INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 3

Fouling Crude Preheat Train Simulator for Override Control on Throughput

Pages 50–80

Edited by John D. Hedengren

How to cite this chapter. "Fouling Crude Preheat Train Simulator for Override Control on Throughput," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 50–80.

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

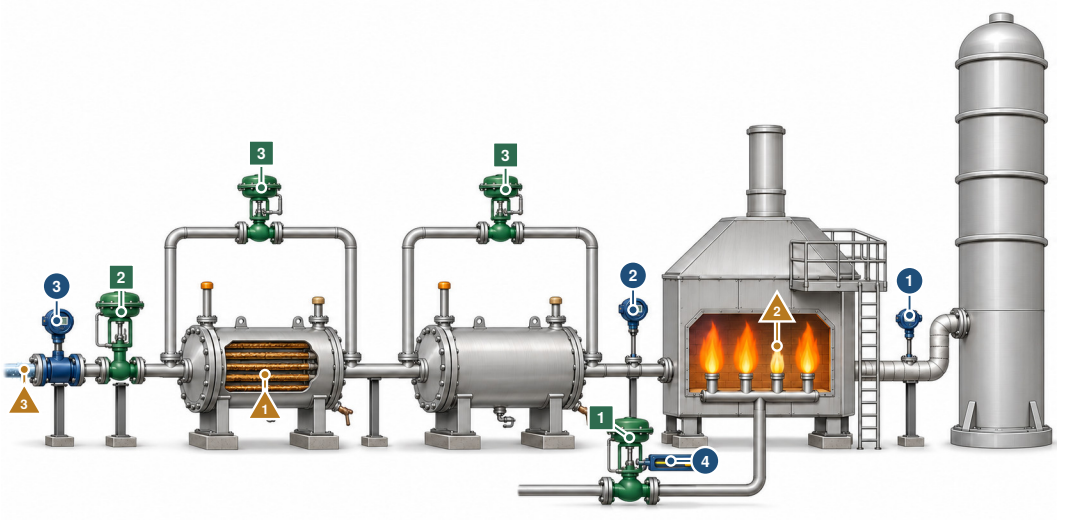
The app for this chapter
https://apmonitor.com/apps/crude-preheat/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v11-refining/#crude-preheat

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

CHAPTER 3

Fouling Crude Preheat Train Simulator for Override Control on Throughput



The controller moves 1 fuel gas valve of the furnace 2 crude charge valve 3 cleaning bypass of each exchanger

The controller reads 1 column inlet temperature 2 furnace inlet temperature 3 crude charge flow 4 fuel valve position

What disturbs it 1 fouling inside the tubes 2 leaner fuel gas 3 colder crude from the tanks

Once fouling has driven the fuel gas valve under the furnace wide open, only the charge valve at the far left can still hold the temperature read at the column inlet on the right. Illustration, not to scale.

IN PLAIN TERMS Slowing the Plant to Hold a Temperature

The main idea. Crude oil must be hot before a refinery can split it into fuels. Most of the heat comes free from hot products in a row of heat exchangers, and a gas furnace adds the rest.

The challenge. Dirt builds up inside the exchangers, so the furnace burns more gas until its fuel valve is fully open. Then the only way to keep the oil hot is to send less oil through, like a driver who slows on a steep hill.

Why it matters. Every tonne of oil held back is product the refinery does not make, so the cut must be kept small.

THREE WAYS TO READ THIS CHAPTER

Explore	9 min	This page, the Did you know? boxes in Sections 3.1 and 3.5.3, the History boxes in Sections 3.2 and 3.4, the Try it box in Section 3.5.2, and Key points in Section 3.9. No equations.
Study	65 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Sections 3.5.1 to 3.5.3 with Figure 3.3, Section 3.5.6, Section 3.6, Section 3.7 and the Open challenge box in Section 3.8.

The control question Fouling and cleaning outages take the fuel valve to its stop. What does holding the column inlet then cost in throughput, and how should an override that cuts the charge be set?

Learning objectives

- Explain why a loop held on setpoint through weeks of fouling hides the lost margin.
- Compute the fouled coefficient and the time to the cleaning threshold at two charge rates.
- Compute the heat-balance firing and the duty a cleaning outage removes.
- Choose a valve position target from both temperature and crude processed.
- Explain why overlapping cleanings cost far more than separate ones.

Prerequisites Sections 1.4.2, 1.4.5, 1.6.2, 1.6.3 and 1.6.4.

The app https://apmonitor.com/apps/crude-preheat/

See also Volume 4, Chapter 9 for valve position control, Volume 4, Chapter 4 for throughput as the override's handle, and Volume 3, Chapter 4 for a heater's feedforward and limits.

3.1 Introduction

The usual evidence that a control loop is working is a measurement sitting on its setpoint. On the unit of this chapter that evidence is incomplete. Three weeks of exchanger fouling leave the column inlet temperature unchanged, hold the integrated absolute error at 0.0 °C h over 504 h, and raise no alarm, while the fuel gas valve opens from 81.3 % to 97.5 % and the share of the crude duty recovered from product streams falls from 60.1 % to 52.1 %. Three weeks of fuel gas costs 76.7 TJ, about $460k at the $6/GJ carried in the parameter list, and roughly a tenth of that is the fouling. A performance report limited to control error would rate the three weeks as flawless. The only variable that records what happened is the position of the valve, which is why valve position is a monitored quantity on every constrained unit in a refinery.

> **DID YOU KNOW? The Largest Energy User in the Refinery**
>
> A 2005 guide for refinery energy managers calls crude distillation the largest energy consuming process of a refinery, and its preheater especially susceptible to fouling. It quotes a federal estimate that the cost penalty for fouling could be as much as $2 billion a year in material and energy [1].

The arrangement ends when the valve reaches its limit. When one exchanger is taken out of service for a sixteen hour cleaning at full charge rate, the fired heater saturates, sits on its 40.6 MW limit for 96.5 % of the outage, and delivers 340.3 °C to a column with a 360 °C setpoint. No feedback tuning recovers duty the furnace cannot supply. One manipulated variable remains, the crude charge rate. A lower charge rate requires less furnace duty and gains more temperature rise per tonne across the same 1620 m^2 of surface, so a small rate reduction produces a large temperature recovery. Reducing the charge rate also reduces refinery throughput directly, and that trade is the subject of this chapter. A valve position controller allowed to write a charge rate holds 360.0 °C through the identical outage at a cost of 555 t of crude out of 19 200 t.

Two general classes of problem sit inside this one plant. The first is override control, in which a second controller on a second manipulated variable is selected only when the primary actuator runs out of authority. Valve position control is its most common industrial form, and the same structure holds the most-open valve in a parallel network near full open, keeps a pump on its curve, and protects a compressor from surge [2, 3]. Volume 4, Chapter 2 treats the surge case, and Volume 4, Chapter 4 uses throughput as the override's handle, as here, with the ground speed of a harvester cut against engine and rotor limits. Constraint switching of exactly this kind has been given a systematic design procedure in terms of the classical structures that implement it [4], and the multivariable version of the same decision, in which throughput is pushed against whichever unit is currently the

bottleneck, is the standard justification for a plantwide coordinating controller [5]. The second class is the slowly moving process gain with a maintenance outage attached. Reverse osmosis membranes, boiler economizers, air cooler finning, pasteurizers and deactivating catalyst all degrade over weeks and are all managed by a schedule rather than by a controller, and the engineering question in every case is when to stop and clean.

However, a first-order transfer function identified from a step test represents neither class. Fouling does not appear in a transfer function because the loop absorbs it, and a firing limit does not appear because a linear model has no saturation. A trend sitting on setpoint establishes only that one loop still has actuator authority in reserve. This chapter works with a crude preheat train and its fired heater, described in enough detail to be rebuilt. It quantifies the throughput cost of the override when it is selected, measures what the valve position target buys on either side of the point where the constraint binds, and shows why two cleaning cycles of unequal period produce their largest excursion when they coincide. The analysis rests on two measured outcomes. Holding temperature through a cleaning outage costs 555 t of crude for the upstream exchanger and about three times that for the downstream one. Moving one of two colliding cleanings 32 h later divides the integrated error by almost twelve while changing the crude processed by 17 t in 199 000. The plant has twelve states on two time scales separated by a factor of five thousand, four control architectures and fourteen scenarios, each with its own completion requirement. The performance report keeps temperature error and throughput separate so that a trade between them is visible. The fouling law, the exchanger model and the override arithmetic are all established practice.

The control ideas are override and valve position control, feedforward design from a steady-state energy balance, actuator saturation and anti-windup, and the time-scale separation between a regulatory loop and a slowly varying process gain, all of which are standard material in process dynamics and control texts [6, 7]. Section 1.4.5 introduces the same three ideas on the running example R-1: the preheat train, the fouling that the outlet loop hides and the valve position override of Equation (1.10). The model itself rests on energy conservation over lumped thermal capacities, a log-mean temperature difference driving force, and a fouling resistance in series with the clean film resistances, and the closed-loop behavior follows from the feedback fundamentals collected in [8]. The behavior matters industrially because a crude unit operates at its most profitable point only when throughput is pushed against the currently active constraint [5], and the engineer who cannot distinguish a capacity limit from a tuning problem retunes a loop that has nothing left to give.

Section 3.2 reviews how a train of recovery exchangers and a cleaning schedule became the front end of every crude unit. Section 3.3 gives the twelve-state model, with fouling as a series resistance, the exchanger and furnace balances, the cleaning event, the parameters and the assumptions. Section 3.4 sets out the temperature loop, the heat-balance feedforward, the valve position override and the predictive option. Section 3.5 measures three weeks of fouling (Section 3.5.1), the cleaning outages (Section 3.5.2), two colliding cleaning cycles (Section 3.5.3), three disturbances that resemble fouling (Section 3.5.4), the graded challenge (Section 3.5.5) and the predictive controller against the override (Section 3.5.6). Section 3.6 sets the loops beside the hardware of a real unit, and Sections 3.7 to 3.9 state the limitations, discuss the scheduling argument and conclude.

3.2 Background

The preheat train exists because fired heater duty is expensive. Crude arriving at tank temperature has to reach the flash zone of the atmospheric column, and every degree supplied in a fired heater is fuel gas burned at a thermal efficiency below unity, while every degree recovered from a hot product stream leaving the same column carries no fuel cost. Refinery energy accounting puts crude distillation among the largest single consumers on the site, and the fired heaters that serve it among the largest emitters [9]. The engineering answer was to place as much surface as the economics allow between tankage and the heater, and the question of how much surface, where, and against which hot stream became a design discipline when heat exchanger network synthesis was posed as a problem

with a thermodynamic target rather than as a matter of judgment [10]. A modern crude unit therefore carries twenty or more shells in series and parallel, with a recovery fraction that falls below the design value in operation.

The shortfall is caused by fouling, and the History box below tells how fouling became a rate equation, a deposition term less a removal term, with an asymptote where the two balance. That model, its limitations and the fifty years of work that followed it are reviewed by Müller-Steinhagen [11], and the mechanisms, velocity dependence and cleaning practice are collected by Bott [12]. The application uses that law directly, because it is the smallest description that reproduces the two behaviors an operator recognizes, a fouling rate that is fastest at the start of a run and a resistance that eventually stops growing.

> ⧖ HISTORY **1959 and 1995: Fouling Becomes a Rate Equation**
>
> Fouling became a quantity that could be modeled, where before it was a design margin, when Kern and Seaton wrote its net rate in 1959 as deposition minus removal. Removal in their model is proportional to the thickness of the deposit, so the resistance approaches an asymptote [13]. Crude oil then acquired a feature of its own. In 1995 Ebert and Panchal reanalyzed tests on crude drawn from a refinery preheat exchanger. Fouling was significant at film temperatures of 370 to 400 °C with velocities of 1.2 and 2.5 m/s, and none was detected at similar wall temperatures with 3.8 and 5.2 m/s. Their rate law sets formation, which rises with film temperature, against removal by wall shear. Where removal exceeds formation, negligible fouling is expected, and the line between is the fouling threshold [13]. The model of this chapter keeps the first idea and leaves out the second.

For crude oil the box also describes a threshold surface condition below which a given exchanger does not foul at all, and the useful design and operating question is which shells in the train sit above it. The threshold picture was reduced to a correlation fitted against laboratory and plant data [14] and then used to decide retrofits, since moving an exchanger to a different place in the train changes both its wall temperature and its velocity [15]. Distributed dynamic models that carry the deposit layer along the tube followed, and they show how far a lumped resistance is from the real spatial picture [16]. The economic weight of the problem is large. Fouling in process heat exchangers is a recurring energy penalty with combustion emissions attached to it, large enough to be studied as an environmental problem of its own [17], and crude preheat carries the largest share of it because it carries the largest surface [18].

Because the resistance grows monotonically, the response is a schedule rather than a controller. Optimal cleaning policies for a fouling exchanger network were posed as a mixed integer problem with the cleaning times as decisions [19], and applied to refinery preheat trains where the objective is the fuel saved against the cost of the outage [20]. A later addition to that literature is the term this application is built around. A fouled train does not merely burn more fuel. It eventually takes the fired heater to its limit, at which point the unit loses throughput continuously without an operating decision, and cleaning schedules computed with throughput in the objective are not the same schedules as those computed on fuel alone [18].

The control side of the same problem developed separately and much earlier, as the History box in Section 3.4 relates. A temperature controller that has saturated its actuator has no authority left, and the practice that answers it is to give the constraint a controller of its own, acting on a second manipulated variable, selected by a limit or a signal selector so that the loop reverts to normal service the moment the constraint releases. Shinskey set out override and valve position control in this form, including the rule that the constraint controller is tuned slower than the loop it protects [2], and the same structures with feedforward designed from a steady-state balance rather than from an identified model are standard textbook material [3]. Two developments since have changed what surrounds them rather than what they are. Constraint switching has been given a systematic design procedure, so the choice of which controller takes which variable in which region is no longer left to heuristics [4]. Model predictive control absorbed the same decision into an optimization

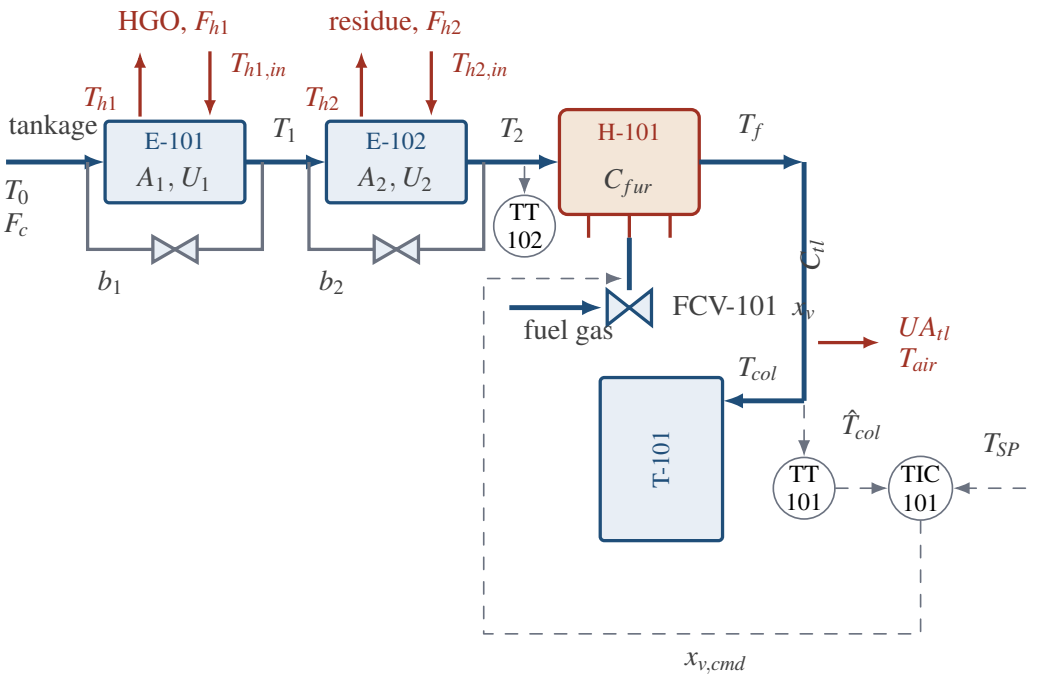


Figure 3.1 Crude preheat train E-101 and E-102, fired heater H-101 and fractionator T-101. Crude enters at tankage temperature T_0 and charge rate F_c, gains heat from heavy gas oil in E-101 and from atmospheric residue in E-102, and leaves the furnace coil at T_f into a transfer line of capacitance C_{tl} that loses $UA_{tl}(T_{col} - T_{air})$ to ambient. TT-101 supplies the measured column inlet to TIC-101, which drives fuel gas valve FCV-101 to position x_v, and TT-102 supplies the furnace inlet T_2 to the feedforward heat balance. Bypass fractions b_1 and b_2 take an exchanger out of service for cleaning, scaling its crude flow, its hot flow and its area in service together. The valve position controller that writes F_c appears in Figure 3.2.

with the constraints written in, which is how a modern crude unit handles it, and the survey of industrial practice records constraint handling rather than trajectory tracking as the primary reason for adoption [21]. The regulatory structure remains in service underneath, because a valve position controller requires one function block and no model.

3.3 Process Model

Crude leaves tankage at 25 °C and a design rate of 400 t/h, passes through two recovery exchangers in series, is fired the rest of the way in heater H-101, and arrives at the flash zone of the atmospheric column at 360 °C. That last temperature is the controlled variable, and the overflash of the column and the yield of every product drawn off it depend on it, so the deviation alarm band is a narrow 3 °C. E-101 is 620 m^2 of crude against heavy gas oil supplied at 285 °C. E-102 is 1000 m^2 of crude against atmospheric residue supplied at 345 °C, and its outlet is the furnace inlet. Each exchanger carries a bypass that takes it out of service for cleaning. The arrangement, the instrument tags and the symbols used below are all in Figure 3.1.

At the design condition with the fouling the application opens on, E-101 recovers 22.87 MW and lifts the crude to 118.6 °C, E-102 recovers 26.79 MW and takes it to 228.1 °C, and H-101 supplies the remaining 32.98 MW with its fuel valve at 81.3 %. The crude absorbs 82.63 MW in all, so 60.1 % of the duty arrives free. The important property of the fuel valve is that it stops: at 100 % the furnace absorbs 40.56 MW and there is no more.

Table 3.1 Controlled, Manipulated and Disturbance Variables With the Twelve Integrated States. Nominal Values Are the Design Condition Every Scenario Opens On.

Symbol	Role	Nominal	Range
T_{col} (TT-101)	CV, measured	360.0 °C	330–375 °C SP
x_v (FCV-101)	MV	81.3 %	0–100 %
$F_{c,cmd}$	MV, override or MPC	400 t/h	240–F_{dem}
F_{dem}	DV, measured	400 t/h	200–440 t/h
T_0 tankage	DV, measured	25 °C	0–45 °C
HV heating value	DV, unmeasured	100 %	80–115 %
$T_{h2,in}$ residue	DV, measured	345 °C	300–360 °C
F_{h2} residue flow	DV, measured	230 t/h	140–300 t/h
b transmitter bias	DV, unmeasured	0 °C	±10 °C
T_{c1}	state	118.6 °C	—
T_{h1}	state	155.8 °C	—
T_{c2}	state	228.1 °C	—
T_{h2}	state	183.8 °C	—
T_f	state	363.0 °C	$\leq$ 400 °C
T_{col}	state	360.0 °C	—
$R_{f,1}$	state	0.320	≥ 0
$R_{f,2}$	state	0.050	≥ 0
F_c	state	400 t/h	—
x_v	state	81.3 %	0–100 %
b_1, b_2	state	0	0–1

Table 3.1 lists the controlled, manipulated and disturbance variables together with the integrated states. Three annunciators bound normal operation, a column inlet deviation outside 3 °C, the fuel valve at or above 99.5 %, and a coil outlet above the 400 °C coking limit. All three annunciate without stopping the run, so a severe excursion can be observed to its end.

3.3.1 Working Units

Time is in hours, because fouling develops over three weeks while the temperature loop settles in minutes. The rest of the unit set follows refinery practice: mass in tonnes, flow in t/h, energy in megajoules, duty in MJ/h displayed as megawatts, heat capacity in MJ/t/K, UA in MJ/h/K, area in m^2, U in MJ/h/m^2/K and R_f in hm^{2}K/MJ. The set is self-consistent, so $Q = F c_p \Delta T$ carries no conversion factor anywhere. Two conversions relate the working units to familiar values. One megawatt is 3600 MJ/h, so the 40.56 MW furnace is 146 000 MJ/h, and 1 W/m^2/K is 3.6×10^{-3} in these units, so the 350 W/m^2/K clean coefficient of E-101 is 1.26.

3.3.2 Fouling as a Series Resistance

A deposit layer is a conduction resistance in series with the two clean film resistances, so the coefficients add reciprocally, as shown in Equation (3.1).

$$\frac{1}{U_i} = \frac{1}{U_{clean,i}} + R_{f,i} \qquad \Longleftrightarrow \qquad U_i = \frac{U_{clean,i}}{1 + U_{clean,i} R_{f,i}} \tag{3.1}$$

where $U_{clean,i}$ is the clean overall coefficient of exchanger i and $R_{f,i}$ is its fouling resistance, floored at zero. Nothing else about the exchanger changes. The area, the two flows and the inlet driving force are unchanged, and the duty falls because UA falls. At its cleaning threshold of 0.42 hm^{2}K/MJ the E-101 coefficient is down from 1.26 to 0.824, a loss of 34.6 %, and at its own threshold of 0.60 hm^{2}K/MJ E-102 is down from 1.152 to 0.681, a loss of 40.9 %, both in MJ/h/m^2/K. Every result in this chapter is a consequence of (3.1) and of the fact that R_f grows monotonically between cleanings.

The resistance follows a deposition term minus a removal term, both dependent on crude velocity, as shown in Equation (3.2).

$$\frac{\mathrm{d}R_{f,i}}{\mathrm{d}t} = (1-b_i)\left[\frac{R_{f\infty,i}}{\tau_{f,i}}\,v^{-n_d} - \frac{v^{n_r}}{\tau_{f,i}}R_{f,i}\right], \qquad v = \frac{F_c}{F_{c,D}} \tag{3.2}$$

where $R_{f\infty,i}$ is the asymptotic resistance at design velocity, $\tau_{f,i}$ is the fouling time constant at design velocity, n_d and n_r are the deposition and removal velocity exponents, v is velocity as a fraction of design and b_i is the bypass fraction, which stops the clock while an exchanger is isolated. Both sides carry units of $\mathrm{m^2K/MJ}$. The steady solution of (3.2) is $R_{f\infty,i}\,v^{-(n_d+n_r)}$ reached with time constant $\tau_{f,i}/v^{n_r}$, so with $n_d = 0.4$ and $n_r = 1.0$ the asymptote scales as $v^{-1.4}$ and operation at reduced rate raises the asymptote while lengthening the approach to it. At 300 t/h the E-101 asymptote rises from 0.55 to $0.55 \times 0.75^{-1.4} = 0.823$ and the time constant stretches from 420 to 560 h, so the threshold is reached at hour 124 instead of hour 240. Deposition falls with velocity because a faster stream spends less time at the wall, while shear removal rises faster than the deposition falls.

3.3.3 Exchanger Energy Balances

Each exchanger is two lumped thermal capacities with a log-mean driving force between them, as shown in Equations (3.3) to (3.5).

$$C_{ic}\frac{\mathrm{d}T_{ci}}{\mathrm{d}t} = (1-b_i)F_c c_{p,c}\,(T_{c,in,i} - T_{ci}) + Q_i \tag{3.3}$$

$$C_{ih}\frac{\mathrm{d}T_{hi}}{\mathrm{d}t} = (1-b_i)F_{hi} c_{p,hi}\,(T_{h,in,i} - T_{hi}) - Q_i \tag{3.4}$$

$$Q_i = (1-b_i)\,U_i A_i\,\mathrm{LMTD}\big(T_{h,in,i} - T_{ci},\; T_{hi} - T_{c,in,i}\big) \tag{3.5}$$

where C_{ic} and C_{ih} are the cold and hot side capacitances in MJ/K, T_{ci} and T_{hi} are the two outlet temperatures, $T_{c,in,i}$ and $T_{h,in,i}$ are the two inlets, and LMTD is the log mean of the two end temperature differences. Where either end is non-positive, which a transient produces while a bypassed exchanger is coming back on line, the arithmetic mean takes over so that the right-hand side stays finite. The crude inlet of E-101 is the tankage temperature and the crude inlet of E-102 is the E-101 station outlet. At steady state (3.3) to (3.5) reproduce the counterflow design equation $Q = UA \cdot \mathrm{LMTD} = \varepsilon C_{min}\Delta T_{max}$ exactly, which one of the application's own verification tests checks to one part in ten thousand.

The bypass fraction b_i appears three times, scaling the crude flow through the tubes, the hot flow through the shell and the area in service by the same factor. Scaling all three together is what a bypass around a fixed exchanger does, and it has two useful consequences. The exchanger's own steady outlet temperature does not depend on b_i, while the station duty falls linearly with it, as shown in Equation (3.6).

$$\begin{aligned} T_{station,i} &= (1-b_i)T_{ci} + b_i T_{c,in,i} \\ Q_{s,i} &= F_c c_{p,c}\,(T_{station,i} - T_{c,in,i}) \end{aligned} \tag{3.6}$$

At $b_i = 1$ every derivative inside that exchanger is zero, so an isolated exchanger holds still instead of drifting toward an unphysical stagnation temperature, and its fouling clock stops by the same factor in (3.2).

3.3.4 Fired Heater and Transfer Line

The furnace adds a duty proportional to valve position and to fuel gas heating value, and the transfer line carries the crude to the column while losing heat to ambient, as shown in Equations (3.7) and (3.8).

$$C_{fur}\frac{\mathrm{d}T_f}{\mathrm{d}t} = F_c c_{p,c}(T_2 - T_f) + \frac{x_v}{100}\,Q_{fur,max}\,\frac{HV}{HV_0} \tag{3.7}$$

$$C_{tl}\frac{\mathrm{d}T_{col}}{\mathrm{d}t} = F_c c_{p,c}(T_f - T_{col}) - UA_{tl}(T_{col} - T_{air}) \tag{3.8}$$

where T_2 is the furnace inlet, $Q_{fur,max}$ is the absorbed duty at full firing, HV/HV_0 is the heating value as a fraction of design, C_{fur} is the coil and firebox capacitance and C_{tl} is the transfer line capacitance. Equation (3.7) is where the constraint of the whole application lives, because x_v is clamped to 100 % and $Q_{fur,max} = 146\,000\,\mathrm{MJ/h}$ is a constant. The loss term in (3.8) is about 0.75 MW at the design condition, which is why the coil has to leave about 3 °C hotter than the column wants and why UA_{tl} appears explicitly in the feedforward of Section 3.4. The heating value ratio multiplies the duty and appears nowhere in any controller, which makes it the disturbance the feedforward is structurally blind to.

The four actuators are first-order lags on commands the controller has already clamped, as shown in Equation (3.9).

$$\begin{aligned} \tau_F \frac{\mathrm{d}F_c}{\mathrm{d}t} &= F_{c,cmd} - F_c, \quad \tau_{fire} \frac{\mathrm{d}x_v}{\mathrm{d}t} = x_{v,cmd} - x_v \\ \tau_{byp} \frac{\mathrm{d}b_i}{\mathrm{d}t} &= b_{i,cmd} - b_i \end{aligned} \tag{3.9}$$

No actuator can therefore walk past its stop and none needs a blocking test at the limit. A bypass commanded shut is snapped onto its command once the remaining travel falls below 10^{-5}, because a first-order response otherwise approaches the command only asymptotically.

3.3.5 Cleaning as a Discrete Event

A cleaning request opens the bypass of one exchanger, holds it open for $t_{clean} = 16\,\mathrm{h}$, then sets that exchanger's resistance to zero in one instant and commands the bypass shut. The sixteen hours represent the outage rather than the cleaning alone, covering isolation, steam out, hydroblasting, box up and line out, and the outage duration is what determines the throughput cost. Nothing else about the exchanger is changed, so a cleaning is exactly a reset of one state in (3.2), and the fouling that follows starts again from zero with a run length set by the charge rate.

3.3.6 Parameters, Initialization and the Instrument Chain

Table 3.2 lists the constants. Every parameter value is exposed and adjustable in the application, so every result below can be reproduced or perturbed.

The initial state is solved in closed form rather than by simulation to steady state, because with both bypasses shut the train is a forward chain: the counterflow effectiveness relation fixes E-101 from the crude and gas oil inlets, then E-102 from the E-101 outlet, and then the firing that puts the column inlet exactly on setpoint, as shown in Equation (3.10).

$$\begin{aligned} T_f^* &= T_{SP} + \frac{UA_{tl}(T_{SP} - T_{air})}{F_c c_{p,c}} \\ x_v(0) &= \frac{100\,F_c c_{p,c}(T_f^* - T_2)}{Q_{fur,max}} \end{aligned} \tag{3.10}$$

At the design condition (3.10) returns $T_f^* = 363.05\,^\circ\mathrm{C}$ and $x_v(0) = 81.32\,\%$, and manual operation provides the same heat balance calculation on request. Where the firing that (3.10) asks for exceeds 100 %, the initializer takes the steady state the clamped valve actually produces rather than the one that was requested.

Only two measurements reach the controller. The column inlet transmitter adds a constant bias and a noise sequence, built by filtering Gaussian draws and rescaling the result to a specified standard deviation, then passes the sum through a thermowell lag of $\tau_{TT} = 0.01\,\mathrm{h}$ (36 s). The furnace inlet transmitter carries no bias, a noise sequence of 60 % of the same amplitude and a faster lag of $\tau_{FT} = 0.004\,\mathrm{h}$, faster because the value of the feedforward depends on the quality of that measurement. Each filter is advanced by $\hat{T}_k = \alpha \hat{T}_{k-1} + (1-\alpha)T_k$ with $\alpha = \exp(-\Delta t_c/\tau)$, the exact solution over one controller period, so no choice of period can destabilize it. There is no lag and no noise on the charge flow indication, the duties or the fouling resistances, which are reported exactly.

Table 3.2 Model Parameters at the Design Condition

Symbol	Meaning	Value	Unit
$F_{c,D}$	design charge rate	400	t/h
$c_{p,c}$	crude heat capacity	2.2	MJ/t/K
$F_{c,min}$	minimum stable charge rate	240	t/h
A_1	E-101 area	620	m^2
$U_{clean,1}$	E-101 clean coefficient	1.26	MJ/h/m^2/K
F_{h1}	heavy gas oil flow	260	t/h
$c_{p,h1}$	heavy gas oil heat capacity	2.45	MJ/t/K
$T_{h1,in}$	heavy gas oil supply	285	°C
C_{1c}, C_{1h}	E-101 capacitances	90, 70	MJ/K
A_2	E-102 area	1000	m^2
$U_{clean,2}$	E-102 clean coefficient	1.152	MJ/h/m^2/K
$c_{p,h2}$	residue heat capacity	2.60	MJ/t/K
C_{2c}, C_{2h}	E-102 capacitances	250, 170	MJ/K
$R_{f\infty,1}$	E-101 asymptote	0.55	hm^2K/MJ
$R_{f\infty,2}$	E-102 asymptote	0.85	hm^2K/MJ
$\tau_{f,1}$, $\tau_{f,2}$	fouling time constants	420, 340	h
$R_{f,clean,1}$	E-101 cleaning threshold	0.42	hm^2K/MJ
$R_{f,clean,2}$	E-102 cleaning threshold	0.60	hm^2K/MJ
n_d, n_r	velocity exponents	0.4, 1.0	—
t_{clean}	cleaning outage	16	h
$Q_{fur,max}$	absorbed duty at full firing	146000	MJ/h
η	furnace thermal efficiency	0.88	—
C_{fur}	coil and firebox capacitance	100	MJ/K
C_{tl}	transfer line capacitance	85	MJ/K
UA_{tl}	transfer line heat loss	8	MJ/h/K
T_{air}	site ambient	25	°C
$T_{f,max}$	coil outlet coking limit	400	°C
—	fuel gas price	6.0	$/GJ
τ_{fire}	firing response lag	0.05	h
τ_F	charge rate response lag	0.05	h
τ_{byp}	bypass swing time	0.20	h
τ_{TT}	column inlet thermowell lag	0.01	h
τ_{FT}	furnace inlet thermowell lag	0.004	h
T_{band}	column inlet alarm band	3	°C

3.3.7 Time Scales and Numerical Integration

The simulation holds a temperature loop that settles in minutes and a fouling cycle that develops over three weeks inside one integration step and one session, so the step size is chosen against the fast end while the session length is chosen against the slow one. The fastest thermal mode in the plant is the transfer line, whose time constant is given by Equation (3.11).

$$\tau_{tl} = \frac{C_{tl}}{F_c c_{p,c} + UA_{tl}} = \frac{85}{880 + 8} = 0.0957\,\text{h} \tag{3.11}$$

which is 5.74 min. The step, the controller period and the log interval below are the three clocks of Section 1.6.1. The fixed integration step is $\Delta t = 0.01$ h (36 s), so (3.11) is resolved 9.6 times over, which is comfortable for classical fourth-order Runge-Kutta on a smooth right-hand side. The controller period is $\Delta t_c = 0.02$ h (72 s), a realistic scan for a furnace outlet loop, and the loop it closes is dominated by the coil at $C_{fur}/(F_c c_{p,c}) = 100/880 = 0.114$ h, which gives 5.7 controller executions per time constant. Neither ratio is generous. Halving either would double the cost of the three-week run for no visible change in the trajectory.

The slow end fixes the session. Three weeks is 504 h, or 50 400 integration steps, each one four evaluations of a twelve-state right-hand side. Measured on this model at 2.4 µs per step, a full three-week run takes 122 ms of processor time, so at the top acceleration of 100× the whole three weeks replays in about five seconds. That replay time is what keeps the fouling scenarios observable within one session, because the lesson of those scenarios is a trajectory shape that develops over the full three weeks. The recorded history holds one sample every third integration step, 0.03 h (1.8 min), or 16 801 rows for a three-week run.

3.3.8 Assumptions

Model detail is concentrated on the constraint. Each simplification below is stated together with the conditions under which it fails.

Two lumped capacities per exchanger with a log-mean driving force. The formulation reproduces the counterflow design equation at steady state and gives the right dominant time constant, but it carries no spatial temperature profile, so a temperature cross, shell-side bypassing and the difference between one shell and two in series cannot appear. A distributed model that tracks the deposit along the tube shows how much of the real behavior that costs [16].

Fouling is one uniform resistance per exchanger driven only by velocity. Real crude fouling depends strongly on wall temperature and on the asphaltene stability of the particular blend, is worse at the hot end of the train than the cold, and is not uniform over the surface. The threshold behavior that decides where in a real train deposit forms at all is absent, and that behavior is the whole basis of fouling-aware design [14, 15]. The time scale is also compressed for teaching, because a real crude preheat exchanger runs six months to two years between cleanings, so no interval reported here predicts a real cleaning schedule.

Cleaning is perfect and instantaneous at the end of a fixed duration. Mechanical cleaning typically recovers 85 to 95 % of the clean coefficient, and the residue left behind makes each successive run shorter.

The hot streams are boundary conditions. In the real unit the heavy gas oil and the atmospheric residue come off the same column the crude is being fed to, so a change in preheat changes the column, which changes the hot streams, which changes the preheat. That loop is cut here, and it is the largest structural omission in the model. Chapter 4 ties its preheat to the pumparound duties that recover the tower's heat.

Furnace absorbed duty is proportional to valve position with a fixed efficiency. There is no draft control, no excess oxygen, no pass flow balancing and no radiant and convective split, so nothing can be said about furnace efficiency optimization or pass balancing, which is where much of the real work on a fired heater goes. Volume 3, Chapter 4 treats the air side of such a heater, with oxygen trim and draft control.

Nothing vaporizes and crude properties are constant. A real crude heater has 30 to 50 % vaporization at the coil outlet, so the coil is two-phase over its last third. Between the tank farm and the flash zone the specific heat rises about 15 % and the viscosity falls by two orders of magnitude, which moves both the duty and the film coefficients by more than the fouling does over a short run.

The column is a temperature target and nothing else. There is no flash calculation, no overflash, no product draw and no tray hydraulics, and the minimum stable charge rate of 240 t/h stands in for the whole of the column's own constraint set.

Perfect repeatability is a teaching convenience. A fixed step and a fixed scan give one trajectory everywhere, which no real control system does.

3.4 Controller Design

Three control architectures of increasing capability are provided, as shown in Figure 3.2, with a model predictive controller as a fourth option beside them (Section 3.4.4). Feedback alone is the first, feedback plus a furnace heat balance is the second, and the third adds the valve position controller

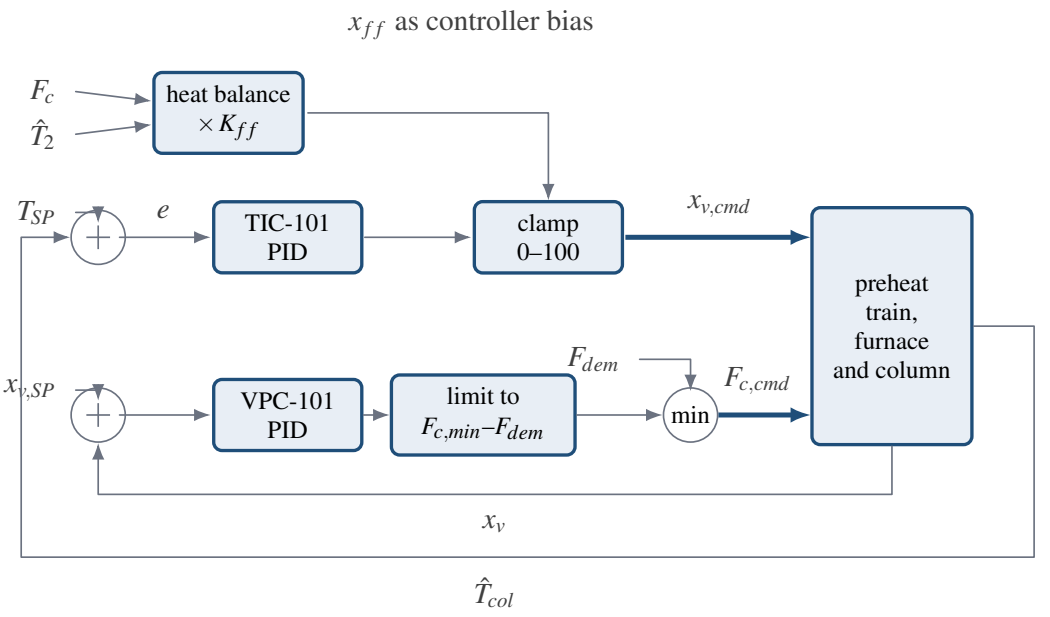


Figure 3.2 The three control architectures. TIC-101 alone is the *Single* architecture. Adding the furnace heat balance, scaled by K_{ff} and entered as the controller bias, gives + *Feedforward*. Adding VPC-101, whose process variable is the fuel valve position itself and whose output is a charge rate low-selected against the operator demand, gives *Override*. The upper output limit of VPC-101 is F_{dem}, so an override that is not selected rests on that clamp and cannot wind up.

that is the subject of this chapter. Switching between them is a bumpless transfer rather than a step, because moving from the first to the second puts the whole heat balance into the controller bias and would otherwise arrive as a demanded step of about eighty points on the fuel valve, the 81.3 % the balance asks for at the design condition. For one execution after any architecture change both controllers are made to track the actuator positions the plant is already holding, which is the same mechanism the manual to automatic transfer of Section 1.6.2 uses, and the switch leaves no mark on the trend.

3.4.1 Temperature Controller TIC-101

The column inlet controller is the velocity-form algorithm of Section 1.6.2, written in the loop's own symbols in Equations (3.12) to (3.14). The derivative acts on the measurement, so a setpoint move on

its own puts no kick into the fuel valve.

$$e_k = T_{SP,k} - \hat{T}_{col,k} \tag{3.12}$$

$$\Delta u_k^P = K_c\,(e_k - e_{k-1})$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c$$

$$\Delta u_k^D = -\,K_c\,\tau_D\,\frac{\hat{T}^{\,f}_{col,k} - 2\hat{T}^{\,f}_{col,k-1} + \hat{T}^{\,f}_{col,k-2}}{\Delta t_c} \tag{3.13}$$

$$s_k = s_{k-1} + \Delta u_k^P + \Delta u_k^D$$

$$c_k = x_{ff,k} + s_k + I_{k-1}$$

$$u_k = \mathrm{sat}(c_k, 0, 100)$$

$$I_k = I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \tag{3.14}$$

Here s_k is the running sum S_k of Equation (1.17), c_k is the valve position the algorithm demands, the v_k of Equation (1.18) with the feedforward x_{ff} as its bias, and the integral used at scan k is the one accumulated through scan $k-1$. The action is $a = +1$ in Equation (1.16), reverse acting in the ISA convention, because more fuel valve raises the column inlet temperature. The derivative increment is a second difference on a measurement first passed through a filter of time constant τ_f defaulting to 0.02 h, and τ_D defaults to zero because derivative action against thermowell measurement noise is rarely justified on a loop of this speed. Reset is limited by back-calculation, with a tracking time T_t that the application passes explicitly as $T_t = \tau_I$, so at the default tuning each scan removes $\Delta t_c/\tau_I = 5.7\,\%$ of the clipped excess. With $T_t = \tau_I$ the integral comes to rest where $x_{ff} + I$ equals the limit, so the proportional term alone carries the demand past it and the valve leaves its stop as soon as the error changes sign. An integral time at its upper bound of 4 h is treated as infinite, and the integral update, back-calculation included, is then skipped. Output limits are 0 and 100 % with no rate limit on top, since the firing lag τ_{fire} of (3.9) is already the rate limit. In manual the untouched valve station is initialized to the solved starting valve position (81.3 % at design conditions), and selecting Manual before Run holds that lineup. A valve position the learner sets survives initialization and Reset. The integral state of (3.14) is seeded so that u_k equals the manual output, which is why no dedicated transfer logic exists. The adjustable ranges are $0 \le K_c \le 6\,\%/^\circ\mathrm{C}$, $0.05\,\mathrm{h} \le \tau_I \le 4\,\mathrm{h}$ and $0 \le \tau_D \le 0.5\,\mathrm{h}$, and the default tuning is $K_c = 1.5\,\%/^\circ\mathrm{C}$, $\tau_I = 0.35\,\mathrm{h}$, $\tau_D = 0$.

The gain limit follows from the process gain and the total lag. The process gain from valve to column inlet follows from (3.7) and (3.8) at steady state as $(Q_{fur,max}/100)/(F_c c_{p,c} + UA_{tl}) = 1460/888 = 1.64\,^\circ\mathrm{C}$ per percent of valve, and the lags between the valve and the thermowell total 0.269 h once the firing response (0.05 h), the coil (0.114 h), the transfer line (0.096 h) and the thermowell (0.01 h) are added. A gain of 6 therefore puts a loop gain near ten around that chain, which is why the measured tuning sweep in Section 3.5 degrades above 1.5 rather than continuing to improve.

3.4.2 Heat Balance Feedforward

The second architecture computes the firing the furnace needs directly from a heat balance on H-101 and the transfer line, as shown in Equation (3.15), the balance of Equation (1.2) with the transfer line loss added.

$$x_{ff} = \frac{100\,K_{ff}}{Q_{fur,max}}\,F_c c_{p,c}\left(T_{SP} + \frac{UA_{tl}(T_{SP} - T_{air})}{F_c c_{p,c}} - \hat{T}_2\right) \tag{3.15}$$

where $\hat{T}_2$ is the measured furnace inlet temperature and K_{ff} is a detuning gain with a default of 1.0. Both inputs are instruments a crude unit already has, the charge flow transmitter and the thermocouple at the train outlet, and there is no step test and no model fit behind (3.15) because the heat balance is the model. At $K_{ff} = 1$ it is algebraically the same expression as the steady-state firing requirement

of (3.10), so the feedback trim on a settled plant is zero to the resolution of the trend, and measured in scenario 3 just before its rate cut at 2 h the feedforward asks for 81.45 % of valve where the plant holds 81.45 %, a trim of 0.004 points. The result is clamped to 0–100 % and entered as the controller bias x_{ff} in (3.14), which is the whole of the implementation. Because the bias passes through the clamp, the back-calculation and the manual tracker, no term had to be added anywhere else. Equation (3.15) is written at the design heating value, because the heating value is unmeasured, and that omission is why the fuel gas scenario defeats the feedforward. Chapter 14 feeds the Wobbe index of the gas forward to its burners instead.

Pause. Try Exercise 3.2 before reading on.

3.4.3 Valve Position Override VPC-101

The third architecture adds a second PID whose process variable is the fuel valve position, the override of Equation (1.10), and whose output is a crude charge rate in t/h, combined with the operator demand as shown in Equation (3.16).

$$\begin{aligned} F_{c,cmd} &= \min\left(F_{dem},\ \mathrm{PID}_{VPC}(x_{v,SP}, x_v)\right) \\ u_{max}^{VPC} &= F_{dem} \end{aligned} \tag{3.16}$$

Three properties of (3.16) make an override inexpensive to implement. It is a high limit rather than a selector chain, since the override's own upper output limit is the operator's demand, so whenever the furnace has margin the controller rests on that clamp and the operator demand sets the charge rate. Its anti-windup is therefore free, because the saturated output that the back-calculation in (3.14) drives the integral toward is the demand itself, and an override that is not selected cannot wind up. The tracking time is again the integral time, so the integral comes to rest on F_{dem} and the proportional term holds the unclamped output $K_{c,VPC}(x_{v,SP} - x_v)$ above it, which means the override is selected only once the valve crosses its target. A recursion that absorbed the whole clipped excess in one scan would leave the integral at $F_{dem} - K_{c,VPC}(x_{v,SP} - x_v)$ and start cutting as soon as the valve began to rise, and in a counterfactual run of that form the E-102 outage of Section 3.5 dips to 326.9 °C instead of 317.5 °C. While VPC-101 is not selected, under another architecture, in manual, or on the execution that switches architecture (the first execution of every run among them), it tracks the demand, and with the valve below its target its integral is set to F_{dem}, the state in which back-calculation holds it. A target lowered before the valve reaches it therefore cuts nothing. It is deliberately slow, with a default integral time of 0.5 h against the 0.35 h of the temperature loop, following the rule that the constraint controller is tuned slower than the loop it protects [2]. Reset is retained all the same, because an override with proportional action alone would settle a fixed distance below its valve position target and surrender throughput the furnace could carry. The default gain is 4 t/h of charge cut per percent of valve above target, the valve position target is 95 %, and the lower output limit is the minimum stable charge rate of 240 t/h (or the demand, when that is lower), which is a fractionator constraint rather than a furnace one.

The order inside one control execution never varies. The tankage temperature and its diurnal term are evaluated first, then both transmitters with their bias, noise and filters, then the feedforward, then the tuning is re-applied, then TIC-101 or its manual tracker, then VPC-101 or its tracker, then the low select of (3.16). The two temperatures the controllers use therefore come through transmitters with their bias, noise and lag, while the charge flow in the feedforward and the valve position that VPC-101 controls are read exactly.

> **HISTORY 1945: Limit Control Enters the Textbooks**
>
> Process control grew up with the large continuous plants, refineries among them, that appeared around 1920. They were first run by hand, with one operator for each valve [22]. Proportional control came in the 1920s, automatic reset followed, and cascade and ratio control were in wide use in the 1930s. Selective

control, also called limit or override control, came into use to handle changes of constraint, and a 1945 book on industrial process control already described it beside ratio, cascade and split range control. By the 1960s the process industry worked with a standard set of such elements, valve position control among them [22]. Predictive control became common in refining at the end of the 1980s, and the review that gives these dates finds the older elements still thriving in practice beside it [22].

3.4.4 Model Predictive Control

The fourth architecture hands the fuel valve and the charge rate to one receding-horizon quadratic program [23] in automatic mode, and the limit VPC-101 exists for becomes a hard bound on x_v inside the optimization. Only scenario 13 opens on it, the three architectures above remain the application's default, and the comparison of Section 3.5.6 treats the predictive controller as an alternative to them rather than as a further rung. The option follows the general form of Section 1.6.3, with a filtered output bias in place of a Kalman filter and the plan's own objective below.

Controller model. The steady state is the feedforward balance of (3.15) solved for the column inlet on the measured furnace inlet, as shown in Equation (3.17).

$$T_{ss} = \frac{F_c c_{p,c} \hat{T}_2 + (x_v/100)\, Q_{fur,max} + UA_{tl} T_{air}}{F_c c_{p,c} + UA_{tl}} \tag{3.17}$$

Its four constants are read where the feedforward reads them, so a learner's change of furnace capacity reaches both structures or neither, and the heating value is the design value, as in (3.15). The slope K_v is 1.644 °C per percent of valve at 400 t/h and 2.14 at 306.6 t/h. The dynamics were identified from valve and charge steps at the design condition with noise off. Three equal lags of $\tau_c = 0.0921$ h from the valve command to the measured column inlet leave a residual of 0.54 % of the response, against 1.42 % for a first-order lag with dead time. At fixed $\hat{T}_2$ the balance carries only 50.7 % of the charge's steady effect, −0.327 of −0.644 °C per t/h, because a slower crude also leaves E-102 hotter. The remainder is the identified furnace inlet response $\delta T_2 = g_{T2}\, \delta F_c / (\tau_{T2} s + 1)$ with $g_{T2} = -0.3238$ °C per t/h and $\tau_{T2} = 0.3783$ h, with which a charge step on the column inlet is predicted within 1.06 % of the response. The duty part of (3.17) enters at the first lag, and the furnace inlet part enters 41 % at the first lag and the rest at the second, a split that keeps the model within 0.54 °C of the measurement while an E-102 bypass with the valve held swings $\hat{T}_2$ by 112 °C. An output bias, the measured column inlet less the model output filtered over 0.06 h, makes the prediction offset-free [24]. Fouling therefore reaches the plan through the furnace inlet thermocouple, while a lean fuel gas, a biased thermocouple and every model error reach it through the bias. The four-state model is relinearized at the applied commands at every decision.

Objective and constraints. Every $T_s = 0.06$ h, three controller scans, the plan covers $N = 33$ steps (1.98 h) in ten move blocks of 1, 1, 1, 2, 2, 3, 4, 5, 6 and 8 steps, 20 decision variables and one slack, applies the first move and holds it, minimizing Equation (3.18).

$$\begin{aligned} J = \sum_{k=1}^{N} \left[q_T \left(\hat{T}_{col,k} - T_{SP} \right)^2 - p F_{c,k} T_s \right] + \rho \left(\varepsilon^2 + \varepsilon \right) \\ + \sum_{j} \left[S_v \Delta x_{v,j}^2 + S_F \Delta F_{c,j}^2 \right] \end{aligned} \tag{3.18}$$

Here $\hat{T}_{col,k}$ is the predicted column inlet, $q_T = 10$ per °C^2 and step, $p = 1$ per tonne prices the charge as a linear cost, $S_v = 0.05$ per %^{2} and $S_F = 0.005$ per (t/h)^2 price the moves Δ at each block boundary, and $\rho = 10^3$. The valve is held hard in $0 \le x_v \le x_{v,max}$, with the ceiling $x_{v,max} = 100$ % as shipped, and the charge hard in $\min(F_{c,min}, F_{dem}) \le F_c \le F_{dem}$, the range VPC-101 clamps to. The one slack $\varepsilon \ge 0$ softens a floor $\hat{T}_{col,k} \ge T_{SP} - T_{band}$ over the whole horizon. The price keeps the charge at the demand while the valve has travel, and once the valve sits on its bound it holds the cut where the full furnace carries the setpoint, 306.612 t/h at hour 10 of scenario 13 against 306.616 t/h from the closed-form

Table 3.3 The Fourteen Scenarios, With the Mechanism Each Is Built Around

#	Scenario	What it exercises
1	Fire the furnace by hand	charge rate as the dominant disturbance
2	Temperature loop against a rate change	feedback distance, the gain limit
3	Feedforward from the heat balance	a model written rather than identified
4	Three weeks of fouling, nothing else	a perfect loop hiding a lost margin
5	Clean E-101 at full rate on feedforward	a saturated actuator ends the loop
6	The same clean under override	throughput as the second MV
7	Clean E-102, the expensive one	where the exchanger sits in the train
8	Two cleaning cycles out of phase	a schedule as a control decision
9	The same two cleanings colliding	two constraints binding at once
10	Fuel gas heating value to 88 %	the disturbance feedforward cannot see
11	Thermocouple reading 6 °C high	feedforward as a soft sensor
12	Three weeks at 300 t/h	velocity dependence of fouling
13	MPC with the whole valve as a hard limit	a gain traced to its sources
14	Graded challenge over three weeks	architecture before tuning

balance. Per degree of column inlet a charge move is priced below a valve move, $S_F/K_F^2 = 0.012$ against $S_v/K_v^2 = 0.019$ per °C^2 with K_F the charge's steady gain, so a fast upset also dips the charge while valve remains. The weights were set once on scenario 13 and are not retuned for any other case.

Acceptance, fallback and cost. A plan is applied only if the solver reports it optimal, every planned command lies inside its box, the plan's first command is the move sent, and its cost, slack and predictions are finite. Otherwise *Override* takes both handles, its TIC-101 and VPC-101 having tracked the applied commands exactly, a fallback mark goes into the historian, and the first accepted solve after it hands both handles back to the plan. Forced at six instants of scenario 13 with noise off, the hand-over moves the valve by at most 0.003 % and the charge by 0.05 t/h, except during the restore at hour 18.3, where the override's own first moves are 0.99 % and 0.48 t/h. The shipped rule that rests the VPC-101 integral at F_{dem} is suspended under the plan, because at hour 6 it would turn a fallback into a charge step of 73 t/h. In Chrome a plan takes 183 to 191 µs on average over the 800 plans of scenario 13, about 0.3 % of one core at a playback rate of one plant hour per second.

Three models kept apart. The plant is the twelve-state model with its true fouling, heating value and thermocouple bias, the identified model is the three lags and the furnace inlet response, and the controller's model is (3.17) with those dynamics and the bias, fed only the two transmitters, the setpoint, the applied commands, the demand and the floor.

3.5 Results

Fourteen scenarios are arranged in instructional order, listed in Table 3.3. Each one fixes the opening mode, the architecture, the three tuning constants, the disturbance settings, any parameter overrides and any timed events, and supplies a briefing, a prediction question, a procedure, one or two hints, a worked explanation and a completion requirement that the app judges from the run (Section 1.6.4). The app lists scenario 14 last, as the challenge, without a number.

3.5.1 Fouling Under Closed-Loop Control

Scenario 4 runs three weeks with no other disturbance. The results are in the first column of Table 3.4, and the principal observation is that the temperature never moves. The integrated absolute error over the whole 504 h is 0.0 °C h and the column inlet sits on 360.0 °C, while the entire disturbance goes into the fuel valve, which opens from 81.3 % to 97.5 % without a step and without an alarm. Figure 1.4 draws the same hidden loss on R-1. The E-101 duty falls from 22.9 to 21.1 MW and E-102 from 26.8 to 21.9 MW, the furnace makes up all of it by going from 33.0 to 39.5 MW, and the recovered share of the duty falls from 60.1 % to 52.1 %.

Table 3.4 Three Weeks (504 h) of Fouling With No Other Disturbance, With $K_c = 1.5$ and $\tau_I = 0.35$ h. Every Figure Is Read at 504 h. The First Two Columns Run on + *Feedforward*. The Third Runs on *Override* and Cleans Each Exchanger at Its Own Threshold, at Hours 240 and 396.

	400 t/h no clean	300 t/h no clean	400 t/h two cleans
IAE (°C h)	0.0	—	57
Fuel valve, final (%)	97.5	—	—
$R_{f,1}$ final	0.481	0.618	—
$R_{f,2}$ final	0.668	0.870	—
E-101 threshold reached (h)	239.6	124.2	240
E-102 threshold reached (h)	395.5	—	396
Recovery, final (%)	52.1	—	—
Crude processed (kt)	201.6	151.2	199.1
Fuel fired (TJ)	76.7	—	72.8

Table 3.5 A Sixteen-Hour Cleaning Outage Started at $t = 2$ h, Every Figure Read at $t = 48$ h. Columns 1 to 3 Take E-101 Out, Column 4 Takes E-102 Out. All Four Carry the Same Feedback Tuning and $K_{ff} = 1.0$. The Worst Column Inlet Is Read From $t = 2$ h and the Lined-Out Value at the End of the Outage.

	feedfwd only	override $x_{v,SP} = 95$	override $x_{v,SP} = 85$	override E-102
Worst column inlet (°C)	340.3	358.6	358.8	317.5
Column inlet lined out (°C)	340.3	360.0	360.0	360.0
IAE over 48 h (°C h)	287	1.3	0.9	59.4
Minimum charge rate (t/h)	400	363	340	295
Average charge rate (%)	100.0	97.1	95.2	91.7
Crude processed (kt)	19.200	18.645	18.279	17.616
Time at firing limit (% of outage)	96.5	0	0	13.3

Pause. Try Exercise 3.1 before reading on.

The remaining margin is the concern. The valve stops at 97.5 % with 2.5 % left, while every measure a temperature loop reports indicates success. Raising the charge demand to 420 t/h puts the same three weeks on the stop at hour 215, and at 440 t/h at hour 32. Fuel over the three weeks comes to 76.7 TJ, about $460k at $6/GJ, and the fouling accounts for roughly a tenth of it.

The second column of Table 3.4 is a counterintuitive result. Turning the unit down to 300 t/h fouls it *faster*. After three weeks the E-101 resistance is 0.618 against 0.481 and E-102 is 0.870 against 0.668, which in threshold units is 1.47× and 1.45× against 1.14× and 1.11×, and E-101 crosses its cleaning threshold at hour 124 instead of hour 240. The mechanism is in (3.2). Deposition falls as $v^{-0.4}$ because a faster stream spends less time in contact with the wall, removal rises as $v^{1.0}$ because shear is what scrubs a soft deposit off, and the asymptote is the ratio, which goes as $v^{-1.4}$. The cost appears on two accounts. The turned-down month processed 151.2 kt rather than 201.6 kt, which was the intended reduction, and it also leaves the next run starting from a higher resistance, which was not.

3.5.2 Cleaning Outages Under Override Control

Scenario 5 opens the E-101 bypass at full charge rate on the feedforward architecture, and scenario 6 repeats it with the override selected and nothing else changed. Table 3.5 carries both, together with the same event on E-102 and the effect of moving the valve position target. Figure 3.3 draws three of its rows as a chart.

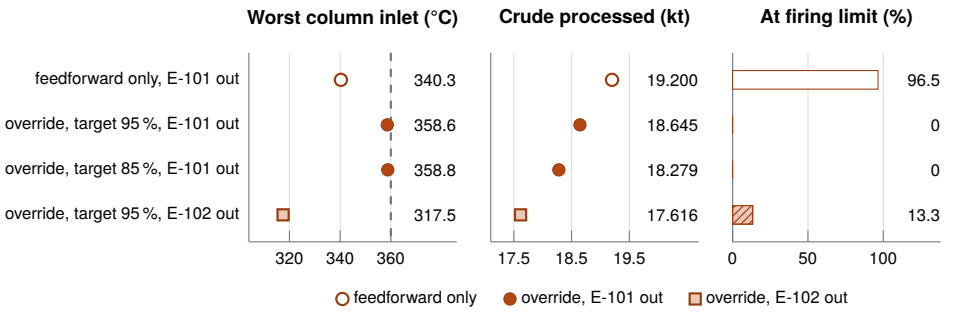


Figure 3.3 Three rows of Table 3.5 as a chart, the last as a share of the outage. The dashed line is the setpoint. The open circle and white bar are feedforward only, filled circles and solid bars the override with E-101 out, the square and hatched bar the override with E-102 out. Notice that the 85 % target gives up more crude than the 95 % target for the same temperature, and that the E-102 outage costs most on all three.

The duty accounting explains the first column. E-101 goes from 22.9 MW to zero, but E-102 goes *up*, from 26.8 to 37.6 MW at hour 6, because its crude inlet has just fallen from 118.6 °C to 25 °C and its driving force has risen by the same amount. The train therefore loses 12.0 MW rather than 22.9 MW. The furnace covers 7.6 of that by going from 33.0 MW to its 40.56 MW ceiling, and the remaining 4.4 MW is the 18 °C the column does not get at hour 6, widening to 19.7 °C by hour 18 as E-102 warms through and gives part of the gain back. The column inlet parks at 340.3 °C and recovery falls from 60.0 % to 48.1 %. The valve reaches its stop 0.57 h after the bypass starts to swing and stays there until 0.54 h after E-101 returns, 96.5 % of the outage, and for that time the temperature loop has no authority at all, because back-calculation holds the integral where it reproduces the saturated output and more gain only raises a demand the valve cannot follow.

The second column is the same event with VPC-101 allowed to write the charge rate. The column inlet lines out on 360.0 °C with a worst deviation of 1.4 °C while the bypass is swinging, the integrated error over the 48 h window falls from 287 to 1.3 °C h, and the fuel valve never reaches its stop. The recovery is obtained by cutting the charge rate to 363 t/h, 90.8 % of design, for the sixteen hours the exchanger is away, and the price is 555 t of crude out of 19 200 t, or 2.9 %. Two effects combine in that cut. A lower rate requires less furnace duty, and the crude also gains more temperature rise per tonne crossing the same 1620 m^2 of surface, which is why a 9 % rate reduction recovers about twenty degrees.

The third column shows why the valve position target is a design decision rather than a default, because the relationship is not monotone. At 100 % the override has nothing to act on once the valve is on its stop, and the run repeats scenario 5 exactly, 19.200 kt with a worst column inlet of 340.3 °C. At 95 % the temperature holds and 18.645 kt goes through. At 85 % the worst column inlet during the outage is no better, 358.8 °C against 358.6 °C, while the charge falls to 340 t/h and 18.279 kt goes through. Below the point at which the constraint actually binds, a tighter target buys no temperature and costs tonnes, and the temperature trend does not show the loss.

🕮 **Pause.** Try Exercise 3.3 before reading on.

The fourth column is the asymmetry between the two exchangers, and it is structural rather than a matter of size. When E-101 is bypassed, E-102 sits downstream of the gap and picks up more than half of the loss by itself. When E-102 is bypassed there is nothing downstream of it except the furnace, so the crude arrives at the coil at 124.6 °C instead of 211.3 °C and the furnace is asked for 58.3 MW, which is 144 % of what it has. The override cuts to 295 t/h, 74 % of design, and the temperature still dips to 317.5 °C because the override cannot cut faster than the duty disappears with a bypass swinging over twelve minutes. Setting the demand to 300 t/h before the run, so that the unit

Table 3.6 Two Cleaning Outages Over Three Weeks, Read at 504 h. The Third Column Moves the Second Clean of Scenario 9 to Hour 280 in a Counterfactual Run. The Fourth Shortens Both Outages to 8 h, the Comparison Scenario 9 Asks For.

	separated (scen. 8)	8 h apart (scen. 9)	moved to 280 h	8 h outages
IAE (°C h)	57	559	48	14
Worst column inlet (°C)	319.8	262.6	326.6	344.9
Minimum charge rate (t/h)	298	240	307	290
Crude processed (kt)	199.061	199.056	199.039	200.315

is already lined out at the lower rate *before* the bypass opens at hour 2, which is what an operator does, turns that dip into 354.4 °C and the integrated error over the window from 59.4 into 6.3 °C h, and it processes 14.334 kt against the 17.616 kt the override gave, so the planned cut processes 3.3 kt less crude than the override. A planned reduction trades throughput for temperature margin, and the trade has a measurable price.

> **TRY IT IN THE APP** **7 · Clean E-102, the expensive one**
>
> Load the scenario and press Run. Nothing has to be set. At hour 2 the bypass opens and E-102 leaves for sixteen hours. Watch the crude charge rate. The override cuts it to 295 t/h, 74 % of design, and the column inlet still dips to 317.5 °C, because the duty disappears faster than the override can cut.

3.5.3 Interaction of Two Cleaning Cycles

Scenario 8 cleans each exchanger when it reaches its own threshold, E-101 at hour 240 and E-102 at hour 396, so only one is ever away at a time. The third column of Table 3.4 carries the outcome: 199.1 kt at 98.7 % of design, an integrated error of 57 °C h over three weeks, the charge rate cut for 6.5 % of the run, and fuel of 72.8 TJ against 76.7 TJ for the same three weeks with no cleaning at all. Two cleans therefore cost 2.5 kt of throughput and buy back 3.9 TJ of fuel gas, about $23k at $6/GJ, plus the whole of the valve margin. Whether that is a good trade requires a crude margin the application does not carry, which is the form the scheduling argument takes in practice.

The periods are different by construction. Started from clean at the design rate, E-101 takes 606 h to reach its threshold and E-102 takes 416 h, and the application opens with E-101 0.76 of the way through its run and E-102 nearly clean, so the two thresholds arrive 156 h apart. Two cycles of unequal length drift in and out of phase forever, and scenario 9 is the in-phase half. Raising the E-102 starting resistance from 0.05 to 0.36 brings E-102 to its threshold at hour 229 and E-101 to its own at hour 239, and the schedule cleans them at hours 240 and 248, so the two sixteen-hour outages overlap for eight hours. The comparison in Table 3.6 is the strongest result in the application.

The reason the tonnes did not move is the reason the temperature did. The override cuts to 240 t/h, exactly the minimum stable charge rate, and then stops, because below that the fractionator trays weep. With the override on its own floor and the fuel valve on its stop, both manipulated variables are gone at once, and the column inlet goes wherever the heat balance puts it, which is 262.6 °C. Moving the second clean to hour 280 costs 17 t out of 199 000 and changes the integrated error by a factor of almost twelve, from 559 to 48 °C h. That run is a counterfactual, because the scripted clean at hour 248 always fires in the application, so scenario 9 poses the same comparison by shortening both outages to 8 h, which lets E-102 leave as E-101 returns. The worst column inlet is then 344.9 °C, the override reaches only 290 t/h and the integrated error is 14 °C h, while 10 h outages that still overlap by two hours fall to the same 262.6 °C and reach 184 °C h. The overlap rather than the length of the outages is what the column cannot carry. The consequence of two constraints binding together is not the sum of the two events, because the constraint set is not additive, and separating two outages that could collide costs essentially nothing.

Table 3.7 Feedforward Against Two Disturbances. The Charge Rate Case Is a 400–340–400 t/h Sequence Read at 16 h, the Fuel Gas Case a Step to 88 % Heating Value Read at 12 h. Identical Feedback Tuning in Every Column.

	charge rate		fuel gas	
	single	+ ff	single	+ ff
Peak deviation (°C)	+10.3	+3.0	−7.14	−7.12
IAE (°C h)	10.7	2.5	2.80	2.61

+ DID YOU KNOW? The Tonnes Stayed, the Temperature Did Not

From the simulator (Table 3.6). Two cleaning outages that overlap for eight hours process 199.056 kt of crude in three weeks, and the same two kept apart process 199.061 kt. The integrated temperature error is 559 °C h against 57 °C h, and the worst column inlet 262.6 °C against 319.8 °C.

3.5.4 Disturbances That Resemble Fouling

Anything that slowly loses duty resembles fouling in the duty accounting, and the application carries three such disturbances. Scenario 10 drops the fuel gas heating value to 88 %, so the same valve position delivers less heat. Table 3.7 puts the two architectures side by side on that disturbance and on the charge rate change of scenarios 2 and 3.

Against the charge rate the feedforward removes 77 % of the accumulated error, cutting the peak deviation from 10.3 to 3.0 °C and the integrated error to 2.5 °C h from 10.7. On the fuel gas it removes almost nothing. The two architectures reach the same peak to within 0.02 °C and differ by 7 % in integrated error. Feedforward acts on measurements and there is no heating value measurement in this plant, so (3.15) returns an unchanged output and the entire correction comes from the feedback trim. The valve settles about twelve points higher at 93.4 %, close to the $1/0.88$ that lean gas costs, and the signature is a valve that moved while every reported duty held still. Gain cannot substitute for either. A sweep on the charge rate case gives 32.0 °C h at $K_c = 0.5$, 10.7 °C h at 1.5, then the wrong way to 82.0 °C h at 3.0 and 142.6 °C h at 6.0, where the loop rings continuously.

The second is a hot stream that has cooled. Dropping the residue supply from 345 to 315 °C takes 30 °C of driving force out of E-102 and produces the same duty loss as fouling. The two are told apart because U is reported beside the duty, and fouling lowers U while a cold supply lowers the duty at unchanged U.

The third is a biased instrument, the only failure a feedback loop cannot detect on its own. Scenario 11 puts 6 °C of bias on the column inlet thermocouple, together with 0.6 °C of noise, so averaged over hours 12 to 24 the true column inlet is 354.0 °C while the measurement reads 360.0 °C and the loop performs perfectly throughout. The evidence exists outside the loop. Over the same hours the fuel valve averages 78.91 % where the same run without the bias puts it at 82.56 %, and 3.65 % of valve at the 1.64 °C per percent process gain is exactly the 6.0 °C. More directly, the feedforward term of (3.15) averages 82.55 % in both runs because it is computed from instruments that have not drifted, while the feedback trim moves from +0.01 to −3.64 %. A feedforward built from independent measurements is also a soft sensor, and a persistent disagreement between it and the controller that trims it indicates a failed instrument. Measurement noise has a separate cost, paid entirely in valve travel. At 0.6 °C of transmitter noise the fuel valve command travels 645 % of full stroke in a day with no disturbance on the plant, and at 2 °C it travels 2149 %, as the scorecard's valve travel reports.

3.5.5 Graded Challenge

Scenario 14 applies a fuel gas step at hour 100, an E-101 cleaning at hour 180, a cold front at hour 250 and an E-102 cleaning at hour 290 over the full three weeks, and asks for the column inlet held within ±3 °C once both exchangers are back, at least 190 kt processed, and the fuel valve off its stop at

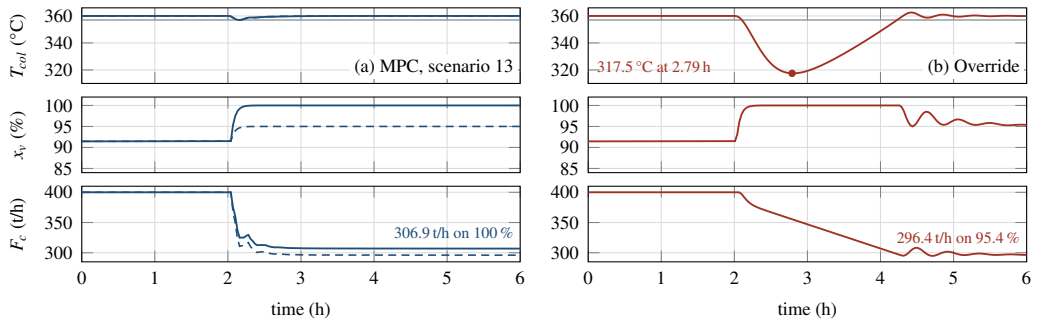


Figure 3.4 Scenario 13, the E-102 outage of scenario 7 beginning at hour 2, over its first six hours, the span the application's 12 h trend window shows at hour 6, under (a) the predictive controller, solid, and the same plan with its valve ceiling at 95 %, dashed, and (b) *Override* with its 95 % valve position target. Top, the column inlet against the plan's floor at 357 °C, the thin gray line. Middle, the fuel valve. Bottom, the charge rate. The plan's lowest column inlet is 357.0 °C at hour 2.16 and the override's 317.5 °C at hour 2.79.

the end. It opens on the single loop with a cautious tuning, and the result establishes that architecture rather than tuning decides this problem. On *Single* at the default gain the run processes 201.6 kt and spends 46.9 % of the three weeks with the fuel valve on its stop, with an integrated error of 3166 °C h. Retuning to $K_c = 1.5$ and $\tau_I = 0.35$ h gives 3091 °C h and adding feedforward gives 3085 °C h, so the two adjustments together are worth 2.5 %, and every one of these runs ends with the valve on its stop and fails. On the shipped settings the challenge's rule reads a worst deviation of 7.7 °C after the cleans and 201.6 kt processed, so the run fails two of its three criteria. Selecting *Override* from the shipped settings and changing nothing else gives 3141 °C h, the same as feedforward at that tuning, because the valve position target opens at 100 % and an override told to keep the valve below 100 % never acts. Moving that one number to 95 %, with the retuned loop, gives 86 °C h, a worst deviation after the cleans of 0.86 °C, reached as the E-102 bypass finishes closing, 2.1 °C inside the band, and 0.7 % of the run at the firing limit, for 194.4 kt against 201.6 kt, 4.4 kt above the floor, and an average charge rate of 96.4 %. The valve ends at 95.0 %, the margin the override was told to keep. With the shipped cautious tuning left in place the same target also passes, at 123 °C h and a worst deviation of 0.52 °C. Fuel falls as well, 69.9 TJ against 72.0 TJ, because a furnace held off its saturation limit does not overfire during recovery.

An integral of absolute error does not stop growing when the transient does, so every IAE above is quoted with the window it was read over, and the comparison carries the information rather than the absolute value. The 1.3 °C h of scenario 6 and the 287 °C h of scenario 5 are both read at 48 h on the identical event, and the three-week figures are all read at 504 h.

3.5.6 The Predictive Controller Against the Override

Table 3.8 runs scenarios 4 to 11 and the challenge, scenario 14, under the shipped structures and the predictive controller with each scenario's own seed, duration, instruments and noise and no preview, the plan at scenario 13's settings throughout, and on the challenge every row but the shipped one from the suggested settings. Three more rows locate the gain. The plan with $x_{v,max} = 95$ % keeps the override's margin. Override with its valve position target raised to 99 % is the shipped loop retuned once in the plan's direction. The same override with the balance of (3.15) as a feedforward on the charge, low-selected at the rate at which the target valve holds the setpoint on $\hat{T}_2$, is the like-for-like classical structure. It exists only in the measurement script, not in the application.

In scenario 13 the plan cuts the charge while the bypass is still swinging, so the lowest column inlet is 357.0 °C against the override's 317.5 °C, and by hour 6 it holds the valve on 100 % at 306.9 t/h

Table 3.8 The Predictive Controller Against the Shipped Structures, Whole Runs Measured on the Engine. Scenario 7's Plant Under the MPC Is Scenario 13.

Case	Structure	Not processed (t)	IAE (°C h)	Worst (°C)	Valve on stop (min)	Fuel (TJ)	Charge travel (t/h)	Plans, fallbacks
4 (504 h)	+ Feedforward, and both rows at 99 %	0	0.00	0.00	0	76.70	0	–
	Override	532	0.00	0.00	0	76.34	6	–
	MPC	0	0.38	0.00	0	76.70	0	8400, 0
	MPC, ceiling 95 %	533	1.05	0.01	0	76.34	6	8400, 0
5, 6 (48 h)	+ Feedforward (scenario 5)	0	286.7	19.7	959	6.82	0	–
	Override (scenario 6)	555	1.27	1.36	0	6.69	78	–
	MPC	386	1.55	0.77	929	6.82	54	800, 0
	MPC, ceiling 95 %	564	1.53	0.79	0	6.68	75	800, 0
	Override 99 % + charge feedforward	419	1.24	1.36	27	6.79	57	–
	Override 99 %	393	16.8	9.42	185	6.80	68	–
7, 13 (48 h)	Override	1584	59.4	42.5	128	6.81	271	–
	MPC	1497	2.57	3.02	954	6.92	240	800, 0
	MPC, ceiling 95 %	1676	2.53	3.14	0	6.79	269	800, 0
	Override 99 % + charge feedforward	1525	4.33	7.93	40	6.90	192	–
	Override 99 %	1003	427.2	65.1	707	6.92	204	–
8 (504 h)	Override	2539	57.4	40.2	184	72.83	421	–
	MPC	2280	4.51	3.21	1895	73.07	387	8400, 0
	MPC, ceiling 95 %	2642	4.42	3.19	0	72.80	445	8400, 0
	Override 99 % + charge feedforward	2340	5.92	7.86	77	73.02	297	–
	Override 99 %	1705	508.2	62.6	1082	73.04	320	–
9 (504 h)	Override	2544	559.1	97.4	572	73.73	816	–
	MPC	2427	485.3	62.6	1384	73.87	369	8400, 0
	MPC, ceiling 95 %	2623	595.2	76.2	0	73.66	367	8400, 0
	Override 99 % + charge feedforward	2443	485.4	62.6	574	73.85	321	–
	Override 99 %	1747	1098.7	133.4	997	73.85	522	–
10 (12 h)	+ Feedforward, and both rows at 99 %	0	2.61	7.11	0	1.63	0	–
	Override	0	2.61	7.11	0	1.63	2	–
	MPC, at either ceiling	3.6	3.00	7.63	0	1.62	28	200, 0
11 (24 h)	+ Feedforward, Override, both rows at 99 %	0	141.1	7.18	0	3.13	0	–
	MPC, at either ceiling	10.2	142.8	6.95	0	3.12	288	400, 0
14 (504 h)	Single, shipped	0	3165.7	97.6	14180	72.00	0	–
	Override, suggested	7194	86.4	52.2	205	69.94	629	–
	MPC	4131	13.6	8.13	14423	72.03	475	8400, 0
	MPC, ceiling 95 %	7312	13.8	8.13	0	69.91	512	8400, 0
	Override 99 % + charge feedforward	4168	426.4	25.0	1381	71.71	455	–
	Override 99 %	3677	827.6	82.5	1381	71.71	454	–

Not processed is $\int (F_{dem} - F_c)\, dt$. IAE is the scorecard's integral on the true column inlet and worst its largest deviation from setpoint. Valve on stop counts the minutes with the fuel valve at or above 99.5 %, the scorecard's firing-limit test. Charge travel sums the changes of the charge command over the logged rows. The MPC rows decide every 0.06 h at scenario 13's settings, and on scenarios 10 and 11 its two ceilings give identical runs.

and the override 95.4 % at 296.4 t/h, as shown in Figure 3.4. Over 48 h the plan processes 17.703 kt against 17.616 kt at an integrated error of 2.57 against 59.35 °C h.

The extra tonnes are the last five points of valve rather than the optimization. With its ceiling at 95 % the plan leaves 1676 t unprocessed on scenario 13 against the override's 1584 t, so those five points are worth 179 t over the outage, and at equal authority the plan gives 92 t back by cutting at once rather than letting the column run 42.5 °C cold. The three runs of the scenario's completion rule, read at hour 26, give 8.903 kt, 8.816 kt and 8.724 kt processed under the plan, the override and the plan at a 95 % ceiling, the same 87 t gain made of the same 179 t less 92 t. At that ceiling it processes less than the override on every case with a clean, 564 against 555 t unprocessed on scenarios 5 and 6, 2642 against 2539 on scenario 8, 2623 against 2544 on scenario 9 and 7312 against 7194 on the challenge. The shallow dip comes from acting on the measured furnace inlet. The override given the balance as a charge feedforward at its 95 % target dips 4.64 °C on scenario 13 against the plan's 3.02 °C, and at a 99 % target it comes within 28 t and 1.8 °C h of the plan. On the challenge that structure comes within 37 t but loses 413 °C h, because a balance written at the design heating value cannot see the gas that turned lean at hour 100, while the plan's output bias carries it. With that event removed from the challenge the gap is 59 t and 1.5 °C h.

Override at 99 % alone beats the plan on tonnes on scenarios 7, 8, 9 and the challenge, though not on scenarios 5 and 6 (393 against 386 t), with 1003 against 1497 t unprocessed on scenario 13 and 3677 against 4131 t on the challenge, by letting the column run as much as 65.1 °C cold at an

integrated error of 427 against 2.57 °C h. It even passes the challenge, with 197.9 kt, a worst deviation of 0.5 °C after the cleans and the valve at 99.0 %, because the temperature clause is read only once both exchangers are back. It is a third structure rather than a recommendation, its tonnes paid for in column inlet temperature that a single score would conceal.

The plan loses scenario 10, at 3.00 against 2.61 °C h and a worst deviation of 7.63 against 7.11 °C, and gives up 3.6 t by dipping the charge to 386.2 t/h with the valve at 86 %. An unmeasured heating value step reaches the plan only through the bias, filtered over 0.06 h and acted on every 0.06 h, where TIC-101 acts every 0.02 h. Counterfactual copies isolate the cause: a pinned charge gives 3.05 °C h, a decision every scan 2.35, a 0.02 h bias filter 2.28 and both together 1.63. On scenario 11 both structures hold the biased reading, at 142.8 against 141.1 °C h, and the plan gives up 10.2 t, of which 8.3 t answer the 0.6 °C noise with the charge (1.9 t without noise).

On the challenge the plan meets the temperature clause, at 1.0 °C worst after the cleans, and the tonnage clause, at 197.5 kt against the suggested override's 194.4 kt, but fails the requirement that the fuel valve finish off its stop, at 100.0 %, because the plan spends the whole valve by design. At a 95 % ceiling it passes with 194.3 kt, 0.9 °C and 95.0 %. The scorecard's firing-limit fraction reads 33.1 % of scenario 13 under the plan against 4.4 % under the override, and 47.7 % of the challenge against 0.7 %. Under the plan that time on the stop is chosen, not authority lost. None of the 71 600 plans in the runs behind Table 3.8 was rejected.

3.5.7 Implementation

Manual and automatic modes are provided, with operator helpers that start a cleaning outage (Clean E-101, Clean E-102), evaluate the heat balance of (3.10) (Balance the furnace), and solve for the highest feasible charge rate in the current exchanger condition (Feasible charge rate). Seeds, Reset and the exported file work as Sections 1.6.1 and 1.6.4 describe, so every run repeats exactly. The application also runs thirty-three internal verification checks, seventeen of them on the predictive controller, covering the design steady state, agreement between $UA \cdot \mathrm{LMTD}$ and effectiveness-NTU to one part in ten thousand, the closed-form fouling integral over 504 h, the $v^{-1.4}$ scaling, sign conventions and the cleaning state machine. Figure 3.5 shows the recovery after a cleaning outage under override control.

3.6 Technology in Practice

Figure 3.6 draws the override structure of Section 3.4 with the hardware a refinery installs in place of the model's ideal instruments.

What Measures It

Both temperatures that the controllers use come from thermocouples in thermowells. The lags τ_{TT} and τ_{FT} of Table 3.2 stand for the wells, and the bias of scenario 11 for an error in the measuring circuit. The process variable of VPC-101 is no process measurement. In the general scheme it is the input that the first controller writes, held near a resting value on a slow time scale [22], and the model reads the valve position exactly.

> **TECHNOLOGY** **Sensor: Thermocouple in a Thermowell**
>
> A thermocouple is two wires of unlike metals joined at the hot end. The junction produces a voltage roughly proportional to temperature, but the connection to the instrument forms a second junction that opposes it, so the receiver must compensate for the temperature at its own terminals [25]. The wrong kind of extension wire moves that second junction, and the compensation is then made for the wrong temperature [25]. The result is a standing offset, and scenario 11 shows what a loop does with one: it holds the wrong temperature exactly. The sensor sits in a thermowell, a pressure-tight sheath that lets it be replaced with the process under pressure. The mass of the well adds a first-order lag [25], which the model carries as 0.01 h on the column inlet (Table 3.2).

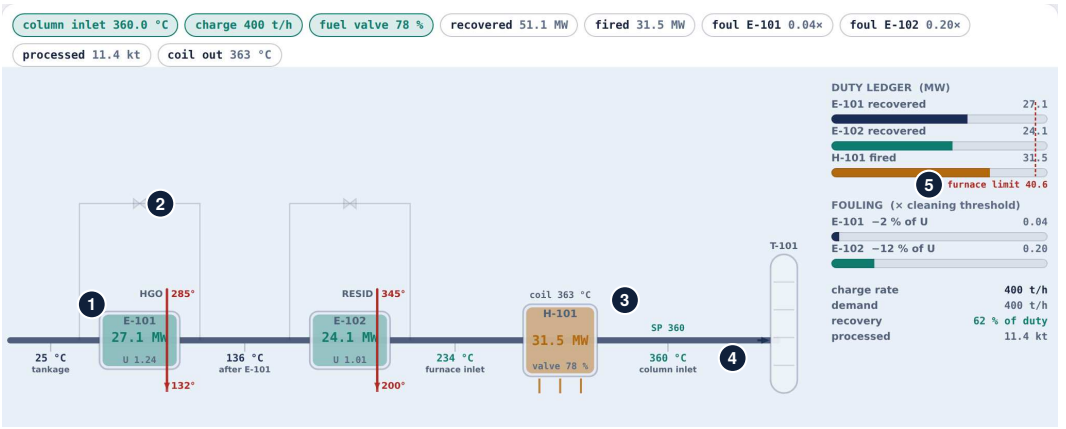


Figure 3.5 Scenario 6, recovery from a sixteen-hour E-101 cleaning outage under override control with the shipped settings, at $t = 30$ h and twelve hours after E-101 returned clean, cropped to the plant view with numbered callouts. E-101 (1) is back in service with its bypass (2) shut, the fired heater (3) supplies the rest of the duty, the column inlet (4) is on its setpoint, and the duty ledger (5) marks the furnace limit. In the full window, outside this crop, the 48 h trends show the charge rate cut to 363 t/h while the exchanger is out of service and back at 400 t/h by hour 18.6, the column inlet within 1.4 °C of 360 °C throughout, and the fuel valve settling at 77.7 % against the 81.3 % it opened on because the cleaning recovered 3.6 points of valve.

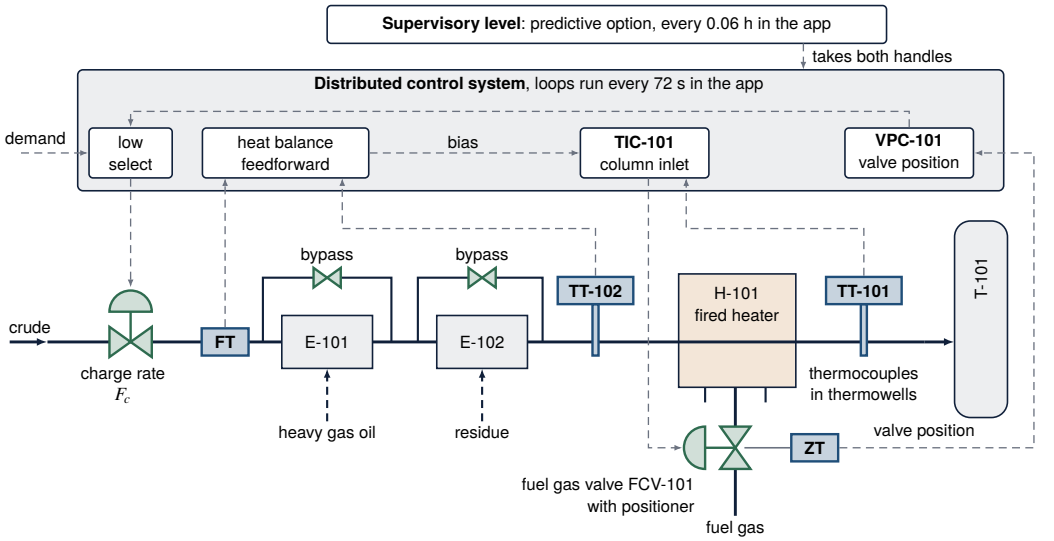


Figure 3.6 The loops of the override structure with their hardware. Lettered boxes are sensing devices, valve symbols are final elements and dashed arrows are signals. Notice that the measurement of VPC-101 is the fuel valve itself, read by a position transmitter (ZT), and that its output reaches the charge valve only through a low select against the operator's demand. Not to scale, and not a piping or safety diagram.

What Moves It

Three kinds of final element carry the chapter, the fuel gas valve FCV-101, the charge rate actuator and the two exchanger bypasses. The model gives each a first-order lag and no dead band or friction,

and its one nonlinearity is the one that matters here, a valve that stops at 100 %. Over 60 % of all fuel used in a refinery is burned in furnaces and boilers, and the average thermal efficiency of furnaces is estimated at 75 to 90 % [1], a range that holds the 0.88 of Table 3.2.

> **TECHNOLOGY** **Final Element: Control Valve With a Positioner**
>
> In a pneumatic control valve, air pressure on a diaphragm or piston works against a spring. Packing friction and the force of the fluid on the plug keep the stem from following the air pressure exactly, so a positioner is added, a device that compares stem position with the command and adjusts the air until they agree [25]. The spring sets the position on loss of air. An air-to-open valve fails closed, and the choice is to be made from the risk of the process, not from habit [25]. A burner management system has shutoff valves of its own to stop the fuel [25]. An electronic positioner also reports the stem position it measures and the total stem travel, like an odometer [25]. A valve position controller is a different thing. It uses a second input to keep the first from saturating [22].

Where the Controller Runs

A refinery runs such loops in a distributed control system, in which a supervisory level sends set points to distributed controllers that act through the actuators from the feedback of the sensors, in cycles from milliseconds to minutes [26]. A survey of industrial practice puts basic control at every second and constraint control at every minute, and builds the conventional constraint layer from PID algorithms, lead-lag blocks and high and low select logic [21]. The structure of this chapter is of that kind, executed in the app every 0.02 h. A controller with integral action that a selector has de-selected tends to wind up, so an override needs limits on its integral [25]. Equation (3.16) gets them by making the operator's demand the upper output limit of VPC-101.

> **IN PRACTICE** **What the Operator Watches: Fouling Resistance and Furnace Inlet**
>
> The model reports its fouling resistances exactly. A refinery has to infer them from temperatures and flows, and only key exchangers carry all four stream temperatures [27]. The most popular indicator is, by far, the fouling resistance of Equation (3.1), computed through a lumped log-mean model like the one of this chapter, and typical measurement errors can put errors of the order of 20 % into it [27]. For the whole train the main indicators are the furnace inlet temperature and the furnace duty, the first often normalized to standard feed conditions. Cleaning needs crews, equipment and cranes, usually contracted at long notice, so operators ask for a stable schedule [28].

What Is Modernizing It

The first of three developments is predictive control, which replaces the blocks of the constraint layer with one controller [21]. Applications typically carry the valve positions of the loops below as extra controlled variables held inside their saturation limits [21], which is valve position control written as a constraint, the form of Section 3.4.4. The second is monitoring with detailed models. A fouling resistance cannot tell the growth of a deposit from a change in its properties, and dynamic thermal and hydraulic models of the exchangers have been proposed to separate them [27]. The third is the electronic positioner, whose sensors and digital communication give the maintenance staff and the control system diagnostic data on the valve [25].

3.7 Limitations

No result here transfers to a real preheat train. Two exchangers stand in for twenty, the fouling law depends only on velocity, nothing vaporizes, the column is a temperature target, and there is no crude margin, so a cleaning interval, an energy saving or the cost of one dirty exchanger are all outside what the model can state, and the fouling literature it draws on exists precisely because those questions are hard [14–16]. The economics are a single fuel gas price, so a decision can be reported in tonnes and gigajoules but the application cannot state which way the decision should go. No limit trips. The

coking limit, the firing limit and the deviation alarm all annunciate while the run continues, which is what lets a severe excursion be read to its end.

One temperature loop, one feedforward term and one constraint controller are the entire shipped control system, and the predictive option beside them plans the same two handles on one heat balance. A modern crude unit runs a multivariable controller with a real-time optimizer above it, and the constraint this application resolves with a valve position controller is exactly the kind that a predictive controller handles by writing it into the optimization, which is the reason that technology was adopted in refining in the first place [21]. Real constraint control on this unit manages firing, flooding, pump capacity, product specifications and tube skin temperatures at once, and the plantwide version pushes throughput against whichever unit is the current bottleneck rather than against one furnace [5]. What survives the simplification is the part a controller has to be designed around, and there are four such features: a process gain that moves with charge rate, an actuator with a hard end, a second manipulated variable that is expensive rather than merely available, and a disturbance that is entirely predictable and therefore answered by a schedule. Section 3.3.8 states where each assumption stops holding.

The documented model also extends beyond the implemented PID, feedforward and override structure to advanced control methods. The predictive option of Section 3.4.4 coordinates the fuel valve and the charge rate as one constrained problem instead of a low select, with the firing limit written in as a bound, and a controller built on the full twelve-state model could add the coking limit as a constraint [29]. The two fouling resistances are unmeasured states that moving horizon estimation or Kalman filtering can reconstruct from the measured temperatures and duties, which turns the soft-sensing argument of Section 3.5 into a formal estimator. An economic objective carrying a crude margin beside the fuel price would let the same controller weigh throughput against fuel continuously, which is the cleaning schedule argument posed as an optimization. The model equations transfer directly to the APMonitor MPC and MHE framework [29] and to the Gekko Optimization Suite in Python [30].

The predictive controller’s model leaves out the heating value, the change in the train’s charge response with fouling or a bypassed exchanger, the slowing of the coil and transfer line at lower charge, the thermocouple bias and the coking limit. The output bias absorbs the first three, and its 0.06 h filter with the 0.06 h decision is what loses scenario 10, at 3.00 against 2.61 °C h. One slack covers the whole horizon, so the floor is not a per-step guarantee. Its advantage over the override is not the optimization as such, because a valve position controller given the same valve range and the same balance as a charge feedforward comes within 28 t and 1.8 °C h on scenario 13, a result that leaves valve position control, whose design and tuning are established in industrial practice [31], competitive for this one constraint.

3.8 Discussion

The quantities the application reports are the ones the real scheduling argument is conducted in, and there are three of them rather than two. The fuel argument favors frequent cleaning, because a fouled exchanger burns gas continuously and carries the emissions that go with it [17], and the two cleanings of scenario 8 recover 3.9 TJ over three weeks. The throughput argument favors rare cleaning, because margin on crude is worth far more per tonne than gas is per gigajoule, and those same two cleanings cost 2.5 kt. Neither argument settles the schedule, because the deciding term is the constraint. An exchanger left too long takes the fuel valve to its stop, and from that point the throughput loss occurs without an operating decision: scenario 4 is that condition approaching with 2.5 % of valve left, and scenario 14 shows the condition after it arrives, at 46.9 % of three weeks on the stop. Schedules computed with throughput in the objective are not the schedules fuel alone produces [18, 20].

⚠ OPEN CHALLENGE Scheduling the Cleaning and Running the Loops Together

The cleaning schedule and the flow control of a preheat train have usually been worked out apart, by heuristics. The operation then reaches a limit such as the firing limit of the furnace, and an emergency cleaning or a cut in production follows [28]. Joining the two is hard. The time scales run from weeks to years, crude blends change within hours or days, and a schedule that changes too often cannot be resourced [28]. A 2020 study closed both loops with predictive control on a model validated against 1240 days of refinery data, but its plant was that model [28]. Progress would be the same loop on a running unit, with a schedule stable enough for the crews to follow.

What the control engineer brings to that decision is a set of measured quantities rather than a controller: the duty the exchanger is still recovering, the fuel valve position, the tonnes the override gave up the last time it was selected, and the length of the next run at the rate the unit is actually being asked to hold. The last of those carries a penalty the application measures, because a unit turned down to 75 % of design fouls toward an asymptotic resistance 1.50 times higher and reaches its threshold in 124 h instead of 240 h, so throughput given up in one month raises the fouling cost of the next.

The scorecard is built around the same argument. It reports the column inlet error, the crude processed, the average charge rate, the fuel fired with its cost, the heat recovered over the run and at the instant, the fraction of the run with the charge rate cut and the fraction at the firing limit, and it never combines them into one number. A run that scores well on temperature and badly on average charge rate has abandoned the problem rather than solved it, and the only way to make that visible is to report both halves and refuse to weight them. What a controller spends on a constraint is invisible unless something counts it, which is the reasoning behind reporting throughput beside the thermal objective in cleaning schedule work [18]. The exported file holds exactly what the trends were drawn from, so a claim about tonnes can be checked in a spreadsheet.

3.9 Conclusion

This chapter studied a crude preheat train and its fired heater as a laboratory for override control, where the manipulated variable of last resort is the plant's own throughput. The simulator integrates twelve states, four exchanger temperatures, the coil outlet, the column inlet, two fouling resistances and four actuators, with fouling entering as a series resistance whose asymptote scales as crude velocity to the -1.4 power. At the design point E-101 recovers 22.87 MW, E-102 recovers 26.79 MW and the furnace supplies 32.98 MW of a 40.56 MW capacity, so 60.1 % of the 82.63 MW duty is recovered from product streams and the fuel valve sits at 81.3 %. Three weeks of fouling never move the column inlet off 360.0 °C and leave the integrated error at 0.0 °C h, while the valve opens to 97.5 % and recovery falls to 52.1 %, which is the case for a performance report that carries more than error. Taking E-101 out for cleaning at full rate delivers 340.3 °C with the valve on its stop for 96.5 % of the outage, while the valve position override holds 360.0 °C for 555 t of crude out of 19 200 t, and losing E-102 instead costs about three times as much because nothing downstream of it picks up the loss. Moving one of two colliding cleaning outages 32 h later changes the integrated error over three weeks from 559 to 48 °C h and the crude processed by 17 t in 199 000, and shortening both outages to 8 h so that they no longer overlap brings it to 14 °C h, which is more than any tuning in the application can offer. A predictive controller offered as a fourth option processes 87 t more than the override on scenario 13 at an integrated error of 2.57 against 59.35 °C h, yet at the override's 95 % valve ceiling it processes 92 t less, so its throughput is the valve's last five points and its temperature the measured furnace inlet acting on the charge. A fuller model would add a column model that closes the loop between preheat and the hot streams, a wall-temperature-dependent fouling law, a crude margin beside the fuel price, and the coking limit as a constraint of the predictive controller. The simulator remains a place to study constraint control where the second manipulated variable is the plant throughput.

KEY POINTS

- A temperature that sits on its setpoint for three weeks can hide a furnace that is running out of room. Only the position of the fuel valve records it.
- With the fuel valve on its stop, no feedback tuning recovers duty the furnace cannot supply. One manipulated variable remains, the crude charge rate.
- Below the point at which the constraint binds, a tighter valve position target buys no temperature and costs tonnes.
- Separating two cleaning outages that could collide costs essentially nothing. The overlap is what the column cannot carry.
- The predictive controller's extra tonnes are the last five points of valve, not the optimization.

Where else this applies. Valve position control returns in Volume 4, Chapter 9, an override that protects a compressor from surge in Volume 4, Chapter 2, throughput as the override's handle in Volume 4, Chapter 4, and a heater's feedforward and limits in Volume 3, Chapter 4.

References

[1] E. Worrell and C. Galitsky, "Energy efficiency improvement and cost saving opportunities for petroleum refineries: An ENERGY STAR guide for energy and plant managers," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-56183, 2005.

[2] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York: McGraw-Hill, 1996.

[3] T. E. Marlin, *Process Control: Designing Processes and Control Systems for Dynamic Performance*, 2nd ed. New York: McGraw-Hill, 2000.

[4] A. Reyes-Lúa and S. Skogestad, "Systematic design of active constraint switching using classical advanced control structures," *Industrial & Engineering Chemistry Research*, vol. 59, no. 6, pp. 2229–2241, 2020.

[5] E. M. B. Aske, S. Strand, and S. Skogestad, "Coordinator MPC for maximizing plant throughput," *Computers & Chemical Engineering*, vol. 32, no. 1-2, pp. 195–204, 2008.

[6] 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.

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

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

[9] A. Szklo and R. Schaeffer, "Fuel specification, energy consumption and CO2 emission in oil refineries," *Energy*, vol. 32, no. 7, pp. 1075–1092, 2007.

[10] B. Linnhoff and J. R. Flower, "Synthesis of heat exchanger networks: I. systematic generation of energy optimal networks," *AIChE Journal*, vol. 24, no. 4, pp. 633–642, 1978.

[11] 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.

[12] T. R. Bott, *Fouling of Heat Exchangers*. Amsterdam: Elsevier, 1995.

[13] W. Ebert and C. B. Panchal, "Analysis of Exxon crude-oil-slip stream coking data," Argonne National Laboratory, Argonne, IL, Tech. Rep. ANL/ES/CP-92175, 1995, presented at Fouling Mitigation of Industrial Heat Exchangers, San Luis Obispo, CA, June 18–23, 1995. OSTI identifier 453433.

[14] G. T. Polley, D. I. Wilson, B. L. Yeap, and S. J. Pugh, "Evaluation of laboratory crude oil threshold fouling data for application to refinery pre-heat trains," *Applied Thermal Engineering*, vol. 22, no. 7, pp. 777–788, 2002.

[15] B. L. Yeap, D. I. Wilson, G. T. Polley, and S. J. Pugh, "Mitigation of crude oil refinery heat exchanger fouling through retrofits based on thermo-hydraulic fouling models," *Chemical Engineering Research and Design*, vol. 82, no. 1, pp. 53–71, 2004.

[16] F. Coletti and S. Macchietto, "A dynamic, distributed model of shell-and-tube heat exchangers undergoing crude oil fouling," *Industrial & Engineering Chemistry Research*, vol. 50, no. 8, pp. 4515–4533, 2011.

[17] H. Müller-Steinhagen, M. R. Malayeri, and A. P. Watkinson, "Heat exchanger fouling: Environmental impacts," *Heat Transfer Engineering*, vol. 30, no. 10-11, pp. 773–776, 2009.

[18] E. M. Ishiyama, W. R. Paterson, and D. I. Wilson, "The effect of fouling on heat transfer, pressure drop, and throughput in refinery preheat trains: Optimization of cleaning schedules," *Heat Transfer Engineering*, vol. 30, no. 10-11, pp. 805–814, 2009.

[19] M. C. Georgiadis, L. G. Papageorgiou, and S. Macchietto, "Optimal cleaning policies in heat exchanger networks under rapid fouling," *Industrial & Engineering Chemistry Research*, vol. 39, no. 2, pp. 441–454, 2000.

[20] F. Smaïli, V. S. Vassiliadis, and D. I. Wilson, "Mitigation of fouling in refinery heat exchanger networks by optimal management of cleaning," *Energy & Fuels*, vol. 15, no. 5, pp. 1038–1056, 2001.

[21] 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.

[22] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[23] 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.

[24] K. R. Muske and T. A. Badgwell, "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.

[25] 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.

[26] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[27] E. Diaz-Bejarano, F. Coletti, and S. Macchietto, "A model-based method for visualization, monitoring, and diagnosis of fouling in heat exchangers," *Industrial & Engineering Chemistry Research*, vol. 59, no. 10, pp. 4602–4619, 2020.

[28] F. Lozano Santamaria and S. Macchietto, "Stability of optimal closed-loop cleaning scheduling and control with application to heat exchanger networks under fouling," *Processes*, vol. 8, no. 12, p. 1623, 2020.

[29] 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.

[30] 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.

[31] B. J. Allison and S. Ogawa, "Design and tuning of valve position controllers with industrial applications," *Transactions of the Institute of Measurement and Control*, vol. 25, no. 1, pp. 3–16, 2003.

3.10 Exercises

3.1 Concept check · 2 min

In scenario 4 the column inlet sits on 360.0 °C for three weeks and the integrated absolute error is 0.0 °C h, while E-101 and E-102 foul. Which reading on the panel shows that the unit is losing its margin?

(a) The fuel valve position, which opens without a step from about 81 to about 97 %

(b) The column inlet error tile, which grows a little each day as the two exchangers foul

(c) The integrated absolute error, which counts the duty the fouled exchangers no longer recover

(d) The average charge rate, which the feedforward trims down as the duty falls off

3.2 Concept check · 2 min

Equation (3.7) is the energy balance on the furnace coil. Which quantity in it does the feedforward of Equation (3.15) have no measurement of, and which scenario shows what that costs?

(a) The furnace inlet T_2, so the rate cut of scenario 3 defeats the feedforward

(b) The heating value ratio HV/HV_0, so the lean fuel gas of scenario 10 defeats it

(c) The charge rate F_c, so the three weeks of fouling in scenario 4 defeat it

(d) The coil capacitance C_{fur}, so the E-101 outage of scenario 5 defeats it

3.3 Concept check · 2 min

In scenario 6 the override holds the column inlet through the E-101 outage with its valve position target at 95 %. The target is set to 85 % before Run and the 48 h are run again. What changes?

(a) The worst column inlet improves by several degrees, and the crude processed is unchanged

(b) The override never acts, and the run repeats scenario 5 with the valve on its stop

(c) The worst column inlet is about the same, and roughly 370 t more crude is given up

(d) The override starts cutting at once, because the valve at 81.3 % is near the target

3.4 Theory and calculation · 7 min

Use Equations (3.1) and (3.2) with the E-101 values of Table 3.2: $U_{clean,1} = 1.26$, $R_{f\infty,1} = 0.55$, $\tau_{f,1} = 420\,\mathrm{h}$ and a cleaning threshold of 0.42, in the working units. At constant velocity the resistance relaxes as $R_f(t) = R_{f\infty}v^{-1.4} - \left(R_{f\infty}v^{-1.4} - R_f(0)\right)\mathrm{e}^{-tv/\tau_f}$.

(a) Find the coefficient of E-101 at its threshold and the share of the clean coefficient it has lost. *Checkpoint:* 0.824 and 34.6 %.

(b) The run opens with $R_{f,1} = 0.320$ at the design rate. Find the hour E-101 reaches its threshold. *Checkpoint:* 239.6 h.

(c) Repeat (b) at 300 t/h, finding the asymptote and the time constant first. *Checkpoint:* 0.823, 560 h and 124.2 h.

3.5 Theory and calculation · 6 min

The design condition has $F_c = 400\,\mathrm{t/h}$, $c_{p,c} = 2.2\,\mathrm{MJ/t/K}$, $T_{SP} = 360\,^\circ\mathrm{C}$, $UA_{tl} = 8\,\mathrm{MJ/h/K}$, $T_{air} = 25\,^\circ\mathrm{C}$, $Q_{fur,max} = 146\,000\,\mathrm{MJ/h}$ and a furnace inlet of 228.14 °C.

(a) Use Equation (3.10) to find the coil outlet the column needs, the transfer line loss and the starting valve position. *Checkpoint:* 363.05 °C, about 0.74 MW and 81.3 %.

(b) Find the steady gain from valve position to column inlet. *Checkpoint:* 1.644 °C per percent.

(c) In scenario 11, averaged over hours 12 to 24, the valve sits at 78.91 % with the thermocouple bias and at 82.56 % without it. Find the bias in degrees, and say where the true column inlet is. *Checkpoint:* 6.0 °C, true inlet near 354 °C.

3.6 Theory and calculation · 6 min

Before the E-101 outage of scenario 5, E-101 recovers 22.87 MW, E-102 26.79 MW and the furnace 32.98 MW of its 40.56 MW. At hour 6, with E-101 bypassed and the valve on its stop, E-102 recovers 37.64 MW.

(a) Find the duty the train loses, the share the furnace covers and the shortfall left over. *Checkpoint:* 12.02 MW, 7.58 MW and 4.44 MW.

(b) Find the column inlet the shortfall leaves at hour 6, with the flowing heat capacity and the transfer line loss coefficient together 888 MJ/h/K. *Checkpoint:* 18.0 °C short, about 342 °C.

(c) In the E-102 outage the crude reaches the coil at 124.6 °C. Find the charge rate at which the full furnace holds 360 °C if the furnace inlet stayed at that value, and compare it with the 295 t/h the override settles at. *Checkpoint:* about 277 t/h.

3.7 App challenge · 8 min

Load scenario 4, “Three weeks of fouling, nothing else”, at the default seed, set the speed to 100 and the trend window to 504 h.

(a) Run the full 504 h. Read the fuel valve position at 504 h (CSV column `fuel_valve_pct`), the Fuel fired tile, and the hour at which E-101 first reaches 1.0 on the fouling trend (column `foul_e101_frac`). *Checkpoint:* 97.5 %, 76.7 TJ and 239.6 h (±0.1 h).

(b) Press Reset, set Crude charge demand to 420 t/h before Run, and run the 504 h again. Find the hour the fuel valve first reaches its stop (FURNACE AT FIRING LIMIT on the trend, or the first 1 in `fuel_saturated`) and the column inlet at 504 h. *Checkpoint:* 215.2 h (±0.1 h) and 352.1 °C (±0.1 °C).

(c) In the 420 t/h run E-101 reaches its threshold at 290.5 h, later than in (a). Use Equation (3.2) to explain why, and why the valve still reaches its stop sooner. *Checkpoint:* $v = 1.05$: asymptote 0.514, time constant 400 h.

3.8 App challenge · 9 min

Compare the E-101 outage without and with the override, at the default seed, running each case 48 h at speed 10. Read the lowest column inlet from the trend or from the minimum of the CSV column `col_inlet_C`, and the crude processed from `crude_total_t` at 48 h.

(a) Run scenario 5, "Clean E-101 at full rate, with feedforward only". Read the lowest column inlet, the crude processed and the At firing limit tile. *Checkpoint:* 340.3 °C, 19 200 t and 33.3 % of the run.

(b) Run scenario 6, "The same clean under override control". Read the same values, the Rate cut tile and the lowest charge rate (`crude_rate_tph`). What did holding the temperature cost? *Checkpoint:* 358.6 °C, 18 645 t (±2 t), 363 t/h, so 555 t or 2.9 %.

(c) Press Reset, set Valve position target (Advanced) to 85 % before Run, and run again. Then do the same at 100 %. Which of the three targets would you commission? *Checkpoint:* 358.8 °C and 18 279 t at 85 %. At 100 % the run repeats (a).

3.9 App challenge · 8 min

Load scenario 9, "The same two cleanings, colliding", at the default seed, set the speed to 100 and run the full 504 h.

(a) Read the scorecard IAE, the lowest charge rate (`crude_rate_tph`), the lowest column inlet (`col_inlet_C`) and the crude processed (`crude_total_t`) at 504 h. *Checkpoint:* 559.1 °C h, 240 t/h, 262.6 °C and 199 056 t.

(b) Press Reset, set Cleaning duration in Show Model, Parameters, to 8 h before Run, and run the 504 h again. Read the same four values. *Checkpoint:* 13.9 °C h, 290.5 t/h, 344.9 °C and 200 315 t.

(c) The cleans are scripted at hours 240 and 248. Find how long the two outages overlap in each run, and explain why the integrated error falls by a factor of about forty. *Checkpoint:* 8 h of overlap, then none.

3.10 Application area · 4 min

Scenario 8 cleans both exchangers at their thresholds and processes 199.1 kt with 72.8 TJ of fuel. Scenario 4, the same three weeks with no cleaning, processes 201.6 kt with 76.7 TJ. Fuel gas costs $6/GJ.

(a) Find the fuel the two cleans save and its value. *Checkpoint:* 3.9 TJ, about $23k.

(b) Find the crude margin per tonne at which the fuel saved just pays for the crude given up. *Checkpoint:* about $9.4 per tonne.

(c) Name two things this break-even leaves out that would change the decision on a real unit.

3.11 Application area · 5 min

The model cleans perfectly. Mechanical cleaning of a real exchanger typically recovers 85 to 95 % of the clean coefficient.

(a) E-101 ($U_{clean,1} = 1.26$, threshold 0.42) is cleaned to 90 % of its clean coefficient. Find the residual resistance and its share of the threshold. *Checkpoint:* 0.088 and 21 %.

(b) Find the hours from that clean to the next threshold at the design rate ($R_{f\infty,1} = 0.55$, $\tau_{f,1} = 420\,\text{h}$), against 606 h from a perfect clean. *Checkpoint:* about 532 h.

(c) What does this do to a cleaning schedule over several runs, and to the chance that two cycles collide?

3.12 Review the agent · 5 min

After running scenario 13 for 48 h, an AI assistant writes: “The MPC processed 17.703 kt against the override’s 17.616 kt and held the column inlet near 357 °C where the override dipped to 317.5 °C. The optimization finds 87 t the override cannot. Replace VPC-101 with the MPC, and for safety keep the override’s 5 % of valve in hand by setting the plan’s valve ceiling to 95 %.” Find the flaw, and say how you would check it in the app.