INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 8

A Diesel Hydrotreater With Two Inlet Handles and a Half-Hour Sulfur Analyzer

Pages 184–207

Edited by John D. Hedengren

How to cite this chapter. "A Diesel Hydrotreater With Two Inlet Handles and a Half-Hour Sulfur Analyzer," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 184–207.

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

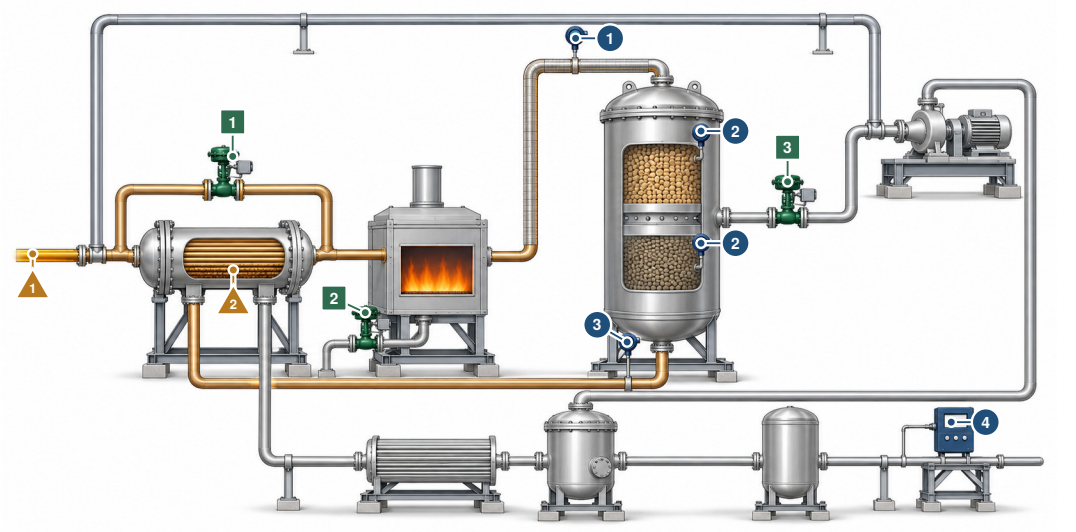
The app for this chapter
https://apmonitor.com/apps/ulsd-hydrotreater/

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

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

CHAPTER 8

A Diesel Hydrotreater With Two Inlet Handles and a Half-Hour Sulfur Analyzer



The controller moves 1 valve that lets feed bypass the exchanger 2 fuel valve of the charge heater 3 cold gas valve between the beds

The controller reads 1 reactor inlet temperature 2 temperature of each bed 3 reactor outlet temperature 4 product sulfur, half an hour late

What disturbs it 1 a sourer feed 2 a fouled exchanger ▲ cracked stock in the feed

The hot oil leaving the bottom of the reactor runs back to the exchanger at the left to heat the feed, and only then through the cooler and two drums to the sulfur analyzer at the lower right. Illustration, not to scale.

IN PLAIN TERMS A Reactor Graded Half an Hour Late

The main idea. A diesel hydrotreater removes sulfur from diesel by passing it with hydrogen over hot grains that help the two react. The hot oil leaving the reactor hands most of its heat to the cold oil coming in.

The challenge. The heat goes round in a loop, so a small change in temperature comes back and grows for hours. The instrument that tests the product for sulfur reports thirty minutes late, when the oil it tested has long left.

Why it matters. Road diesel may hold only a few parts of sulfur in a million, and running hotter to be safe wears out the grains sooner.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	This page, the History boxes in Sections 8.2 and 8.5, the Did you know? boxes in Sections 8.1 and 8.6, the Try it boxes in Sections 8.4 and 8.6, and Key points in Section 8.9. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 8.6 with Figure 8.3, Section 8.7 and Section 8.8 with its Open challenge box.

The control question A diesel hydrotreater heats its feed with its own effluent and learns its product sulfur from an analyzer half an hour late. Which handle should hold the reactor inlet, and what can a sulfur loop, a Smith predictor or a predictive controller do before the analyzer reports?

Learning objectives

- Compute the loop gain of the feed and effluent exchanger and explain why a heater step keeps moving the inlet for hours.
- Explain the wrong-way rise of the reactor outlet after an inlet cut.
- Design mid-ranging so the fast bypass holds the inlet while the slow heater returns it to mid-range.
- Compare a PI, a Smith predictor and an MPC on a late analyzer, and state what a wrong predictor dead time costs.
- Judge a sulfur setpoint back-off against catalyst life and bed-rise limits.

Prerequisites Sections 1.4.4, 1.4.5, 1.4.6, 1.6.2 and 1.6.3. Chapter 7 for beds with interbed quench.

The app https://apmonitor.com/apps/ulsd-hydrotreater/

See also Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging and valve-position control. Volume 5, Chapter 5 for a Smith-type predictor and Volume 5, Chapter 3 for a loop on a slow analyzer. Volume 7, Chapter 2 for an exothermic reactor near instability.

8.1 Introduction

A diesel hydrotreater must remove nearly all of the sulfur it is fed. In the generic unit simulated here, 200 t/h of diesel at 0.8 wt% sulfur leaves the reactor with 8.0 wppm (parts per million by weight), so 999 of every thousand sulfur atoms are taken out. Most of the heat the feed needs comes back from the reactor itself, since the feed–effluent exchanger returns 21.6 MW to the charge against the 7.34 MW fired in the charge heater. The product, finally, is judged by an online analyzer whose every report describes oil that left the reactor 30 min earlier.

The first of those facts is set by the regulations that created ultra-low-sulfur diesel (ULSD). In the United States the highway diesel limit fell from the 500 ppm in force since 1993 to 15 ppm for refiners from June 1, 2006 [1], and in the European Union diesel of at most 10 mg/kg had to be marketed from January 1, 2005 and to be the only diesel sold from January 1, 2009 [2]. In the simulator the room between the 8 wppm target and the 10 wppm limit is 2 wppm, and the evidence that the product is inside it arrives 30 min late.

> ✦ DID YOU KNOW? **The Largest Process After Crude Distillation**
>
> On January 1, 2003 the refineries of the United States could hydrotreat 11,247,745 barrels per calendar day, 67.7 % of their crude distillation capacity and more than any other major process downstream of the crude unit [3]. The unit of this chapter treats 200 t/h of diesel.

Three beliefs that a block diagram never challenges are corrected by the simulator. The first is that a fired heater moves the reactor inlet by what it adds. In case 1, 1 MW more in the coil is worth 5.36 K at once, yet the inlet ends 12.83 K higher after 400 min, 2.40 times as much, because the warmer effluent preheats the feed. The second is that cutting the inlet cools the reactor. In case 2 a 10 K cut raises the reactor outlet by 1.10 K for 24.5 min before it falls 13.3 K. The third is that a dead-time compensator removes the analyzer's delay. In case 4 the Smith predictor saves about a fifth of the integrated error against the best proportional-integral (PI) tuning found, 243 against 298 wppm min over 480 min, none of it before the first report, and with its assumed delay 15 min too long it costs 699 wppm min.

Each behavior belongs to a class. An adiabatic reactor preheated by its own effluent closes a positive energy loop whose design and control were analyzed early [4, 5], and whose bypass and furnace give two handles of different speed on one temperature, the pair that dynamic controllability

favors [6]. A packed bed whose composition answers faster than its catalyst temperature shows a transient of the wrong sign when its feed is cooled [7]. A fast valve of small range working beside a slow actuator of large range is the setting of mid-ranging control [8, 9], and a quality loop closed on a late analyzer is the setting of the Smith predictor [10]. A catalyst that deactivates in service [11] adds a far slower time scale. The same shapes recur in any reactor with heat recovery, in any trim valve beside a large actuator, and in any loop whose measurement is a laboratory or an analyzer.

These pieces usually reach a course separately: a heat-integrated reactor as a stability analysis, mid-ranging as a block diagram, dead time as a delay inside a transfer function with a perfectly known output. This chapter studies all three in the plant where they meet, a simulated ULSD hydrotreater in which the bed temperatures, the exchanger, the heater, the hydrogen loop, the sulfur analyzer and the catalyst activity are computed together. It measures how an inlet loop, mid-ranging, a quench override, a bed temperature loop and three sulfur controllers hold a 10 wppm specification. The lumped kinetics are teaching values and the controllers are established practice. Two of the measurements set bounds on the others. The wrong-way rise grows by about 0.4 K with every doubling of the cells and has not converged at 192 cells per bed, and the graded day needs a sulfur setpoint back-off once an expanded noise sweep exposes a missed specification. The model has sixty-four states whose sulfur, hydrogen and energy balances close over disturbed runs. It is the second user of the series' shared packed-bed cell module and sits behind a measurement boundary that keeps the true sulfur away from every controller. Nine cases each carry a completion rule.

Section 8.2 traces hydrotreating from cobalt molybdate to ten parts per million, Section 8.3 builds the model, Section 8.4 examines the plant without feedback, Section 8.5 sets out the controllers, and Section 8.6 reports the nine cases, the graded day among them. Section 8.7 sets the hardware of a real unit beside the model, Section 8.8 states what the model leaves out, and Section 8.9 draws the results together.

8.2 From Cobalt Molybdate to Ten Parts per Million

The catalyst came first. In 1943 Byrns, Bradley and Lee described the catalytic desulfurization of gasolines by a cobalt molybdate process [12], and four decades later the catalysts used for hydrodesulfurization and the related hydrotreating reactions were described as molybdenum, or tungsten, on high-surface-area alumina with cobalt or nickel added as promoters, although no general agreement had yet been reached on the structures present in the active sulfided catalyst, the Co–Mo–S type structures of that review's title, or on the origin of promotion [13]. By 1996 the sales of hydrotreating catalysts, which serve the catalytic hydrogenation processes that saturate unsaturated hydrocarbons and remove sulfur, nitrogen, oxygen and metals from refinery streams, made up close to 10 % of the world market for catalysts [14].

The specification came from engines, and the History box below gives its dates.

> ⧖ HISTORY **2001: The Rule That Set Fifteen Parts per Million**
>
> The United States Environmental Protection Agency set fuel quality standards that limited sulfur and aromatics in highway diesel beginning in 1993, at 500 ppm. Its rule of January 2001 tied new heavy-duty engine standards from model year 2007 to exhaust catalysts that sulfur poisons, and so required refiners to produce highway diesel with at most 15 ppm of sulfur beginning June 1, 2006, a reduction of 97 %, with retail stations supplied by September 1, 2006. A temporary option let up to 20 % of highway diesel be produced at 500 ppm through 2009 [1]. Directive 2003/17/EC required each member state of the European Union to market diesel of at most 10 mg/kg by January 1, 2005, and to allow only such diesel from January 1, 2009, a date the directive asked to be confirmed by the end of 2005 [2].

The hardware followed. The agency projected that every refiner could meet 15 ppm by extending the conventional hydrotreating it already ran, which would likely mean a second stage with an additional reactor and supporting hydrogen, sulfur recovery, amine and sour water units, and it

reported that the vendors' projections of hydrogen consumption and reactor volume varied widely [1]. For the deep desulfurization of gas oils, the apparent reaction order and the role of hydrogen sulfide carry consequences for the configuration of an industrial reactor [15]. An order above one for a lumped mixture of first-order species is the subject of Ho and Aris's note on apparent second-order kinetics [16], and the simulator borrows that reading (Section 8.3.1). Hydroprocessing catalysts also deactivate in service [11], and in the simulator the temperature needed for the same product climbs until a limit ends the cycle.

The control of heat-integrated reactors has its own history. The design and control of feed–effluent exchanger and reactor systems was analyzed in 1962 [4], and in 1982 the effect of design on the stability and control of fixed-bed reactors with heat feedback was set out as a set of concepts [5]. Later work on an exchanger, furnace and adiabatic reactor in series noted that the positive feedback of energy can make the process open-loop unstable, that the reactor inlet is usually held by bypassing the exchanger, and that dynamic controllability favors using both the bypass and the furnace, particularly when the reactor gain $\Delta T_{out}/\Delta T_{in}$ is large [6]. A pseudo-homogeneous plug-flow analysis in 1981 named the wrong-way behavior, a transient rise of temperature after a sudden reduction of the feed temperature, and found its size set by the adiabatic temperature rise, the activation energy, the size of the drop and the reactor length [7]. Axial dispersion of heat was later shown to decrease that excursion and prolong it [17], a result that matters here because the cells of the simulator are themselves a dispersion (Section 8.4).

The analyzer's delay has the oldest answer, the predictor that carries the name of Smith, from his 1957 paper on closer control of loops with dead time [10]. Mid-ranging, controlling one variable with a fast and a slow manipulated variable while the slow one returns the fast one toward a target, has its own design analysis [8] and is regarded as well established, with bumpless transfer and saturation the practical issues in a distributed control system [9]. Above those loops, model predictive control became the industrial standard, with refining among its main fields of application [18].

8.3 The Unit as Simulated

The simulated unit is a generic 200 t/h diesel hydrotreater at 60 bar, as shown in Figure 8.1. Diesel at $T_f = 180\,^{\circ}\mathrm{C}$ joins treat gas (recycle gas rich in hydrogen) at 4.0 mol/s per t/h, about 275 $\mathrm{Nm^3/m^3}$ at 850 $\mathrm{kg/m^3}$. The combined feed splits between the tubes of the feed–effluent exchanger and a bypass whose valve opening b sends a fraction $f = b f_{max}$ of it around the exchanger, with $f_{max} = 0.30$. The two streams mix, a fired charge heater raises the mix to the reactor inlet temperature (RIT), and the oil and gas flow down through bed 1 (70 t of catalyst), receive cold recycle gas as quench, and pass bed 2 (130 t). The effluent returns through the exchanger's shell, an air cooler and a separator, whose gas is purged, made up with hydrogen and recompressed as treat gas and quench, and whose liquid is stripped to product. Time in the engine runs in minutes, while bed flows are in mol/s and duties in MW. The engine, its clocks and its seeded noise are those of Section 1.6.1.

The variables are classified in Table 8.1. The state vector holds sixty-four quantities in use, the 48 cell temperatures, eleven other physical states (the heater duty Q, the coil outlet RIT, the exchanger outlet T_X, the bypass and quench valve positions, the recycle purity y, the feed sulfur w_{S0} and cracked-stock share x_{LCO} reaching the reactor, the feed rate F and the two bed activities), and five ledgers that integrate catalyst age, feed, product sulfur, fuel and hydrogen consumed.

8.3.1 Two Lumped Reactions

Hundreds of sulfur compounds and aromatic rings are folded into two reactions with teaching constants. Hydrodesulfurization (HDS), $\mathrm{S} + 3\,\mathrm{H_2} \rightarrow \mathrm{H_2S}$, consumes organic sulfur in the liquid, and the saturation of aromatics, $\mathrm{A} + 3\,\mathrm{H_2} \rightleftharpoons \mathrm{N}$, turns saturable aromatics into naphthenes up to an

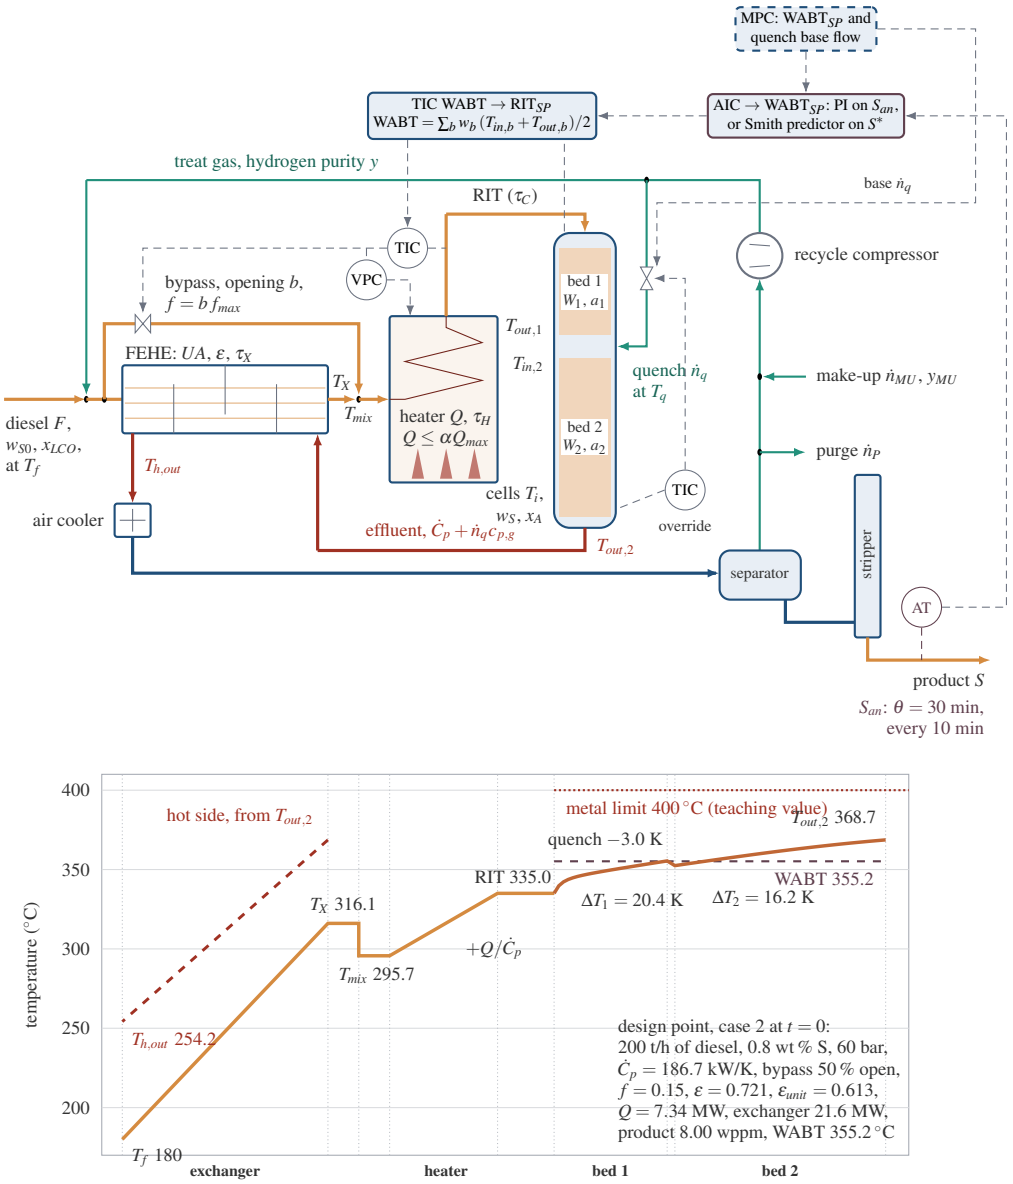


Figure 8.1 ULSD hydrotreater as simulated, drawn after the application's wide Unit view. Above, the loop. Feed at T_f and treat gas of purity y split between the exchanger's tubes and the bypass ($f = b\,f_{max}$), mix at T_{mix}, gain Q in the heater and reach the reactor at RIT. Each bed is 24 cells at T_i with activity a_b, the quench $\dot{n}_q$ at T_q joins between them, and the effluent at $T_{out,2}$ heats the feed and leaves the shell at $T_{h,out}$. The TIC on RIT moves the bypass, the VPC moves the heater to bring the bypass back to its target, the override TIC adds quench on $T_{out,2}$, the WABT controller sets the RIT setpoint, the AIC (a PI on the analyzer reading S_{an} or the Smith predictor on S^*) sets the WABT setpoint, and the MPC (dashed) can set the WABT setpoint and the quench base flow. Below, on the same stations, the design temperatures drawn to scale, with the 48 cell temperatures, the hot side dashed and the 400 °C metal limit dotted. It is a drawing of the model and not of piping or safeguards.

Table 8.1 Hydrotreater Variables, Their Roles and Design Values (Case 2 at $t = 0$)

Variable	Role	Design	Range
S	CV, product sulfur (true)	8.00 wppm	SP 4–15
S_{an}, S^*	measured, estimated sulfur	8.00 wppm	—
WABT	CV and MV, bed temperature	355.2 °C	SP 320–400
RIT	CV and MV, reactor inlet	335.0 °C	SP 300–395
$T_{out,2}$	CV of the override	368.7 °C	SP 360–400
b	MV, bypass valve	50 % open	0–100 %
Q	MV, heater duty	7.34 MW	1–14
$\dot{n}_q$	MV, quench gas	60 mol/s	0–400
F	DV, feed rate	200 t/h	150–240
$w_{S,f}$	DV, feed sulfur (tank)	0.8 wt%	0.3–1.6
x_{LCO}	DV, light cycle oil	0	0–40 %
ϕ	DV, exchanger UA	100 % of clean	40–100 %
α	DV, heater available	100 %	30–100 %
y_{MU}, bias	DV, make-up purity, analyzer	97 mol%, 0	85–100, ±3
T_i	states, 48 cells	Figure 8.1	—
T_X, T_{mix}	exchanger outlet, mix	316.1, 295.7 °C	—
ΔT_1, ΔT_2	bed rises	20.4, 16.2 K	limit 30
y	state, recycle purity	87.3 mol%	—
a_1, a_2	states, activity	1.000, 1.000	—

equilibrium. Equation (8.1) gives both rates on a catalyst-mass basis.

$$\begin{aligned} r_S &= a\,k_S\,\kappa_S\,e^{-\gamma_L x_{LCO}}\,\frac{w_{S0}}{w_r}\left(\frac{w_S}{w_{S0}}\right)^n\frac{p_H/p_r}{1+K_{H_2S}\,p_{H_2S}} \\ r_A &= a\,k_A\,\kappa_A\,\frac{p_H}{p_r}\left(x_A-\frac{x_N}{K_A\,p_H^3}\right) \\ K_A &= \exp\left(\Delta S_A/R-\Delta H_A/RT\right) \end{aligned} \tag{8.1}$$

where $\kappa_j = \exp[-E_j/R\,(1/T - 1/T_r)]$ with $T_r = 350\,°\mathrm{C}$, a is the bed's activity, w_S the sulfur mass fraction of the liquid and w_{S0} its value at the reactor inlet, x_A and x_N the mole fractions of aromatics and naphthenes in the liquid, $w_r = 10^{-3}$, $p_r = 45\,\mathrm{bar}$, and p_H and p_{H_2S} are partial pressures, P times each flow over the gas, which counts hydrogen, H_2S, the light impurities of the recycle and a fixed vaporized share of 0.30 of the oil's moles. Light cycle oil carries the more hindered sulfur, through $e^{-\gamma_L x_{LCO}}$ with $\gamma_L = 2$, and more saturable aromatics, 0.65 against 0.25 mole fraction for straight-run diesel. The same ΔH_A enters the equilibrium and the energy balance.

The order $n = 1.65$ is written on the fraction of the feed's sulfur left, not on the absolute sulfur, and the reason is the reading of an order above one. For a mixture of first-order species whose rate constants k follow a gamma distribution of shape m and mean $\bar{k}$, the fraction left after a contact time τ is as shown in Equation (8.2).

$$\begin{aligned} X(\tau) &= \int_0^\infty g(k)\,e^{-k\tau}\,dk = \left(1+\frac{\bar{k}\tau}{m}\right)^{-m} \\ \frac{dX}{d\tau} &= -\bar{k}\,X^{1+1/m} \end{aligned} \tag{8.2}$$

which is an n-th order lump with $n = 1 + 1/m$, here $m = 1.54$, and an exponential distribution ($m = 1$) gives the apparent second order of Ho and Aris [16]. The fraction left depends on the reactor and not on the feed's level, so a sourer feed makes a proportionally sourer product. With the order written on the absolute sulfur and calibrated to the same rate at 0.8 wt%, a feed 50 % sourer came out cleaner at a fixed 335 °C inlet, 7.32 wppm against 8.00, where the fraction form gives 10.92 wppm and needs a 5.1 K hotter inlet. At design the beds release 6.87 MW, 5.07 MW of it from saturation and 1.80 MW from HDS, and bed 1 removes 99.17 % of the sulfur. The equilibrium saturates 92, 61 and 38 % of

the aromatics at 335, 368 and 385 °C under 45 bar of hydrogen, so a hotter late-cycle bed releases less heat.

8.3.2 Thermal Cells With Algebraic Composition

Each bed is a train of 24 cells of equal catalyst mass from the shared packed-bed module of the series, one train per bed so the quench can add gas and heat capacity between them. The liquid crosses a bed in minutes while the catalyst holds its heat for tens of minutes, so each cell's composition is algebraic and its temperature is a state, as shown in Equation (8.3).

$$\begin{aligned} F_i &= F_{i-1} + \nu^{\mathsf{T}}\xi_i, \qquad \xi_{i,j} = W_c\, r_j(T_i, F_i) \\ C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} &= \dot{C}_{p,b}\,(T_{i-1} - T_i) - \textstyle\sum_j \xi_{i,j}\,\Delta H_j \end{aligned} \tag{8.3}$$

where F holds the molar flows of S, H_2S, A, N and H_2, ν is the stoichiometric matrix, $W_c = W_b/24$, $C_i = W_c c_{cat} + C_{v,b}/24$ the catalyst plus a share of the liquid held up and the vessel, T_0 the bed inlet, and $\dot{C}_{p,b}$ the heat-capacity flow, held along a bed, $\dot{C}_p = \dot{m}c_{p,L} + \dot{n}_g c_{p,g} = 186.7\,\mathrm{kW/K}$ in bed 1 and $\dot{C}_{p,2} = \dot{C}_p + \dot{n}_q c_{p,g} = 188.6\,\mathrm{kW/K}$ in bed 2. The bed heat capacities are 114.2 and 212.2 MJ/K, thermal time constants of 10.19 and 18.75 min. Each cell is solved by Newton's method on the extents with a fixed four iterations from zero extent, converged to 2.6×10^{-9} of a forty-iteration solution from 300 to 395 °C, so the derivative depends on the state alone and every run replays bit for bit.

8.3.3 Quench and the Weighted Average Bed Temperature

The quench is recycle gas at $T_q = 65\,^{\circ}\mathrm{C}$, so bed 2's inlet and the weighted average bed temperature (WABT) are as shown in Equation (8.4).

$$\begin{aligned} T_{in,2} &= \frac{\dot{C}_p\, T_{out,1} + \dot{n}_q c_{p,g} T_q}{\dot{C}_p + \dot{n}_q c_{p,g}} \\ \mathrm{WABT} &= \textstyle\sum_{b=1}^{2} w_b\,(T_{in,b} + T_{out,b})/2 \end{aligned} \tag{8.4}$$

where $w = (0.35, 0.65)$ are the catalyst shares and $T_{in,1} = \mathrm{RIT}$. At design 60 mol/s of quench (4.8 $\mathrm{kNm^3/h}$) cools bed 2's inlet by 2.96 K, and the valve passes at most 400 mol/s.

8.3.4 Exchanger, Bypass and Charge Heater

The exchanger is counter-current, with the cold flow $(1-f)\dot{C}_p$ in the tubes and the effluent $\dot{C}_{p,2}$ in the shell. Its effectiveness, its cold outlet through a lag for the metal and holdup, and the mix are as shown in Equation (8.5).

$$\begin{aligned} \varepsilon &= \frac{1 - e^{-N_t(1-C_r)}}{1 - C_r e^{-N_t(1-C_r)}}, \qquad N_t = \frac{\phi\, UA}{C_{min}} \\ q_X &= \varepsilon\, C_{min}\,(T_{out,2} - T_f) \\ \tau_X \frac{\mathrm{d}T_X}{\mathrm{d}t} &= T_f + \frac{q_X}{(1-f)\dot{C}_p} - T_X \\ T_{mix} &= (1-f)\,T_X + f\,T_f, \qquad T_{h,out} = T_{out,2} - \frac{q_X}{\dot{C}_{p,2}} \end{aligned} \tag{8.5}$$

where $C_r = C_{min}/C_{max}$, $UA = 344\,\mathrm{kW/K}$ clean, ϕ the fouling factor, q_X the duty returned to the feed and $\tau_X = 3\,\mathrm{min}$. At design with the bypass half open ($f = 0.15$), $\varepsilon = 0.7212$, the exchanger returns 21.60 MW, the mix is at 295.67 °C and the effluent leaves at 254.20 °C. A first-order lag carries the heater's command to its absorbed duty, and a second lag, the coil, carries the duty to the inlet, as shown in Equation (8.6).

$$\tau_H \frac{\mathrm{d}Q}{\mathrm{d}t} = \mathrm{sat}(Q_c) - Q, \qquad \tau_C \frac{\mathrm{dRIT}}{\mathrm{d}t} = T_{mix} + \frac{Q}{\dot{C}_p} - \mathrm{RIT} \tag{8.6}$$

where sat clamps the command Q_c between $\min(Q_{min}, \alpha Q_{max})$ and αQ_{max}, with $Q_{max} = 14\,\text{MW}$, $Q_{min} = 1\,\text{MW}$, α the share of burners in service, $\tau_H = 6\,\text{min}$ and $\tau_C = 1\,\text{min}$. The bypass and quench valves follow their commands through lags of 0.1 and 0.15 min. The design duty is 7.343 MW.

8.3.5 Recycle Hydrogen

The hydrogen loop is one well-mixed inventory whose pressure is held, so the make-up replaces every mole that leaves, as shown in Equation (8.7).

$$\begin{aligned} N_L \frac{\mathrm{d}x}{\mathrm{d}t} &= \dot{n}_{MU}(1 - y_{MU}) - \dot{n}_P x, \qquad x = 1 - y \\ \dot{n}_{MU} &= \dot{n}_{H_2,used} + \dot{n}_P \end{aligned} \tag{8.7}$$

where $\dot{n}_P = 36.4\,\text{mol/s}$ is the purge, $N_L/\dot{n}_P = 60\,\text{min}$, and $\dot{n}_{H_2,used}$ the hydrogen the beds consume. At design the beds use 117.6 mol/s, the make-up is 12.4 $\text{kNm}^3\text{/h}$ and the purity 87.3 mol%. The recycle carries no H_2S, as if scrubbed.

8.3.6 Feed Line, Activity and a Compressed Cycle

The feed's sulfur and cracked-stock share reach the reactor through a 20 min surge drum, $\tau_{feed}\,\dot{w}_{S0} = w_{S,f} - w_{S0}$ with $w_{S,f}$ the tank's sulfur, and its rate through a 1 min lag. Each bed loses activity at its mean temperature $\bar{T}_b$, as shown in Equation (8.8).

$$\frac{\mathrm{d}a_b}{\mathrm{d}t} = -k_d \exp\!\left[-\frac{E_d}{R}\left(\frac{1}{\bar{T}_b} - \frac{1}{628.15\,\text{K}}\right)\right](1 + \beta x_{LCO})\,\frac{c}{43\,200} \tag{8.8}$$

per minute, where $k_d = 0.019$ per month at 355 °C, $E_d = 100\,\text{kJ/mol}$, $\beta = 2$, 43 200 min make a month, and c is a declared cycle compression, 1 in every case but the cycle case, where $c = 720$ passes a month of catalyst life every simulated hour. Five kelvin hotter ages the catalyst 16 % faster and 20 % of light cycle oil 40 % faster. The run ends when the hottest catalyst reaches the 400 °C metal limit.

8.3.7 What the Controllers Can Read

All the controllers share one packet of measurements, refreshed at each 10 s execution. The inlet, bed and outlet thermocouples lag 6 s with 0.1 K of noise, the duty, feed and quench meters carry 0.3, 0.2 and 1 %, the recycle purity analyzer lags 1 min, and the assay of the feed tank lined up supplies its sulfur and its cracked-stock share. The product sulfur analyzer draws a sample every 10 min and reports it 30 min after the sample left the reactor, a 20 min product path and sample line followed by a 10 min analysis, with 0.1 wppm of noise and a calibration bias. It is the sampled, delayed analyzer of Section 1.4.6. Out of service, its reading holds and is marked stale after two missed reports. The true sulfur, the cell profiles and the activities never enter the packet, and each execution draws the same normal deviates whatever the settings, so runs differing in a single setting meet identical noise.

8.3.8 Parameters, Initial State and Integration

Table 8.2 lists the parameters. Each case starts lined out, the inlet at its setpoint (or, with a sulfur loop in Auto, at the inlet that holds the sulfur setpoint), the bypass at its 50 % target, the heater at the duty that holds the inlet, the quench at its base flow and the purity at the fixed point of Equation (8.7). A cell profile is lined out bed by bed, each cell bracketed from its inlet and bisected. The state advances by fourth-order Runge–Kutta steps of 5 s, the controllers execute every 10 s, and a row is logged every 30 s. Halving the step leaves the wrong-way rise of case 2 at 1.1035 K and the true sulfur of case 4 at 240 min at 8.056 14 wppm. Over four disturbed runs the sulfur entering the reactor equals the product plus the H_2S made to better than 10^{-12} of itself, the hydrogen loop closes to round-off, and the energy around the exchanger, heater, beds and quench closes to 3×10^{-7} of the heat entering, with the stored heat of the beds, coil and exchanger metal counted.

Table 8.2 Parameters of the Simulated Hydrotreater by Access (A Adjustable in Advanced, H Hidden, F Written Into the Source)

Symbol	Meaning	Value	Unit		Symbol	Meaning	Value	Unit	
P	reactor pressure	60	bar	A	UA	exchanger, clean	344	kW/K	H
W_1, W_2	catalyst per bed	70 000, 130 000	kg	H	f_{max}	bypass share at full open	0.30	—	A
c_{cat}	catalyst heat capacity	1000	J/(kg K)	H	τ_X	exchanger metal lag	3	min	H
$C_{v,1}, C_{v,2}$	holdup and vessel per bed	44.2, 82.2	MJ/K	H	τ_V, τ_q	bypass, quench valve lags	0.1, 0.15	min	H
N_c	cells per bed	24	—	H	Q_{max}, Q_{min}	heater rated, minimum	14, 1	MW	A, H
k_S	HDS at 350 °C	1.104×10^{-3}	mol/(s kg)	H	τ_H, τ_C	firebox, coil	6, 1	min	A, H
E_S, n	HDS activation, order	120, 1.65	kJ/mol, —	H	η_H	heater efficiency	0.85	—	H
K_{H_2S}	H_2S inhibition	0.2	1/bar	H	T_f	combined feed	180	°C	H
γ_L	LCO sulfur refractoriness	2	—	H	T_q	quench gas	65	°C	H
ΔH_S	heat of HDS	−130	kJ/mol S	H	τ_{feed}	surge drum	20	min	A
k_A	saturation at 350 °C	6.364×10^{-4}	mol/(s kg)	H	$\dot{n}_P$, $N_L/\dot{n}_P$	purge, loop time	36.4, 60	mol/s, min	H
$E_A, \Delta H_A$	saturation activation, heat	80, −200	kJ/mol	H	θ, period	analyzer delay, period	30, 10	min	A
ΔS_A	saturation entropy	−403	J/(mol K)	H	k_d	activity lost at 355 °C	0.019	1/month	H
$M_L, c_{p,L}$	diesel molar mass, heat capacity	230, 2900	g/mol, J/(kg K)	H	E_d, β	deactivation activation, LCO	100, 2	kJ/mol, —	H
$c_{p,g}$	recycle gas heat capacity	32	J/(mol K)	H	age, c	starting age, compression	0, 1	months, —	A
—	vaporized oil share	0.30	mol/mol	H	ΔT_{max}	bed rise limit	30	K	A
—	treat gas to oil	4.0	mol/s per t/h	H	T_{metal}	metal limit	400	°C	A
x_A	saturable aromatics, SR, LCO	0.25, 0.65	mol/mol	H	w_b	catalyst shares	0.35, 0.65	—	F
$\dot{n}_{q,max}$	quench valve capacity	400	mol/s	H	T_r, w_r, p_r	reference point	350, 10^{-3}, 45	°C, —, bar	F

8.3.9 Assumptions

Two lumped reactions stand for hundreds, so olefin saturation, denitrogenation and mild cracking are folded into the HDS heat, and the constants are teaching values chosen to give 8 wppm at a 335 °C inlet with rises of 20.4 and 16.2 K, not fitted to a catalyst. Each bed is 24 stirred cells, which act as axial dispersion with a Péclet number near 48 per bed, so the wrong-way rise depends on the cell count (Section 8.4). The liquid and gas hold no inventory, the reactor has one pressure and no pressure drop, the wetting and maldistribution of a trickle bed are absent, and the vaporized share of the oil is fixed. The exchanger has one lag, and its stored heat is not a fixed capacity when the bypass moves. The heater has no tube metal, draft or fuel system, the separator, cooler, scrubber and stripper are not modeled beyond the analyzer's delay, and the surge drum is a fixed 20 min lag whatever the flow, which leaves 0.026 % of the sulfur unbalanced in the drum when the feed rate steps. Activity is one number per bed with no start-of-run decline, metals or nitrogen poisoning. Sulfur in wppm is reckoned on the feed's liquid mass, which the product undercuts by 0.42 %, so the design's 8.00 wppm would read 8.03 on the product's own mass. The bed rise limit of 30 K, the 400 °C metal limit and the 14 MW rating are teaching values, where real limits come from the licensor's operating guidelines, the reactor's design temperature under the ASME pressure vessel code and the heater's design.

8.4 Three Behaviors Before Any Feedback

8.4.1 The Heat That Comes Back

The exchanger and its bypass return $\varepsilon_{unit} = (T_{mix} - T_f)/(T_{out,2} - T_f) = 0.613$ of the effluent's temperature above the feed, and at design a kelvin more at the inlet makes 1.187 K more at the outlet,

so a kelvin added at the inlet returns as 0.728 K and the series converges slowly. That return is the loop gain L of Equation (1.11), and Equation (1.8) gives the amplification $1/(1-L)$ it implies for a small move (Section 1.4.5). The gain of that loop moves with the operating point, 0.899 at a 320 °C inlet and 0.431 at 350 °C, because a hotter bed brings its aromatic saturation nearer equilibrium and releases less extra heat. In case 1 in Manual, 1 MW more firing at 10 min raises the inlet 5.23 K in the next 30 min, 95 % of that within 19 min, then 8.65, 11.69 and 12.83 K at 60, 120 and 400 min after the step. Opening the bypass from 50 to 60 % at 40 min takes the inlet from 340.23 to 337.67 °C in 2 min, 81 % of it in the first minute, because only the coil stands between the mix and the inlet. The two handles act on two clocks, and both are then multiplied by the loop. A unit with a larger exchanger or a hotter exotherm could pass a loop gain of one, where the positive feedback of energy makes the process open-loop unstable [6].

8.4.2 Wrong Way, and Its Size Under Refinement

Case 2 lowers the inlet setpoint from 335 to 325 °C at 10 min. The reactor outlet rises from 368.69 to 369.79 °C, peaks at 31.5 min, stays above its start from 10.5 to 35.0 min, and only then falls, to 355.9 °C at 60 min and 13.3 K below its start at 120 min. The first cells cool and react less, so unsaturated aromatics and sulfur reach catalyst still at its old temperature and release their heat there, while the cooling front moves at the pace of the catalyst's heat capacity. The mechanism was isolated with the beds alone. For an ideal 10 K inlet step at 24 cells the rise is 1.173 K whether the heat capacity is scaled by 1, 0.1 or 0.01, and only its time moves (20.7, 2.07 and 0.21 min), while with the inlet approaching over 2 min the rise shrinks with the heat capacity, 1.082, 0.800, 0.427, 0.159 and 0.057 K at scales of 1, 0.3, 0.1, 0.03 and 0.01. The aromatic saturation carries the effect, 0.685 K with its heat alone against 0.018 K with the HDS heat alone, because bed 1's first cells remove nearly all the sulfur in any case.

> **TRY IT IN THE APP** **2 · Cut the inlet 10 °C**
>
> Load the case and press Run. At 10 min lower the RIT setpoint from 335 to 325 °C, after Run and not before. Watch the Reactor outlet trend. The outlet first rises, from 368.69 to 369.79 °C at 31.5 min, and stays above its start until 35.0 min. By 120 min it is 13.3 K below its start.

The size of the rise has not converged with the cells. Case 2 as shipped gives 0.69, 1.10, 1.50, 1.88 and 2.24 K at 12, 24, 48, 96 and 192 cells per bed, about 0.4 K for each doubling, and the peak moves later, from 28.5 to 35.0 min between 12 and 96 cells, as shown in Figure 8.2a. Stirred cells in series are an axial dispersion, and dispersion is known to reduce the excursion [17], so the 24-cell figure is the size this model's dispersion allows, below what a plug-flow bed would show. The sign, the timing within a few minutes and the cause survive every refinement, while the size belongs to the dispersion the cells stand for. A real bed has its own dispersion and interphase transfer, which neither the cells nor the plug-flow limit represent. The converter beds of Volume 13, Chapter 4 show the same dependence of the wrong-way rise on the cell count.

Pause. Try Exercise 8.1 before reading on.

8.4.3 A Sourer Feed Brings Its Own Heat

At a fixed 335 °C inlet a feed of 1.2 wt% sulfur leaves 10.92 wppm, and 340.11 °C at the inlet is needed for 8 wppm. The extra sulfur also releases heat in bed 1. Under the Smith predictor at 240 min of case 4, with the inlet at 340.06 °C, bed 1 rises 27.6 K on 5.09 MW of reaction heat and bed 2 only 12.6 K, the outlet reaches 377.13 °C, the exchanger returns 22.59 MW, and the heater fires 7.30 MW, slightly less than at design, for an inlet 5 K hotter (Figure 8.4). A rise measured as outlet less inlet can also exceed the reaction's share. In case 8 as shipped, at 60 min, bed 2 shows 34.58 K while its reaction heat is worth 15.33 K of its heat-capacity flow, since the collapsing inlet draws the catalyst's stored heat out through the outlet.

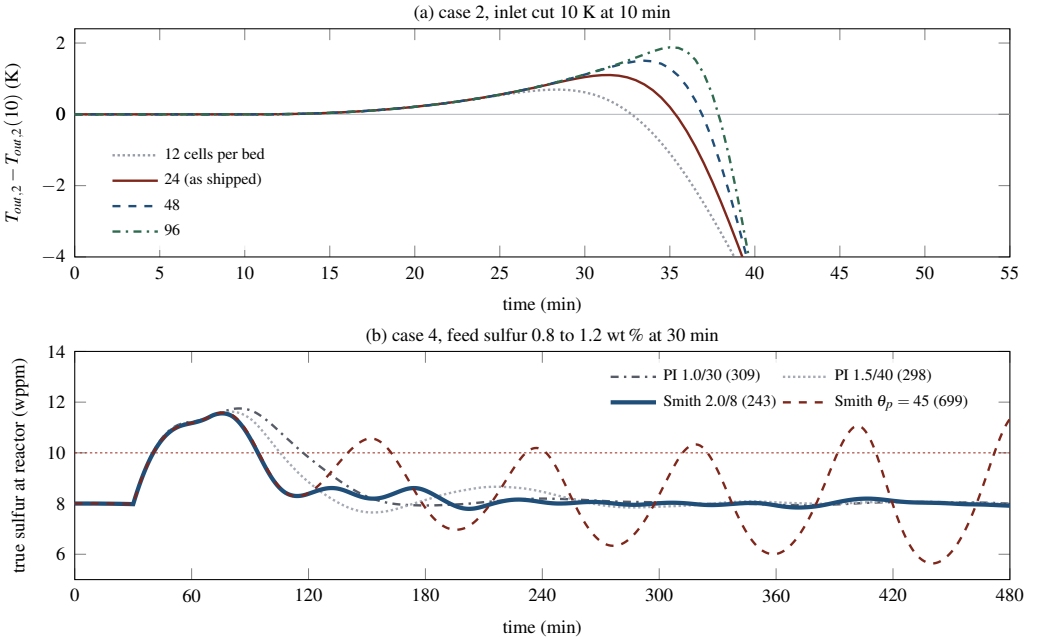


Figure 8.2 Two responses from the application's engine. (a) Case 2, the reactor outlet after a 10 K inlet cut at 10 min, less its value at the cut, at four cell counts. The rise grows with refinement and has not converged. (b) Case 4, the true product sulfur after the feed sulfur steps from 0.8 to 1.2 wt% at 30 min, with the IAE over 480 min in parentheses and the 10 wppm specification dotted.

8.5 The Control Layers in the Code

At each 10 s execution the code first forms the true state, then the instrument readings and the sulfur estimates, and then runs the sulfur layer or the MPC, the WABT loop, the inlet loop, mid-ranging and the quench override from the top down. Only then are the valve and heater commands written. A single Manual and Auto switch in the header covers the bypass, heater and quench stations, and while in Manual each loop follows its station, so switching back to Auto causes no bump.

8.5.1 The Shared PID

Every loop uses the series' velocity-form PID routine of Section 1.6.2, Equations (1.16) to (1.19). The direction $\sigma = \pm 1$ given for each loop below is the action a of those equations. Anti-windup back-calculates the integral only, with τ_I as its tracking time, and any integral time from 600 min upward leaves the loop proportional-only. No loop uses derivative action. In Manual, or when a layer above holds, the routine solves for the integral that makes its output equal the value imposed, which is how every transfer stays bumpless.

8.5.2 The Inlet Loop and Mid-Ranging

The inlet loop moves the bypass between 0 and 100 % at $K_c = 3\,\%/\mathrm{K}$, $\tau_I = 1$ min and $\sigma = -1$, so a hot inlet opens the bypass. Mid-ranging is a valve-position controller on the bypass command against its 50 % target, moving the heater at 0.04 MW per percent with $\tau_I = 20$ min and $\sigma = +1$ between $\min(Q_{min}, \alpha Q_{max})$ and αQ_{max}. With mid-ranging off the heater stays at its manual station, which is how every case but 4 to 7 ships. A power block with three such layers is the subject of Volume 4, Chapter 10.

8.5.3 The Quench Override

A PI on the measured reactor outlet adds quench above the base flow when the outlet passes its override setpoint, 385 °C by default and 374 °C in case 6, at 6 mol/s per kelvin with $\tau_I = 20$ min and $\sigma = -1$, limited to the valve's capacity less the base. In development, a faster tuning oscillated against bed 2's own wrong-way response. The quench loops of Chapter 7 hold each bed's inlet instead.

8.5.4 WABT and Sulfur Layers

A PI on the measured WABT sets the inlet setpoint between 300 and 395 °C at 0.5 K/K with $\tau_I = 10$ min. Above it, the sulfur controller sets the WABT setpoint between 320 and 400 °C with $\sigma = -1$. As a Sulfur PI it acts once per analyzer report on the reading, integrating over the time since the last report. As a Smith predictor the same PI acts every 10 s on the estimate S^* of Equation (8.9). Both layers hold their integrals when the inlet cannot move further, with the bypass shut and the heater fixed or at its limit (or the bypass full open and the heater at its minimum), and the sulfur layer holds while the analyzer is stale or out of service. A WABT schedule, when on, learns a ramp from the sulfur controller's own moves, its rate changing by the trim's last move over a 30 min learning time at each report, and adds the integrated ramp to the WABT setpoint ahead of the trim.

8.5.5 The Smith Predictor

The predictor models the logarithm of the sulfur answering the WABT setpoint that the sulfur controller itself sets, and moves the last report forward to the present, as shown in Equation (8.9).

$$\begin{aligned} z &= \frac{g}{(\tau_1 s + 1)(\tau_2 s + 1)} \left(\mathrm{WABT}_{SP} - \mathrm{WABT}_0\right) \\ S^* &= S_{rep} \exp\left[z(t) - z(t_{rep} - \theta_p)\right] \end{aligned} \tag{8.9}$$

where $g = -0.057$ per kelvin and $\tau_1 = \tau_2 = 10$ min were identified on the commissioned plant from ±5 K setpoint steps, S_{rep} is the last report, t_{rep} the time it arrived, and θ_p the predictor's dead time, a setting of 30 min. Written in logarithms, the correction is exact for a gain that scales with the sulfur. Because the model's input is the controller's own output, no disturbance moves it, so a feed change still arrives one delay late. With θ_p equal to the true delay the delayed model cancels each move when the analyzer catches up, and with θ_p wrong the controller sees its own correction appear and vanish. In RIT only and Manual, where no controller sets the WABT, the estimate shown runs a second model of 2.75 min from the measured WABT. In development, a predictor driven by the measured WABT read a sourer feed's own exotherm as lower sulfur and moved the wrong way, which is why the controller's output drives this one. Volume 5, Chapter 5 uses a predictor of the same family on the delay of a forming line.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.5.6 Model Predictive Control

The MPC is the series' linear MPC of Section 1.6.3, here on two decisions, the WABT setpoint W and the quench base flow q, every 5 min over 36 steps (3 h) with moves blocked as 1, 1, 2, 2, 4, 6, 8 and 12 steps. Its model is fixed, identified on the commissioned plant with the WABT loop closed and mid-ranging on. The logarithm of the sulfur answers W at −0.0562 per kelvin through lags of 10.5 and 10 min, the outlet answers W at 0.952 K/K and q at −0.0191 K per mol/s, each bed rise answers both, and the feed tank's sulfur assay enters the sulfur one-for-one in logarithm through the surge

drum, a feedforward the PI and the Smith predictor lack. The program is Equation (8.10).

$$\begin{aligned} \min\ & \textstyle\sum_k \left[w_S (\ln S_k - \ln S_{SP})^2 + c_W \mathrm{WABT}_k \right] \\ & \quad + \textstyle\sum_j \Delta u_j^{\mathsf{T}} \Lambda \Delta u_j + \text{slack penalties} \\ \text{s.t. } & 320 \le W \le 400, \quad 0 \le q \le 400 \\ & |\Delta W| \le 4, \quad |\Delta q| \le 100 \\ & T_{out,2,k} \le T_{metal} - 5, \quad \Delta T_{b,k} \le \Delta T_{max} - 1 \end{aligned} \tag{8.10}$$

where $w_S = 100$ on every step and on the last, the catalyst-life cost is $c_W = 0.2$ per kelvin of WABT, $\Lambda = \mathrm{diag}(1, 10^{-4})$, a weight of 10^{-6} holds the quench decision, and the outlet and rise limits are soft, priced at 50 per squared unit and 5 per unit of slack. Each output carries an offset, the sulfur's formed at each report against the model one predictor dead time earlier, as the Smith predictor's is, and the temperatures' filtered over 5 min. If the solver returns a move outside its bounds, the previous decisions stay in force. Over the challenge day the MPC solves 288 times with no fallback. The inlet loop and mid-ranging stay under it, so the bypass keeps its 10 s scan.

> ⧖ HISTORY **Late 1980s: A Hydrotreater Tests Predictive Control**
>
> A hydrotreater was a test case in the history of predictive control. In the late 1980s engineers at a research center in France developed a controller that they described as a bridge between state-space and predictive algorithms, and to motivate it they discussed a hydrotreater with four reactor beds in series [18]. The control system had to hold the average bed temperature at a setpoint, keep the temperature differences between the beds close to a profile and prevent any violation of the maximum temperature in each reactor. It moved the first reactor inlet temperature and the quench flows between the beds. A disturbance in the first bed finally affects every bed downstream, which a model of independent outputs does not carry [18]. The MPC of this section decides the same two kinds of move, a bed temperature setpoint and a quench flow.

8.5.7 Three Models Kept Apart

The plant is Section 8.3. The identified model is whatever a learner estimates from step tests, for example the 5.36 K/MW direct gain of the heater or the gain of −0.057 per kelvin in the logarithm of the sulfur. The controllers' models are the predictor's two lags and the MPC's fixed map, which know neither the activity, the feed's true sulfur in the beds nor the fouling. The end of run reads the true hottest cell and latches.

8.6 Nine Cases Measured

The cases are listed in Table 8.3, each from the defaults of Section 8.5 (RIT only, the inlet at 335 °C, the slow Sulfur PI at 0.4 K/wppm and 90 min ready, mid-ranging off, the override at 385 °C) with only its differences listed. The figures below come from the engine advanced without a display at the default seed, with each learner move at the time the text names. Every case runs identically when loaded first, from the dropdown after another case and after Reset, and the 211 numbers the application's own texts quote reproduce on it. The outcomes are collected in Table 8.4.

8.6.1 Two Handles, Mid-Ranging and a Fouled Exchanger

In case 3 the bypass alone holds an 8 K rise of the inlet setpoint, but it closes fully by 27.5 min and stays shut until 41 min while the cooler effluent of the first minutes comes back, the inlet sagging to 341.7 °C, and then parks at 32.3 %, with 32 % of its travel left for the next upset. Turned on at 120 min, mid-ranging raises the heater from 7.33 to 7.89 MW and returns the bypass within 3 % of its target by 175 min, the inlet within 0.5 K of 343 °C throughout. With mid-ranging on from the start the bypass dips only to 23.3 % and is back by 42.5 min.

Case 8 shows what a bypass with nothing behind it does. Fouled to 60 % at 30 min, the exchanger would return 17.70 MW instead of 21.6 at the design inlet. The bypass shuts by 34.5 min, and the

Table 8.3 The Nine Cases

#	Preset and events	Teaching point
1	Manual, learner raises the heater 1 MW, then opens the bypass to 60 %	two handles, two clocks
2	Learner cuts the RIT setpoint 335 to 325 °C	wrong-way rise
3	Learner raises RIT to 343 °C with mid-ranging off, then on	mid-ranging
4	Sulfur PI 1.0/30 min and mid-ranging, feed 0.8 to 1.2 wt% at 30 min	PI against Smith
5	Mid-ranging, feed 1.2 wt% at 30, 230 t/h at 240 min	feed changes
6	Mid-ranging, override 374 °C, 20 % LCO at 30 min	quench and heat
7	Sulfur PI, mid-ranging, compression 720	catalyst cycle
8	Exchanger fouled to 60 % at 30 min	a bypass out of range
9	Challenge: feed 1.0 wt% at 4, analyzer out 10–12, fouling 60 % at 16 h	the whole problem

Table 8.4 Measured Outcomes by Case, Default Seed

#	Setting	Outcome
1	heater +1 MW	RIT +5.23 in 30 min, +12.83 K in 400 min
	bypass 50 to 60 %	RIT −2.57 K in 2 min
2	RIT SP −10 K	outlet +1.10 K at 31.5 min, −13.3 K at 120
3	mid-ranging off	bypass shut 27.5–41 min, parks at 32.3 %
	on at 120 min	heater 7.33 to 7.89 MW, bypass back by 175 min
4	PI 1.0/30	IAE 309, 77 min over 10 wppm
	Smith 2.0/8	IAE 243, 54 min over
	Smith, $\theta_p = 45, 15$	IAE 699, 399
5	RIT only	440 min over, peaks 11.1 and 15.1 wppm
	Smith	92 min over, 7.91 wppm at 480 min
6	RIT only, override	outlet held at 374.2 °C, 12.3 wppm
	Smith from 240 min	quench to 32.3 kNm^3/h, bed 1 rise 29.5 K, 8.1 wppm
7	schedule off	over 10 wppm from 11.6 months, end of run at 32.8
	schedule on	7.9–8.9 wppm, end of run at 25.6 months
8	mid-ranging off	RIT 282.5 °C, 295 wppm at 240 min
	mid-ranging on	heater 7.35 to 11.26 MW, RIT back at 51 min
	heater at 60 %	RIT 299.8 °C, 90 wppm at 240 min

inlet keeps falling, 319.3, 298.2 and 282.5 °C at 45, 120 and 240 min, the sulfur reaching 295 wppm, because each kelvin lost comes back through the exchanger and a cooler bed releases less heat. With mid-ranging the heater climbs from 7.35 to 11.26 MW, the inlet sags to 329.9 °C at 40 min and is back at 335 °C by 51 min, the bypass at 49.5 %. With the heater limited to 60 % (8.4 MW) the inlet falls to 299.8 °C and the sulfur to 90 wppm. That case is infeasible by design. Lined out with the bypass shut, holding 335 °C at 200 t/h needs 10.39 MW, the limited heater holds only 294.8 °C, and only at 180 t/h does it reach 335.5 °C, so the remedy is the feed rate and a cleaning.

> **TRY IT IN THE APP** **8 · A fouled exchanger**
>
> Load the case and press Run. Touch nothing. At 30 min the exchanger fouls to 60 % of clean. Watch the bypass valve and the reactor inlet. The bypass is shut by 34.5 min and the inlet keeps falling, to 282.5 °C at 240 min. Then turn Mid-ranging On, press Reset and run again. The heater climbs to 11.26 MW and the inlet is back at 335 °C by 51 min.

8.6.2 Sulfur on a Late Analyzer

Case 4 steps the feed sulfur from 0.8 to 1.2 wt% at 30 min, and no controller can act before the analyzer's first reading over 9 wppm, at 70 min in every run, so the true sulfur peaks between 11.48 and 11.76 wppm whatever the controller, as shown in Figure 8.2b. The IAE of the true sulfur against its setpoint over 480 min is 309 for the shipped PI (1.0 K/wppm, 30 min) with 77 min over 10 wppm, and 243 for the Smith predictor at 2.0 K/wppm and 8 min with 54 min over. Both tunings are the best of their own grids, the predictor's of 36 and the PI's of 42 (1.5 K/wppm, 40 min, 298 and 65 min), each at an interior point. Against the best PI the predictor saves 18 % of the IAE and 11 min over

specification, all of it in the recovery. At seeds 1, 2 and 3 the shipped PI gives 304, 299 and 306, the best PI 297, 284 and 291, and the predictor 244, 238 and 241.

The predictor is sensitive to its dead time. With $\theta_p = 45$ min, 50 % long, the IAE is 699, over twice the PI's, and the true sulfur still swings between 5.6 and 11.3 wppm in the last hour, because each move's modeled effect is taken back 15 min after the analyzer has shown it. With $\theta_p = 15$ min it is 399. Both detuned predictors lose to the plain PI at every seed tried (583 to 724 and 379 to 412), and the curve is not symmetric. The predictor run with the PI's own tuning gives 487, so its advantage needs its own tuning.

> **+ DID YOU KNOW? A Wrong Delay Costs More Than No Predictor**
>
> From the simulator (Table 8.4). On a sourer feed the shipped PI scores an integrated error of 309 and the Smith predictor 243. Give the predictor a dead time of 45 min where the analyzer's is 30 min and the same controller scores 699, more than twice the PI. Set 15 min too short, it scores 399.

8.6.3 Feeds and Cracked Stock

Case 5 runs RIT only through a sourer feed at 30 min and 230 t/h at 240 min. The product passes 10 wppm at 40 min, reaches 11.1 and then 15.1 wppm, and spends 440 min over specification. With the Smith predictor the sourer feed takes the true sulfur to 11.56 wppm at 75.5 min, back to 8.15 by 230 min, the rate step, which arrives within a minute instead of through the drum, peaks at 10.70 wppm at 263 min, and the product is at 7.91 wppm by 480 min, 92 min over in all. The WABT climbs from 355.3 to 364.8 °C for the sulfur and to 369.8 °C for the rate, and mid-ranging takes the heater to 10.43 MW. The Sulfur PI with the predictor's gains, acting on the raw analyzer, is unstable, reaching 70.7 wppm before the rate step and 136.9 wppm after it.

Case 6 blends 20 % light cycle oil at 30 min with the inlet fixed. The aromatics of the blend rise from 25 to 33 mol%, the rises from 20.4 and 16.2 to 24.0 and 20.0 K lined out, and the outlet to 375.6 °C at 93.5 min. The override adds quench, 4.8 to 8.2 kNm^3/h by 240 min, and holds the outlet at 374.2 °C while the cooler bed 2 and the harder sulfur leave the product at 12.3 wppm (11.5 wppm with the override off). Chosen at 240 min, the Smith predictor raises the WABT, the outlet overshoots to 382.0 °C at 291 min, the quench reaches its 32.3 kNm^3/h maximum by 556.5 min and brings the outlet back to 375.2 °C, but the WABT loop puts the heat back through the inlet (349.6 °C) and bed 1's rise climbs to 29.5 K. The sulfur returns to 8.1 wppm. Under a WABT loop, quench moves heat from bed 2 to bed 1 rather than removing it.

Pause. Try Exercise 8.3 before reading on.

8.6.4 A Catalyst Cycle in a Day

Case 7 compresses two years into a day with the slow Sulfur PI. Without the schedule the product passes 10 wppm at 11.6 months and stays over for 21.2 months, 26.6 wppm at the end, and the run ends at 32.8 months only because the beds ran cooler than the product needed. With the schedule the sulfur stays between 7.9 and 8.9 wppm, the WABT climbs from 355.2 to 364.0 °C in the first year and to 383.6 °C by month 24, bed 2 falls to an activity of 0.20 against bed 1's 0.43, and bed 1's outlet reaches the 400 °C metal limit at 25.6 months, with the quench at 22.6 kNm^3/h and the heater at 13.1 MW. The compression is a teaching device, and a real unit sees the same climb over years.

8.6.5 A Graded Day and Its Margin

The challenge runs 24 h with a feed of 1.0 wt% sulfur from 4 h through the surge drum, the analyzer out of service from 10 to 12 h and the exchanger fouled to 60 % at 16 h. Its target judges the true sulfur at the reactor every 10 min, analyzer in service or not, at or under 10 wppm at every one of 144 readings and at or under 8.5 wppm over all the day's diesel, both rises at or under 30 K at every step, catalyst life used at or under 1.25 days per day, and the heater never at its limit. The variants are listed in Table 8.5. Figure 8.3 draws three columns of the table against their limits.

Table 8.5 The Challenge Over 24 h, Default Seed. Worst Reading, Readings Over 10 wppm, Mean Sulfur, Peak Rise and Catalyst Life

Settings	Worst (wppm)	Over	Mean (wppm)	Rise (K)	Life (d/d)	Verdict
shipped (slow PI, MR off)	386.97	53	86.27	35.2	0.81	fails
slow PI, MR on	10.14	5	8.83	28.4	1.11	fails
Smith 2.0/8, MR off	386.81	47	85.08	36.5	0.85	fails
PI 1.0/30, MR on	9.88	0	8.10	28.9	1.17	passes
Smith 2.0/8, MR on	9.68	0	8.08	28.8	1.17	passes
Smith 2.0/10, MR on	9.70	0	8.08	28.9	1.17	passes
Smith 2.0/10, SP 7.5 wppm[a]	9.11	0	7.58	29.2	1.21	passes
Smith 2.0/10, SP 6.5 wppm	7.94	0	6.57	30.05	1.30	fails
MPC, MR on	9.26	0	8.18	28.8	1.16	passes

[a]The suggested settings. Other rows use SP 8 wppm, except where stated. MR is mid-ranging.

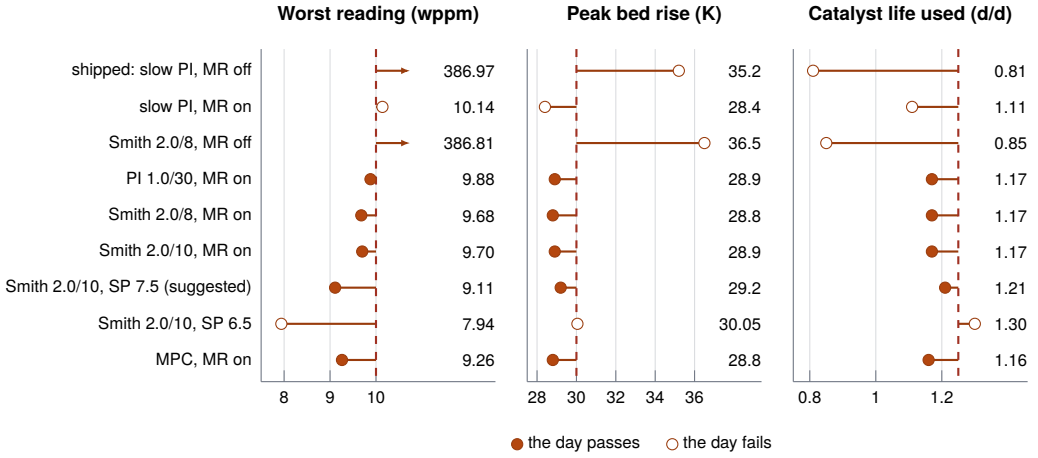


Figure 8.3 Three columns of Table 8.5 against their limits (dashed): the worst of the day's readings against 10 wppm, the peak bed rise against 30 K and the catalyst life used against 1.25 days per day. A filled circle marks a day that passes and an open circle a day that fails. Notice that the two rows without mid-ranging fail by far, that every passing row sits close under the sulfur limit, and that the 6.5 wppm setpoint fails on the two limits at the right and not on sulfur. The worst readings of 386.97 and 386.81 wppm lie off the scale. The mean sulfur and the count of readings over the limit are not drawn.

As shipped, the fouling at 16 h takes the inlet down to 281.8 °C with the heater fixed at 7.34 MW, 53 readings exceed 10 wppm, and the rise peaks at 35.2 K at 16.2 h, the transient of case 8 rather than a runaway. Mid-ranging saves the inlet, and the slow PI then lets the feed switch through, five readings over and 8.83 wppm on average. With mid-ranging, the Smith predictor at 2.0 K/wppm and 10 min passes the default 8 wppm setpoint day with a worst reading of 9.70 wppm, mean 8.08 wppm, rise 28.9 K and catalyst use 1.17 days per day. Expanding the sweep to the default seed and seeds 0–29 exposes seed 26: two readings exceed specification, peaking at 10.05 wppm at both step sizes. The suggested settings therefore use a 7.5 wppm sulfur setpoint. Their default worst reading is 9.11 wppm, mean 7.58 wppm, rise 29.2 K and catalyst use 1.21 days per day. Heater duty peaks at 12.08 MW of 14. The back-off buys quality margin with catalyst life, within the same limits. All 62 trials pass across those 31 seeds and both integration steps. Worst readings span 9.006–9.448 wppm, peak rise stays below 29.736 K and catalyst use below 1.212 days per day.

At the common 8 wppm setpoint, case 4's 8 min integral time passes the default day (9.68 wppm) but fails seed 6 (10.27 wppm, two readings over). Its setpoint-driven predictor cannot see WABT

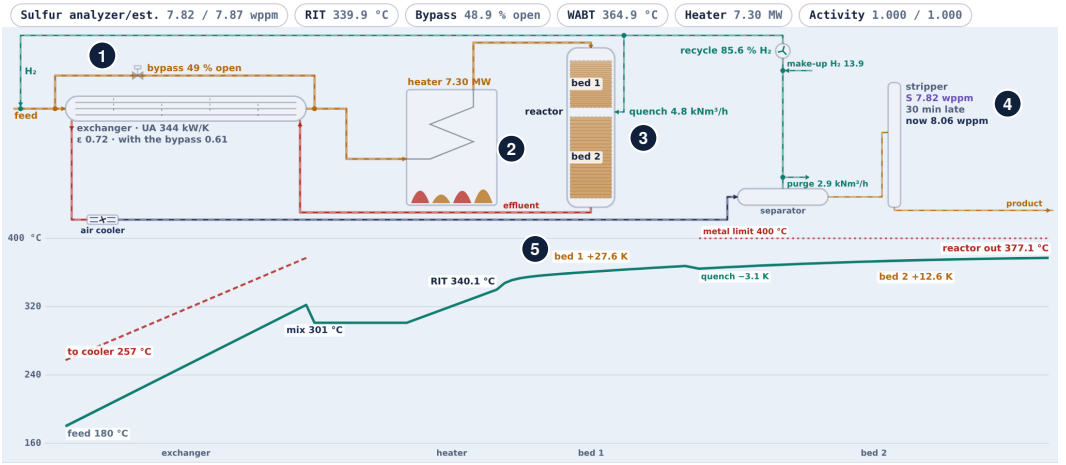


Figure 8.4 Case 4 at $t = 240$ min on the Unit view, cropped to the plant view with numbered callouts, the Smith predictor at 2.0 K/wppm and 8 min chosen before Run. The feed carries 1.2 wt% sulfur since 30 min. The bypass (1) is 48.9 % open and the heater (2) fires 7.30 MW, with the inlet at 339.9 °C and the WABT at 364.9 °C. The quench (3) enters between the beds. At the stripper (4) the analyzer reads 7.82 wppm, 30 min late, the predictor's estimate is 7.87 wppm and the true sulfur at the reactor 8.06 wppm. The temperature ladder (5) reads a mix of 301 °C, an inlet of 340.1 °C, a bed 1 rise of 27.6 K, a quench step of −3.1 K, a bed 2 rise of 12.6 K, a reactor outlet of 377.1 °C and 257 °C to the cooler, and the exchanger's effectiveness is 0.72, or 0.61 with the bypass. The recycle reads 85.6 mol% hydrogen with 13.9 kNm^3/h of make-up. In the full window, outside this crop, the trends show a true peak of 11.56 wppm at 75.5 min that the analyzer reported as 11.38 wppm at 110.5 min, and the scorecard gives a true sulfur IAE of 229 wppm min so far and a run average of 8.93 wppm.

fall short during fouling recovery. The MPC gives 9.25 to 9.45 wppm over seven seeds with the tank assay in its model. Comparing it with the 7.5 wppm suggestion would mix a setpoint change with an architecture change. At 6.5 wppm, the default worst reading falls to 7.94 wppm, but catalyst use 1.30 and rise 30.05 K fail both budgets.

The graded feed switch is 25 %. At 50 %, the original 8 wppm Smith 2.0/10, case 4 predictor, PI 1.0/30 and slow PI with mid-ranging reach 11.40, 11.37, 11.75 and 12.33 wppm before the analyzer reports it, with rises of 31.7–32.6 K. Even MPC reaches 10.08 wppm on two readings. Certification covers the stated day, not an arbitrary sour-feed envelope. The state of case 4 at 240 min under the Smith predictor is shown in Figure 8.4.

8.7 Technology in Practice

Figure 8.5 draws the loops of Section 8.5 with the hardware a refinery installs.

What Measures It

The inlet, the beds and the outlet are read by thermocouples, which are less accurate than resistance thermometers but more rugged and wider in span [19]. Their 6 s lag in Section 8.3 is short beside every other clock of the unit. The measurement that sets the pace is the product sulfur. In one industrial diesel hydrotreater it was known from a laboratory analysis made once a day [20]. The analyzer of this chapter, with a report every 10 min, is the faster instrument that a sulfur loop needs.

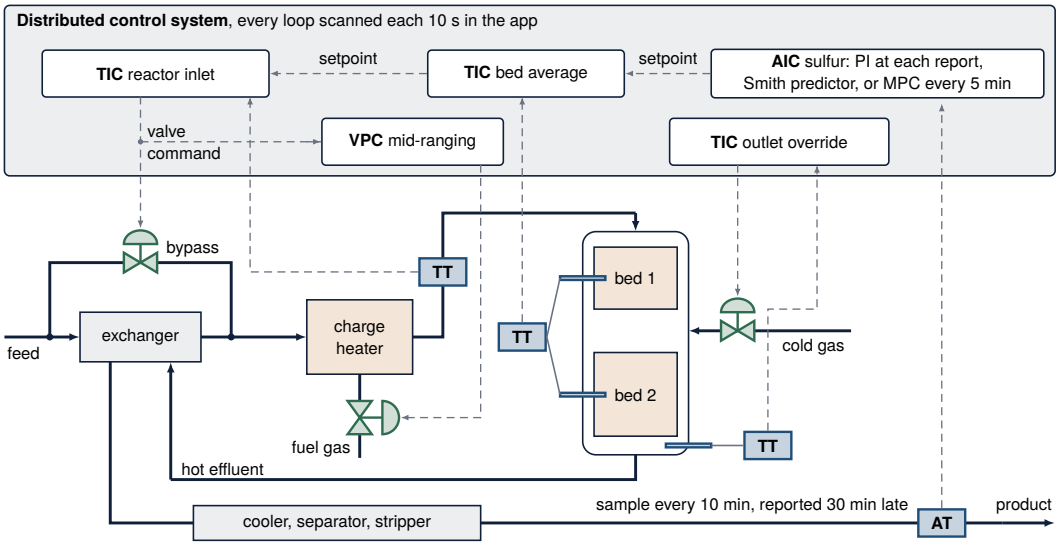


Figure 8.5 The loops of the hydrotreater with their hardware. Lettered boxes are sensing devices, thermocouples with their transmitters (TT) and the sulfur analyzer (AT). Valve symbols are final elements and dashed arrows are signals. Notice that the three controllers of the top row form one cascade from the late analyzer down to the fast bypass valve, and that the heater is moved only by the position of that valve. Not to scale, and not a piping or safety diagram.

> **TECHNOLOGY** **Sensor: Product Sulfur Analyzer**
>
> Product sulfur is defined by laboratory methods. Under the federal fuel rule every batch of ultra-low-sulfur diesel is tested before it leaves the refinery, and the referee procedure is ASTM D2622, wavelength dispersive X-ray fluorescence [21]. ASTM D5453 finds total sulfur by ultraviolet fluorescence [1]. Sulfur dioxide, formed when sulfur burns, fluoresces under ultraviolet light, and a photomultiplier tube measures the light emitted [19]. Online analyzers for product sulfur are mostly expensive, with high maintenance costs [22]. An analyzer away from the process needs a sample system, and analyzer problems quite often stem from it. With a plugged sample line the analyzer measures old, stale sample and has no way of knowing [19]. The 20 min product path and 10 min analysis of Section 8.3 stand for this sensor.

What Moves It

Three final elements carry the control. The bypass valve is the fast one, because only the heater coil stands between the mix and the inlet. The fired heater is the slow one. Such a heater may not be required in normal operation but is always needed for startup [6]. In the audited unit of the box below, every temperature controller of the reactor section had a very slow action and was retuned, one of them to more than three times its former gain [23].

> **TECHNOLOGY** **Final Element: Bypass and Quench Control Valves**
>
> Bypassing the feed and effluent exchanger is the typical way to control a reactor inlet temperature [6], and in a hydrotreater cold gas injected as quench controls the reactor temperature [23]. A control valve with excess stem friction moves in jerks, which is called stiction, and with integral action the controller output then ramps up and down as the valve slips and sticks [19]. An audit of an industrial diesel hydrotreater found this in a quench flow loop. Stiction of about 0.4 % of the valve's travel made the flow swing by 2.5 % of its range, fast enough that the reactors filtered it out of their temperatures [23]. The model's bypass and quench valves are lags of 0.1 and 0.15 min with no friction.

Where the Controller Runs

A refinery runs these loops in a distributed control system. Its loops are sometimes nested, the set point of one being set by another, and their cycle times range from milliseconds to minutes [24]. The app scans every loop each 10 s and solves its MPC every 5 min, inside that range. The sulfur, WABT and inlet controllers are three PID blocks in cascade, and the override is a fourth. Mid-ranging is a well established technique, but in a distributed control system it must handle bumpless transfer, saturation and operation with one of its two handles out of service [9]. Above the loops, a predictive controller replaces a combination of PID, lead-lag and select blocks with one controller [18].

> **IN PRACTICE** **A Leak That No Temperature Fixes**
>
> A real exchanger has a failure that the model does not carry. The feed side stands at the higher pressure, by the pressure drop of the reactor and the exchanger, so a leaking tube passes raw feed into the effluent. At a refining industry panel such leaks were called a big problem for meeting the product specification [25]. One refiner tracks a sulfur species that the reactor converts completely. Found in the product from the exchanger, it shows a leak. A quick field test raises the reactor temperature. Product sulfur that does not answer successive temperature moves points to a problem such as a leak. Tubes are seal welded to the tube sheet and inspected at turnarounds [25].

What Is Modernizing It

The first development is the soft sensor. In one refinery a neural network inferred the product sulfur of a diesel hydrotreater in real time from thirteen online measurements of flow, temperature and pressure, and was meant for the unit's advanced control [20]. The second is the audit of the loops themselves. Software that grades each loop from routine data, with the minimum variance controller as its reference, was applied to the 56 loops of an industrial diesel hydrotreater and found a tuning problem in 27 of them [23]. The third is predictive control. A 2003 survey counted more than 4600 applications, and refining held 67 % of those classified [18]. The app's MPC adds what the PI and the Smith predictor lack, the feed tank's assay as a feedforward.

8.8 Limitations

Every size, rate constant, heat and limit is a generic teaching value for a 200 t/h unit, chosen so the design point gives 8 wppm at 335 °C with rises of 20.4 and 16.2 K, and none was fitted to a catalyst or a plant. The HDS order, its activation energy and the H_2S constant were not checked against a published fit, and the hydrogen consumed, 40 Nm^3/m^3 at design, is a teaching value as well. Two lumped reactions replace the chemistry, so nitrogen, olefins and cracking appear only in the HDS heat. Twenty-four stirred cells set the size of the wrong-way rise, which has not converged, and the trickle bed's wetting, maldistribution and interphase transfer are absent. The exchanger has a single lag, the heater no tube metal, and the separator, scrubber and stripper exist only as the analyzer's delay. The catalyst cycle is compressed 720 times in its case.

The Smith predictor's and the MPC's models are fixed at the commissioned design, so their offsets learn the activity and the feed only through the analyzer, one delay late, and a wrong dead time misleads both. The WABT schedule reads a step as a ramp and overshoots it. The MPC's limits are soft and it holds its last moves on a failed solution. The challenge requires a sulfur setpoint back-off for margin across the expanded noise sweep, and its feed switch was softened to 25 % so that feedback alone can pass. The end of run reads the truth and stops the run rather than modeling a shutdown.

OPEN CHALLENGE **A Sulfur Estimate That Ages With the Catalyst**

An analyzer that reports half an hour late invites an estimate from faster measurements, a soft sensor, and such models have been built for the product sulfur of hydrotreaters [20, 22]. The open problem is keeping one true. A review of soft sensors finds that after launch the prediction usually deteriorates because the process drifts, with catalyst deactivation among the causes, until the model must be retrained or rebuilt. Most soft sensors have no automated mechanism for their own maintenance, and telling a process drift from a sensor drift takes expert knowledge [26]. The predictor and the MPC of this chapter carry fixed models too. Progress would be an estimate that follows the catalyst through its cycle and reports its own error.

8.9 Conclusion

A ULSD hydrotreater is a reactor that heats its own feed and is graded half an hour late, and this simulator computes its 48 cell temperatures, its exchanger, heater and hydrogen loop, its analyzer and its catalyst activity together. At design the exchanger returns 21.6 MW against 7.34 MW of firing, and each kelvin at the inlet comes back as 0.73 K, so 1 MW of heater ends as 12.8 K at the inlet while the bypass moves it within minutes. A 10 K inlet cut first raises the outlet 1.10 K, a size that grows with the cell count while its sign and cause do not. On a sourer feed the Smith predictor saves about a fifth of the IAE against the best PI, 243 against 298 wppm min, and loses to it with a dead time 50 % wrong either way. Mid-ranging keeps a fouled exchanger from taking the inlet from 335 to 282.5 °C, a WABT schedule holds the product within 7.9 to 8.9 wppm across a compressed cycle, and the graded day needs a 7.5 wppm sulfur setpoint to preserve margin over the expanded noise sweep while respecting temperature and catalyst-life budgets. The kinetics are teaching values and the beds are stirred cells. Further work would add a plug-flow bed with axial dispersion fitted to a real one, kinetics fitted to a published ULSD catalyst, a separator and stripper model, and a laboratory sample that corrects the analyzer. The simulator shows that a hydrotreater's inlet, its beds and its sulfur loop have to be designed as one loop.

KEY POINTS

- A hydrotreater heats its own feed, so each kelvin at the reactor inlet comes back as 0.73 K and a heater move grows for hours.
- The bypass moves the inlet within minutes, and mid-ranging lets the slow heater return that valve to mid-range.
- A 10 K inlet cut first raises the outlet 1.10 K, a size that grows with the cell count while its sign and cause do not.
- No controller acts before the analyzer's first report, and a Smith predictor with a dead time 50 % wrong loses to the PI.
- The graded day needs a 7.5 wppm sulfur setpoint, a margin bought with catalyst life.

Where else this applies. Mid-ranging returns in Volume 4, Chapter 10 and Volume 4, Chapter 9, a predictor of the Smith family in Volume 5, Chapter 5, and a loop on a slow analyzer in Volume 5, Chapter 3.

References

[1] U.S. Environmental Protection Agency, "Control of air pollution from new motor vehicles: Heavy-duty engine and vehicle standards and highway diesel fuel sulfur control requirements," *Federal Register*, vol. 66, pp. 5002–5193, Jan. 2001, document 01-2, published January 18, 2001.

[2] European Parliament and Council of the European Union, "Directive 2003/17/EC of 3 March 2003 amending directive 98/70/EC relating to the quality of petrol and diesel fuels," Official Journal L 76, pp. 10–19, Mar. 2003, 22 March 2003.

[3] 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.

[4] J. M. Douglas, J. C. Orcutt, and P. W. Berthiaume, "Design and control of feed-effluent, exchanger-reactor systems," *Industrial & Engineering Chemistry Fundamentals*, vol. 1, no. 4, pp. 253–257, 1962.

[5] J. L. Silverstein and R. Shinnar, "Effect of design on the stability and control of fixed bed catalytic reactors with heat feedback. 1. concepts," *Industrial & Engineering Chemistry Process Design and Development*, vol. 21, no. 2, pp. 241–256, 1982.

[6] F. Reyes and W. L. Luyben, "Steady-state and dynamic effects of design alternatives in heat-exchanger/furnace/reactor processes," *Industrial & Engineering Chemistry Research*, vol. 39, no. 9, pp. 3335–3346, 2000.

[7] P. S. Mehta, W. N. Sams, and D. Luss, "Wrong-way behavior of packed-bed reactors: 1. the pseudo-homogeneous model," *AIChE Journal*, vol. 27, no. 2, pp. 234–246, 1981.

[8] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5-6, pp. 469–474, 1998.

[9] N. Alsop, "Implementing mid ranging in a DCS environment," *IFAC-PapersOnLine*, vol. 49, no. 7, pp. 550–555, 2016.

[10] O. J. M. Smith, "Closer control of loops with dead time," *Chemical Engineering Progress*, vol. 53, no. 5, pp. 217–219, 1957.

[11] E. Furimsky, "Deactivation of hydroprocessing catalysts," *Catalysis Today*, vol. 52, no. 4, pp. 381–495, 1999.

[12] A. C. Byrns, W. E. Bradley, and M. W. Lee, "Catalytic desulfurization of gasolines by cobalt molybdate process," *Industrial & Engineering Chemistry*, vol. 35, no. 11, pp. 1160–1167, 1943.

[13] H. Topsøe and B. S. Clausen, "Importance of Co-Mo-S type structures in hydrodesulfurization," *Catalysis Reviews*, vol. 26, no. 3-4, pp. 395–420, 1984.

[14] H. Topsøe, B. S. Clausen, and F. E. Massoth, "Hydrotreating catalysis," in *Catalysis: Science and Technology*, J. R. Anderson and M. Boudart, Eds. Berlin, Heidelberg: Springer, 1996, vol. 11, pp. 1–269.

[15] S. T. Sie, "Reaction order and role of hydrogen sulfide in deep hydrodesulfurization of gas oils: consequences for industrial reactor configuration," *Fuel Processing Technology*, vol. 61, no. 1-2, pp. 149–171, 1999.

[16] T. C. Ho and R. Aris, "On apparent second-order kinetics," *AIChE Journal*, vol. 33, no. 6, pp. 1050–1051, 1987.

[17] V. Pinjala, Y. C. Chen, and D. Luss, "Wrong-way behavior of packed-bed reactors: II. impact of thermal dispersion," *AIChE Journal*, vol. 34, no. 10, pp. 1663–1672, 1988.

[18] 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.

[19] 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.

[20] L. Salvatore, M. B. de Souza Júnior, and M. C. M. M. Campos, "Design and implementation of a neural network based soft sensor to infer sulfur content in a Brazilian diesel hydrotreating unit," *Chemical Engineering Transactions*, vol. 17, pp. 1389–1394, 2009.

[21] U.S. Environmental Protection Agency, "Regulation of fuels, fuel additives, and regulated blendstocks," Code of Federal Regulations, Title 40, Part 1090, sections 1090.95, 1090.305, 1090.1310, 1090.1350 and 1090.1360, edition of July 1, 2024, 2024, govinfo.gov.

[22] S. Shokri, M. T. Sadeghi, M. A. Marvast, and S. Narasimhan, "Improvement of the prediction performance of a soft sensor model based on support vector regression for production of ultra-low sulfur diesel," *Petroleum Science*, vol. 12, no. 1, pp. 177–188, 2015.

[23] L. G. S. Longhi, A. S. Reginato, H. C. G. Teixeira, C. A. C. Cortez, L. P. Lusa, S. S. Gonzalez, and T. D. Fleck, "Control loop performance assessment and improvement of an industrial hydrotreating unit and its economical benefits," *Sba: Controle & Automação Sociedade Brasileira de Automática*, vol. 23, no. 1, pp. 60–77, 2012.

[24] 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.

[25] American Fuel & Petrochemical Manufacturers, "Question 15: ULSD reactor feed/effluent heat exchanger leaks can be a big problem meeting product specifications. what are best practices for detecting and preventing leaks?" AFPM Q&A Database, panel answers by a refiner, a licensor and suppliers, undated, afpm.org, question and answer search.

[26] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

8.10 Exercises

8.1 Concept check · 2 min

In case 2 the inlet loop cuts the reactor inlet from 335 to 325 °C at 10 min. What does the reactor outlet do over the next twenty minutes?

(a) It rises by about 1 K, because the cooler first cells pass unreacted aromatics and sulfur to catalyst still at its old temperature

(b) It falls by about 10 K within minutes, because the oil and gas cross each bed in only a few minutes

(c) It rises by about 1 K, because the inlet loop opens the bypass and the cold feed starves the exchanger of heat

(d) It holds still for twenty minutes, because the catalyst's heat capacity delays every change at the outlet

8.2 Concept check · 2 min

In case 4 the feed sulfur steps from 0.8 to 1.2 wt% at 30 min. The shipped PI, the best PI found and the Smith predictor all let the true sulfur peak between 11.48 and 11.76 wppm. Why does the predictor not lower this peak?

(a) Its gain of 2.0 K/wppm is too low to answer a step of half the feed's sulfur

(b) Its model's input is the controller's own output, so the feed change reaches it only with the analyzer

(c) The surge drum spreads the step over 20 min, longer than the predictor's two 10 min lags

(d) The sulfur layer holds its integral while the analyzer is stale, so the predictor's moves are lost

8.3 Concept check · 1 min

In case 6 the quench override holds the reactor outlet after a light cycle oil blend, and at 240 min the Smith predictor takes over the WABT to bring the sulfur back. Where does the extra heat end up?

(a) In the quench gas, which carries it out of the reactor, so both bed rises fall

(b) In the product, whose sulfur settles above the 10 wppm specification

(c) In bed 1, because the WABT loop raises the inlet as the quench cools bed 2

(d) In the exchanger, which returns it to the feed and lowers the heater duty

8.4 Theory and calculation · 6 min

At design the charge carries $\dot{C}_p = 186.7\,\mathrm{kW/K}$, the exchanger and its bypass return $\varepsilon_{unit} = 0.613$ of the outlet's temperature above the feed, and a kelvin more at the inlet makes 1.187 K more at the outlet.

(a) What inlet rise does 1 MW more in the heater give at once, before the exchanger returns any of it? *Checkpoint:* 5.36 K.

(b) Compute the loop gain L and the amplification $1/(1-L)$ of Equation (1.8). What final inlet rise do they predict for the 1 MW step? *Checkpoint:* $L = 0.728$, 3.67, 19.7 K.

(c) Case 1 ends 12.83 K higher after 400 min. The loop gain is 0.899 at a 320 °C inlet and 0.431 at 350 °C. Explain the gap, and bound the amplification over the range the inlet crosses. *Checkpoint:* 2.40 actual, between 1.76 and 3.67.

8.5 Theory and calculation · 6 min

The HDS rate of Equation (8.1) is written on the fraction $X = w_S/w_{S0}$ of the feed's sulfur left, with order n.

(a) Using Equation (8.2), compute n for a gamma shape $m = 1.54$, and show that $m = 1$ gives a second-order lump. *Checkpoint:* $n = 1.65$.

(b) Hold the bed temperature and the hydrogen and H_2S pressures fixed. Show that X at the outlet does not depend on w_{S0}, and predict the product sulfur for a 1.2 wt% feed if a 0.8 wt% feed gives 8.00 wppm. *Checkpoint:* 12.0 wppm.

(c) At a fixed 335 °C inlet the simulator gives 10.92 wppm, and the order written on the absolute sulfur gave 7.32 wppm. Explain both departures from part (b). *Checkpoint:* The extra sulfur's heat. The absolute form inverts the trend.

8.6 Theory and calculation · 7 min

Bed 1 enters at 335 °C and rises 20.4 K, bed 2 rises 16.2 K, the catalyst shares are 0.35 and 0.65, $\dot{C}_p = 186.7\,\mathrm{kW/K}$, and the quench gas enters at 65 °C with $c_{p,g} = 32$ J/(mol K). Use Equation (8.4).

(a) Compute bed 2's inlet with the design quench of 60 mol/s, and the step the quench makes. *Checkpoint:* 352.44 °C, 2.96 K.

(b) Compute the WABT at design. *Checkpoint:* 355.17 °C.

(c) The quench opens to its 400 mol/s maximum. Holding both rises fixed, compute bed 2's inlet at the design inlet, and the reactor inlet a WABT loop needs to restore the WABT of part (b). *Checkpoint:* 336.77 °C. An inlet of 345.6 °C.

8.7 App challenge · 8 min

Load scenario 1, "Two handles: heater and bypass (Manual)", at the default seed. Before pressing Run, set Heater duty to 8.35 MW. The control moves in steps of 0.05 MW, so this is the setting nearest to 1 MW above the 7.34 MW the case starts from.

(a) Run to 400 min at speed 60 and export the CSV. Read `truth_rit_C` at 0, 30 and 400 min. How far has the inlet risen at 30 min and at 400 min, and what multiple of the heater's direct effect, $\Delta Q/\dot{C}_p$, is the final rise? *Checkpoint:* 335.00, 340.27 and 347.90 °C: 5.27 K and 12.90 K, 2.39 times 5.39 K (±0.02 K).

(b) Press Reset (the heater stays at 8.35 MW). Run at speed 5, pause at 30.0 min, set Bypass valve to 60 % and run to 32 min. Read `truth_rit_C` at 30, 31 and 32 min. How far does the inlet fall in 2 min, and what share of that comes in the first minute? *Checkpoint:* 340.27, 338.18 and 337.71 °C: 2.56 K, about 81 % in the first minute.

8.8 App challenge · 8 min

Load scenario 4, "Sulfur on a late analyzer: PI against Smith", at the default seed and speed 120. Each run is 480 min. After each run, read the Sulfur IAE, true tile and export the CSV.

(a) Run the case as shipped (Sulfur PI, 1.00 K/wppm, 30 min). Then set the Sulfur controller gain to 1.50 K/wppm and its integral time to 40 min, press Reset and run again. Then choose Smith predictor with 2.00 K/wppm and 8 min, press Reset and run a third time. For each run, read the IAE and count the minutes in which `truth_S_reactor_ppm` is above 10 (rows are 0.5 min apart). *Checkpoint:* 309 and 77 min, 298 and 65 min, 243 and 54 min (±2 and ±1 min).

(b) What share of the better PI's IAE does the predictor save? When does the true sulfur first exceed 10 wppm in each run, and why is that time the same? *Checkpoint:* About 18 %. At 40.5 min in all three.

8.9 App challenge · 9 min

Load scenario 9, "Challenge: a day on the hydrotreater", at the default seed and speed 240. The run stops at 24 h. Read the three So far tiles (the sulfur tile's hover text gives the readings over 10 wppm).

(a) Run it as shipped. Which limits fail, and what happened at 16 h? *Checkpoint:* Worst 386.97 wppm, mean 86.27 wppm, life 0.81 days per day, peak rise 35.2 K.

(b) Turn Mid-ranging On, press Reset and run the day again. Which limits fail now? *Checkpoint:* Worst 10.14 wppm, mean 8.83 wppm, life 1.11, rise 28.4 K: five readings over.

(c) Choose Smith predictor, set Sulfur setpoint to 7.5 wppm, the gain to 2.00 K/wppm and the integral time to 10 min, press Reset and run the day again. *Checkpoint:* Worst 9.11 wppm, mean 7.58 wppm, life 1.21, rise 29.2 K: passes.

8.10 Application area · 4 min

A unit's feed and effluent exchanger has fouled to 60 % of its clean heat transfer, and burners out of service limit the charge heater to 60 % of its 14 MW rating. Holding the inlet at 335 °C at 200 t/h with the bypass shut needs 10.39 MW. What should the console do now, and why should exchanger cleaning be planned from the heater's duty rather than from the bypass position?

8.11 Application area · 4 min

The challenge's suggested settings hold a 7.5 wppm sulfur setpoint against a 10 wppm limit. Using Table 8.5 and the seed sweep, say what the back-off from 8 to 7.5 wppm buys and costs, why 6.5 wppm is not better, and what in the plant sets the size of the back-off.

8.12 Review the agent · 5 min

An AI assistant reviewing the hydrotreater writes: "The sulfur analyzer's delay is 30 min, but sample lines foul and slow down. Set the Smith predictor's dead time to 45 min to be safe. Overestimating a dead time is conservative, just as detuning a PI is, so the worst case is a loop a little slower than the PI." Find the flaw and say how you would check it in the app.