25.66	849.9	217.1	0

t_{TAH} is the time at or above TAH-102 and Sat the time at a valve limit. Band $-h_N$ is the nominal band moved in by the tube's end half-width. IMC-PI gains 6.3, 4.4 and 3.0 %/°C, $\tau_I = 3$ min. Each MPC row made 290 plans, none rejected. No row reached TAHH-102, TAL-102 or FAL-102.

temperature, which the feedforward reads at TT-101, reaches the MPC only through TT-102 and $\hat{d}$.

On the other shipped upsets the MPC lies between the single loop and the structure built for each (IAE 4.01 against 6.60 and 0.78 °C min for the single loop and the cascade on the header sag, 7.19 against 11.67 and 4.81 °C min for the single loop and the feedforward on the colder feed), it beats the single loop under stiction (5.88 against 10.03 °C min). Outside case 17 no bound was active in any plan, so there the MPC acted as a fixed linear law and its results come from the model, the preview and the estimator rather than from constraint handling. All 3260 of its plans in Table 4.7 were accepted, at about 85 µs each in Node.js where no bound was active and 150 to 250 µs on case 17, and at 0.25 to 0.30 ms in Chrome, or 0.3 % of one processor core in real time (wall-clock values on one desktop computer).

4.7.3 Alarm Band From the Learner's Fit

The tube case starts in manual operation, steps FCV-102 by 5 % at 1 min as the step test, hands the valve to the MPC at 11 min, raises the setpoint from 85 to 89 °C, one degree under TAH-102, at 14 min, and drops the feed from 20 000 to 17 000 kg/h at 28 min, which heats the outlet toward the alarm. Every design is judged from 11 to 40 min with the same seed, instruments, $\hat{\tau}$ and $\hat{\theta}$, with $\hat{K}_p$ entered at 0.7, 1.0 and 1.5 times the 0.36 °C/% of the step test, as reported in Table 4.8. The PI rows use Equation (4.17) with $\tau_c = \hat{\theta}$ on the same entry, the feedforward row $K_{FF} = 1$ with that PI as the trim, and the MPC rows the settings of Section 4.5.4. The last row of each group is the nominal MPC retuned once by the tube's own measure, its band moved in at every decision by the half-width that the tube plan reaches at its end.

With the shipped entry of 0.355 °C/% the misfit is 1.03 %, the tube is 0.02 to 0.04 °C wide, and both bands give the same run (peak 89.91 °C, IAE 19.22 °C min), as they do at the step test's gain. At 0.7 times neither band reaches TAH-102, and the tube lowers the peak by 0.12 °C and the IAE by 2.6 % for 2.3 % more travel. At 1.5 times, δ is 34.3 %, the tube widens to ±1.49 °C at the end of the plan, and TT-102 crosses TAH-102 about 2.1 min after the feed drop under both bands. The tube shortens the time above it from 1.88 to 1.14 min with 3.4 % less IAE and 22 % more travel, as shown in Figure 4.4, but does not remove it.

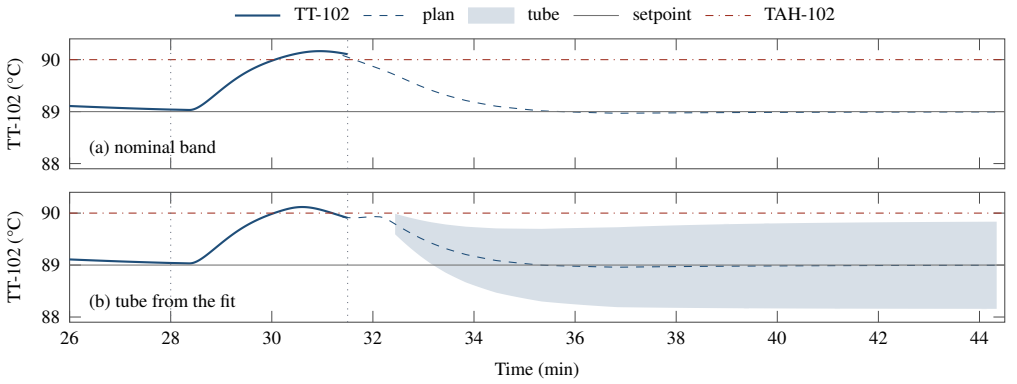


Figure 4.4 The tube case with $\hat{K}_p = 0.54$, 1.5 times the step test's gain, paused at $t = 31.5$ min (dotted vertical lines, with the feed drop at 28 min), with TT-102 to that instant (thick solid) and the reading expected by the plan made at 31.4 min (dashed). The setpoint is the thin gray line, TAH-102 the dash-dotted line and the tube the shaded band. (a) Nominal band, TT-102 at 90.10 °C with TAH-102 active and FCV-102 at 44.0 %. (b) Tube from the learner's fit, TT-102 at 89.91 °C, last at 90 °C at 31.18 min, FCV-102 at 40.9 % and a tube widening from ±0.20 °C to ±0.84 °C.

The remaining error is the feed path's. Applied open loop to a copy of the exchanger, the plans made at the feed drop (28.0 to 28.4 min) under-predict TT-102 by 0.26 °C 2.1 min ahead and by 0.57 °C 3.4 min ahead even with the step test's gain, inside tubes of 0.02 to 0.03 °C, because the feed path of Equation (4.19), identified at 20 000 kg/h, is off at 17 000 kg/h. With the right gain the feedback corrects this in time (peak 89.92 °C). At 1.5 times every correction is a third too small, the same plans miss by 0.50 and 1.11 °C, and their tubes of 0.34 and 0.72 °C cover the gain's share but not the feed path's. The crossing is not decided in the dead time, since closing FCV-102 at any plan up to 29.4 min would have held TT-102 at or below 89.97 °C, and a tube twice as wide removes it (peak 89.87 °C) while one three times as wide crosses again for 1.04 min.

The nominal band moved in by the tube's end half-width does remove the crossing at 1.5 times (peak 89.26 °C), at 3.5 times the tube's travel and 18 % more IAE, and at 0.7 times it saturates the valve for 0.30 min and raises the IAE from 16.36 to 23.58 °C min. The tube wins on IAE and travel and the constant back-off on the alarm. The PID structures, the cascade included (5.52 to 6.92 min, not tabulated), cross TAH-102 at every entry, and to stay under it the IMC-PI needs a setpoint of 88.00 to 88.50 °C and the feedforward 87.25 to 88.00 °C. Over 23 faults at their extremes in place of the feed drop, the nominal band and the tube spend 8.78 and 8.82 min above TAH-102 in all at 1.0 times the gain and 14.78 and 12.44 min at 1.5 times, against 26.78 to 100.24 min for the PID structures. No MPC run reaches TAHH-102, and the tube spends longer than the nominal band at a limit in 4 of 46 runs, by 0.04 to 0.10 min.

A wide tube works the valve hard. Its travel is 109 against 55 % for the nominal band at a misfit of 53 % and 539 against 76 % at 65 %, and at 70 % its plans hold FCV-102 at a stop for 1.90 min and set off FAL-102 twice, although it never spends longer above TAH-102. The MPC of Section 4.5.4 has a limit of its own. With $\hat{K}_p = 2.0$, 5.6 times the step test's gain, the nominal band moves FCV-102 through 1665 % of travel in the window, 6.90 min of it at a stop, with five FAL-102 alarms, in a cycle of plan and estimator with a period of about 2.4 min (the IMC-PI from that entry moves 15.8 % but spends 9.50 min above TAH-102). Around a linear plant of the model's structure the same controller cycles at a gain ratio of 5.56 (729 % of travel) but not at 5.0 (64 %), nor at 5.56 with $\kappa = 5$ or a factor of 0.1 on q_d (43 and 38 %), so the margin against an overestimated gain, about a factor of five, belongs to the design and not to the exchanger. The tube lowers it, with 2422 % of travel at $\hat{K}_p = 1.5$ where the nominal band moves 70 %.

4.7.4 Numerical Implementation

The application runs on the engine of Section 1.6.1, with its three clocks set as follows. The physics advances by classical fourth-order Runge-Kutta on a fixed 0.005 min step, the controller executes every 0.02 min, and the historian records every fourth physics step, 50 samples per simulated minute, the rate the application declares. The trajectory does not depend on frame rate and the noise is seeded, so two machines produce identical exports and noisy scenarios replay exactly.

Within a control execution the order is fixed, instruments first with the transport delay ahead of the bias, the noise and the transmitter lag, then the available-duty refresh, then the feedforward calculation, then the estimator of the predictive controller, then the controllers, then the stiction model, then integration. Placing the delay line ahead of everything else preserves the measurement chain, because no controller reads a state variable directly (the feedforward takes its latent heat at the design shell pressure, which no transmitter reads). The exported record carries eighty-six columns sampled every 0.02 min (1.2 s), including the true and the measured temperature, the commanded and the actual valve position, the separate proportional, integral and derivative contributions, the duty ledger, every disturbance, the estimate and plan of the predictive controller, and the misfit and widths of its tube, which is sufficient for an offline identification or grading task in any language. Forty self-checks run inside the application. Eighteen cover the steam-table fits, the log-mean fallback, the zero-derivative condition at the design steady state, the sign of every disturbance, the flow dependence of gain and dead time, the fouling and flooding relations, and the two tuning rules. Seven test the predictive controller, including its first move against a least-squares closed form, a forced rejection handed to TIC-101 without a bump, and a check of its source for any read of the plant state. Fourteen test the tube, including δ against a replay of the logged step test and each half-width against a recomputation, and one checks that the feedforward reads neither the shell pressure nor the steam quality.

4.8 Technology in Practice

The loop of Section 4.3 reads TT-102 through a first-order lag and moves FCV-102 through another. Figure 4.6 draws it with the devices an exchanger in a plant would carry in the place of the schematic's symbols.

What Measures It

An outlet temperature in a plant is read by a resistance temperature detector or a thermocouple, held in a thermowell. Thermocouples are the most rugged and wide-ranging of the contact sensors, and resistance detectors need no reference compensation and are more linear. Their accuracy varies with the type and the quality of the installation [26], so no single figure applies. The lag τ_{TT} of Table 4.2 stands for the element and its well, and the bias b for a lead-wire or calibration error that no controller can see.

> **TECHNOLOGY** **Sensor: A Resistance Thermometer in a Thermowell**
>
> A resistance temperature detector is a winding of pure metal wire, usually platinum, whose resistance always rises with temperature. It is relatively insensitive but very linear. The resistance of its lead wires adds to the reading and is taken as a falsely high temperature, a steady bias on the trend. The thermowell that protects the element adds a first-order lag through its own mass and heat capacity. An air gap between element and well adds a second lag, and multiple lags raise the phase shift of a feedback loop and may lead to oscillation [26]. How fast the reading answers is therefore set mostly by the installation, not by the element.

What Moves It

FCV-102 of Equation (4.9) is a first-order stem with a lag of 0.20 min, an installed characteristic x^a and an optional deadband S. A plant valve in this service is pneumatic, with a positioner that closes an inner loop on the stem [5], and the deadband stands for the friction that the positioner cannot

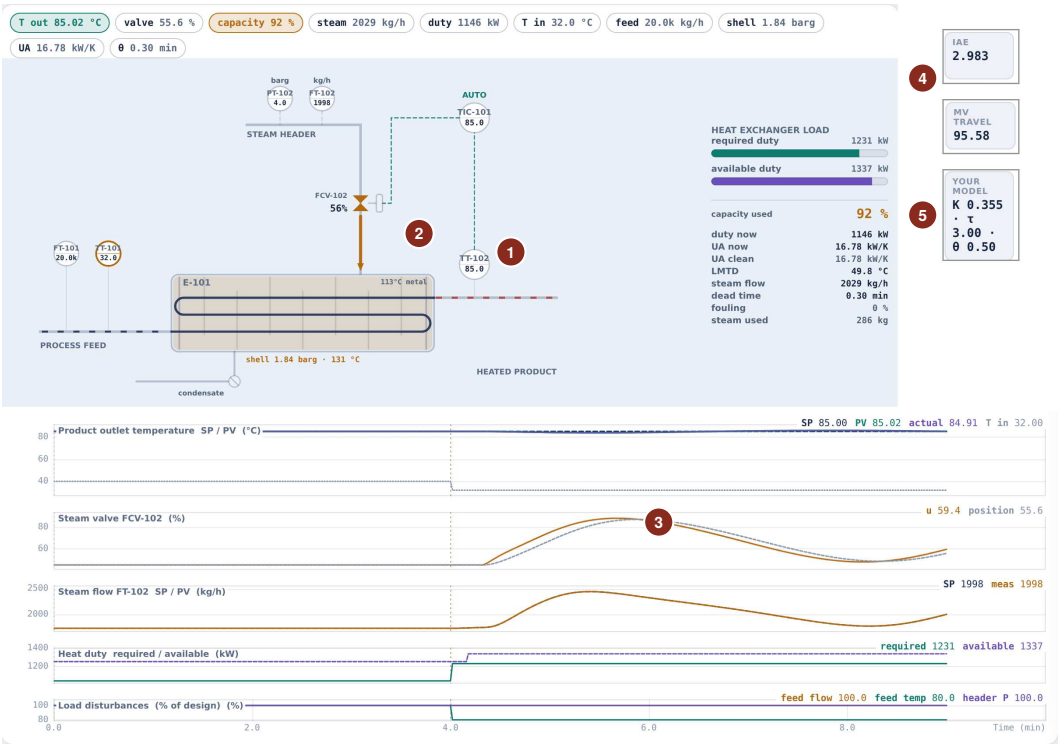


Figure 4.5 Tuning comparison after an inlet-temperature disturbance, at $t = 9$ min. A Ziegler-Nichols PID computed from (4.18) with the dead time read as $\hat{\theta} = 0.5$ min instead of the fitted 0.94 min ($K_c = 20.3\,\%/°\text{C}$, $\tau_I = 1.0$ min, $\tau_D = 0.25$ min) meets a feed 8 K colder at $t = 4$ min. The outlet stays within 1.24 °C of setpoint (Table 4.6) while the valve and the steam flow ring about the new resting position, the steam flow rising 386 kg/h above its final 2056 kg/h and falling 288 kg/h below it, where the SIMC tuning makes a single monotone valve move. IAE is 2.983 °C min and valve travel 95.6 % of span at this instant, reaching the 3.43 °C min and 112 % of Table 4.6 at $t = 10$ min. The identified constants are $\hat{K}_p = 0.355$, $\hat{\tau} = 3.00$ and $\hat{\theta} = 0.50$, which is where the error entered. The picture is a composite: the readouts and plant view, three scorecard tiles moved into the empty band beside them, and the trends. The callouts mark (1) TT-102, (2) FCV-102, (3) the ringing valve trend, (4) the IAE and valve travel tiles and (5) the identified model with its short dead time.

overcome. The heat-balance feedforward of Equation (4.16) also has its plant form. A feed flow transmitter and a summing block change the steam at once with the load, and because the steam heats more metal than the liquid does, the steam path is the slower one and the feedforward needs dynamic compensation [26], which is the role of τ_{lead} and τ_{lag}.

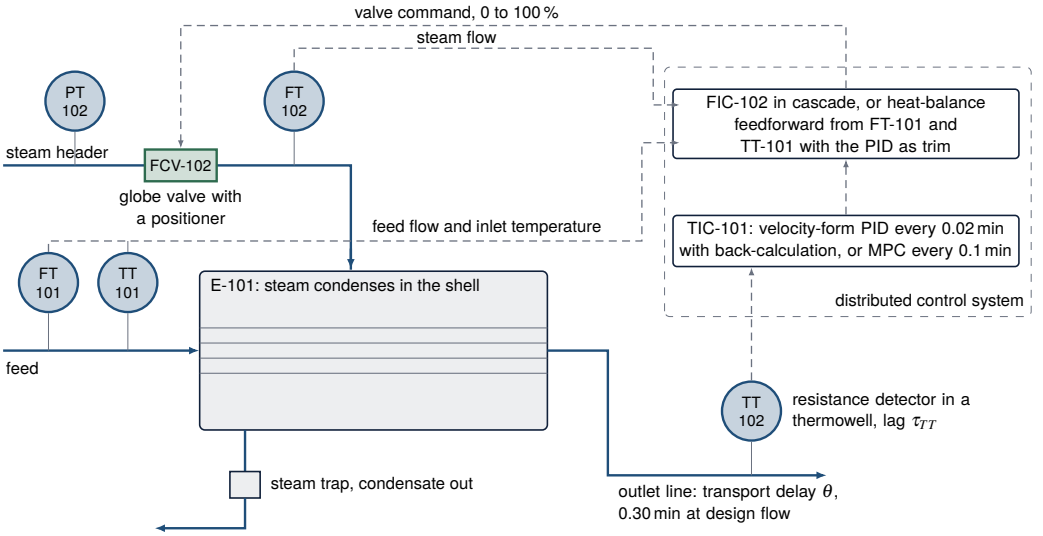


Figure 4.6 The loop of the chapter with plant devices in place of the schematic's symbols. A resistance detector in a thermowell reads TT-102 at the end of the outlet line, and the controllers on a distributed control system move valve FCV-102 through its positioner, either directly, through the steam flow loop FIC-102, or through the heat-balance feedforward. Notice that the transport delay and the thermowell both lie between the bundle and the controller, so the controller sees every move of the valve late.

> **TECHNOLOGY** **Final Element: The Steam Valve and Its Positioner**
>
> Steam, not the process liquid, is the stream to throttle, since changing the liquid's flow would upset the units around the exchanger [26]. A pneumatic valve is chosen for its failure mode first. An air-to-open valve fails closed on loss of air, and the risk of the process, not habit, should decide that position [26]. A positioner closes an inner loop on the stem to reduce the effects of friction and nonlinearity [5]. Friction it cannot overcome makes the stem stick and jump, and integral action then hunts above and below the setpoint [26].

Where the Controller Runs

A distributed control system runs production at one site in chemical manufacturing, refining and similar industries. Its loops are often nested, with the setpoint of one taken from another, and they run with cycle times from milliseconds to minutes [27]. TIC-101 and FIC-102 execute every 0.02 min (1.2 s) and the predictive controller decides every 0.1 min, both well inside that range. The cascade of the chapter is a remote setpoint from one PID block to another. Industrial practice places model predictive control one level above the PID loops that move the valves [20], while the MPC of Section 4.5.4 moves FCV-102 directly and hands it to TIC-101 when a plan is rejected.

IN PRACTICE The Plant Step Test

Before a model-based controller is commissioned, engineers run a plant test. A pre-test steps each manipulated variable to adjust the instruments and the existing PID loops, finds each output's time to steady state and gives data for a first model. The formal test runs around the clock with engineers watching. Each variable is stepped 8 to 15 times with a signal-to-noise ratio of at least six, which can take 5 to 15 days. No PID configuration or tuning change is allowed during the test, and operators may intervene only to avoid critical situations [20].

What Is Modernizing It

Three developments touch this loop. The first is the smart positioner. It records actuator pressure against stem position, and this valve signature identifies excessive packing friction, trim interference and poor seating [26], the faults behind the stiction of case 15. The second is the soft sensor for fouling. Regression and neural-network models estimate the fouling resistance of a steam heater from routine readings of flows and temperatures [16], the role the Conclusion gives to an online estimate of R_f. The third is predictive control on identified models. It originated in the process industries, where the processes are often slow enough to allow the computation [5], and its identification packages rely mainly on step-like or pseudo-random binary test signals [20], the industrial form of the step test of Section 4.4.

4.9 Limitations

The simulator is a teaching instrument and not a design tool. The lumped formulation cannot represent axial profiles, shell-side maldistribution, baffle effects or two-phase behavior on the tube side, so it must not be used to size equipment. Disturbances are deterministic steps and ramps, scripted by a scenario or set by the learner, rather than stochastic plant behavior. Two of them are stylized. A fouling ramp of 2.5 %/min would not occur in a physical system, and it serves to show the drift of a controller output under a slowly changing plant. A step in header pressure of 20 % is larger than most plants would see, and it separates the disturbance the inner flow loop can reject from the two it cannot. Sensors are ideal apart from the explicit lag, bias and noise that a user enables. The condensate system is reduced to one flooding parameter, and no equipment fault beyond fouling, flooding, stiction and transmitter bias is represented. The steam correlations are valid over 1 to 10 bar absolute and nowhere else, and the film correlation is extrapolated below the turbulent range. The fixed-step integrator and the fixed control period make every run reproducible, which a real distributed control system is not. Each simplification that was retained, the log-mean driving force, the flow-dependent film coefficient, the square-root valve law and the transport delay, was retained because removing it removes a behavior the loop must be tuned around. Section 4.3.5 states where each assumption fails.

The predictive controller plans on a fixed linear model, a FOPDT valve path from wherever the learner identified it and a feed path identified once at the design point. It omits the gain that falls with load and outlet temperature (0.56 to 0.22 °C/% between 16 000 and 24 000 kg/h, a secant gain of 0.277 °C/% between 85 and 90 °C), the dependence of a fitted gain on the direction of the step (0.355 and 0.440 °C/% for +5 and −5 % steps), the dead time that falls with feed rate, the available duty and stiction, and it sees the inlet temperature and the header pressure only through $\hat{d}$. Nothing re-identifies it on line, and its random-walk disturbance trails a ramp.

The tube inherits the reach of one valve step at one operating point and in one direction. It bounds a relative error of the valve path near that point, not an error in the declared feed path, a dead time that moves with flow or an unmeasured load, and for a gain error it is exact only once the estimator has settled. A misfit of 100 % or more builds no tube, and from about 70 % the tube drives the valve into its stops.

⚠ OPEN CHALLENGE Knowing When a Step Test Is Enough

A survey of industrial practice found that identification rests on step-like test signals that need the close attention of experienced engineers. No tool tells whether the data represent the process dynamics well enough for the controller, so plants are generally over-tested, which raises the cost of the project. Adaptive predictive controllers that keep their models current were proposed, but few reached the market, which reflects how hard adaptive control is in practice [20]. Progress would be a test that stops when the model is good enough and a controller that follows the load.

4.10 Discussion

Heat exchange accounts for most of the energy transferred within a process plant and for a large share of the energy lost. Pinch analysis established how strongly the recovery achievable in a network depends on how its exchangers are matched [28], and that recovery is realized only if each service holds its duty during upsets, which is a control problem. Fouling turns the same argument into a direct cost. An industrial survey of New Zealand plants quantified the combined capital and operating penalty of exchanger fouling across a national industrial base [29], and much of that penalty accrues through lost duty and extra fuel rather than through cleaning. Reading fouling from a drifting controller output, the task of scenario 14, is a maintenance-planning skill before it is a control skill, and the rise in balancing valve position from 44.9 % to 71.9 % across a full fouling ramp is the signal a monitoring system would trend.

Automation of these services also changes what an operator is asked to do. A cascade rejects the header-pressure upset almost completely, holding TT-102 within 0.16 °C of setpoint where the single loop dips 1.22 °C, and a heat-balance feedforward added on top answers measured load changes before they reach the outlet. That is desirable until the assumed heat capacity is wrong and the feedback trim has been absorbing the resulting bias over a long period. A heat-capacity error of $\pm 12\,\%$ is available as a disturbance for that reason, and (4.16) uses the nameplate value so the error is not silently corrected. Increasing automation does not remove the need to understand the loop. It moves the required understanding from moment-to-moment control to recognizing when the automation is compensating for a fault that ought to be corrected.

⚠ OPEN CHALLENGE Seeing Fouling Before It Costs Duty

Fouling remains the main unresolved problem in heat transfer and an almost universal one in the design and operation of exchangers. The deposit adds a thermal resistance, narrows the flow passage and raises the cost of maintenance. Neural-network and regression models now estimate the fouling resistance from routine plant readings to plan cleaning schedules [16], and online determination of that resistance in shell-and-tube exchangers is under study [30]. Progress would be an estimate trusted enough to schedule a cleaning by.

4.11 Conclusion

This chapter studied a steam-heated shell-and-tube exchanger whose plant constants are hidden, so the controller is built from an experiment. The simulator integrates five states with a log-mean driving force, a tube-side film coefficient that scales as $F_p^{0.8}$, saturated-steam correlations and a transport delay of 0.30 min at design flow that varies inversely with production rate. A 5 % valve step at the design point identifies $\hat{K}_p = 0.355\,^{\circ}\mathrm{C}/\%$, $\hat{\tau} = 3.02$ min and $\hat{\theta} = 0.94$ min, three lumped surrogates for five dynamic elements, and the same test at 16 000 and 24 000 kg/h moves the identified gain by a factor of 2.6. A dead time read 0.44 min short raises the reaction-curve gain from the 10.8 %/°C that the fitted dead time warrants to the 20.3 %/°C the scenario loads, where the SIMC rule on the same fit returns 4.5 %/°C. After an 8 K inlet drop both reaction-curve tunings reach an IAE near 3.4 °C min at $t = 10$ min against 8.61 °C min for SIMC, so the error advantage belongs to the rule. The misreading adds nothing to it, multiplies the valve travel by 4.5 and leaves a loop that grows into a valve-limited oscillation when the feed falls to 18 000 kg/h. Around that comparison the application arranges load

upsets that motivate cascade and feedforward, and four fault conditions whose symptoms resemble a tuning problem without being one.

A linear MPC built on the same three identified numbers, with FT-101 as a preview and the valve range as a hard limit, reaches 14.05 °C min on the added comparison case against 20.64 for the IMC-PI and 17.21 for the feedforward, and it loses the setpoint step alone (16.21 against 13.22 °C min) and the fouling ramp. A tube sized from the learner's own step-test misfit costs nothing with a correct fit and, with the gain entered 1.5 times too high, shortens the time above TAH-102 from 1.88 to 1.14 min where a band moved in by its end width removes it at 3.5 times the valve travel, because the error that remains lies in the feed path and not in the gain.

The model supports control methods beyond the four architectures. Equations (4.1) to (4.11) are a small differential-algebraic system with measured inputs, so they can be posed directly as the prediction model of a nonlinear model predictive controller [31, 32]. The linear MPC of Section 4.5.4 plans against the 0–100 % valve span and coordinates the steam valve with the measured feed rate, but it knows neither the available duty of 1252 kW at design flow nor the gain that falls with load, and a nonlinear predictive controller on the full model would add both. Moving horizon estimation or an extended Kalman filter driven by TT-102, FT-101 and FT-102 would infer the fouling resistance R_f and the tube-metal temperature T_w, neither of which is measured. An online estimate of R_f turns the drifting controller output of scenario 14 into a soft sensor for cleaning schedules. A fuller model would add a distributed tube-side model for axial profiles, a condensate system with stall dynamics, a deposition-minus-removal fouling law, and scenarios at more than one production rate to make gain scheduling measurable. The lesson of the exchanger is that the three identified numbers belong to an operating load as much as to the plant, and that the experiment behind them decides more than the tuning rule applied to them.

KEY POINTS

- A small step of the steam valve gives three numbers for tuning, but they belong to one feed rate as much as to the exchanger.
- The fitted dead time of 0.94 min exceeds the 0.30 min transport delay because it lumps the valve, the metal and the transmitter.
- A dead time read 0.44 min short raises the reaction-curve gain from 10.8 to 20.3 %/°C, adds nothing to the error advantage and multiplies the valve travel by 4.5.
- Fouling shows in the controller output, which drifts upward while the temperature stays at setpoint.

Where else this applies. The same step test and tuning rules meet a tank in Volume 1, Chapter 2 and a heated device in Volume 1, Chapter 4.

References

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

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

[3] F. W. Dittus and L. M. K. Boelter, "Heat transfer in automobile radiators of the tubular type," *International Communications in Heat and Mass Transfer*, vol. 12, no. 1, pp. 3–22, 1985.

[4] E. N. Sieder and G. E. Tate, "Heat transfer and pressure drop of liquids in tubes," *Industrial & Engineering Chemistry*, vol. 28, no. 12, pp. 1429–1435, 1936.

[5] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[6] D. J. Correa and J. L. Marchetti, "Dynamic simulation of shell-and-tube heat exchangers," *Heat Transfer Engineering*, vol. 8, no. 1, pp. 50–59, 1987.

[7] H. Müller-Steinhagen, "Heat transfer fouling: 50 years after the Kern and Seaton model," *Heat Transfer Engineering*, vol. 32, no. 1, pp. 1–13, 2011.

[8] W. Miao and B. Xu, “Application of feedforward cascade compound control based on improved predictive functional control in heat exchanger outlet temperature system,” *Applied Sciences*, vol. 13, no. 12, p. 7132, 2023.

[9] A. Vasičkaninová, M. Bakošová, A. Mészáros, and J. J. Klemeš, “Neural network predictive control of a heat exchanger,” *Applied Thermal Engineering*, vol. 31, no. 13, pp. 2094–2100, 2011.

[10] S. Bennett, “The past of PID controllers,” *Annual Reviews in Control*, vol. 25, pp. 43–53, 2001.

[11] J. G. Ziegler and N. B. Nichols, “Optimum settings for automatic controllers,” *Transactions of the ASME*, vol. 64, no. 8, pp. 759–765, 1942.

[12] D. E. Rivera, M. Morari, and S. Skogestad, “Internal model control: PID controller design,” *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

[13] S. Skogestad, “Simple analytic rules for model reduction and PID controller tuning,” *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

[14] B. W. Bequette, “Process control practice and education: Past, present and future,” *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[15] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, “A survey of good practice in control education,” *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[16] R. Jradi, C. Marvillet, and M. R. Jeday, “Estimation and sensitivity analysis of fouling resistance in phosphoric acid/steam heat exchanger using artificial neural networks and regression methods,” *Scientific Reports*, vol. 13, p. 17889, 2023.

[17] M. A. A. S. Choudhury, N. F. Thornhill, and S. L. Shah, “Modelling valve stiction,” *Control Engineering Practice*, vol. 13, no. 5, pp. 641–658, 2005.

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

[19] 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.

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

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

[22] 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.

[23] L. Chisci, J. A. Rossiter, and G. Zappa, “Systems with persistent disturbances: predictive control with restricted constraints,” *Automatica*, vol. 37, no. 7, pp. 1019–1028, 2001.

[24] D. Q. Mayne, M. M. Seron, and S. V. Raković, “Robust model predictive control of constrained linear systems with bounded disturbances,” *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

[25] M. V. Kothare, V. Balakrishnan, and M. Morari, “Robust constrained model predictive control using linear matrix inequalities,” *Automatica*, vol. 32, no. 10, pp. 1361–1379, 1996.

[26] 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.

[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] B. Linnhoff and E. Hindmarsh, “The pinch design method for heat exchanger networks,” *Chemical Engineering Science*, vol. 38, no. 5, pp. 745–763, 1983.

[29] R. Steinhagen, H. Müller-Steinhagen, and K. Maani, “Problems and costs due to heat exchanger fouling in New Zealand industries,” *Heat Transfer Engineering*, vol. 14, no. 1, pp. 19–30, 1993.

[30] M. Markowski and P. Trzcinski, “Comparison of the analytical model with neural network model on the case of heat exchanger behavior under fouling,” *Scientific Reports*, vol. 15, p. 42078, 2025.

[31] 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.

[32] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, “GEKKO optimization suite,” *Processes*, vol. 6, no. 8, p. 106, 2018.

4.12 Exercises

4.1 Concept check · 1 min

TIC-101 can drive the steam valve directly, send a steam-flow setpoint to FIC-102, or trim a heat-balance feedforward. Which variable is the manipulated variable of the exchanger in every one of these structures?

(a) The steam-flow setpoint of FIC-102, in percent of the FT-102 span
(b) The position of the steam valve FCV-102, in percent of its travel
(c) The pressure of the steam header at PT-102, in bar gauge
(d) The process feed flow measured at FT-101, in kilograms per hour

4.2 Concept check · 2 min

Equation (4.1) is the energy balance on one of the five inventories of the model. Which inventory is it, and where does its heat come from?

(a) The tube metal, which receives the duty Q_s from the condensing steam
(b) The steam in the shell, which is charged by the flow through FCV-102
(c) The process fluid in the bundle, which receives the duty Q_p from the tube metal
(d) The process fluid in the bundle, which receives the duty Q_s straight from the steam

4.3 Concept check · 2 min

The tubes foul slowly while the loop holds 85 °C in automatic at a constant feed rate and inlet temperature. Which signal shows the fouling?

(a) TT-102 drifts below its setpoint as the conductance of the bundle falls
(b) The transport delay to TT-102 grows as the deposit builds in the tubes
(c) TT-101 at the feed inlet falls as less heat reaches the product
(d) The valve position drifts upward while TT-102 stays on its setpoint

4.4 Theory and calculation · 5 min

Exchanger E-101 heats 20 000 kg/h of process liquid with $C_p = 4.18\,\mathrm{kJ/kg/K}$ from 40 to 85 °C. The shell settles at 2.60 bar absolute.

(a) Find the duty in kilowatts. *Checkpoint:* 1045 kW.

(b) Find the latent heat from Equation (4.8) and the steam flow that supplies the duty. *Checkpoint:* 2178 kJ/kg and 1727 kg/h.

(c) Find the transport delay to TT-102 at this feed rate and at 30 000 kg/h, with $V_{\mathrm{trans}} = 0.1\,\mathrm{m}^3$ and $\rho = 1000\,\mathrm{kg/m}^3$. *Checkpoint:* 0.30 and 0.20 min.

4.5 Theory and calculation · 8 min

A valve step at the design point is fitted as $\hat{K}_p = 0.355\,^{\circ}\mathrm{C}/\%$, $\hat{\tau} = 3.0\,\mathrm{min}$ and $\hat{\theta} = 0.94\,\mathrm{min}$.

(a) Find the SIMC PI settings of Equation (4.17) with $\tau_c = \hat{\theta}$. *Checkpoint:* $K_c = 4.5\,\%/^{\circ}\mathrm{C}$ and $\tau_I = 3.0\,\mathrm{min}$.

(b) Find the reaction-curve PID settings of Equation (4.18). *Checkpoint:* $K_c = 10.8\,\%/^{\circ}\mathrm{C}$, $\tau_I = 1.88\,\mathrm{min}$ and $\tau_D = 0.47\,\mathrm{min}$.

(c) The dead time is misread as 0.50 min. Find both sets of settings again, with $\tau_c = \hat{\theta}$ as before, and say what changes in each. *Checkpoint:* SIMC 8.5 and 3.0 min. Reaction curve 20.3, 1.0 min and 0.25 min.

4.6 Theory and calculation · 6 min

The clean conductances of Table 4.2 are $UA_{p0} = 1380\,\mathrm{kJ/min/K}$ and $UA_{s0} = 3726\,\mathrm{kJ/min/K}$, and they act in series.

(a) Find the overall conductance of the clean exchanger at the design feed rate, in kW/K. *Checkpoint:* 16.78 kW/K.

(b) Find it again at 100 % fouling, where $R_f = 2.6 \times 10^{-4}\,\mathrm{min\,K/kJ}$, and the loss in percent. *Checkpoint:* 13.30 kW/K, a loss of 20.7 %.

(c) Find it for the clean exchanger at 24 000 kg/h from Equation (4.10), and compare its rise with the rise in duty. *Checkpoint:* 18.63 kW/K, a rise of 11 % against 20 % in duty.

4.7 App challenge · 9 min

Load scenario 3, Step test, and change no setting. The run starts in Manual at a steady state with FCV-102 at 44.9 %. Use the default seed. This scenario has no measurement noise. Press Run, let about one minute pass and press Step valve +5%. Run until 20 min after the step and export the CSV file. Measure times from the last row in which `valve_cmd_pct` still reads 44.9.

(a) Read TT-102, column `temp_pv_C`, before the step and 20 min after it, and compute the process gain. *Checkpoint:* 85.01 and 86.78 °C, so 0.354 °C/%, within 0.005.

(b) Find the times at which TT-102 has covered 28.3 and 63.2 % of that change, and apply Equation (1.9). *Checkpoint:* 2.00 and 3.96 min, within 0.05 min, so $\hat{\tau} = 2.94\,\mathrm{min}$ and $\hat{\theta} = 1.02\,\mathrm{min}$.

(c) Press Fit FOPDT to my step test and read the Regression tile. Compare the three numbers with the two-point fit and with the transport delay of 0.30 min. *Checkpoint:* 0.355, 3.02 and 0.94, within 0.002, 0.06 min and 0.02 min.

4.8 App challenge · 8 min

Load scenario 8, Ziegler-Nichols, and an optimistic dead time. The loop is in Auto with the reaction-curve tuning from a dead time read as 0.5 min, and the feed inlet falls from 40 to 32 °C at 4 min. Use the default seed and run each case to 24 min. For each run read the IAE and the valve travel on the scorecard and find the lowest TT-102 in the exported file.

(a) Run the scenario as shipped. *Checkpoint:* 6.05 °C min, 187 % and 83.76 °C.

(b) Press Reset. Before Run, set Identified dead time to 0.94 min and select ZN PID under Tune from my model, which loads $K_c = 10.8\,\%/°\mathrm{C}$, $\tau_I = 1.88\,\mathrm{min}$ and $\tau_D = 0.47\,\mathrm{min}$. Run again. *Checkpoint:* 3.67 °C min, 25.8 % and 83.64 °C.

Accept an IAE within 0.1 °C min, a travel within 5 % and a temperature within 0.03 °C. State what the misread dead time bought and what it cost.

4.9 App challenge · 8 min

Load scenario 17, Capacity limit. The feed is 27 000 kg/h, the tubes are 70 % fouled and the header is at 3.2 barg. At 20 min the exchanger is cleaned and the header returns to 4.0 barg. Use the default seed, run each case to 28 min and export the CSV file.

(a) Run the scenario as shipped. Read TT-102 in the row at 19.9 min, and the required and the available duty in the columns `duty_required_kW` and `duty_available_kW`. *Checkpoint:* 79.82 °C, 1411 kW and 1248 kW.

(b) In the same file find the time at which `valve_cmd_pct` first falls below 100 and the highest TT-102 after it. *Checkpoint:* 22.66 min, within 0.05 min, and 86.85 °C.

(c) Press Reset, set Anti-windup to Off at the Advanced level and run again. Read the valve command and TT-102 at 28 min, and compare the column `I_term` at 19.9 min in the two files. *Checkpoint:* 100 % and 87.70 °C. The integral term is 254.6 % against 100.0 %.

Accept a temperature within 0.03 °C and a duty within 2 kW. Say why no tuning removes the offset before 20 min, and what the anti-windup changes after it.

4.10 Application area · 3 min

At the design load a ramp to 100 % fouling moves the balancing valve position from 44.9 to 71.9 % and the overall conductance from 16.78 to 13.30 kW/K, while the loop holds 85 °C.

(a) Say what has happened to the duty over the ramp, and why the outlet temperature is of no use as a fouling monitor. *Checkpoint:* The duty is unchanged at 1045 kW.

(b) Name two other faults of the model that raise the valve position at a constant load, and the instrument or check that tells each from fouling. *Checkpoint:* For example a sagging header, seen at PT-102.

4.11 Application area · 3 min

In the cascade TIC-101 writes the setpoint of FIC-102 in percent of the 4000 kg/h span of FT-102, limited to 0–100 %. The steam header delivers at most 2600 kg/h.

(a) Express the header capacity and the design steam flow of 1727 kg/h in percent of span. *Checkpoint:* 65 % and 43.2 %.

(b) During a large upset the setpoint from TIC-101 passes 65 %. Say what each integral does, and what an engineer would configure to prevent it. *Checkpoint:* The outer integral winds.

4.12 Review the agent · 5 min

An AI assistant is asked to tune TIC-101 at the design point without a step test and replies as follows. "The dead time of this loop is the transport delay from the bundle to TT-102, $\theta = V_{\text{trans}}\rho/\dot{m}_p = 0.1 \times 1000/333.3 = 0.30$ min. This is exact physics, so it is better than a dead time read from a trend. With $K_p = 0.355\,^{\circ}\text{C}/\%$ and $\tau = 3.0$ min the reaction-curve rule gives $K_c = 1.2 \times 3.0/(0.355 \times 0.30) = 33.8\,\%/^{\circ}\text{C}$, $\tau_I = 0.60$ min and $\tau_D = 0.15$ min, and the loop will reject loads about three times better than with a dead time of 0.94 min." Find the flaw in the proposal and say how it could be checked.