INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 9

Interheater Coordination and Inferred Octane on a Continuous Catalytic Reformer

Pages 208–232

Edited by John D. Hedengren

How to cite this chapter. "Interheater Coordination and Inferred Octane on a Continuous Catalytic Reformer," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 208–232.

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

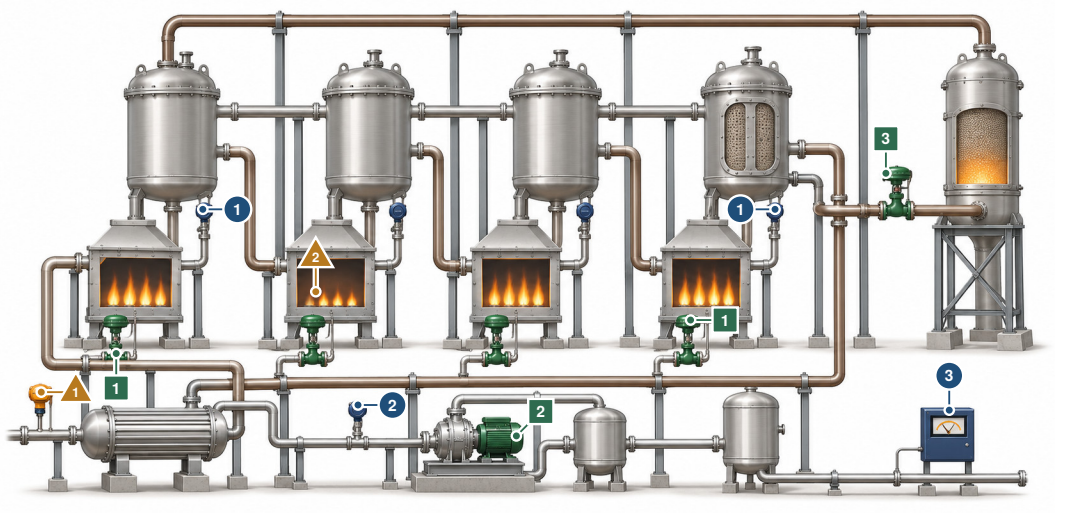
The app for this chapter
https://apmonitor.com/apps/catalytic-reformer/

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

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

CHAPTER 9

Interheater Coordination and Inferred Octane on a Continuous Catalytic Reformer



The controller moves 1 fuel valve of each of the four heaters 2 drive of the recycle compressor 3 catalyst circulation to the regenerator

The controller reads 1 temperature at each bed inlet 2 recycle gas flow 3 reformate octane analyzer

What disturbs it 1 a leaner naphtha 2 a heater short of burners ▲ an analyzer that reads high

Each heater warms the gas on its way into the reactor above it, and the one octane analyzer stands at the far end of the product line at the lower right, after the compressor and the two drums. Illustration, not to scale.

IN PLAIN TERMS Four Heaters and One Late Octane Test

The main idea. A reformer rebuilds the molecules of a light oil so that it resists knocking in an engine, a quality called octane. Its four beds of hot grains do work that soaks up heat, so a heater must warm the stream before each bed.

The challenge. The first bed swallows most of any extra heat, so a push on one heater fades along the line. The octane is tested late and far downstream, and the tester can read high without anyone knowing.

Why it matters. Running too hot turns good gasoline into gas, and running too cool makes fuel that falls short of its grade.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	This page, the History boxes in Section 9.2, the Did you know? boxes in Sections 9.1 and 9.6, the Try it boxes in Sections 9.4 and 9.6, and Key points in Section 9.9. 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 9.6 with Figure 9.4, Section 9.7 and Section 9.8 with its Open challenge box.

The control question A reformer's beds absorb the heat its interheaters add, and its octane is measured late, rarely or by inference. Which layers of control hold that octane through a leaner naphtha, a heater short of duty and an analyzer that drifts, and at what cost in yield?

Learning objectives

- Compute each bed's reaction heat, each heater's tube metal and the regenerator's burn from the drops, the duties and the catalyst circulation.
- Explain why the first bed absorbs most of a heater step and passes only part of it on.
- Price an octane number in C5+ yield and find the WAIT above which more severity stops paying.
- Compare an octane trim on a late analyzer with one on an inferred octane that validates the analyzer.
- Choose the feedforward, inlet distribution and octane measurement that carry the unit through a graded day of upsets.

Prerequisites Sections 1.4.2, 1.4.4, 1.4.5, 1.4.6, 1.6.2 and 1.6.3, and Chapter 7 for adiabatic beds in series.

The app https://apmonitor.com/apps/catalytic-reformer/

See also Volume 5, Chapter 10 for a soft sensor corrected by a delayed laboratory and Volume 5, Chapter 3 for composition control through a slow analyzer. Volume 9, Chapter 4 for offset-free MPC and Volume 9, Chapter 8 for an optimum that sits on a moving constraint.

9.1 Introduction

A catalytic reformer cools its own feed. In the generic unit simulated here, 150 t/h of naphtha carrying 21 vol% of naphthenes and 20.5 vol% of aromatics, with recycle hydrogen at 3 mol per mole of hydrocarbon, enters the first of four catalyst beds at 500 °C and leaves it 70.4 K colder. The second, third and fourth beds take 45.1, 11.8 and 7.2 K, so four fired heaters must put back 21.6 MW and the gas temperature along the train traces a staircase. Lifting every bed inlet by 10 K raises the research octane number (RON) of the reformate by 2.01 and lowers its liquid yield of molecules with five carbons and more (C5+) by 1.51 points of the feed's mass once the unit lines out.

The unit matters to the whole refinery. Reforming exists to upgrade low-octane naphtha into gasoline blending stock, to make aromatics for petrochemical plants and to supply hydrogen [1], and a commercial reformer is a set of heater and reactor pairs in series working between about 400 and 500 °C, followed by a flash drum whose gas is partly recycled and a stabilizer whose bottoms is the reformate [2]. A catalyst that cokes, a hydrogen ratio that slows the coking, a regenerator that burns the coke off and an octane that the laboratory confirms only a few times a day make the reformer a plant whose operation is governed as much by measurement and constraints as by its chemistry [3, 4].

> ✦ DID YOU KNOW? **A Third of the Gasoline**
>
> On January 1, 2003 the refineries of the United States could reform 3,512,237 barrels per calendar day, 21.1 % of their crude distillation capacity, and the reformer produced around 30 to 40 % of all the gasoline made in the country [5]. The unit of this chapter reforms 150 t/h of naphtha.

The simulator corrects three beliefs that a block diagram leaves intact. The first is that a heater step reaches the bed it feeds. In case 1 heater 1 receives 1.6 MW more, about 10 K at its outlet, and by 120 min its outlet has risen 15.7 K because the warmer effluent comes back through the feed–effluent exchanger, yet bed 1 absorbs 9.3 K of that in a larger drop and hands only 6.3 K to bed 2. The second is that an octane loop on an online analyzer holds the octane being made. The near-infrared (NIR) analyzer here reports the stabilizer bottoms about 20 min after the reactors, and when its calibration shifts by 0.8 RON a trim that trusts it lowers the true octane to 94.24 and keeps it more than 0.3 under target for 3.2 h. The third is that more severity always buys more octane. Lined out at a hydrogen ratio of 3, the regenerator reaches its burn capacity at a weighted average inlet temperature (WAIT)

of 523.5 °C, after which coke returns to the first bed and the octane gained per kelvin falls from 0.23 to about 0.1.

Each of these behaviors has relatives in other plants. A train of adiabatic beds with heating between them is a distributed temperature problem in which each bed's outlet is the next heater's disturbance and the inlet profile, not one temperature, sets the product. Endothermic beds damp what they receive, where the exothermic beds of Chapter 7 amplify it (Section 1.4.4). A quality that is sampled slowly and late while the inputs move fast is the dual-rate problem that inferential control addresses [4], met again in the milk powder moisture of Volume 5, Chapter 10, and a supervisory optimizer that trades a product property against yield inside equipment limits is the refinery setting in which model predictive control is most widely applied [6]. A slowly deactivating catalyst whose activity a regenerator restores adds a slow time scale, months between regenerations in a semiregenerative unit [3, 7] and hours in this compressed simulator, that the regulatory loops never see.

However, the reforming literature treats these pieces apart. Kinetic models of industrial reformers lump the naphtha into paraffins, naphthenes and aromatics, or into a few dozen pseudocomponents, and fit the reactor train to plant data [3, 8], optimization studies follow a semiregenerative unit over its catalyst life [3], and control studies address the heaters, the octane or the optimizer one at a time [4, 9], while a reactor temperature loop usually reaches a student as a transfer function whose output is measured exactly. This chapter works on a dynamic simulator in which the bed temperatures, the coke on each bed, the analyzer, the laboratory and an inferred octane are computed together. It shows how heater loops, a WAIT coordinator, an octane trim with a feed-quality feedforward and a model predictive controller (MPC) hold a reformer's octane through a leaner naphtha, a heater short of duty, a recycle compressor turndown, a stopped catalyst lift and a drifting analyzer, and it records each equation and constant of the reformer model so that the model can be reproduced.

The lumped kinetics and the controllers are established practice, and the octane correlation is an illustrative teaching value. The results are measurements made on a single generic unit with the default random seed, repeated at three further seeds wherever noise could change the outcome. Two of them carry the chapter. One is the burn-capacity ceiling on octane above, and the other is the finding that every strategy on the analyzer fails the graded challenge, because the shifted reading persists for longer than the challenge allows whatever the controller does. The model has forty-three states, and its mass, carbon and hydrogen balances close at the reactor boundary to 10^{-15}. Its bed module carries no chemistry and is shared with the bed reactors of other chapters, among them Chapter 8, and a measurement boundary keeps the true octane, coke and activity away from every controller. Each of the nine cases has a completion rule that grades it (Section 1.6.4).

Section 9.2 follows platinum reforming from its patents to the moving bed. Section 9.3 gives the model, Section 9.4 its behavior before any controller, and Section 9.5 the controllers, the estimator and the optimizer as they are coded. Section 9.6 measures the nine cases, Section 9.7 sets the hardware of a real unit beside the model, Section 9.8 states what the model leaves out, and Section 9.9 draws the results together.

9.2 Platinum, Rhenium and the Moving Bed

The process took its modern form in a patent filed in November 1947 and granted in August 1949, whose content the History box below gives. Two years later Haensel and Sterba compared the new Platforming process with thermal reforming, then in commercial operation at a Pennsylvania refinery, by the relation between yield and octane number, the product properties and the hydrocarbon types [10]. The process crossed the Atlantic quickly, and in 1952 CEPSA became the first European company to license Platforming, for its refinery on Tenerife [11].

> ⧖ HISTORY **1947: Platinum Reforming Is Patented**
>
> Vladimir Haensel filed the patent in November 1947 and received it in August 1949 [12]. It defines reforming as the treatment of gasoline fractions to improve their antiknock rating and lists what the catalyst must do, which is to dehydrogenate naphthenes to aromatics, to close straight-chain paraffins into aromatics and to crack selectively. It gives the high cost of the catalyst as the reason platinum had found limited commercial acceptance and describes catalysts with as little as 0.01 to 1 wt% of platinum on a halogen-containing alumina. It also names the problem a regenerator exists for, coke that deposits on the catalyst, shortens the processing cycle and forces either a burn or a new charge. One test raised a Pennsylvania naphtha of motor octane 41.2 to 64.2 at 872 °F and 500 psi. In 1997 the National Academy of Engineering awarded Haensel the Charles Stark Draper Prize for catalytic reforming [13].

The second catalyst came from Chevron. Kluksdahl's patent, granted in December 1968, claimed platinum with rhenium for reforming sulfur-free naphtha [14]. Its account of the prior art is a compact statement of the trade that runs through the present chapter. Higher octane comes from higher temperature or lower space velocity while the gasoline yield suffers, reforming catalysts are rated by their yield at the desired octane, and hydrocracking, being exothermic where reforming is endothermic, can drive a bed toward runaway. In one of its examples the temperature was adjusted throughout the run to hold a 100-octane product, and between 200 and 845 h on a sulfur-free feed the temperature required rose by about 0.024 °F per hour, while in a simulated life test of a platinum catalyst the yield at the desired octane fell steadily with time. That constant-octane schedule is what case 8 reproduces on a compressed clock.

The regeneration mode separates the generations of reformers, as the History box below tells.

> ⧖ HISTORY **Late 1960s: Regeneration Without a Shutdown**
>
> A semiregenerative unit keeps its catalyst in fixed beds and shuts down for an in-place regeneration every six months to two years. A cyclic unit regenerates on a cycle of a few weeks to a few months and suits poor feeds. A continuous unit regenerates a fixed amount of catalyst at a time, at a rate set by its catalyst circulation [7]. The continuous reformer was introduced in the late 1960s. Its catalyst moves down through the reactors as a column of particles, and small quantities are withdrawn, carried to a regeneration unit and returned [13]. In 1976 this was described as the most recent innovation in catalytic reforming, valued for continuous and stable production of gasoline and hydrogen [15]. Case 8 of this chapter stops the catalyst lift, and the WAIT for constant octane then climbs as it does in a semiregenerative unit.

By 2004 the technology had moved from semiregenerative fixed beds toward continuous units for their efficiency, reliability and operating cost [1]. Optimization of a semiregenerative unit over a whole catalyst life, at constant inlet temperatures or at constant octane, is the problem that continuous regeneration largely removes [3].

Control followed the computers. A reformer controller using nonlinear programming with a penalty-function gradient method for constraints, first run in a process computer, was configured in 1983 in the function blocks of a microprocessor distributed control system [9]. The octane itself resisted automation. The knock engine is slow, and online analyzers infer octane from spectra instead, as near-infrared features were shown to do for gasolines in 1989 [16]. In a pilot study of a platinum–rhenium catalyst the RON of the reformate was linear in its aromatics content [17], the relation the simulator's correlation borrows. As late as 2003 octane quality control on Shell Canada's continuous reformers was typically manual because the RON was measured so seldom, and a dual-rate inferential MPC built for those units cut the octane variance by 40 % [4].

9.3 Four Beds, Four Heaters and a Moving Catalyst

The simulated unit is a generic low-pressure continuous reformer, as shown in Figure 9.1. Naphtha at 150 t/h joins the recycle gas, is heated by the reactor effluent in the feed–effluent exchanger, and

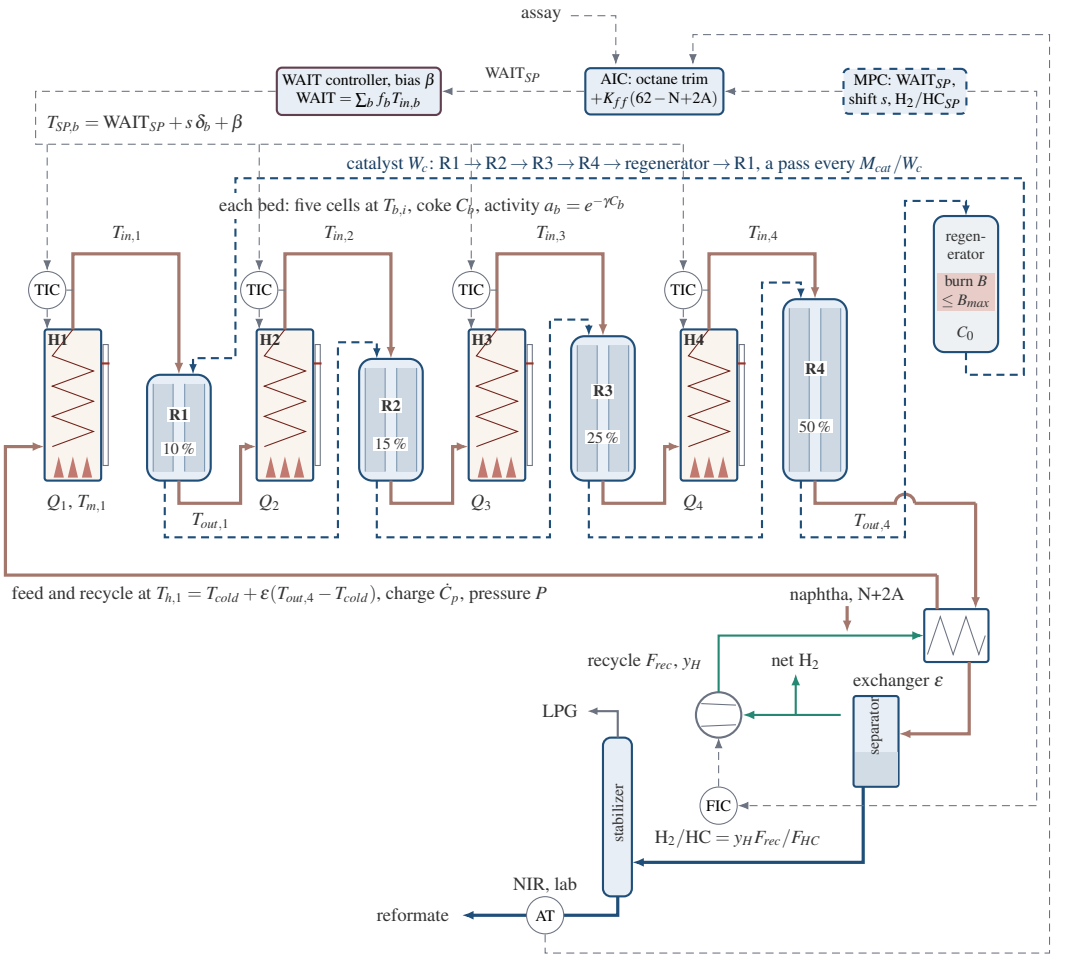


Figure 9.1 Continuous catalytic reformer as simulated, drawn after the wide layout of the application's Unit view, with the product section moved under the last beds. The exchanger delivers feed and recycle at $T_{h,1}$ to heater 1, each heater b with duty Q_b and tube metal $T_{m,b}$ (gauge, with a tick at the 620 °C limit) raises the gas to the bed inlet $T_{in,b}$, and each bed of five cells at $T_{b,i}$ returns it at $T_{out,b}$ to the next heater. Catalyst (dashed) moves from bed to bed at W_c, carrying coke C_b, to a regenerator that burns $B \le B_{max}$ and returns catalyst with coke C_0 to bed 1. The compressor sets the recycle F_{rec} of purity y_H and so H_2/HC against the hydrocarbon feed $F_{HC} = F_P + F_N + F_A$. Four TIC loops hold $T_{SP,b}$ from the WAIT controller, the AIC octane trim on the NIR, laboratory or inferred octane with an N+2A feedforward sets WAIT_{SP}, and the MPC (dashed block) can set WAIT_{SP}, the severity shift and the H_2/HC setpoint. Signals are thin dashed lines. The drawing represents the model, not piping or safety design, and vessel heights follow the cube root of the catalyst.

passes heater 1 (the charge heater), bed 1, heater 2, bed 2, heater 3, bed 3, heater 4 and bed 4 in series at 4.5 bar. Each bed is a radial-flow reactor holding 10, 15, 25 or 50 % of 66 t of catalyst. The effluent gives up its heat in the exchanger and reaches a separator, whose gas the recycle compressor returns to the feed while the net hydrogen leaves, and whose liquid a stabilizer turns into reformate. The catalyst moves down through the four beds in turn at $W_c = 8.25\,\mathrm{t/h}$, is lifted to a regenerator that burns its coke, and returns to bed 1. The clock is the minute, flows inside the beds are moles per second, duties megawatts and coke weight percent of the catalyst.

Each variable in Table 9.1 has a role in the loops. Forty-three states are integrated. Thirty-eight

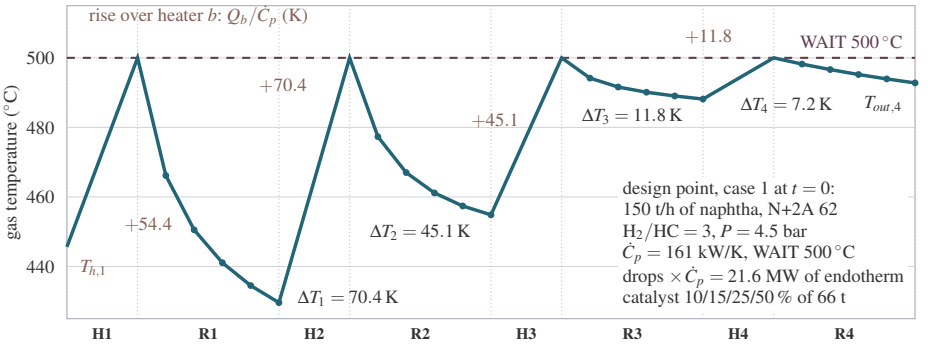


Figure 9.2 The gas temperature along the train at the design point, drawn to scale on the horizontal slots of the heaters and beds of Figure 9.1, with the five cell temperatures of each bed as dots and the WAIT dashed. Each heater raises the gas by $Q_b/\dot{C}_p$ and each bed cools it by its drop ΔT_b.

Table 9.1 Reformer Variables, Their Roles and Design Values (Case 1 at $t = 0$)

Variable	Role	Design	Range
RON	CV, reformate octane	95.0	SP 90–102
WAIT	CV and MV, severity	500.0 °C	SP 470–545 °C
$T_{in,b}$	CV, bed inlets	500 °C each	—
H_2/HC	CV, recycle ratio	3.00 mol/mol	SP 1.5–4.5, floor 2.5
Q_b	MV, heater duties	8.72, 11.30, 7.25, 1.90 MW	0–12, 15, 10, 5
N	MV, compressor speed	84.7 %	0–100 %
W_c	MV, catalyst circulation	8.25 t/h	0–16
N+2A	DV, feed quality	62 (N 21, A 20.5)	48 after a switch
$\dot{m}_f$	DV, feed rate	150 t/h	90–170
a_{H2}	DV, heater 2 available	100 %	30–100 %
N_{cap}	DV, compressor limit	100 %	40–100 %
b_{NIR}	DV, NIR calibration	0	±3 RON
$T_{b,i}$	states, 20 cells	see Figure 9.1	—
C_1–C_4	states, coke	0.21, 0.49, 1.35, 2.99 wt%	—
a_1–a_4	outputs, activity	0.98, 0.94, 0.85, 0.70	—
C_0	state, regenerated coke	0.10 wt%	—
y_H	state, purity	88.9 mol%	—
ΔT_b	outputs, drops	70.4, 45.1, 11.8, 7.2 K	—
$T_{m,b}$	outputs, tube metal	580, 583, 580, 542 °C	limit 620, trip 640
B	output, burn	239 kg/h	capacity 400
Y_{C5+}	output, C5+ yield	94.7 wt% (91.7 vol%)	—
A_{vol}	output, aromatics	44.2 vol%	—

are physical, the temperatures of the twenty cells (five per bed), the four heater duties Q_b, the coke on each bed C_b and on the regenerated catalyst C_0, the compressor speed N, the recycle purity y_H, the naphthenes x_N and aromatics x_A of the feed reaching the reactors, the catalyst circulation W_c, and two lags each for the octane and the yield of the reformate. Five are ledgers that integrate the feed, the C5+ liquid, the net hydrogen, the fuel and the coke burned.

9.3.1 Four Lumped Reactions

The chemistry is Smith's four-reaction lumping on paraffins P, naphthenes N and aromatics A, as reproduced in reformer simulations that represent each group by a single compound [8], with every lump a C7 and a light-ends lump L of equal moles of C1 to C5. The rates per kilogram of catalyst are

as shown in Equation (9.1).

$$\begin{aligned} \mathrm{N} \rightleftharpoons \mathrm{A} + 3\mathrm{H_2} &: r_1 = k_1 a\,(p_N - p_A p_H^3 / K_1) \\ \mathrm{N} + \mathrm{H_2} \rightleftharpoons \mathrm{P} &: r_2 = k_2 a^m (p_N p_H - p_P / K_2) \\ \mathrm{P} + \tfrac{4}{3}\mathrm{H_2} \rightarrow \tfrac{7}{3}\mathrm{L} &: r_3 = k_3 a\, p_P p_H / P \\ \mathrm{N} + \tfrac{7}{3}\mathrm{H_2} \rightarrow \tfrac{7}{3}\mathrm{L} &: r_4 = k_4 a\, p_N p_H / P \end{aligned} \tag{9.1}$$

where $p_i = P F_i / F_t$ is the partial pressure in bar at the reactor pressure P, a the bed's activity and $m = 2$. Dehydrocyclization of a paraffin runs as r_2 backward and then r_1, and ring closure needs both the metal and the acid function, so it alone takes the square of the activity. The rate and equilibrium constants follow Arrhenius and van 't Hoff with the same heats that enter the energy balance, as shown in Equation (9.2).

$$\begin{aligned} k_j &= k_{j,r} \exp\Big[-\frac{E_j}{R}\Big(\frac{1}{T} - \frac{1}{T_r}\Big)\Big] \\ K_j &= \exp\Big(A_j - \frac{\Delta H_j}{RT}\Big) \end{aligned} \tag{9.2}$$

where $T_r = 773.15\,\mathrm{K}$, $E_4 = E_3$ and $\Delta H_4 = \Delta H_2 + \Delta H_3$ by Hess's law. The molar masses are built from toluene ($M_A = 92.14\,\mathrm{g/mol}$) and hydrogen, $M_N = M_A + 3M_H$, $M_P = M_N + M_H$ and $M_L = \frac{3}{7}(M_P + \frac{4}{3}M_H) = 44.10\,\mathrm{g/mol}$, so every reaction conserves mass and atoms exactly. At 500 °C $K_1 = 7.44 \times 10^5\,\mathrm{bar}^3$ and $K_2 = 0.400\,\mathrm{bar}^{-1}$. The constants are teaching values chosen so that the design point gives the bed drops of the plan. Against Smith's constants as reproduced in [8], k_1 is 8.5 times his at 500 °C with his activation energy of about 160 kJ/mol, his equilibrium constants are within a factor of 1.6 of these between 480 and 520 °C, and the two hydrocracking steps are first order in hydrogen where Smith wrote them zero order. At a hydrogen pressure near 3 bar they run at about a fifth of his rate, and in this model more hydrogen cracks more.

9.3.2 Stirred Cells and the Bed Heat Balance

Each bed is five cells of equal catalyst mass W along the radial gas path. The gas crosses a cell in well under a second, so a cell's composition is the stirred-tank solution at its temperature, while its temperature is a state on the heat capacity of the catalyst and a share of the vessel, as shown in Equation (9.3).

$$\begin{aligned} F_{out} &= F_{in} + \nu^{\mathsf{T}} \xi, \qquad \xi_j = W\, r_j(T, F_{out}) \\ C_{cell} \frac{\mathrm{d}T}{\mathrm{d}t} &= \dot{C}_p\,(T_{prev} - T) - \sum_j \xi_j \Delta H_j \end{aligned} \tag{9.3}$$

where ν is the stoichiometric matrix, ξ the extents in mol/s, $C_{cell} = W c_{cat} + C_{ves}/5$, T_{prev} the previous cell's temperature or, for a bed's first cell, the bed inlet $T_{in,b}$, and $\dot{C}_p = \sum_i F_i c_{p,i}$ the heat-capacity flow of the charge, held along the train. Holding $\dot{C}_p$ leaves out the small change of heat capacity by reaction and makes every bed's drop times $\dot{C}_p$ its reaction heat exactly, 11.30, 7.25, 1.90 and 1.16 MW at design with $\dot{C}_p = 160.5\,\mathrm{kW/K}$. The cell composition is found by Newton's method on the extents, four iterations from zero extent with each step damped so no flow falls below a tenth of its value, which makes the right-hand side a pure function of the state. The thermal time constants $C_{bed}/\dot{C}_p$ are 1.55, 1.91, 2.63 and 4.43 min.

9.3.3 Heaters, Tube Metal and the Exchanger

Each heater is first order on its absorbed duty and adds that duty to the gas, and the charge heater's inlet is set by the exchanger, as shown in Equation (9.4). The tube metal is Equation (1.3) without

the coke term, and the exchanger returns heat from the last bed to the first heater as in Section 1.4.5.

$$\begin{aligned} \tau_H \frac{\mathrm{d}Q_b}{\mathrm{d}t} &= \min(Q_{cmd,b}, Q_{cap,b}) - Q_b \\ T_{in,b} &= T_{h,b} + \frac{Q_b}{\dot{C}_p}, \qquad T_{m,b} = T_{in,b} + \theta_m \frac{Q_b}{Q_{max,b}} \\ T_{h,1} &= T_{cold} + \varepsilon\,(T_{out,4} - T_{cold}) \end{aligned} \tag{9.4}$$

where $T_{h,b}$ is the heater inlet, equal to $T_{out,b-1}$ for $b \geq 2$, $Q_{cap,b}$ the rated duty $Q_{max,b}$ times the share of burners in service (zero after a trip), and $\theta_m = 110\,\mathrm{K}$ the rise of the hottest tube metal over the outlet at rated duty. At design the exchanger returns the feed at 445.7 °C and the tube metals sit at 580, 583, 580 and 542 °C. A heater at its rated duty runs 110 K above its outlet, so with an outlet above 530 °C it passes the 640 °C trip. Fuel is the sum of the duties over an efficiency of 0.88, 33.1 MW at design.

9.3.4 Recycle Hydrogen and the Feed

The compressor speed sets the recycle flow, and its purity relaxes toward that of the net gas the reactions make, as shown in Equation (9.5).

$$\begin{aligned} F_{rec} &= \frac{N}{100} F_{rec,max}, \qquad \mathrm{H_2/HC} = \frac{y_H F_{rec}}{F_P + F_N + F_A} \\ \tau_s \frac{\mathrm{d}y_H}{\mathrm{d}t} &= \frac{h}{h+g} - y_H \\ \tau_N \frac{\mathrm{d}N}{\mathrm{d}t} &= \min(N_{cmd}, N_{cap}) - N \end{aligned} \tag{9.5}$$

where h is the net hydrogen made, g the C1 to C3 and half the C4 of the light ends ($0.7F_L$), and the feed flows of P, N and A come from the feed rate and the assay through molar volumes of 150, 128 and 107 $\mathrm{cm^3/mol}$. At design the recycle is 1440 mol/s at 84.7 % speed and 88.9 mol% hydrogen. The naphthenes and aromatics of a new tank reach the reactors through a first-order lag of 6 min.

9.3.5 Coke, Activity and the Regenerator

One coke state per bed treats each bed's catalyst as well mixed. Coke is laid down in proportion to the aromatics and to the inverse cube of the hydrogen pressure, and moves down the train with the catalyst, as shown in Equation (9.6).

$$\begin{aligned} \frac{\mathrm{d}C_b}{\mathrm{d}t} &= r_{c,b} - \frac{W_c}{M_b}(C_b - C_{b-1}), \qquad a_b = e^{-\gamma C_b} \\ r_c &= k_c\, a\, e^{-\frac{E_c}{R}\left(\frac{1}{T} - \frac{1}{T_r}\right)} \frac{p_A}{p_{A,r}} \left(\frac{p_{H,r}}{p_H}\right)^3 \\ B &= \min\left[W_c(C_4 - C_r),\ B_{max}\right] \\ \tau_{reg} \frac{\mathrm{d}C_0}{\mathrm{d}t} &= \max\left(C_r,\ C_4 - B_{max}/W_c\right) - C_0 \end{aligned} \tag{9.6}$$

where $r_{c,b}$ is the mean of the cell rates in bed b, M_b the bed's catalyst, $p_{A,r} = 0.4\,\mathrm{bar}$, $p_{H,r} = 3\,\mathrm{bar}$, $C_r = 0.1\,\mathrm{wt\%}$ the coke a full burn leaves, and the units are reconciled in the code (weight percent, tonnes per hour and kilograms per hour). Below its capacity $B_{max} = 400\,\mathrm{kg/h}$ the regenerator burns the catalyst back to C_r, and above it the catalyst returns to bed 1 carrying coke. At design the beds carry 0.21, 0.49, 1.35 and 2.99 wt% and the regenerator burns 239 kg/h, 60 % of its capacity. The pass through the reactors takes $M_{cat}/W_c = 8.0\,\mathrm{h}$, a deliberate compression (the application's texts put a commercial unit's cycle at 3 to 6 days, about twelve times longer) so that coke moves within a case, and the laydown, the burn and the capacity are compressed by the same factor. The coke in weight percent is therefore of a realistic size while the kilograms per hour are not. The capacity equals the coke on spent catalyst at 4.95 wt% burned back to C_r at the design circulation.

9.3.6 Reformate, Octane and Yield

The reformate is the C5+ liquid of the bed 4 effluent, the P, N and A lumps plus the pentane fifth of the light ends, and its octane is an illustrative correlation on its aromatics and paraffins, as shown in Equation (9.7).

$$\begin{aligned} \mathrm{RON} &= R_0 + b_A A_{vol} - b_P P_{vol} + b_T(\mathrm{WAIT} - T_R) \\ \mathrm{WAIT} &= \textstyle\sum_b f_b T_{in,b} \\ Y_{C5+} &= 100\,\dfrac{\dot m_P + \dot m_N + \dot m_A + \dot m_{C5}}{\dot m_f} \end{aligned} \tag{9.7}$$

where A_{vol} and P_{vol} are volume percent of the liquid, f_b the catalyst shares 0.10, 0.15, 0.25 and 0.50, and $R_0 = 86.205$ anchors the design point at 95.0. The linear dependence on aromatics follows pilot measurements [17], but the coefficients are teaching values and the correlation holds only for this reformate between WAIT 480 and 535 °C, where it gives 91.8 to 101.2 on the design naphtha and 87.2 to 96.6 on the paraffinic one. The octane and yield of the effluent reach the stabilizer bottoms through two first-order lags of 10 min, the separator and the stabilizer, so the reformate there is on average the reactor effluent of 20 min before.

9.3.7 What the Instruments Report

Every controller reads one measurement packet, built every 10 s. The bed inlet and outlet thermocouples lag 12 s with 0.3 K of noise, the tube-metal thermocouples carry 1 K of noise, the duties, feed and recycle meters 0.5, 0.3 and 0.5 %, and the recycle hydrogen analyzer lags 1 min with 0.002 of noise in mole fraction. The distributed control system computes H_2/HC from the purity, the recycle meter and the feed meter over the assay's molar mass. The NIR analyzer on the stabilizer bottoms draws a sample every 2 min through a 1 min sample line, reports one minute later with 0.1 RON of noise and its own calibration bias b_{NIR}, and is flagged stale after four missed cycles. A knock engine tests a sample every 8 h, reports an hour later with a repeatability of 0.1 RON, and is the reference. Chapter 15 holds the octane of the finished gasoline with the same pair of instruments. The regenerator's burn is read from its air with 2 % of noise and reads the capacity itself at the oxygen limit, the catalyst flow meter carries 1 %, and the assay of the tank in service gives N, A and N+2A. Every execution draws the same 23 normal deviates whatever the settings, so runs that differ in a single setting see identical noise. In the exported file every instrument column is matched by a column of the true value it stands for, whose name begins with `truth`.

9.3.8 Parameters, Starting State and Integration

Table 9.2 lists every parameter with its default. Each case starts lined out at its own operating point. A steady-state solve finds the cell temperatures for the bed inlets the heater loops are asked for, iterating the recycle purity and the coke to a fixed point (the coke bed by bed by bisection, with the burn-limited regenerated coke solved when the capacity binds) and capping any heater short of duty. The state advances in fixed fourth-order Runge–Kutta steps of 5 s. Halving the step moves case 5's octane, yield, WAIT and bed 1 drop at 120 min by 2.6×10^{-8} of their values and leaves cases 1, 6 and 8 unchanged to four decimals. Ten and twenty cells per bed in place of five enlarge bed 1's drop by 2.5 and 3.9 K and raise the octane by 0.15 and 0.23, with every conclusion unchanged. Over the challenge day with the suggested settings, mass, carbon and hydrogen balance at the reactor boundary to 9.2×10^{-16} of the throughput, and the integrated heat balance of the beds closes to 0.04 MJ against 2.5 TJ of heater duty. Over case 6 the 10 509 kg of coke laid down equals the 9425 kg burned plus the 1085 kg stored on the beds and in the regenerator. A simulated minute costs about 2 ms in a JavaScript engine, so the 24 h challenge runs in about 3 s.

9.3.9 Assumptions

Four C7 lumps stand for a naphtha of many carbon numbers, so the ring expansion of cyclopentanes, benzene precursors, heavier aromatics and isomerization are folded into four reactions, and the model

Table 9.2 Parameters of the Simulated Reformer by Access (A Adjustable, H Hidden, F Written Into the Source)

Symbol	Meaning	Value	Unit		Symbol	Meaning	Value	Unit	
P	reactor pressure	4.5	bar	A	θ_m	tube metal over outlet at rated duty	110	K	H
M_{cat}	catalyst, all four beds	66 000	kg	H	$T_{m,max}$	tube-metal limit (teaching value)	620	°C	A
f_b	catalyst shares	0.10, 0.15, 0.25, 0.50	—	F	T_{trip}	heater trip	640	°C	A
c_{cat}	catalyst heat capacity	1050	J/(kg K)	H	η_H	heater efficiency	0.88	—	H
C_{ves}	vessel heat capacity per reactor	8	MJ/K	H	ε	exchanger effectiveness	0.88	—	H
$k_{1,r}$	dehydrogenation at 500 °C	0.40	mol/(s kg bar)	A	T_{cold}	feed and recycle into the exchanger	100	°C	H
E_1	its activation energy	160	kJ/mol	H	k_c	coke laydown at 500 °C, fresh	0.76	wt %/h	A
$k_{2,r}$	ring opening and closing	6.4×10^{-4}	mol/(s kg bar^2)	A	E_c	its activation energy	120	kJ/mol	H
E_2	its activation energy	250	kJ/mol	H	β	laydown order in hydrogen	3	—	H
$k_{3,r}$	paraffin hydrocracking	2.64×10^{-3}	mol/(s kg bar)	A	m	activity exponent of ring closure	2	—	H
$k_{4,r}$	naphthene hydrocracking	6.6×10^{-4}	mol/(s kg bar)	H	γ	activity lost per wt % coke	0.12	1/wt %	A
$E_3 = E_4$	hydrocracking activation	250	kJ/mol	H	C_r	coke left by a full burn	0.1	wt %	H
ΔH_1	dehydrogenation (endothermic)	+210	kJ/mol	H	B_{max}	regenerator burn capacity	400	kg/h	A
ΔH_2, ΔH_3	ring closing, cracking	−40, −60	kJ/mol	H	τ_{reg}	regenerator catalyst residence	30	min	H
A_1, A_2	ln K intercepts	46.19, −7.14	—	H	τ_{lift}	catalyst lift response	5	min	H
$c_{p,i}$	P, N, A, H, L, G vapor	290, 260, 210, 29.7, 70, 65	J/(mol K)	F	$F_{rec,max}$	recycle at full speed	1700	mol/s	H
M_G	recycle light gas	20	g/mol	F	τ_N, τ_s	compressor speed, purity lag	0.3, 3	min	H
$Q_{max,b}$	rated duties	12, 15, 10, 5	MW	A	τ_{feed}	feed line to the reactors	6	min	A
τ_H	heater firing response	2	min	A	τ_{prod}	separator, stabilizer (each)	10	min	A
R_0	octane constant	86.205	—	H	b_A, b_P	octane per vol % A, P	0.38, 0.15	1/vol %	H
b_T, T_R	octane per K of WAIT, reference	0.07, 500	1/K, °C	H	$p_{A,r}$, $p_{H,r}$	coke reference pressures	0.4, 3	bar	F

gives the order of the beds' drops and the slope of the severity trade rather than a particular catalyst's yields. Each bed is five stirred cells with no gas holdup and no pressure drop, which underestimates bed 1's drop by a few kelvin against plug flow. The heaters are first order on their duty with a tube metal proportional to it, leaving out the firebox, the draft and the fuel header. One coke state per bed ignores the coke profile along a bed and the difference between coke on the metal and on the support, and the regenerator is one well-mixed vessel with a burn limit rather than burn, chlorination and reduction zones. The catalyst loop is compressed about twelve times, which keeps the coke in weight percent realistic but makes the kilograms per hour and the capacity twelve times a real unit's. The separator pressure and the feed flow are taken as held by fast loops the model does not include, and the separator and stabilizer are lags on the octane and yield. The octane correlation is illustrative. Because the plan's bed drops at $H_2/HC = 3$ limit the endotherm the model can spend, its 95-octane reformate carries 44.2 vol% of aromatics and its C5+ yield sits about 7 volume points above a published correlation (Section 9.4).

9.4 What the Plant Does Before Any Controller

9.4.1 The First Bed Takes the Heat

Bed 1 meets the naphthenes first (Figure 9.2). Across it the naphthenes fall from 22.0 to 9.4 mole percent of the hydrocarbon and the aromatics rise to 38.4, and the dehydrogenation cools the bed until its own rate slows, 70.4 K at design. Bed 2 finishes the naphthenes, down to 1.4 mole percent,

with a 45.1 K drop, and beds 3 and 4 are left with the slow conversion of paraffins, which takes heat, and hydrocracking, which gives some back, for 11.8 and 7.2 K. The consequence for control is the one the introduction stated. A hotter bed 1 reacts faster and cools itself more, so it passes on only a part of any heat given to it. In case 1, in Manual, heater 1 receives 1.6 MW more at 10 min. By 120 min its outlet is 15.7 K higher, 6.4 K of it from the exchanger, whose outlet has risen as the warmer bed 4 effluent comes around. Of that 15.7 K, bed 1 absorbs 9.3 K in a drop that grows to 79.7 K and passes 6.3 K to bed 2, the WAIT rises 10.5 K, the octane of the reformate goes from 95.0 to 97.3, and heater 1's tube metal from 580 to 610 °C. In Auto the heater loops break the loop through the exchanger.

Pause. Try Exercise 9.1 before reading on.

9.4.2 Severity Buys Octane With Yield

The WAIT moves octane and yield in opposite directions, as shown in Figure 9.3. Lined out with the coke at its new steady state, 10 K of WAIT at design buys 2.01 RON and costs 1.51 points of C5+ yield by mass. With the coke frozen at design the same step gives 2.30 RON and 1.67 points, the difference being the extra coke a hotter catalyst carries. In case 2, stepped from 500 to 510 °C at 10 min, the octane at the reactors is 97.3 by 20 min, the inferred octane passes 96 at 11 min and the analyzer's reading at 27.5 min, and by 4 h the reformate is at 97.1 with the yield down from 94.7 to 93.1 wt%, the aromatics up 2.4 vol%, the hydrogen up from 2.20 to 2.31 t/h and the fuel from 33.1 to 34.7 MW. The coke on the spent catalyst follows over hours, from 3.0 to 3.7 wt% by 10 h. One octane number of giveaway is about three quarters of a point of yield here.

> **TRY IT IN THE APP** **2 · Severity against yield: WAIT +10 °C**
>
> Load the case and press Run. At 10 min raise the WAIT setpoint from 500 to 510 °C. Watch the octane and the C5+ yield. The octane at the reactors is 97.3 by 20 min, and the analyzer's reading passes 96 only at 27.5 min. By 4 h the reformate is at 97.1 and the yield has fallen from 94.7 to 93.1 wt%.

The regenerator puts a ceiling on that curve. Lined out at $H_2/HC = 3$ and the design circulation, the coke made per hour reaches the 400 kg/h capacity at a WAIT of 523.5 °C on the design naphtha, where the octane is 100.1, and at 524.5 °C on the paraffinic naphtha, at 96.1. Above it the regenerated catalyst returns to bed 1 with coke, 3.36 wt% of it at 535 °C on the design naphtha, and five more kelvin buy about 0.5 RON where below the limit they bought 1.1 to 1.2. The laydown and the capacity carry the same compression of the catalyst loop, so the WAIT at which this ceiling appears does not depend on it. The paraffinic naphtha, with an N+2A of 48 against 62, needs 520.2 °C for 95 RON instead of 500, and at WAIT 535 °C it reaches only 96.6.

Pause. Try Exercise 9.3 before reading on.

9.4.3 How the Model's Yield Compares

Gary, Handwerk and Kaiser chart the C5+ volume yield of a reformer against the reformate's octane for naphthas of different N+2A [18]. Read by interpolation for N+2A 62, the chart gives about 84.5 volume percent at 95 RON where the model gives 91.7, and over 93 to 99 RON the model sits 7.0 to 7.6 points above the chart, the range spanning two independent readings of it. The slope agrees within the precision of the reading, 0.93 volume percent per octane number from 93 to 97 RON in the model against 0.83 to 0.87 on the chart, and the leaner naphtha costs the model 4.7 points at 95 RON against the chart's 3.8 to 4.0. The level and the hydrogen agree on the cause. The model makes 1.47 wt% of net hydrogen on feed at 95 RON, where the same source's hydrogen chart gives about 2.4 at its yield, and since each aromatic ring formed releases three molecules of hydrogen or more, the model aromatizes less of its naphtha than the units behind the chart. Its reformate at 95 RON carries 44.2 vol% of aromatics. The plan fixed a drop of about 70 K for the first bed and a few kelvin

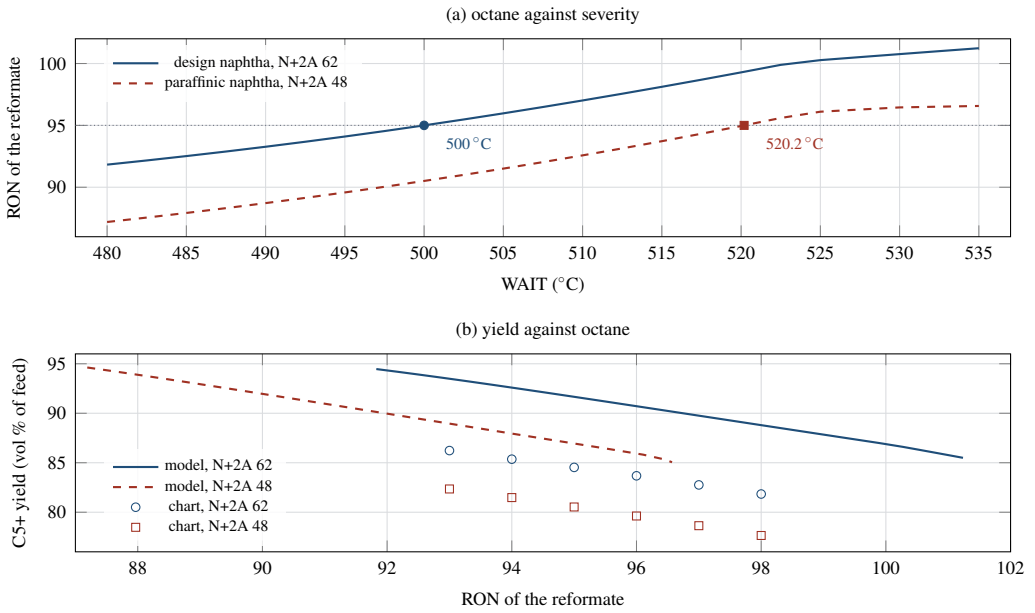


Figure 9.3 The severity trade from the application's own steady state at $H_2/HC = 3$ with the coke lined out at each WAIT from 480 to 535 °C. (a) The octane rises with WAIT on both naphthas (design solid, paraffinic dashed), and the curves bend where the regenerator reaches its burn capacity, from 523.5 °C on the design naphtha and 524.5 °C on the paraffinic one. Dots mark 95 RON. (b) C5+ volume yield against octane, with the yield read from Gary, Handwerk and Kaiser's chart for naphthas of N+2A 62 and 48 (open circles and open squares, interpolated between its curves). The model's slope is close to the chart's, its level about 7 points higher.

for the last at $H_2/HC = 3$, and those drops bound the endotherm, and with it the aromatics, that the model can spend.

9.4.4 Coke Answers in Hours

Coke is the slow variable. Lowering H_2/HC from 3 to 2 at 500 °C raises the lined-out burn demand from 239 to 334 kg/h, still under the capacity, and a cut of the ratio from 3 to 2 at 95 octane leaves the burn demand at 332 kg/h after a day and moves the WAIT the trim needs by only 0.3 K. Only where the regenerator binds does coke reach the octane, and the turndown of case 6 is therefore run at 98 octane. With the lift stopped the coke has no exit at all, and the WAIT for constant octane climbs as it does in a semiregenerative unit (case 8).

9.5 Controllers, Estimator and Optimizer as Coded

The controllers execute every 10 s, two physics steps, and the MPC every 2 min. Each execution builds the truth record, passes it through the instruments, updates the octane estimates and runs the octane layer, the WAIT layer, the heater loops and the H_2/HC loop, in that order, before writing the commands the plant's actuators follow. The header's Manual and Auto switch the four heater duty stations and the compressor station together, and in Manual every loop tracks its station so that the return to Auto is bumpless.

9.5.1 The Shared Velocity-Form PID

All four loop types use the velocity-form PID of Section 1.6.2, with the derivative taken from the measured variable, as shown in Equation (9.8).

$$\begin{aligned} e_k &= \mathrm{SP}_k - \mathrm{PV}_k, \qquad P_k = K_c e_k, \qquad D_k = -K_c \tau_D \dot{\mathrm{PV}}_k \\ s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ \hat{u}_k &= s_k + I_{k-1}, \qquad u_k = \min(u_{max}, \max(u_{min}, \hat{u}_k)) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{\tau_I} \right] \Delta t \end{aligned} \tag{9.8}$$

where the running sum s_k is never clamped and the back-calculation acts on the integral only, with a tracking time equal to τ_I, as in Equations (1.17) to (1.19). Any of these loops given an integral time of at least 600 min runs proportional-only. No loop here uses derivative action. The error is setpoint minus measurement in every loop, so a low reading raises the output.

9.5.2 Heater Loops and the WAIT Layer

Each heater outlet loop moves its duty between zero and the capacity the control system knows, the rated duty times the burners in service, at $K_c = 0.25\,\mathrm{MW/K}$ and $\tau_I = 2.5\,\mathrm{min}$ for all four. The process gain of an outlet to its duty is $1/\dot{C}_p = 6.2\,\mathrm{K/MW}$ through the 2 min firing lag. The setpoints come from one WAIT controller over a declared distribution, as shown in Equation (9.9).

$$T_{SP,b} = \mathrm{WAIT}_{SP} + s\,\delta_b + \beta, \quad \beta = \mathrm{PI}_W(\mathrm{WAIT}_{SP} - \widehat{\mathrm{WAIT}}) \tag{9.9}$$

where $\widehat{\mathrm{WAIT}}$ is the catalyst-weighted mean of the measured inlets, s the severity shift (0 for equal inlets and 1 for the aged distribution), and $\delta = (-6, -4, 0, +2.4)$ K moves severity from beds 1 and 2 to bed 4 with no change in the catalyst-weighted mean. The bias β is limited to $\pm 15\,\mathrm{K}$, its gain is 0.3 K/K with $\tau_I = 6\,\mathrm{min}$, and its integral is held while every heater sits at the limit the error pushes toward, so it cannot wind up behind them. The heater loops do not watch the tube metal. A saturated heater asked for more simply fires at its cap, and with $\theta_m = 110\,\mathrm{K}$ a WAIT setpoint stepped 30 K at once in Auto trips heater 3 3.9 min later. A tube-metal override on each heater loop would prevent that, but it would also hold the charge heater of case 7 under its limit and remove that case's lesson, so the simulator leaves the tube metals to the operator, the inlet distribution and the MPC.

9.5.3 Octane Trim, Feedforward and the Inferred Octane

The octane trim sets the WAIT setpoint from the octane the controllers read together with a feedforward on the assay of the tank in service, as shown in Equation (9.10).

$$\mathrm{WAIT}_{SP} = \mathrm{PI}_R(\mathrm{RON}_{SP} - \widehat{\mathrm{RON}}) + K_{ff}\,(62 - \overline{\mathrm{N{+}2A}}) \tag{9.10}$$

where $\overline{\mathrm{N{+}2A}}$ is the assay passed through the controllers' 6 min model of the feed line and $K_{ff} = 1.4\,\mathrm{K}$ per unit, which supplies 19.6 K for the paraffinic naphtha against the 20.2 K its lined-out octane needs. The trim's gain is 2 K/RON with $\tau_I = 20\,\mathrm{min}$, and its output limits are the WAIT range, 480 to 540 °C, less the feedforward, so its integral never winds up behind the clamp on their sum. On the analyzer the trim executes once per report, over the time since the last, holds when the analyzer is stale, and holds its integral when every heater is at the limit the error pushes toward.

The octane the controllers read is either the NIR corrected by the laboratory or an inferred octane, as shown in Equation (9.11). The NIR's correction is the laboratory bias update of Equation (1.14).

$$\begin{aligned} \mathrm{RON}_{NIR+lab} &= y_{NIR} + \hat{b}, \quad \hat{b} \leftarrow \hat{b} + \lambda\left(y_{lab} - (\bar{y}_{NIR} + \hat{b})\right) \\ \mathrm{RON}_{inf} &= \mathscr{M}(\widehat{\mathrm{WAIT}}, \overline{\mathrm{N{+}2A}}, \widehat{\mathrm{H_2/HC}}, s) + \hat{b}_s \end{aligned} \tag{9.11}$$

where $\bar{y}_{NIR}$ is the mean of the three NIR reports around the laboratory sample, $\lambda = 0.7$, and $\mathscr{M}$ is the controllers' octane map, a seventeen-term polynomial in $(\text{WAIT} - 500)/10$, $(\text{N+2A} - 62)/7$, $\text{H}_2/\text{HC} - 3$ and s (quadratic with the cube of the first and one cross term) fitted at 663 lined-out points of the commissioned plant with the catalyst as designed, with residuals of 0.020 RON root mean square and 0.103 at worst. The map gives the octane being made at the reactors with no delay. The offset $\hat{b}_s$ carries what the map cannot see, chiefly coke, and follows each corrected NIR report, compared with the map's value passed through the controllers' model of the two 10 min product lags, with a factor $1 - e^{-2/30}$ per report. A corrected report more than 0.4 RON from the inferred octane on two reports running is held as suspect, $\hat{b}_s$ stops following the NIR, each laboratory result then corrects $\hat{b}_s$ directly, the result that finds a suspect NIR re-references it with $\lambda = 1$, and the flag clears when the corrected NIR agrees within 0.2 RON. A calibration shift jumps where coke drifts, and the gate tells them apart. The flag is shown whichever octane the trim reads, but only the inferred source acts on it.

9.5.4 Model Predictive Control

The MPC is the linear MPC of Section 1.6.3 on three decisions, the WAIT setpoint W, the severity shift s and the H_2/HC setpoint H. Its model is the controllers' own lined-out maps of octane, C5+ yield, burn, the four duties and the four tube metals, linearized at every execution by central differences around the measured WAIT and H_2/HC and the present shift, made dynamic by first-order lags of 4, 4 and 1.5 min on the decisions and discretized exactly for a zero-order hold. Each predicted output carries the offset between its measurement and its map, filtered over 10 min for the duties and tube metals and taken from the octane estimator for the octane, while the burn has none and is held on the lined-out map. This output bias is the simplest way to remove offset, where Volume 9, Chapter 4 estimates an unmeasured input with a Kalman filter instead. The program over 30 steps of 2 min in blocks of 1, 1, 2, 4, 8 and 14 is Equation (9.12).

$$\begin{aligned} \min\ & \textstyle\sum_k \left[w_R (\text{RON}_k - \text{RON}_{SP})^2 - w_Y Y_k \right] + c_H \sum_j H_j \\ & + \textstyle\sum_j \Delta u_j^{\mathsf{T}} S \, \Delta u_j + \text{slack penalties} \\ \text{s.t. } & 480 \le W \le 540, \ \ 0 \le s \le 1, \ \ 2.7 \le H \le H_{cap} \\ & |\Delta W| \le 8, \ \ |\Delta s| \le 0.25, \ \ |\Delta H| \le 0.25 \\ & B_k \le B_{max} - 10, \ \ T_{m,b,k} \le T_{m,max} - 8 \\ & Q_{b,k} \le Q_{cap,b} - 0.05 \end{aligned} \tag{9.12}$$

where $w_R = 20$ per RON squared, $w_Y = 1$ per weight percent, $c_H = 0.4$ per unit of H_2/HC, $S = \text{diag}(0.5, 50, 50)$, the output limits are soft with quadratic and linear slack prices of 10^4 and 10^3, H_{cap} is the ratio the compressor can reach, and the lower bound of 2.7 keeps a margin of 0.2 over the floor of 2.5. A heater already at its duty limit has that limit dropped, its tube-metal limit kept, and the WAIT layer moves its share of the severity to the others. A solution that is not finite or not inside its bounds is rejected and the last decisions are held. A dual active-set quadratic programming module, shared by the predictive controllers of the series, solves the program. Over the challenge day the MPC solves 721 times with no fallback.

9.5.5 Trips and the Models Kept Apart

Heater protection sits apart from the loops. It reads the true tube metal, trips a heater that reaches 640 °C and stops the run there. The 620 °C limit, the 640 °C trip, the 400 kg/h capacity and the H_2/HC floor of 2.5 are teaching values. Real ones come from the tube alloy and an API 530 heater design, from the regenerator's air and oxygen system and from the licensor's operating guidelines. The application keeps three models apart. The plant is everything in Section 9.3. The identified model is what a learner fits from a step, such as the 6.2 K/MW heater gain. The controllers' model is the set of maps with the catalyst as designed, which knows neither the present coke nor the analyzer's bias.

Table 9.3 The Nine Cases

#	Case, preset and events	Teaching point
1	Manual, heater 1 duty raised by the learner	the staircase
2	WAIT setpoint raised by the learner, 500 to 510 °C	severity against yield
3	Heater loops 0.08/8 min, WAIT controller off, WAIT SP +10 at 20, heater 2 at 60 % at 120 min	heater loops, WAIT layer
4	Trim on NIR, NIR +0.8 RON at 30 min, laboratory at 150, reported 210 min	a biased analyzer
5	Trim on NIR, paraffinic naphtha at 30 min	feedforward from N+2A
6	Trim on NIR at 98 RON, compressor at 62 % at 30 min	H_2/HC, coke, regenerator
7	Paraffinic, trim on inferred with feedforward, heater 2 at 60 % at 30 min	tube metal, distribution
8	Trim on NIR, catalyst lift stops at 30 min	semiregenerative decline
9	Challenge: paraffinic at 4, heater 2 at 10, NIR +0.8 at 16 h	the whole problem

Table 9.4 Measured Outcomes by Case, Default Seed

#	Setting	Outcome
1	+1.6 MW at 10 min	bed 1 drop 79.7, outlet +6.3, heater 1 outlet 515.7 °C at 120 min
2	WAIT 510 °C	RON 97.1, C5+ 94.7 to 93.1 wt% at 4 h, coke 3.0 to 3.7 wt%
3	as shipped	WAIT dips to 501.6, holds 507.0, bed 2 20 K short
	heaters 0.25/2.5 min	dip to 505.3 °C, deficit unchanged
	+WAIT 0.3/6 min	bias 3.4 K, WAIT 510.0 °C by 180 min
4	trim on NIR	true RON 94.24, under 94.7 from 54 to 248 min
	trim on inferred	true RON never under 94.97
5	no feedforward	lowest 92.2, 112 min out of 95 ± 0.3
	feedforward	lowest 94.9, never out of the band
6	as shipped	H_2/HC 2.08, burn limit at 312 min, WAIT +2.2 K by 24 h
	feed 120 t/h at 8 h	H_2/HC 2.56, C_0 back to 0.10 wt%
7	equal inlets	heater 1 tube metal 619 to 626 °C
	aged at 45 min	under 620 °C at once, near 605, RON 95.1
8	as shipped	WAIT 499.6, 507.4, 513.2 °C at 0, 12, 24 h, bed 4 10.5 wt%

9.6 What the Nine Cases Show

The cases appear in Table 9.3. All start from the defaults of Section 9.5, WAIT only at 500 °C with equal inlets, the trim on the NIR and the design naphtha, and each row lists only what differs from them. The values below come from the engine of Section 1.6.1 advanced without a display at the default seed, with each learner move applied at the time the text gives. A case run first, after another case or from Reset produces the same trajectory. The outcomes are collected in Table 9.4.

9.6.1 Heater Loops Against Their Neighbors

Case 3 ships the heater loops detuned, 0.08 MW/K with an 8 min integral time, and the WAIT controller off. The WAIT setpoint step of 10 K at 20 min reaches the beds, but the inlets of beds 3 and 4 overshoot to about 512 °C as each warmer upstream outlet arrives at the next heater as a disturbance. At 120 min heater 2 loses burners and can give 60 % of its rated duty. The WAIT falls to 501.6 °C and stays 3 K short at 507.0 °C, with bed 2's inlet 20 K under its setpoint. Retuned loops at 0.25 MW/K and 2.5 min cut the dip to 505.3 °C and leave the deficit, because no heater loop can make heater 2 bigger. The WAIT controller at 0.3 K/K and 6 min adds a 3.4 K bias to every setpoint, beds 1, 3 and 4 follow it, and the WAIT is back at 510.0 °C by 180 min.

9.6.2 Octane Measured Late, Wrong and Rarely

In case 4 the NIR starts reading 0.8 RON high at 30 min. The trim on the analyzer does what it is told, lowering the WAIT setpoint to 495.8 °C and holding the corrected reading at 95 while the true reformate falls to 94.24, under 94.7 from 54 min. The knock-engine sample drawn at 150 min and reported at 210 min moves the analyzer's correction to −0.83, and the true octane is back over 94.7 at 248 min. On the inferred octane the validation holds the NIR as suspect from 33.1 min, two reports after the shift reaches the sample, to 211.0 min, and the true octane never falls under 94.97. At seeds

1, 2 and 3 the trim on the NIR takes the true octane down to 94.24, 94.22 and 94.19, with laboratory corrections of −1.00, −0.94 and −0.79, while on the inferred octane it never falls under 94.99.

Case 5 switches to the paraffinic naphtha at 30 min. Bed 1's drop starts to shrink as soon as the leaner feed arrives, from 70.3 to 63.9 K by 60 min even as the WAIT climbs, and the octane at the reactors passes under 94.7 at 30.5 min. The reformate follows at 38 min and the analyzer reports it at 41.5 min. The trim on the analyzer then chases a late signal, the true octane falls to 92.2 and stays outside 95 ± 0.3 for 112 min (139, 120 and 133 min at seeds 1, 2 and 3), and the WAIT settles near 519 °C. With the feedforward the WAIT rises with the feed and the lowest octane is 94.9 (94.82 to 94.92 at the four seeds). The leaner naphtha costs about 4 points of yield at the same octane, 94.7 to 90.6 wt%, and at that severity the charge heater's tube metal exceeds its limit for 11 min, reaching 621 °C.

> **TRY IT IN THE APP** **5 · A paraffinic naphtha: the octane falls before the analyzer says so**
>
> Load the case and press Run. Touch nothing. At 30 min the feed changes to a leaner naphtha. Watch the true octane and the analyzer's reading. The octane at the reactors is under 94.7 at 30.5 min, and the analyzer reports it at 41.5 min. The true octane falls to 92.2. Then turn N+2A feedforward On, press Reset and run again. The lowest octane is 94.9.

9.6.3 Hydrogen, Coke and the Regenerator

Case 6 holds 98 RON at a WAIT of 514.1 °C, since at 95 the regenerator never limits, and limits the recycle compressor to 62 % speed at 30 min. The ratio falls to 2.08, and no setpoint or speed can raise it. Lined out with the compressor at 20 to 62 %, H_2/HC is 0.71 to 2.08 and stays at 2.08 for any speed asked above 62. The coke on the spent catalyst rises from 4.1 to 5.2 wt% by 8 h, the burn demand passes the capacity at 312 min, the regenerated catalyst returns with 1.15 wt% of coke by 24 h, and the trim raises the WAIT by 2.2 K to keep 98 RON while the coke still climbs. Lined out, the WAIT needed for 98 RON against this compressor is 6.0 K higher. The feed rate is the honest remedy. Cut to 120 t/h at 8 h, it raises H_2/HC to 2.56, the burn falls under its limit at 677.5 min, and the regenerated catalyst is clean again by 24 h.

Pause. Try Exercise 9.2 before reading on.

Case 8 stops the catalyst lift at 30 min. The trim on the analyzer raises the WAIT steadily, about 0.6 K an hour, from 499.6 °C to 507.4 by 12 h and 513.2 by 24 h, holding the octane within 0.15 of 95. Bed 4, with half the catalyst and the most aromatics, cokes fastest, to 10.5 wt% by 24 h, its activity falls to 0.28, and its drop shrinks from 6.8 to 0.9 K. The C5+ yield hardly moves, 94.7 to 95.0 wt%, because the deactivated catalyst cracks less while its hotter beds crack more, whereas a platinum catalyst lost yield steadily at constant octane in a published life test [14]. Restarted at 12 h, the lift raises bed 4's activity from 0.42 to 0.62 by 24 h while the regenerator, at its limit, burns the stored coke, and the trim brings the WAIT setpoint back from 507.3 to 501.6 °C.

9.6.4 A Heater Short of Duty Near Its Tube-Metal Limit

Case 7 runs the paraffinic naphtha at 95 RON, WAIT 520 °C, with the trim on the inferred octane and the feedforward on, and the charge heater's tube metal already at 619 °C. At 30 min heater 2 drops to 60 % of its duty, bed 2's inlet falls 15.7 K, the WAIT controller adds 2.7 K to every setpoint, and heater 1's tube metal passes 620 °C within 2 min, reaching 626 °C. Selected at 45 min, the aged distribution takes bed 1 6 K and bed 2 4 K cooler and bed 4 2.4 K hotter. Bed 1 holds a tenth of the catalyst, so its cooler inlet costs little octane and spares the most duty, and heater 1's tube metal falls under the limit at once and settles near 605 °C, heater 3's rises to about 605 and heater 4's peaks at 596 °C, with the octane at 95.1.

Table 9.5 The Challenge Over 24 h at the Default Seed. Share of Judged Time Within 95 ± 0.3, Tube-Metal Peak, Burn-Demand Peak, Lowest H_2/HC and Average C5+

Settings	Band (%)	T_m (°C)	B (kg/h)	H_2/HC	C5+ (wt %)	Verdict
shipped (WAIT only, equal)	13.6	587	239	2.95	94.17	fails
trim on NIR	77.9	625	361	2.95	91.36	fails
trim on NIR, ff, aged	84.7	607	360	2.90	91.22	fails
trim on inferred, ff, equal	100	624	366	2.90	91.19	fails
trim on inferred, aged, no ff	87.7	606	364	2.93	91.27	fails
trim on inferred, ff, aged[a]	100	606	364	2.90	91.16	passes
MPC on NIR	86.4	617	383	2.61	91.23	fails
MPC on inferred	100	617	387	2.61	91.16	passes

[a]The suggested settings. Here ff is the N+2A feedforward.

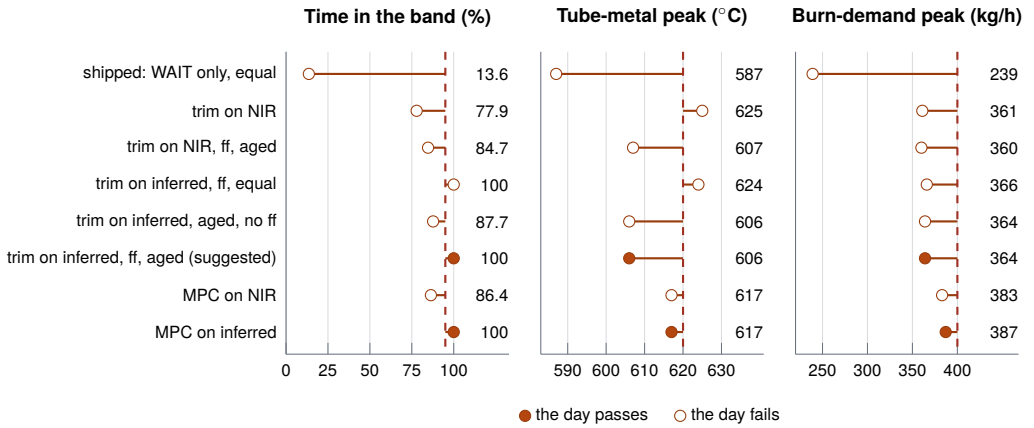


Figure 9.4 Three columns of Table 9.5 against their limits (dashed): the share of judged time within 95 ± 0.3 against 95 %, the tube-metal peak against 620 °C and the burn-demand peak against 400 kg/h. A filled circle marks a day that passes and an open circle a day that fails. Notice that every row on the analyzer stops short of the band target whatever else it has, that equal inlets put two rows over the tube-metal limit, and that the MPC runs closer to both limits at the right than the suggested trim does. The lowest H_2/HC and the C5+ yield, which no row fails, are not drawn.

9.6.5 The Challenge and the Analyzer It Cannot Trust

The challenge runs 24 h with the paraffinic naphtha from 4 h, heater 2 at 60 % from 10 h and the NIR reading 0.8 RON high from 16 h, with knock-engine samples at 2, 10 and 18 h. Its target is the true reformate octane within 95 ± 0.3 for 95 % of the time after the first hour, every tube metal at or under 620 °C, the regenerator's burn demand at or under 400 kg/h, the true H_2/HC at or over 2.5, an average C5+ yield within 0.5 points of the best the unit can make at 95 RON over the day (91.18 wt%, so 90.68 wt% or more), and no heater trip. The variants are listed in Table 9.5. Figure 9.4 draws three columns of the table against their limits.

As shipped the octane falls to 90.4 after the feed switch and never returns, 13.6 % of the time in the band. Each event needs its own layer. The feedforward handles the switch (without it the switch alone costs 169 min out of the band), the aged distribution keeps heater 1 under its limit (with equal inlets it reaches 624 °C after the heater 2 limit), and the inferred octane rides through the analyzer's shift. The suggested trim on the inferred octane with the feedforward and the aged distribution stays in the band all day, at 606 °C, 364 kg/h of burn demand, H_2/HC of at least 2.90 and 91.16 wt% of C5+, and passes at seeds 1, 2 and 3 with 91.17 to 91.20 wt%. The MPC on the inferred octane passes as well, using the room to 617 °C of tube metal and 387 kg/h of burn that the trim leaves unused, with

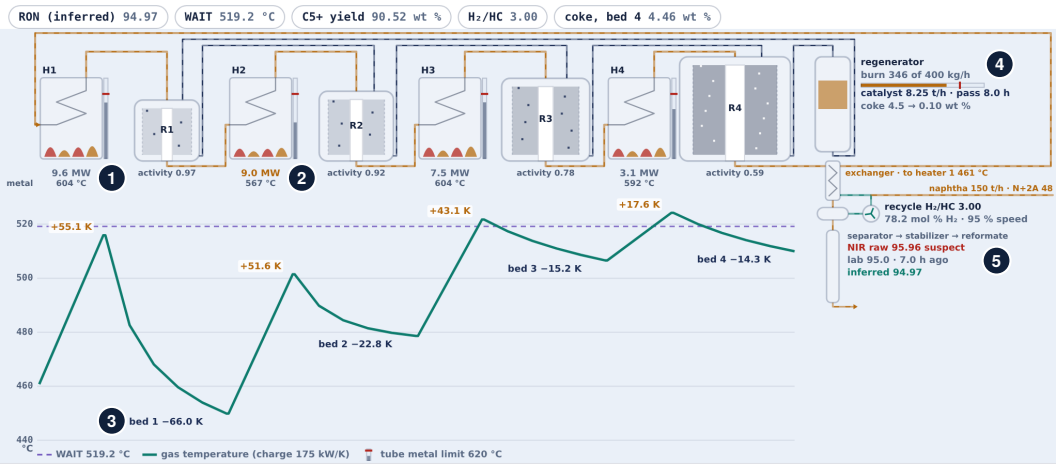


Figure 9.5 The challenge at $t = 18$ h, Unit view, cropped to the plant view with numbered callouts, with the suggested settings (octane trim on the inferred octane, N+2A feedforward, aged distribution). On the paraffinic naphtha with heater 2 held at its 9.0 MW cap (2), the staircase rises 55.1, 51.6, 43.1 and 17.6 K through the heaters and drops 66.0 (3), 22.8, 15.2 and 14.3 K through the beds for a charge of 175 kW/K, the exchanger returns the feed at 461 °C, and the tube metals read 604 (1), 567, 604 and 592 °C. The beds' activities are 0.97, 0.92, 0.78 and 0.59, and the regenerator (4) burns 346 kg/h of its 400 by its air reading, the catalyst passing at 8.25 t/h, a pass every 8.0 h, with coke falling from 4.5 to 0.10 wt%. The NIR's raw report (5), 95.96, is flagged suspect beside a laboratory value of 95.0 reported 7.0 h earlier and an inferred octane of 94.97. The WAIT stands at 519.2 °C.

the same yield.

> ✦ **DID YOU KNOW? A Better Controller Cannot Mend a Wrong Reading**
>
> From the simulator (Table 9.5). On the graded day the MPC that reads the analyzer keeps the octane in its band 86.4 % of the time and fails. A plain octane trim on the inferred octane, with the feedforward and the aged distribution, holds 100 % and passes. The MPC passes too, once it reads the inferred octane.

No strategy that believes the analyzer can pass. The judged time is 1380 min, so the rule allows 69 min outside the band, but the shift at 16 h lasts until the sample drawn at 18 h is reported at 19 h, and the trim on the corrected NIR loses 190.5 min after 16 h with or without the other layers, the MPC on it 188.5 min. The best of them, the MPC on the NIR, keeps the octane in the band 86.4 % of the time and the best trim 84.7 %. This is why the simulator offers the inferred octane as a measurement choice, and why the target is not met with the plan's 8 h laboratory alone. The challenge scores the regenerator's burn demand, which the plant caps at the capacity, so a demand over 400 kg/h is a violation even though the burn never exceeds it. The state at 18 h under the suggested settings is shown in Figure 9.5.

9.7 Technology in Practice

Figure 9.6 draws the loops of Section 9.5 with the hardware a refinery installs.

What Measures It

The bed inlets and outlets and the heater tubes are read by thermocouples, which are less accurate than resistance thermometers but more rugged and wider in span [19]. Their 12 s lag in Section 9.3 is

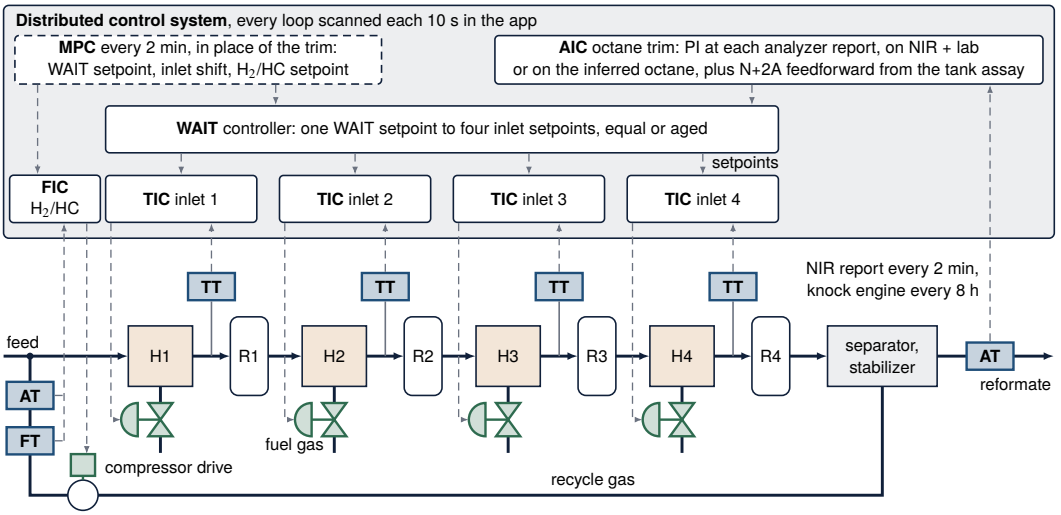


Figure 9.6 The loops of the reformer with their hardware. Lettered boxes are sensing devices, thermocouples with their transmitters (TT), the recycle flow meter (FT) and the analyzers (AT). Valve symbols and the compressor drive are final elements, and dashed arrows are signals. Notice that four fuel gas valves answer to one octane number through three controllers in cascade, and that the only measurement of that octane stands at the far end of the product path. Not to scale, and not a piping or safety diagram.

short beside the 2 min firing lag of a heater. The measurement that sets the pace is the octane, which the analyzer reads on a product that left the reactors about 20 min earlier.

> **TECHNOLOGY** **Sensor: Near-Infrared Octane Analyzer**
>
> Octane is defined by an engine test, and the research method is ASTM D2699 [20]. An online analyzer runs no engine. It passes near-infrared light through the sample, where the overtone and combination bands of the carbon and hydrogen bonds carry information on the types of hydrocarbon, and a multivariate model turns the spectrum into a property. A measurement takes a minute or less. The method is a secondary one: its model is built from samples rated by the reference method, may need hundreds or even thousands of them, and has to be updated frequently as the sample matrix changes [21]. A reading can therefore repeat well and still be wrong, by the error of its model. The bias b_{NIR} of Section 9.3 stands for that error, and the knock engine every 8 h for the reference method.

What Moves It

Four fuel gas valves carry the severity, the drive of the recycle compressor carries the hydrogen ratio, and the catalyst lift carries the circulation. A review of commercial practice describes the pattern the model has [13]. Heaters stand before and between the reactors, the temperature falls fastest in the first reactor, and the first reactor is the smallest while the third and fourth are very large. Its table of typical conditions for three reactors gives temperature drops of 69, 31 and 6 K, beside the 70.4, 45.1, 11.8 and 7.2 K of the model's four beds.

TECHNOLOGY Final Element: Fuel Gas Valve of a Fired Heater

A heater's duty is moved through a control valve on the fuel gas to its burners. At one refinery reactor furnace the inlet temperature controller set the furnace temperature controller, which in turn set the fuel gas pressure [22]. A pneumatic valve that air opens fails closed, and the failure mode is to be chosen from the risks of the process and not from habit. Beside the control valve stands a burner management system. It watches for low fuel pressure, high fuel pressure and loss of flame, and stops the fuel with shutoff valves of its own. Where a master controller can call for a setpoint that is unsafe, a limit function belongs between the two controllers [19]. The model's heater is a 2 min lag on its duty, capped by the burners in service, and its loops carry no tube-metal limit.

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 [23]. The app scans every loop each 10 s and solves its MPC every 2 min, inside that range. The four heater controllers are PID blocks, the WAIT controller sets their setpoints and the octane trim sets the WAIT setpoint, a cascade three deep. Industrial predictive controllers meet the two awkward cases of this chapter with standard features. When an analyzer delivers no new value, the typical answer is to skip the bias update for a number of intervals. When a lower-level controller is saturated, a temporary hard constraint keeps the move from going further the wrong way [6]. The chapter's trim holds on a stale analyzer, and its MPC drops the duty limit of a heater already at its cap.

IN PRACTICE A Regenerator That Burns on a Leash

The model's regenerator is one vessel with a burn limit. A real one works in steps, which a federal rule lists as depressurization, purge, coke burn-off, catalyst rejuvenation with a chloride compound and a final purge [24]. The burn is held back on purpose. Oxygen and air flow are regulated so that the coke burns under control, because higher temperatures cost the catalyst surface area. The screens that hold the moving catalyst plug with fines, which leaves coke unburned [7]. In the United States, 40 CFR Part 63, Subpart UUU limits the hydrogen chloride vented during coke burn-off and rejuvenation, as a percent reduction or a concentration, and requires continuous monitoring and a written operation, maintenance and monitoring plan [24].

What Is Modernizing It

The first development is the spectroscopic analyzer, which replaced a slow engine test by a spectrum read in a minute. Fiber optics let one instrument watch several points of a plant up to 100 m away [21]. The second is the soft sensor, of which the inferred octane is one. A review of such models names a use beyond prediction, the detection of a failed sensor and a back-up for it. The same review warns that after launch the prediction usually deteriorates as the process drifts, catalyst deactivation among the causes, and that telling a process drift from a sensor drift takes expert knowledge. A sensor drift calls for recalibration and for no corrective action on the process [25]. The trim of case 4 moved the process instead. The third is predictive control, of which a 2003 survey counted more than 4600 applications, with refining holding 67 % of those classified [6].

9.8 Limitations

Every size, rate constant, limit and the octane correlation are generic teaching values for a 150 t/h low-pressure continuous reformer, chosen so that the design point gives the plan's bed drops, and none was fitted to a unit. The kinetics depart from Smith's in two declared ways, a dehydrogenation constant 8.5 times his and hydrocracking first order in hydrogen. The model's 95-octane reformate carries less aromatics and its C5+ yield is about 7 volume points higher than a published correlation gives, so the slope of the trade is right and its level is not. One coke state per bed, a well-mixed regenerator and a catalyst loop compressed about twelve times make the coke lessons visible within a

day at the price of coke rates twelve times a real unit's. An aging catalyst here loses almost no yield. Pressure drop, the heater fireboxes, the fuel header, the separator pressure loop and the stabilizer's composition profile are outside the model. No refinery switches its naphtha as a step or shifts an analyzer by exactly 0.8 RON, but clean events let each effect in the cases be attributed to one cause.

The inferred octane uses maps of the commissioned plant, so it learns coke only through its offset, and its validation is a jump test against the analyzer. The MPC linearizes those maps at each execution and holds the burn on the lined-out map, its limits are soft, and it carries no guarantee beyond its hour. The heater loops do not protect the tube metals, which only the MPC, the distribution and the operator do, and the trip reads the truth and stops the run rather than model a restart. The challenge scores the burn demand, not the capped burn, and sets no price on hydrogen or fuel.

> ⚠ **OPEN CHALLENGE** **Octane With Less Benzene**
>
> A reformer makes octane by making aromatics, and the lightest aromatic is regulated. A review of 2003 called the limits on benzene and vapor pressure that followed the Clean Air Act Amendments of 1990 the biggest challenge facing reforming, one that forced a complete rethink of the process. It expected modified catalysts, new processes after the reformer and reluctance to reform the full range of naphtha [13]. The federal fuel rule now holds a manufacturer's gasoline to an average of 0.62 volume percent benzene [26]. The four lumps of this chapter fold benzene precursors and heavier aromatics into one aromatic, so its octane target cannot see the limit. Progress would be a severity target that weighs octane, yield, hydrogen and benzene together.

9.9 Conclusion

A continuous catalytic reformer is a train of endothermic beds whose first bed does most of the work, run on an octane that is never measured where it is made, and this simulator computes the cell temperatures, the coke on each bed, the analyzer, the laboratory and an inferred octane together from a forty-three-state model. At design the beds cool the gas by 70.4, 45.1, 11.8 and 7.2 K, and heat added to the charge heater is mostly absorbed by bed 1, 9.3 K of a 15.7 K rise. Ten kelvin of WAIT buy 2.01 RON for 1.51 points of yield, until the regenerator's capacity, reached near 100 RON on the design naphtha, cuts what each further kelvin buys. A trim that trusts a shifted analyzer holds its reading while the true octane falls to 94.24, a feedforward from N+2A keeps a leaner naphtha's octane above 94.8 where feedback alone loses 112 min, and the aged distribution keeps the charge heater under its tube-metal limit. The graded day passes with the inferred octane and fails on any strategy that believes the analyzer, because the shift outlasts what the target allows. The octane correlation is illustrative and the yield level high. Further work would add plug-flow beds with a coke profile, a regenerator with burn and chlorination zones, a pressure and stabilizer model, and an octane correlation fitted to measured reformates. The simulator shows that a reformer is controlled through its temperature profile and its measurements as much as through its chemistry.

KEY POINTS

- A reformer cools its own feed, so four fired heaters must put the heat back between the beds.
- The first bed absorbs most of a heater step, 9.3 K of a 15.7 K rise.
- Ten kelvin of WAIT buy 2.01 RON for 1.51 points of yield, until the regenerator reaches its burn capacity.
- A trim that trusts a shifted analyzer holds its reading while the true octane falls to 94.24.
- The graded day passes with the inferred octane and fails on any strategy that believes the analyzer.

Where else this applies. A soft sensor corrected by a delayed laboratory returns in Volume 5, Chapter 10, the same pair of octane instruments in Chapter 15, beds in series that amplify where these damp in Chapter 7, and an optimum on a moving constraint in Volume 9, Chapter 8.

References

[1] P. Doolin, D. Zalewski, and S. Oyekan, "Catalyst regeneration and continuous reforming issues," in *Catalytic Naphtha Reforming*, 2nd ed., ser. Chemical Industries, G. J. Antos and A. M. Aitani, Eds. New York: Marcel Dekker, 2004, vol. 100, pp. 414–435.

[2] L. Turpin, "Control systems for commercial reformers," in *Catalytic Naphtha Reforming*, 2nd ed., ser. Chemical Industries, G. J. Antos and A. M. Aitani, Eds. New York: Marcel Dekker, 2004, vol. 100, pp. 470–518.

[3] U. Taskar and J. B. Riggs, "Modeling and optimization of a semiregenerative catalytic naphtha reformer," *AIChE Journal*, vol. 43, no. 3, pp. 740–753, 1997.

[4] D. Li, S. L. Shah, T. Chen, and K. Z. Qi, "Application of dual-rate modeling to CCR octane quality inferential control," *IEEE Transactions on Control Systems Technology*, vol. 11, no. 1, pp. 43–51, 2003.

[5] 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.

[6] 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.

[7] A. Gupta and S. K. Gupta, "Catalyst regeneration techniques in naphtha reforming: Short review," *Chemical and Process Engineering*, vol. 43, no. 2, pp. 101–108, 2022.

[8] K.-m. Liang, H.-y. Guo, and S.-w. Pan, "A study on naphtha catalytic reforming reactor simulation and analysis," *Journal of Zhejiang University-SCIENCE B*, vol. 6, no. 6, pp. 590–596, 2005.

[9] A. J. Lazar and J. H. Klatt, "Modern control of a catalytic reformer in distributed control," in *Proc. 1983 American Control Conference*, Jun. 1983, pp. 7–13.

[10] V. Haensel and M. J. Sterba, "Comparison of Platforming and thermal reforming," in *Progress in Petroleum Technology*, ser. Advances in Chemistry. American Chemical Society, 1951, pp. 60–75.

[11] R. Larraz and R. Arvelo, "Modeling catalytic naphtha reforming," in *Catalytic Naphtha Reforming*, 2nd ed., ser. Chemical Industries, G. J. Antos and A. M. Aitani, Eds. New York: Marcel Dekker, 2004, vol. 100, pp. 519–561.

[12] V. Haensel, "Process of reforming a gasoline with an alumina-platinum-halogen catalyst," U.S. Patent 2 479 110, Aug. 16, 1949.

[13] U. T. Turaga and R. Ramanathan, "Catalytic naphtha reforming: Revisiting its importance in the modern refinery," *Journal of Scientific & Industrial Research*, vol. 62, no. 10, pp. 963–978, 2003.

[14] H. E. Kluksdahl, "Reforming a sulfur-free naphtha with a platinum-rhenium catalyst," U.S. Patent 3 415 737, Dec. 10, 1968.

[15] E. L. Pollitzer, "Continuous Platforming," *Platinum Metals Review*, vol. 20, no. 1, pp. 2–6, 1976.

[16] J. J. Kelly, C. H. Barlow, T. M. Jinguji, and J. B. Callis, "Prediction of gasoline octane numbers from near-infrared spectral features in the range 660–1215 nm," *Analytical Chemistry*, vol. 61, no. 4, pp. 313–320, 1989.

[17] K. Moljord, H. G. Hellenes, A. Hoff, I. Tanem, K. Grande, and A. Holmen, "Effect of reaction pressure on octane number and reformate and hydrogen yields in catalytic reforming," *Industrial & Engineering Chemistry Research*, vol. 35, no. 1, pp. 99–105, 1996.

[18] J. H. Gary, G. E. Handwerk, and M. J. Kaiser, *Petroleum Refining: Technology and Economics*, 5th ed. Boca Raton, FL: CRC Press, 2007.

[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] U.S. Federal Trade Commission, "Automotive fuel ratings, certification and posting," Code of Federal Regulations, Title 16, Part 306, section 306.0, edition of January 1, 2026, 2026, govinfo.gov.

[21] C. Pasquini, "Near infrared spectroscopy: fundamentals, practical aspects and analytical applications," *Journal of the Brazilian Chemical Society*, vol. 14, no. 2, pp. 198–219, 2003.

[22] U.S. Environmental Protection Agency, "EPA chemical accident investigation report: Tosco Avon refinery, Martinez, California," U.S. Environmental Protection Agency, Chemical Emergency Preparedness and Prevention Office, Washington, DC, Tech. Rep. EPA 550-R-98-009, 1998.

[23] 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.

[24] U.S. Environmental Protection Agency, "National emission standards for hazardous air pollutants for petroleum refineries: Catalytic cracking units, catalytic reforming units, and sulfur recovery units," Code of Federal Regulations, Title 40, Part 63, Subpart UUU, sections 63.1566, 63.1567 and 63.1579, edition of July 1, 2025, 2025, govinfo.gov.

[25] 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.

[26] U.S. Environmental Protection Agency, "Regulation of fuels, fuel additives, and regulated blendstocks," Code of Federal Regulations, Title 40, Part 1090, section 1090.210, edition of July 1, 2024, 2024, govinfo.gov.

9.10 Exercises

9.1 Concept check · 2 min

In case 1 the reformer is in Manual and heater 1 receives about 1.6 MW more, enough for about 10 K at its outlet. Two hours later heater 1's outlet has risen 15.7 K, yet bed 1's outlet has risen only 6.3 K. Why does so little reach bed 2?

(a) The hotter bed reacts faster, and its endothermic reactions keep most of the heat in a larger drop

(b) Heater 1 is held at its rated duty, so most of the extra firing never reaches the gas at all

(c) The feed–effluent exchanger takes heat back out of heater 1's inlet as the outlet warms

(d) The recycle hydrogen rises with the temperature and spreads the added heat over more gas

9.2 Concept check · 2 min

In case 6 the recycle compressor is limited to 62 % speed while the unit takes 150 t/h of naphtha, and H_2/HC falls to 2.08, under its floor of 2.5. Which move brings the ratio back over the floor?

(a) Raise the H_2/HC setpoint, so that the loop asks the compressor for more speed

(b) Cut the naphtha feed rate, which lowers the hydrocarbon moles in the ratio

(c) Raise the WAIT, which makes more hydrogen and so raises the recycle purity

(d) Raise the catalyst circulation, so that the regenerator burns more coke

9.3 Concept check · 1 min

Lined out on the design naphtha at $H_2/HC = 3$, five kelvin of WAIT buy 1.1 to 1.2 RON below a WAIT of 523.5 °C and about 0.5 RON above it. What happens at that WAIT?

(a) The octane correlation stops being linear in the aromatics above 100 RON

(b) Heater 4 reaches its rated duty and can no longer lift bed 4's inlet

(c) The coke made reaches the regenerator's burn capacity, and coke returns to bed 1

(d) Dehydrogenation reaches equilibrium in bed 1, so no more aromatics can form

9.4 Theory and calculation · 7 min

At the design point every bed inlet is at 500 °C, the bed drops are 70.4, 45.1, 11.8 and 7.2 K, the charge carries $\dot{C}_p = 160.5\,\mathrm{kW/K}$, and the heaters fire 8.72, 11.30, 7.25 and 1.90 MW of their rated 12, 15, 10 and 5 MW, with $\theta_m = 110\,\mathrm{K}$.

(a) Compute each bed's reaction heat and the total endotherm. *Checkpoint:* 11.30, 7.24, 1.89 and 1.16 MW; 21.6 MW.

(b) Compute the four tube metals from Equation (9.4). *Checkpoint:* 579.9, 582.9, 579.8 and 541.8 °C.

(c) In case 7 all four inlets are at 520.2 °C. What share of its rated duty can a heater fire before its tube metal reaches 620 °C, and what duty is that for heater 1? At $t = 0$ heater 1 fires 10.62 MW. What is its tube metal, and at what outlet would a heater at rated duty reach the 640 °C trip? *Checkpoint:* 0.907 and 10.89 MW; 617.5 °C; 530 °C.

9.5 Theory and calculation · 6 min

The four beds hold 66 t of catalyst, circulated at $W_c = 8.25\,\mathrm{t/h}$. A full burn leaves $C_r = 0.1$ wt% and the regenerator can burn $B_{max} = 400\,\mathrm{kg/h}$.

(a) How long does the catalyst take to pass once through the beds? *Checkpoint:* 8.0 h.

(b) At design the spent catalyst leaving bed 4 carries 2.99 wt%. Use Equation (9.6) to compute the burn and its share of the capacity. *Checkpoint:* About 238 kg/h, 60 %.

(c) At what coke on the spent catalyst does the burn reach its capacity? *Checkpoint:* 4.95 wt%.

(d) In case 6 the spent catalyst carries 6.02 wt% at 24 h. What coke does the regenerated catalyst return to bed 1 with, once the regenerator has lined out? *Checkpoint:* 1.17 wt%.

9.6 Theory and calculation · 6 min

The design naphtha carries 21 vol% of naphthenes and 20.5 vol% of aromatics, the paraffinic one 16 vol% of each. The trim's feedforward is Equation (9.10) with $K_{ff} = 1.4\,\mathrm{K}$ per unit of N+2A.

(a) Compute N+2A for both naphthas and the feedforward after the switch. Lined out, the paraffinic naphtha needs a WAIT of 520.2 °C for 95 RON. How much must the trim's own output supply? *Checkpoint:* 62 and 48; 19.6 K; 0.6 K.

(b) The trim's output limits are the WAIT range, 480 to 540 °C, less the feedforward. Give the trim's own range on the paraffinic naphtha, and say why the limits are written this way. *Checkpoint:* 460.4 to 520.4 °C.

(c) The challenge judges 1380 min and requires the octane in 95 ± 0.3 for 95 % of it. The NIR shifts at 16 h, and the next knock-engine sample is drawn at 18 h and reported an hour later. How many minutes may the run spend outside the band, and for at least how long does a trim that believes the NIR act on the shifted reading? *Checkpoint:* 69 min; 180 min.

9.7 App challenge · 8 min

Load case 1, "The staircase: four drops, four rises (Manual)", at the default seed. Before pressing Run, raise Heater 1 duty (TIC · four heater outlets card) from 8.72 to 10.30 MW. Run to 120 min at speed 20× (about 6 s) and export the CSV.

(a) From the columns `bed1_in_C` and `truth_drop1_K`, read heater 1's outlet and bed 1's drop at 0 and 120 min. How much of the outlet's rise does bed 1 keep in its drop, and how much reaches bed 1's outlet? *Checkpoint:* 500.0 to 515.7 °C (±0.8 K); drop 70.39 to 79.65 K. About 9.3 K is kept and 6.3 K passed on.

(b) The extra 1.58 MW over $\dot{C}_p = 160.5\,\mathrm{kW/K}$ accounts for about 9.8 K. Where does the rest of the outlet's rise come from, and why does the outlet keep climbing for an hour after the step? *Checkpoint:* The feed–effluent exchanger, about 6 K. The heat goes round the train and back.

9.8 App challenge · 9 min

Load case 4, "The octane trim on a late analyzer and an 8 h laboratory", at the default seed and speed 20×. It loads with Severity control on Octane trim and Octane measurement on NIR + lab, and the NIR starts reading 0.8 RON high at 30 min.

(a) Run as shipped to 300 min and export the CSV. In `truth_ron_product`, find the lowest true octane, the first row under 94.7 and the first row back over 94.7. Read `nir_bias_est` after the laboratory result at 210 min. *Checkpoint:* 94.24 (±0.02); 53.5 min; 248 min; −0.83.

(b) Select Inferred + lab in Octane measurement, press Reset and run to 300 min again. What is the lowest true octane now, and how does the WAIT setpoint (`wait_sp_C`) compare with run (a)? *Checkpoint:* 94.97. It stays near 500 °C where run (a) lowered it to 495.8 °C.

(c) Why is the change in (b) followed by Reset, rather than made at 60 min into run (a)? *Checkpoint:* By then the trim has already lowered the WAIT and the octane.

9.9 App challenge · 8 min

Load case 9, the challenge, at the default seed. Before pressing Run, set Severity control to Octane trim, N+2A feedforward to On and Inlet distribution to Aged, and leave Octane measurement on NIR + lab. Set the speed to 300× (the day takes about 5 s) and run to the stop at 24 h.

(a) Read the verdict. Which targets are met and which are missed? *Checkpoint:* Missed: 84.7 % in the band. Met: tube metal 607 °C, burn demand 360 kg/h, H_2/HC 2.89, C5+ 91.22 wt%.

(b) Select Inferred + lab, press Reset and run the day again. Read the verdict, and say which of the day's three events the change dealt with. *Checkpoint:* Passes: 100.0 %, 606 °C, 364 kg/h, 2.90, 91.16 wt%. It deals with the NIR shift at 16 h.

9.10 Application area · 4 min

The heater loops of this simulator do not watch the tube metal, and a WAIT setpoint stepped 30 K at once in Auto trips heater 3 3.9 min later. On a real reformer, name two layers that keep a tube metal under its limit, say where the limit itself comes from, and explain why the simulator leaves a tube-metal override out of its heater loops.

9.11 Application area · 4 min

Near design one octane number of giveaway costs about three quarters of a point of C5+ yield.

(a) On 150 t/h of naphtha, what does a trim that holds the reformate 0.5 RON above target cost in reformate per hour? *Checkpoint:* About 0.56 t/h.

(b) The model's yield sits about 7 volume points above a published chart while its slope agrees. Does the answer to (a) survive that, and why? *Checkpoint:* Yes, because it depends on the slope.

9.12 Review the agent · 5 min

An AI assistant reviews case 4 and writes: "The true octane falls because the NIR reports about 20 min late. Raise the octane trim gain from 2 to 4.0 K per RON and cut its integral time from 20 to 10 min so the trim corrects sooner, and the octane will stay within 0.3 of 95." Find the flaw, and say how you would check the claim in the app.