INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 11

A Delayed Coker Whose Drum Switch Is a Disturbance Known in Advance

Pages 262–283

Edited by John D. Hedengren

How to cite this chapter. "A Delayed Coker Whose Drum Switch Is a Disturbance Known in Advance," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 262–283.

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

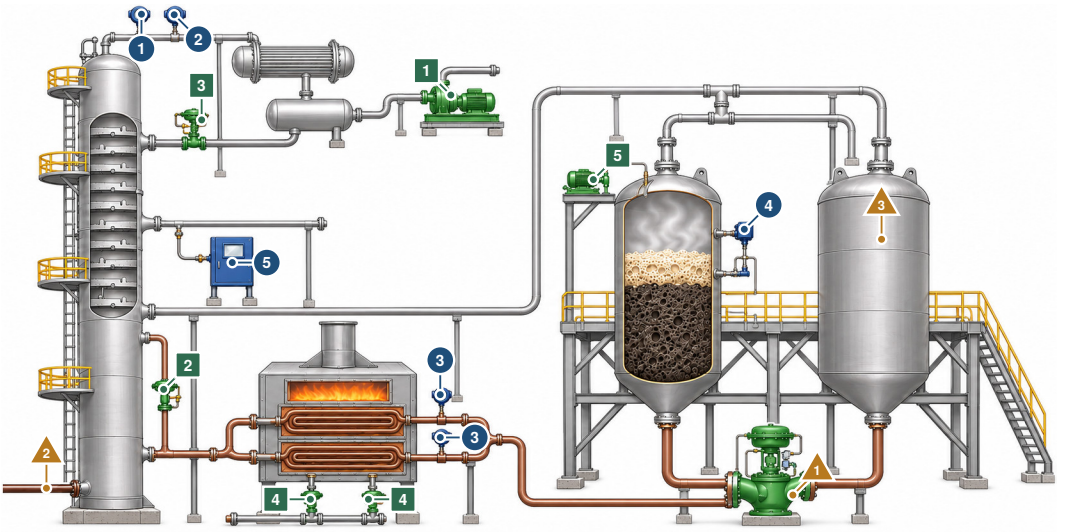
The app for this chapter
https://apmonitor.com/apps/delayed-coker/

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

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

CHAPTER 11

A Delayed Coker Whose Drum Switch Is a Disturbance Known in Advance



The controller moves 1 wet gas compressor 2 valve on the wash recycle 3 top reflux valve 4 fuel to the two furnace passes 5 antifoam pump

The controller reads 1 fractionator pressure 2 overhead temperature 3 two coil outlet temperatures 4 foam level on the live drum 5 analyzer on the heavy gas oil

What disturbs it 1 the scheduled drum switch 2 a heavier, foamier residue 3 a cold drum after a skipped warm-up

Notice the switch valve below the drums: at each switch it moves the charge from the cut-open live drum, where foam rides on the coke bed, to the closed drum beside it, which is colder than the vapor it will receive. Illustration, not to scale.

IN PLAIN TERMS Two Drums That Take Turns

The main idea. A delayed coker heats the heaviest part of crude oil and lets it break apart for hours in a tall drum. Solid coke stays behind, and the vapor goes on to a column that splits it into fuels.

The challenge. When a drum is full, the flow moves to an emptied drum that is still cooler than the oil. That drum soaks up heat and holds back vapor, and the column loses pressure unless its controls act on the known switch time.

Why it matters. The switch is the most predictable upset in the plant, yet only warming the new drum first saves the vapor.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History box in Section 11.2, the Did you know? boxes in Sections 11.1 and 11.5, the Try it boxes in Section 11.3, and Key points in Section 11.8. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 11.5 with Figure 11.4, Section 11.6 and Section 11.7 with its Open challenge box.

The control question A delayed coker's two drums take turns on a fixed cycle while its furnace and fractionator run without a pause, and every switch sends the charge into a drum colder than its vapor. How much of that upset can control remove when the schedule that causes it is known hours ahead, and what can only the drum's warm-up prevent?

Learning objectives

- Compute the coke made, the rise of the bed and the outage left at the end of a fill from the yield correlation and the drum size.
- Calculate the vapor a cold drum holds back from its shell temperature, at the switch and in total.
- Predict the first compressor move of a feedforward built from the schedule and the shell thermocouples.
- Compare loops, schedule feedforward and MPC on a cold switch, and explain why grading the heavy gas oil cut at every instant separates them.
- Weigh the coil outlet temperature against coke yield and furnace run length, and the antifoam dose against foam-over.

Prerequisites Sections 1.4.2 (heaters and tube metal), 1.4.3 (cut points), 1.4.6 (analyzer delay), 1.6.2 (velocity form) and 1.6.4 (scenarios). Section 1.6.3 is needed only for the MPC. Chapter 5 makes the residue this unit cracks.

The app https://apmonitor.com/apps/delayed-coker/

See also Volume 3, Chapter 6 for feedforward from a measured load beside a backwash sequencer, Volume 3, Chapter 7 for a foam limit that moves with pressure and Volume 3, Chapter 4 for heater outlet control. Volume 8, Chapter 6 for a recipe the controller knows in advance and Volume 8, Chapter 13 for a continuous plant that a diversion sequence interrupts.

11.1 Introduction

The coke drum is a batch reactor that sits inside a continuous refinery unit. In the generic coker simulated here, 120 t/h of vacuum residue and 12.0 t/h of recycle leave a two-pass furnace at a coil outlet temperature (COT) of 495 °C, and the charge cracks in whichever of two drums is on line. About 36.6 t/h of coke stays behind, so the bed rises 1.12 m an hour, and after 16 h the switch valve sends the charge to the other drum, which has meanwhile been steamed, quenched, cut out and warmed. A drum whose shell sits 300 K below the vapor it receives holds back 16.2 t/h of vapor at the switch, and the fractionator, which sees only what leaves the drums, finds out first.

> ✦ DID YOU KNOW? **Coking Capacity in 2003**
>
> On January 1, 2003, the coking units of United States refineries could take 2,243,947 barrels per calendar day, 13.5 % of the country's crude distillation capacity [1]. The coker of this chapter takes 120 t/h of vacuum residue.

The simulator is built around one belief it is meant to correct, that a recurring disturbance belongs to feedback. A switch is the most predictable upset in the plant, because the sequencer sets its time and the shell thermocouples measure how cold the new drum is. On a switch onto a drum left cold by a skipped warm-up (case 8), the fractionator pressure falls 0.178 bar under its loops alone, while a schedule feedforward that moves the wet gas compressor with the predicted gas keeps the dip to 0.024 bar. A second belief follows from the first, that a feedforward which holds the pressure has fixed the switch. The cold drum keeps 15.4 t of vapor in the drum over four hours whatever the loops do, and only the warm-up moves that heat to a time the fractionator can spare it. The same shape, a batch step on a known schedule feeding a continuous section, appears in bed regenerations, filter backwashes, dryer and adsorber swaps and cyclic reactors, and the integration of production schedules with process control is reviewed as a field of its own [2]. Delayed coking itself is reviewed in [3, 4], and model predictive control of the main fractionator of a cracking unit in [5].

However, the published material on cokers separates the drum, whose kinetics and bed structure are modeled without a fractionator [6–8], from the furnace, whose coil fouling is treated as a run-length problem [9], and from the drum's mechanical life [10]. This chapter puts the drum, the furnace and the fractionator on one simulated coker whose sequencer, drum shells, foam, furnace coils and fractionator loops act on one another, and it measures what changes when the schedule is handed to the controllers and what does not.

The yield correlation, the velocity-form PI controller, the feedforward and the linear model predictive controller (MPC) are established methods, and the results are measurements on one synthetic unit. One of them carries the chapter. A pressure PI with five times the shipped gain also holds the challenge's pressure band, yet it leaves the heavy gas oil cut 9.9 K off at a cold switch, so the challenge, which grades that cut at every instant, needs the feedforward or the MPC. The model has 41 states, its charge balance over each drum cycle closes to 2.6×10^{-13} of the charge, and its controllers see instruments only, never the pitch, the true foam or the coil coke. Each of the nine cases credits its completion by a rule written against the plant state (Section 1.6.4).

Section 11.2 follows delayed coking and its drum cycle. Section 11.3 sets out the unit as modeled, from the yields and the live drum through the sequencer, the warm-up, the foam and the furnace coils to the fractionator, its pressure and the instruments. Section 11.4 describes the control layers, which are five PI loops, the schedule feedforward and the MPC. Section 11.5 reports the nine cases, from one fill by hand to two graded days on the unit, Section 11.6 sets the hardware of a real unit beside the model, Section 11.7 states what the model leaves out, and Section 11.8 draws the lessons together.

11.2 Delayed Coking and Its Drum Cycle

Delayed coking grew out of the first thermal cracking processes, as the History box below tells. The drums make the unit a batch process. A full drum leaves the line to be steamed, quenched, drained and emptied of its coke, and is warmed before it fills again, the sequence the simulator's sequencer runs in Section 11.3.

> ⧖ HISTORY **1913: From Unwanted Coke to Coke Drums**
>
> The first thermal cracking process was developed around 1913. Distillate fuels and heavy oils were heated under pressure in large drums until they cracked, and the method made large amounts of solid, unwanted coke. That process evolved into visbreaking, steam cracking and coking, which turns heavy residue into lighter products and by-product coke [11]. In delayed coking the furnace is designed and controlled to prevent premature coking in its tubes, and large coke drums give the long residence time the cracking needs. When the coke reaches a set level in one drum, the flow is diverted to another to keep the unit running, and the full drum is steamed, cooled with water and cut out by high-pressure water from a rotating cutter [11]. Fuel demand and ever heavier crudes renewed interest in the bottom of the barrel, and delayed coking, able to take virtually any refinery stream, gained in importance [3].

Yields came to be predicted from the feed. Refining texts give coke, gas and liquid yields as functions of the residue's Conradson carbon residue (CCR) [12], the residue's chemistry sets how it cracks [13], and an industry consortium studied the fundamentals of the process, from feed and operating conditions to coke morphology [14]. Inside the drum, kinetic models describe the deep thermal cracking [6, 7] and the bed that results is far from uniform [8]. Two limits outside the drum shape the cycle. The furnace coils foul with coke, so the furnace runs until its tube metal approaches its limit and is then decoked, and antifoulants are judged by laboratory tests against field experience [9]. The heater of Chapter 5, whose residue this unit takes, cokes its coil at a rate that doubles with every 10 K of film temperature, as the coils of Section 11.3.4 do. The drums themselves fatigue under the temperature swing of every cycle, so throughput is weighed against drum life [10].

Control followed the general history of refinery automation. Each draw, the furnace outlets and the overhead pressure received their own loops, and model predictive control reached fractionators of cracking units by the 1990s [5, 15]. The drum switch remained an operating event for the regulatory

Table 11.1 Variables and Their Roles at the Design Point

Variable	Role	Design	Range
P	CV, overhead pressure	1.000 bar g	SP 0.6–1.6
T_{95}	CV, HCGO 95 % point	480.0 °C	SP 460–500
T_{top}	CV, overhead temperature	130.0 °C	SP 110–150
COT_i	CV, coil outlets	495.0 °C	SP 470–515
h_{out}	CV (on foam level), outage	18.45 m	SP 2–10
u	MV, compressor	75.2 %	0–100
R	MV, recycle	10.0 % of F_f	0–30
L_r	MV, top reflux	40.0 t/h	0–150
Q_i	MV, absorbed duty per pass	10.10 MW	0–13
D	MV, antifoam	3 ppm	0–25
L	set, LCGO draw	22.4 t/h	0–40
F_f	DV, fresh residue	120 t/h	80–140
CCR, φ	DV, carbon, foaming	17.5 wt %, 1	12–26, 0.5–4
—	DV, compressor available	100 %	50–100
t_{cyc}, t_w, f_w	schedule	16 h, 4 h, 8 %	10–24 h
M_k, M_p, T_m, h_f	states, per drum	—	—
$Q_i, \mathrm{COT}_i, \delta_i$	states, per pass	—	—
T_N, T_L, T_R, T_{fz}	states, cut and flash zone	—	—

loops, while the textbook answer to a measured disturbance, feedforward [16], applies directly once the schedule is treated as a measurement, as Volume 3, Chapter 6 computes its coagulant dose from a raw water turbidimeter rather than wait for the settled water. Coordinating a schedule with the loops it disturbs is now studied in general [2], and the coker is one of its clearest cases because the schedule is fixed by the decoking crew and known to the minute.

11.3 The Unit as Modeled

The unit is generic and taken from no plant, as shown in Figure 11.1. Fresh residue F_f enters the fractionator's bottoms, where the wash adds the recycle R, and the charge $F_c = F_f + R$ is pumped through two furnace passes A and B, each taking half of it with no pass flow balancing. The passes join in the transfer line, the switch valve sends the charge to the live drum, and the drum's vapor goes to the flash zone of the main fractionator. The off-line drum follows its own sequence, and for the last 4 h of a 16 h cycle part of the live vapor is drawn through it to warm its shell. The fractionator washes the heaviest ends back into the bottoms as recycle and draws heavy and light coker gas oil (HCGO and LCGO), naphtha and the wet gas that the compressor takes. Twenty-nine physical states and twelve ledgers are integrated, the ledgers being the running totals of charge, vapor, condensate, coke laid and cut, products, recycle, antifoam, fuel, gas and liquid that let conservation be checked. The clock is in minutes, flows in t/h, duties in MW and pressures in bar g.

Table 11.1 classifies the variables and gives their values at the design point, case 3 at $t = 0$, four hours into drum A's fill.

11.3.1 Drum Yields and the Live Drum

The charge cracks in the live drum at yields set by a correlation of the published form [12], written on the residue's Conradson carbon, the COT and the drum pressure, as shown in Equation (11.1).

$$\begin{aligned} y_c &= 0.016\,\mathrm{CCR} - k_T\,\Delta\mathrm{COT} + k_P(p_d - 1.70) \\ y_g &= 0.078 + 0.00144\,\mathrm{CCR} + k_g\,\Delta\mathrm{COT} \\ C &= F_f y_c + R\,y_{c,R}, \qquad V_{gen} = F_c - C \end{aligned} \tag{11.1}$$

where $\Delta\mathrm{COT} = \mathrm{COT} - 495\,^\circ\mathrm{C}$, y_c and y_g are the coke and gas fractions of the fresh feed, clipped to 0.05–0.6 and 0.02–0.3, $k_T = 0.002$ and $k_g = 0.0008$ per K, $k_P = 0.015$ per bar, the recycle cracks a quarter to coke ($y_{c,R} = 0.25$) and 5 % to gas, and C and V_{gen} are the coke and the vapor the

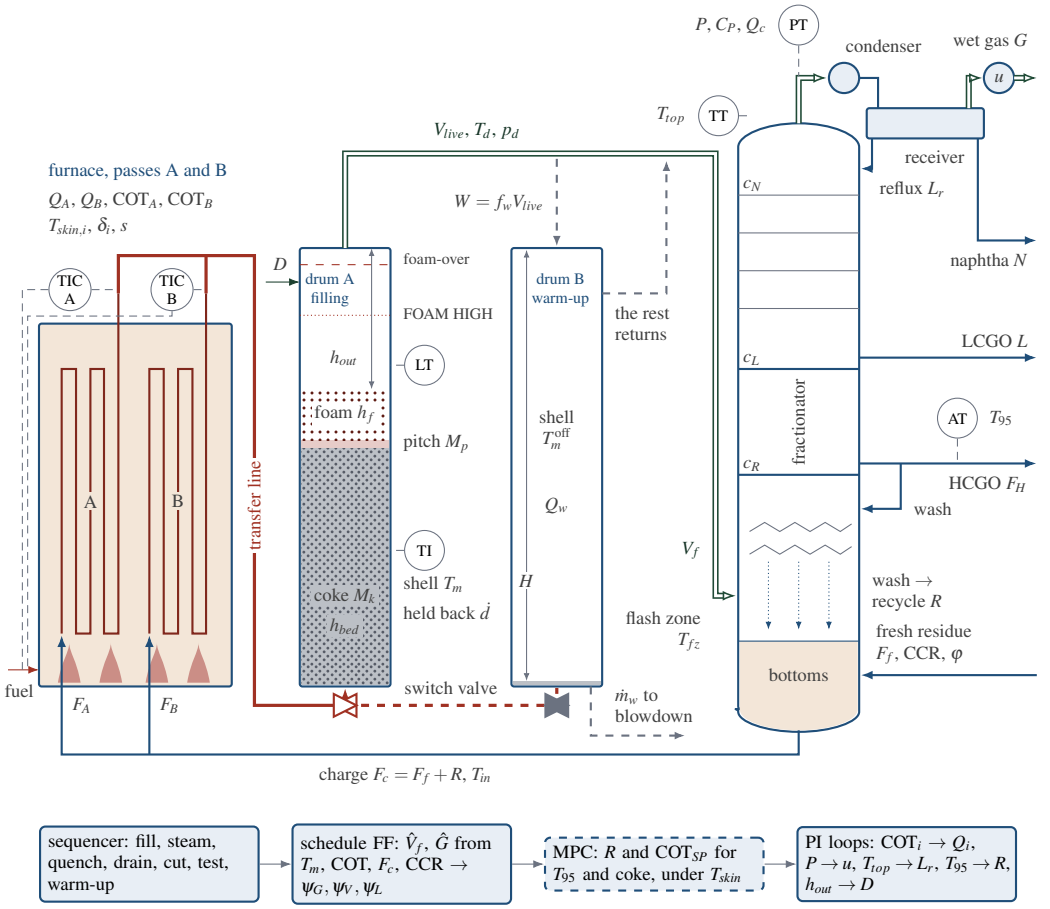


Figure 11.1 The two-drum delayed coker as modeled, redrawn from the Unit view of case 1 at $t = 10$ h. Drum A has filled 13 h and holds a coke bed of 14.6 m under 3.0 m of foam, leaving an outage of 8.4 m below the vapor line, with the FOAM HIGH (4 m) and foam-over (1 m) lines marked. Drum B is in its warm-up, its shell at 223 °C, condensing 4.2 t/h of a 7.6 t/h draw to blowdown while the rest rejoins the overhead line. The fractionator receives 91.3 t/h, its wash returns 12.1 t/h of recycle to the bottoms, and the analyzer AT reads the HCGO 95 % point. Drum heights are to scale, everything else is not. Liquid streams are thick solid lines (blue), the furnace effluent very thick (red, dashed on the shut leg to drum B), vapor and gas hollow double lines (green) and the warm-up gray dashed lines. The strip below lists the control layers, the MPC in a dashed box because it is an option.

charge makes. At 17.5 wt % and 495 °C the coke is 30.5 % of the fresh feed, the steady design value. The drum pressure is algebraic, $p_d = P + \Delta p_\ell (V/V_0)^2$, with $\Delta p_\ell = 0.7$ bar at the design vapor $V_0 = 95.4$ t/h.

A drum whose shell is colder than its vapor takes heat from its contents, and the charge that heat

would have cracked stays behind as pitch, as shown in Equation (11.2).

$$\begin{aligned} C_m \frac{\mathrm{d}T_m}{\mathrm{d}t} &= UA_d (T_{d0} - T_m) \\ \dot{d} &= UA_d \max(0,\, T_{d0} - T_m)/h_{eff} \\ \frac{\mathrm{d}M_p}{\mathrm{d}t} &= \dot{d} - \frac{M_p}{\tau_p}, \quad \frac{\mathrm{d}M_k}{\mathrm{d}t} = C + y_{c,p}\frac{M_p}{\tau_p} \\ V_{live} &= V_{gen} - \dot{d} + (1 - y_{c,p})\frac{M_p}{\tau_p} \end{aligned} \tag{11.2}$$

where $T_{d0} = \mathrm{COT} - 45\,\mathrm{K}$ is the drum vapor, $C_m = 60\,\mathrm{MJ/K}$, $UA_d = 18\,\mathrm{kW/K}$, $h_{eff} = 1200\,\mathrm{kJ/kg}$ is the cracking and vaporization heat per kilogram of vapor not made, $\dot{d}$ is capped at $0.6V_{gen}$, the pitch M_p cracks back with $\tau_p = 60\,\mathrm{min}$, 35 % of it to coke ($y_{c,p}$), and M_k is the coke. The shell's time constant on line is $C_m/UA_d = 55.6\,\mathrm{min}$, and each kelvin of shell below the vapor holds back 0.054 t/h. A quarter of the shell's heat is taken from the vapor's temperature, $T_d = T_{d0} - 0.25\,UA_d(T_{d0} - T_m)/(2.5\,V_{gen})$ with V_{gen} in kg/s, so the flash zone cools at a switch as well. The bed height is $M_k/(\rho_k A) + M_p/(0.95A)$ with $\rho_k = 0.90\,\mathrm{t/m^3}$ and $A = 36.3\,\mathrm{m^2}$ for a 6.8 m drum, 32.7 t of coke per meter.

Pause. Try Exercise 11.1 before reading on.

11.3.2 The Sequencer, the Off-Line Drum and the Warm-Up

The sequencer is plant state, advanced after each step, and exactly one drum is live at every instant. Off line a drum passes through steam-out (1.5 h), water quench (4), drain (1.5), cut and re-head (4) and steam test (1), twelve hours in all, then stands by until its warm-up, scheduled to end at the switch. In each phase its shell relaxes toward a declared temperature (380, 100, 95, 60 and 155 °C), the pitch left at the switch cracks during steam-out and is quenched into the coke, and the coke is cut out at a uniform rate. A switch falls due when the live drum has filled for t_{cyc}, and it is held until the other drum is ready, so a cycle under 12 h cannot be kept. During the warm-up a fraction f_w of the live vapor passes through the off-line drum, as shown in Equation (11.3).

$$\begin{aligned} W &= f_w V_{live}, \qquad Q_w = \varepsilon c\, W (T_d - T_m^{\mathrm{off}}) \\ C_m \frac{\mathrm{d}T_m^{\mathrm{off}}}{\mathrm{d}t} &= Q_w, \qquad \dot{m}_w = \min\left(W(1-g),\; Q_w/\lambda_w\right) \\ V_f &= V_{live} + V_{steam} - \dot{m}_w \end{aligned} \tag{11.3}$$

where $\varepsilon c = 2.16\,\mathrm{kJ/(kg\,K)}$, $\lambda_w = 900\,\mathrm{kJ/kg}$, g is the gas share of the vapor, V_{steam} is the steam-out's cracked pitch and V_f is the vapor reaching the fractionator. Only the condensate $\dot{m}_w$, never the gas, drains to blowdown, and the rest of the draw rejoins the overhead line, so the draw is subtracted from the fractionator once, as its condensate. In case 1 the warm-up starts at 9 h with the shell at 151 °C, draws 7.64 t/h and condenses 5.47 t/h, falling to 4.16 t/h at 10 h and 1.83 t/h by the switch, 13.3 t over the warm-up, and leaves the shell at 350.5 °C.

> **TRY IT IN THE APP** **1 · One fill, by hand**
>
> Load the case and press Run. Every station stays in Manual. Watch drum A: its bed climbs from 3.4 to 18.0 m with the foam riding on it, and the outage shrinks from 19.6 to 5.0 m. At 13.0 h the sequencer switches the charge to drum B, and with no loop to answer it the pressure falls to 0.907 bar g.

11.3.3 Foam and Outage

The foam on the live drum's bed approaches a steady thickness set by the vapor rate, the pressure, the residue's foaming tendency and the antifoam dose, as shown in Equation (11.4).

$$\begin{aligned} h_{f,ss} &= \frac{h_n \varphi \,(V_{live}/V_0)\,\pi_p}{1+\eta D}, \qquad \pi_p = \frac{p_{d0}+p_{atm}}{p_d+p_{atm}} \\ \tau_f \frac{\mathrm{d}h_f}{\mathrm{d}t} &= h_{f,ss} - h_f, \qquad h_{out} = H - h_{bed} - h_f \end{aligned} \tag{11.4}$$

where $h_n = 6\,\mathrm{m}$, $\eta = 0.33$ per ppm, $\tau_f = 20\,\mathrm{min}$, $p_{d0} = 1.70\,\mathrm{bar\,g}$ and $H = 26\,\mathrm{m}$ to the vapor line. At 3 ppm the design foam is 3.0 m, and a 16 h fill ends with 5.0 m of outage. FOAM HIGH below 4 m of outage and a latched FOAM-OVER at 1 m, which stops the run, are teaching values. Real alarm points come from the drum's radiation level detectors and the licensor. Through the factor π_p a fall in drum pressure swells the foam, the lever that Volume 3, Chapter 7 meets as a pressure margin against foaming at a filler.

> **TRY IT IN THE APP** **5 · Foam and antifoam**
>
> Load the case and press Run. The residue foams three times as much as the design residue, and the antifoam is fixed at 3 ppm. The outage reaches 1 m and the run stops with FOAM-OVER at 11.2 h. Set Antifoam to On foam level and press Reset. The loop then holds the outage within 4.39 to 5.44 m and uses 6.2 kg.

11.3.4 The Furnace and Its Coils

Each pass carries an absorbed duty that follows its fuel command through the firebox and a coil outlet that lags its heat balance, as shown in Equation (11.5).

$$\begin{aligned} \tau_{fb} \frac{\mathrm{d}Q_i}{\mathrm{d}t} &= Q_{i,cmd} - Q_i \\ \tau_c \frac{\mathrm{d}\mathrm{COT}_i}{\mathrm{d}t} &= T_{in} + \frac{Q_i}{F_i c} - \mathrm{COT}_i \\ T_{skin,i} &= \mathrm{COT}_i + \theta_f \phi_i \frac{Q_i}{Q_0} + \beta \delta_i \\ \frac{\mathrm{d}\delta_i}{\mathrm{d}t} &= k_f \, 2^{(T_{film,i} - T_{film,0})/10\,\mathrm{K}} \sqrt{F_0/F_i}\; s^{-0.4} \end{aligned} \tag{11.5}$$

where $\tau_{fb} = 2\,\mathrm{min}$, $\tau_c = 3\,\mathrm{min}$, $Q_{i,cmd} \le 13\,\mathrm{MW}$, $F_i = F_c/2$, $c = 3.0\,\mathrm{kJ/(kg\,K)}$ with the coil's vaporization, $\theta_f = 55\,\mathrm{K}$, $Q_0 = 10.10\,\mathrm{MW}$, $\phi_A = 1$ and $\phi_B = 1.06$ because pass B sits nearer the burners, $\beta = 8\,\mathrm{K/mm}$, $k_f = 0.267\,\mathrm{mm/day}$, $T_{film,i} = \mathrm{COT}_i + 0.55\theta_f \phi_i Q_i/Q_0$ with $T_{film,0} = 525.25\,^\circ\mathrm{C}$, $F_0 = 66\,\mathrm{t/h}$ and s the velocity steam in wt % of the charge. The skin is the tube metal temperature of Equation (1.3) with a coke term that grows by the coking law. The charge enters at T_{in}, which lags the mix of fresh residue at 300 °C and recycle at $T_{fz} - 25\,\mathrm{K}$ with 10 min, 311.4 °C at design. On coil day 2 the skins are 554.3 and 558.2 °C. A compression parameter multiplies the coil coking rate so that case 7 runs 100 coil days per simulated day, and the skin limit of 650 °C is a teaching value for the coil alloy's design metal temperature.

11.3.5 The Fractionator and Its Pressure

The vapor reaching the flash zone is cut on its own pseudo true-boiling-point curve, through 15 points from −160 °C at zero to 200 °C at 0.354, 345 °C at 0.589, 495 °C at 0.874 and 620 °C at 1, as shown in Equation (11.6).

$$
\begin{aligned}
c_N &= g_F + n_0\left(\frac{\rho_0}{L_r/V_f}\right)^{0.45}, \quad c_L = c_N + \frac{L}{V_f} \\
c_R &= 1 - \frac{\min(R_w,\, 0.3V_f)}{V_f} \\
F_H &= (c_R - c_L)V_f, \qquad T_{95} = T_R - 15\,\mathrm{K} \\
\tau_N \frac{\mathrm{d}T_N}{\mathrm{d}t} &= \mathrm{TBP}(c_N) - T_N \quad \text{(likewise } T_L,\, T_R\text{)} \\
C_P \frac{\mathrm{d}P}{\mathrm{d}t} &= G - u\,Q_c \frac{P + p_{atm}}{P_0 + p_{atm}}
\end{aligned}
\tag{11.6}
$$

where g_F is the gas share of V_f (all the gas arrives), $n_0 = 0.2175$, $\rho_0 = 0.42$, R_w is the wash's recycle command after a 0.5 min valve lag, the LCGO draw is limited to $(c_R - c_N - 0.02)V_f$, and the cut temperatures lag with 4, 8 and 10 min. The HCGO F_H is the balancing stream between the LCGO and the recycle, so its 95 % point is the HCGO and recycle cut. At design one more percent of recycle lowers it 7.29 K and adds 0.3 points of coke. The overhead temperature is $T_{top} = T_N - 70\,\mathrm{K} + 25\,\mathrm{K/bar}\,(P - P_0)$, the flash zone lags the drum vapor with 5 min, and the recycle reaches the charge through a 15 min bottoms lag. In the pressure balance $G = g_F V_f$ is the wet gas that the compressor takes, $C_P = 0.645\,\mathrm{t/bar}$, $Q_c = 17.3\,\mathrm{t/h}$ and $P_0 = 1.0\,\mathrm{bar\,g}$, which gives an overhead time constant of 6.0 min with the compressor at 75.2 % of its capacity. The condensable vapor is taken as condensed at once.

11.3.6 Instruments, Initialization and Numerics

The controllers read a measurement packet and nothing else. The coil outlets carry a 0.2 min lag and 0.3 K of noise, the skins 1 min and 1 K, the overhead 0.3 min and 0.2 K, the drum shells 2 min and 1 K, the overhead pressure 0.05 min and 0.004 bar, the flows 0.3 %, and the radiation level detectors read each drum's foam top with 1 min and 0.15 m. An online distillation analyzer draws a sample every 10 min, analyzes it for 10 min behind a 5 min sample line and reports T_{95} with 0.7 K of noise, a sampled analyzer whose effective delay follows Equation (1.12). The feed tank's assay gives the CCR, and the sequencer's phases and clocks are known exactly because the control system runs them. Every case starts lined out at its setpoints, with the drums pre-run through two cycles to the hour of drum A's fill the case names. The engine of Section 1.6.1 integrates by fourth-order Runge–Kutta with a 3 s step, the controllers execute every 6 s and the historian logs every 12 s. The exported log sets every reading beside the simulated quantity behind it, the pitch, the true foam, the held-back vapor, the coil coke and the ledgers, in columns whose names begin with `truth_`. Over every full cycle in the challenge and in case 8, the charge equals the vapor to the fractionator, the warm-up condensate and the coke laid to within 2.6×10^{-13} of the charge, and at every step $V_f = V_{live} + V_{steam} - \dot{m}_w$ exactly. The parameters are collected in Table 11.2.

11.3.7 Assumptions

The yields are teaching correlations of a published form [12], and real cokers are characterized on their own residues [13, 14]. A drum is one mass of coke, one pool of pitch and one shell temperature, with no bed structure [8], no liquid level inside the bed and no thermal stress, which is the real limit on how fast a cycle may run [10]. The new drum's deficit is written as vapor not made, while a real drum also condenses vapor on its walls. Both hold material in the drum for a while and both depend on the shell's temperature. The foam is a steady thickness with no ceiling, so on a foaming residue with little antifoam it stands as tall as 18 m, which a real unit would not allow. The fractionator has

Table 11.2 Parameters. A: Adjustable Under Parameters in Show Model. V: Fixed and Stated in the App. H: Fixed and Not Displayed

Symbol	Meaning	Value	
H, D_d, ρ_k	drum height, diameter, bed density	26 m, 6.8 m, 0.90 t/m^3	V
C_m, UA_d	shell capacity, contents to shell	60 MJ/K, 18 kW/K	V
h_{eff}, τ_p, $y_{c,p}$	vapor not made, pitch	1200 kJ/kg, 60 min, 0.35	V
εc, λ_w	warm-up exchange, condensing	2.16 kJ/(kg K), 900 kJ/kg	V
k_T, k_g, k_P	yield slopes, points of coke or gas	0.20/K, 0.08/K, 1.5/bar	V
h_n, τ_f	foam at design vapor, lag	6 m, 20 min	V
η	antifoam efficiency	0.33 per ppm	A
—	FOAM HIGH, FOAM-OVER outage	4 m, 1 m	A
—	antifoam pump capacity	25 ppm	V
Q_{max}	duty limit per pass	13 MW	A
τ_{fb}	firebox	2 min	H
τ_c	coil outlet	3 min	A
c	coil heat capacity	3.0 kJ/(kg K)	V
η_F	furnace efficiency (fuel ledger)	0.85	H
θ_f, β, ϕ_B	film, coke on skin, pass B flux	55 K, 8 K/mm, 1.06	V
k_f	coil coking at design film	0.267 mm/day	V
$T_{skin,max}$	skin limit	650 °C	A
—	coil day at start, compression	2 days, 1×	A
C_P, Q_c, Δp_ℓ	overhead, compressor, drum line	0.645 t/bar, 17.3 t/h, 0.7 bar	V
τ_N, τ_L	naphtha and LCGO cut lags	4, 8 min	V
τ_R	HCGO and recycle cut lag	10 min	A
τ_B, τ_{fz}, τ_{in}, τ_v	bottoms, flash zone, charge, valves	15, 5, 10, 0.5 min	H
—	analyzer period, sample line	10 min, 5 min	A
—	feedforward model error in UA_d	0 %	A

no trays, pumparounds or wash-bed coking, and the condensable vapor condenses at once. Coil coke raises only the skin, without the pressure drop it adds. Every size, yield, rate, alarm and limit is a teaching value, and the real skin limit comes from the coil alloy's design metal temperature under API Standard 530.

11.4 The Control Layers as Implemented

The sequencer owns the drums and no controller moves the switch. A global Manual and Auto switch is the master. In Manual every station drives its actuator and every loop tracks it, as Section 1.6.2 describes, so the return to Auto is bumpless. Each execution, every 6 s, builds the measurement packet, steps the feedforward's model, runs the MPC when it is due, then the loops, and writes the actuator commands.

11.4.1 Five PI Loops

All five loops run the velocity-form routine of Section 1.6.2 with the derivative term set to zero, written in this chapter's symbols in Equation (11.7).

$$\begin{aligned} e_k &= a(\mathrm{SP}_k - \mathrm{PV}_k), \qquad & S_k &= S_{k-1} + K_c(e_k - e_{k-1}) \\ \hat{v}_k &= v_0 + S_k + I_{k-1}, & v_k &= \min(v_{max}, \max(v_{min}, \hat{v}_k)) \\ I_k &= I_{k-1} + \left[(K_c/\tau_I)\, e_k + (v_k - \hat{v}_k)/\tau_I\right]\Delta t \end{aligned} \tag{11.7}$$

where $a = \pm 1$ is the direction, S the unclamped running sum, v_0 the output at the start and the back-calculation tracking time equals τ_I. In the symbols of Equations (1.18) and (1.19), v_0 is u_b, $\hat{v}_k$ is v_k and v_k is u_k. Setting an integral time at or above 600 min leaves a loop with proportional action alone. The COT loops ($a = +1$) move each pass's duty at 0.06 MW/K and 5 min within 0–13 MW. The pressure loop ($a = -1$) moves the compressor at 20 % per bar and 20 min within 0–100 %. The overhead loop ($a = -1$) moves the top reflux at 0.6 t/h per K and 6 min. The HCGO loop ($a = -1$) moves the recycle at 0.04 % of the fresh feed per K and 15 min within 0–30 %, and it executes only

when the analyzer reports, with Δt the time since the last report, holding while the analyzer is stale. Antifoam runs either at a fixed ratio to the charge, 3 ppm as shipped, or on foam level, where a PI ($a = +1$) holds the measured outage at 5 m at 2 ppm per m and 20 min within 0–25 ppm. The LCGO draw has no loop.

11.4.2 Schedule Feedforward

The feedforward carries its own copy of the drum constants, as commissioned, and predicts the vapor and gas the fractionator will get from the measurements and the schedule, as shown in Equation (11.8).

$$\begin{aligned} \hat{V}_f &= \hat{V}_{gen} - \hat{d} + (1 - y_{c,p})\frac{\hat{M}_p}{\tau_p} + \hat{V}_{steam} - \hat{m}_w \\ \hat{G} &= \hat{g}\,(\hat{V}_f + \hat{m}_w), \quad \psi_G = \hat{G}/G_0, \quad \psi_V = \hat{V}_f/V_0 \\ \psi_L &= \psi_V \left(n_0/\hat{n}_v\right)^{1/0.45}, \quad u = \psi_G\, u_{PI}, \quad L_r = \psi_L L_{r,PI} \\ R &= \psi_V (F_{ref}/F_f)\, R_{PI}, \quad L = \psi_V L_{set} \end{aligned} \tag{11.8}$$

where $\hat{V}_{gen}$ and $\hat{g}$ come from the yield correlation at the measured coil outlets, charge and fresh feed and at the assay lagged by its 15 min transport, $\hat{d}$ uses the measured live shell, $\hat{M}_p$ is the model's own pitch, $\hat{m}_w$ uses the off-line shell during the warm-up, $\hat{n}_v = \max(0.03,\ n_0 - (\hat{g}_F - g_0))$ corrects the reflux for the gas share, G_0, V_0 and g_0 are design values and $F_{ref} = 120\,\mathrm{t/h}$. The measured inputs are filtered over 1 min, while the schedule's steps pass at once. Each PI loop keeps running on its own error, with its output limit divided by its multiplier, and turning the feedforward on or off re-bases the loops so the actuators do not jump. In the last 2 h of a fill it adds 4 ppm of antifoam. A parameter scales the model's UA_d to give the feedforward a known model error, zero as shipped.

11.4.3 Model Predictive Control of the Recycle and the COT

The MPC, the option of Section 1.6.3 built here on top of the feedforward, plans the recycle base u_1 (the HCGO loop's output, % of the fresh feed at the design vapor) and the COT setpoint u_2 every 2 min over 120 steps (4 h) in graduated blocks. Its model is relinearized at each solve from the measurements, as shown in Equation (11.9).

$$\begin{aligned} \tau_R \dot{x}_1 &= k_r x_3 + b - x_1, \quad 6\dot{x}_2 = u_2 - x_2, \quad \dot{x}_3 = u_1 - x_3 \\ y_2 &= y_{c0} - k_T(x_2 - \mathrm{COT}_0) + k_{rc}(x_3 - r_0) \\ y_3 &= s_0 + k_s(x_2 - \mathrm{COT}_0) \\ y_4 &= \ell_0 + k_\ell(x_2 - \mathrm{COT}_0) \\ \min\ & \textstyle\sum_k w_{95}(x_{1,k} - T_{95,SP})^2 + w_c(y_{2,k} - y_{c,SP})^2 \\ & \quad + \Delta u_k^\top S\, \Delta u_k \\ \text{s.t. } & y_3 \le T_{skin,max} - 5 + \sigma_1, \ \ y_4 \le \bar{u}_c - 2 + \sigma_2 \\ & 0 \le u_1 \le 30, \ \ 480 \le u_2 \le 510 \\ & |\Delta u_1| \le 2, \ \ |\Delta u_2| \le 3 \end{aligned} \tag{11.9}$$

where x_1 is the HCGO 95 % point with the analyzer bias, moved halfway toward each new analyzer result, compared with what the model held at the instant that sample was taken, x_2 the COT, x_3 the delivered recycle, k_r the slope of the curve at the cut, b the term that places the steady value at the present 95 % point, y_2 the coke yield in percent of the fresh feed with the assay as a measured disturbance, y_3 the hotter skin, y_4 the compressor load against its available capacity $\bar{u}_c$, y_{c0}, s_0 and ℓ_0 the present coke yield, skin and load at the present COT COT_0 and recycle r_0, k_{rc}, k_s and k_ℓ their slopes, $w_{95} = w_c = 1$, $y_{c,SP} = 30.5\,\%$, $S = \mathrm{diag}(4, 2)$, and the slacks σ are priced $10^3\sigma^2 + 10\sigma$. The framework's linear MPC module solves each quadratic program by the Goldfarb–Idnani dual algorithm [17], as industrial MPC does [15], and a solution that is infeasible, non-finite or out of range holds the last decisions. Over the challenge's two days it solves 1440 times with no fallback.

Table 11.3 The Nine Cases

#	Case and events	Teaching point
1	Manual, drum A 3 h into its fill	one fill, the warm-up, the switch
2	Loops, drum B warming, switch by hand	the switch transient
3	COT 495 to 505 °C	yield against coil coking
4	Loops, then schedule FF, switch at 2 h	feedforward from a schedule
5	Residue foaming 3×, fixed 3 ppm	antifoam on foam level
6	CCR 22 wt %, foaming 1.6×	the cycle must shorten
7	Coil clock 100×, from day 2	skin limit, run length
8	Switch warmed at 2 h, warm-up skipped, cold at 18 h	what only the warm-up does
9	Challenge, 48 h: cold switch at 28 h, CCR 20 at 20 h, compressor 80 % at 30 h	the whole unit

11.4.4 Three Models and Six Design Choices

The plant is Section 11.3. A second, identified model is the one a reader builds from the trends, for instance the 6 min overhead or the 56 min shell. The controllers' models are the feedforward's commissioned copy of the drum and the MPC's linearization, neither of which sees the pitch, the true foam or the coil coke.

Six choices depart from the app's original specification and are design decisions of the simulator. First, the HCGO 95 % point is held by the wash, the recycle, with the HCGO draw as the balancing stream. On one boiling curve the 95 % point is the HCGO and recycle cut, so the specification's two loops, one on the HCGO draw and one on the recycle ratio, would act on the same cut point, since each draw moves every cut below it (Section 1.4.3), the interaction that Chapter 4 measures on a crude tower. Second, the challenge holds three switches, the second onto a cold drum, and judges the pressure over the whole run, because with warmed switches only the shipped loops already stay within 0.1 bar (a 0.060 bar dip in case 8). Third, only the warm-up's condensate is taken from the fractionator. Subtracting the whole draw made the start of the warm-up a larger step in gas than the switch itself. Fourth, there is no pass flow balancing, so pass B's higher flux makes it the pass that ends the furnace run. Fifth, the challenge judges the HCGO 95 % point within ±5 K at every instant. An earlier rule asked for the band over 95 % of the run, and a pressure PI alone at 100 % per bar passed it (Section 11.5). Sixth, the foam has no ceiling, the documented simplification above.

11.5 Nine Cases, Measured

The cases appear in Table 11.3 and the principal measurements in Table 11.4. Every number comes from the engine run without a browser at the default seed with the browser's own step, any change a case asks for is made after the run begins unless the text says Reset, and seeds 1 to 3 and a halved step were used to check the challenge. Pressures are the logged transmitter values unless called true.

11.5.1 One Fill and One Switch

Case 1 holds every station still in Manual from 3 h into drum A's fill. The charge leaves 36.6 t/h in the drum, 30.5 % of the fresh feed, and sends 95.7 t/h of vapor on, so the bed climbs steadily from 3.4 to 18.0 m by the switch at 13.0 h, with the foam riding on it at 3.0 m and the outage shrinking from 19.6 to 5.0 m. The warm-up from 9 h takes the fractionator's vapor from 95.46 to 89.98 t/h, by its condensate, and lightens the HCGO to 475.7 °C with the recycle fixed. The pressure barely moves at the warm-up start (0.982 to 1.006 bar g) because the condensate carries no gas, and at the switch, with no loop, it falls to 0.907 bar g.

In case 2 the loops run without the schedule. Switched by hand at 10 min, with the new shell at 325 °C after 3 h 10 min of its warm-up, the drum holds back 6.73 t/h, the pressure falls to 0.923 bar g after 8.8 min, the overhead to 126.2 °C, and the pressure is back above 0.99 bar g for good 39 min after the switch, then overshoots to 1.027 bar g at 72 min as the pitch cracks back. Left to the schedule, the shell reaches 350.5 °C and the dip is 0.941 bar g with 5.37 t/h held back.

Table 11.4 Case Outcomes Measured on the Engine

Case and condition	Result
1, the fill	bed 3.4 to 18.0 m, outage 19.6 to 5.0 m
1, the switch at 13.0 h	P to 0.907 bar g, no loop
2, by hand at 10 min (325 °C)	$\dot{d}$ 6.73 t/h, P 0.923 bar g
2, scheduled at 1 h (350 °C)	$\dot{d}$ 5.37 t/h, P 0.941 bar g
3, COT +10 K	coke 30.60 to 28.72 %, coking ×2.23/2.25
4, loops	P 0.940 bar g at +9 min
4, schedule FF	P 0.988–1.014 bar g
5, fixed 3 ppm	FOAM-OVER at 11.2 h, 4.4 kg
5, fixed 12 ppm	outage ≥ 4.42 m, 20.6 kg to the switch
5, on foam level	outage ≥ 4.39 m, 6.2 kg
6, cycle 16 / 15 / 14 h	outage 2.97 / 3.74 / 4.26 m
6, cycle 14 / 12 h	new shell 268 / 151 °C
7, COT 495 / 490 / 485 °C	skin limit on day 40.0 / 63.7 / 101
8, warmed / cold, loops	dip 0.060 / 0.178 bar, 4.9 / 15.4 t held
8, cold, schedule FF	dip 0.024 bar, T_{95} 479.4–480.9 °C

11.5.2 What the Schedule Is Worth

Case 4 runs the same scheduled switch twice. With the loops, the pressure falls to 0.940 bar g 9 min after the switch while the compressor reaches its lowest, 73.4 %, only after 17 min. With the schedule feedforward the compressor steps with the switch, from 75.0 to 74.3 % in the first 12 s and 70.8 % at 2 min, and the logged pressure stays within 0.988–1.014 bar g, 0.993–1.003 bar g in truth against the transmitter's 0.004 bar noise. Case 8 puts a warmed switch and a cold one on one axis, as shown in Figure 11.2. The cold shell, at 137 °C, holds back 16.9 t/h against 5.38 t/h for the warmed one and 15.4 t against 4.9 t over four hours, 3.1 times as much, and with the loops the dip is 0.178 bar against 0.060 bar, while the true HCGO 95 % point swings from 471.1 to 484.9 °C. The feedforward cuts the cold dip to 0.024 bar and holds the 95 % point within 479.4–480.9 °C, but the vapor lost to the fractionator does not change, because the heat to warm 60 MJ/K of steel by 300 K must come from the charge.

🕮 **Pause.** Try Exercise 11.2 before reading on.

The feedforward is only as good as its copy of UA_d. With that copy 20 % low the cold switch of case 8 dips to 0.959 bar g, with it 20 % high the pressure overshoots to 1.056 bar g, against 0.976–1.018 bar g with the exact copy, and the challenge still passes at 0.040 and 0.052 bar with the 95 % point within 1.5 and 3.4 K. At 50 % low it passes at 0.095 bar with the 95 % point 4.6 K off, so close to its band that seed 1 fails it at 5.3 K, and at 50 % high it fails at 0.121 bar with the 95 % point 7.3 K off. The degradation is graceful and roughly in proportion to the error, and an error that overpredicts the held-back vapor is the more harmful one because it pulls the compressor down further than the gas falls.

11.5.3 Coke Against Tube Life, Foam and Heavier Feed

Raising the COT to 505 °C in case 3 lowers the coke from 30.60 to 28.72 % of the fresh feed and raises the liquid from 58.71 to 59.72 % and the gas from 10.83 to 11.64 %, which loads the compressor from 75.1 to 80.2 %. The duty rises by 1.03 MW, the hotter skin from 558.1 to 571.6 °C, and the coil coking rate by 2.23 and 2.25 times in passes A and B, so the run left falls from 38.1 to 14.5 days. Case 7 follows that trade to its end at 100 coil days per simulated day. Pass B reaches 650 °C on coil day 40.0 at 495 °C, day 63.7 at 490 and day 101 at 485, and the cooler run costs coke, 31.55 against 30.61 % at the start.

In case 5 a residue that foams three times as much stands 9.05 m of foam on the fixed 3 ppm, and as the bed rises under it the outage reaches 1 m at 11.2 h, 14.2 h into the fill, after 4.44 kg of antifoam. A fixed ratio that reaches the switch needs 6 ppm for 2.01 m of outage, under the FOAM HIGH line, and 12 ppm for 4.42 m at 20.6 kg. On foam level the loop doses 0.24 ppm at 1 h, 3.49 at 8 h and 12.9

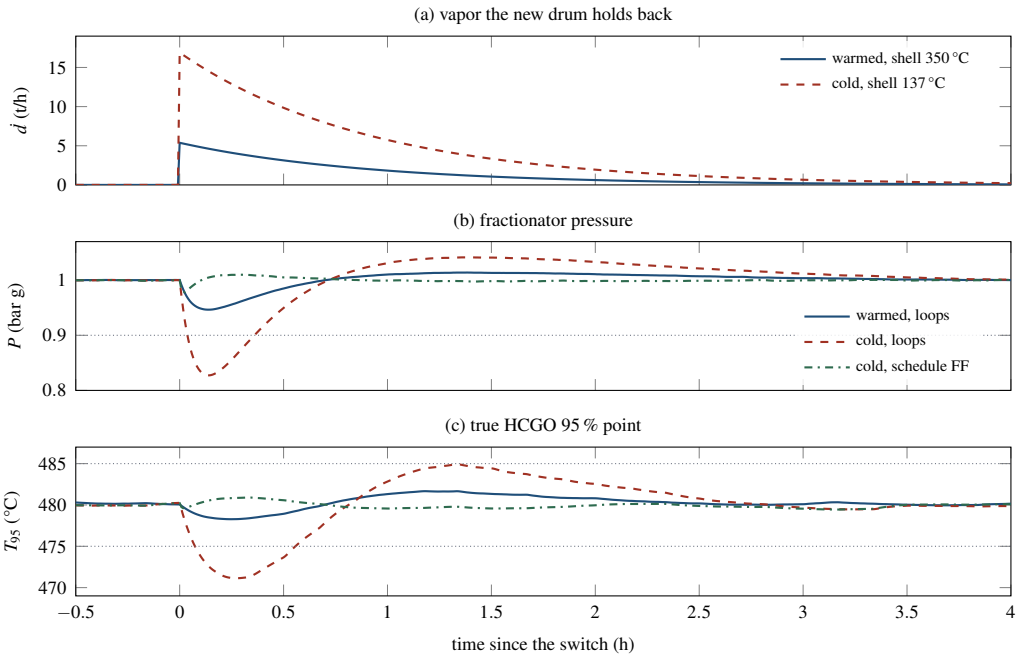


Figure 11.2 The two switches of case 8 on a common axis of time since the switch, true values every 30 s. The warmed switch is drawn solid, the cold switch on the loops dashed and the cold switch with the schedule feedforward dash-dotted. (a) The cold shell holds back 3.1 times the vapor. (b) With the loops the cold switch dips 0.17 bar in truth, and the schedule feedforward keeps it within 0.978–1.010 bar g. The gray dotted line is the challenge's lower band. (c) The HCGO 95 % point leaves its ±5 K band, the gray dotted lines, with the loops and not with the feedforward.

just before the switch, holds the outage within 4.39–5.44 m and uses 6.2 kg. Early in the fill the same loop lets the foam stand 16.1 m tall, and 17.9 m on the new drum after the switch, the consequence of a foam law with no ceiling. In case 6 a residue at 22 wt % CCR leaves 44.7 t/h of coke, the bed rises 1.37 m/h, and on the 16 h cycle the outage falls below 4 m 14.6 h into the fill and reaches 2.97 m at the switch with the antifoam at its 25 ppm limit. Shortening the cycle to 14 h keeps 4.26 m, but the shorter warm-up puts the next drum on line at 268 °C instead of 339 °C, and at 12 h at 151 °C. With 26 wt % at 135 t/h no cycle the decoking allows keeps the outage above 4 m (2.99 m at 12 h), so past that floor only the feed rate is left.

🕮 **Pause.** Try Exercise 11.3 before reading on.

11.5.4 Two Days on the Coker

The challenge runs 48 h from coil day 30, with switches at 12, 28 and 44 h, the second onto a cold drum because its warm-up is skipped at 23 h, a residue at 20 wt % CCR and foaming 1.5 times from 20 h, and the compressor limited to 80 % from 30 h. It is met with the true HCGO 95 % point within ±5 K and the true pressure within ±0.1 bar throughout, no foam-over, the skin under 650 °C and at most 15 kg of antifoam, judged on the truth at every step. The 95 % point is graded on the simulator's value rather than on the analyzer's reports, which arrive 15 to 25 min late. The configurations are compared in Table 11.5, and Figure 11.4 draws three of its columns against their limits. As shipped the pressure leaves its band first, at the cold switch at 28.0 h, the 95 % point 1.9 min later, and the fixed 3 ppm passes the budget at 38.0 h. The feedforward and antifoam on foam level each fix one of those failures, and together they pass at all four seeds and with the step halved (0.024–0.025 bar,

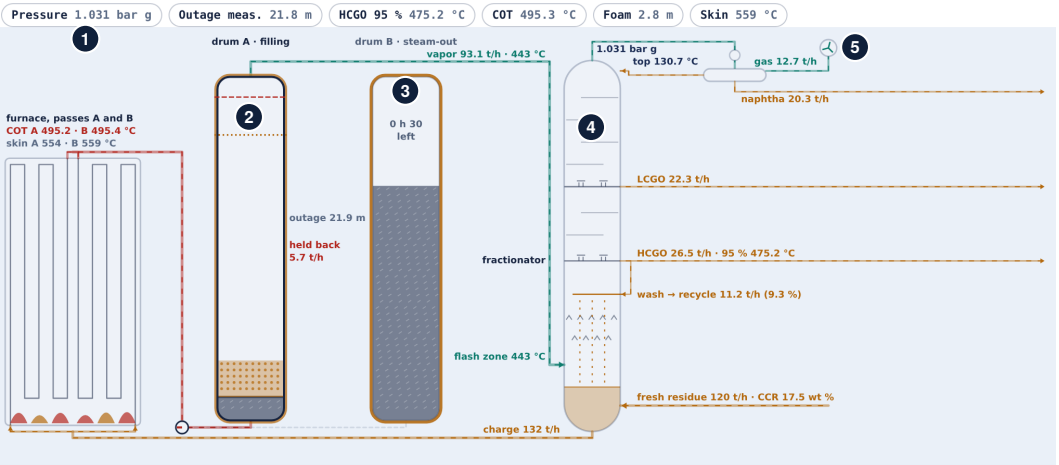


Figure 11.3 Case 8 on the loops alone at $t = 19$ h, one hour after the switch onto the cold drum, on the Unit view, cropped to the plant with numbered callouts. The badges (1) read 1.031 bar g as the pressure overshoots on the pitch cracking back, 21.8 m of measured outage, a COT of 495.3 °C, 2.8 m of foam and a 559 °C skin, and the HCGO analyzer reports 475.2 °C. Drum A (2) is filling, its shell still holding back 5.7 t/h, and drum B (3), full, is in steam-out with 30 min left. The vapor reaches the fractionator (4) at 93.1 t/h and 443 °C, the wash returns 11.2 t/h, and the wet gas compressor (5) takes the gas from the overhead.

Table 11.5 Challenge Configurations Scored at 48 h

Configuration	P worst (bar)	T_{95} in band	T_{95} off (K)	Antifoam (kg)	Verdict
shipped (loops, 3 ppm)	0.186	98.7 %	9.7	18.9	fails
loops, on foam level	0.186	98.7 %	9.7	6.0	fails
schedule FF, 3 ppm	0.024	100 %	1.0	22.1	fails
schedule FF, 2 ppm	0.024	100 %	1.0	15.8	fails
PI 70 %/bar, 6 min, foam level	0.103	98.2 %	9.9	6.0	fails
PI 100 %/bar, 10 min, foam level	0.089	98.2 %	9.9	6.0	fails
suggested (FF, foam level)	0.024	100 %	1.0	6.8	passes
MPC, foam level	0.049	100 %	1.2	6.1	passes

the 95 % point within 1.0–1.3 K, 6.81–6.84 kg). The MPC with antifoam on foam level also passes, raising the COT as high as 508.4 °C to win back liquid on the heavier residue until the skins, 646 °C against its 645 °C soft limit, and the compressor stop it.

> ✦ DID YOU KNOW? **The Tightest Pressure Still Fails**
>
> From the simulator (Table 11.5). On a fixed antifoam ratio of 3 ppm, the schedule feedforward holds the pressure within 0.024 bar and the heavy gas oil cut within 1.0 K, yet it fails the challenge. It spends 22.1 kg of antifoam against a budget of 15 kg. On foam level the same feedforward uses 6.8 kg and passes.

The pressure band by itself does not need the feedforward. Without it a pressure PI at 100 % per bar holds 0.077, 0.089 and 0.092 bar at integral times of 2, 10 and 20 min, and the 100 % per bar, 10 min setting keeps the pressure inside its band at all four seeds and with the step halved (0.089–0.090 bar). That setting still fails the challenge, on the heavy gas oil. It leaves the true HCGO 95 % point 9.6–10.5 K off at the cold switch over the four seeds, against 1.0–1.3 K for the suggested settings and 1.2–1.6 K for the MPC. The excursion lasts under 2 % of the run, so the earlier rule,

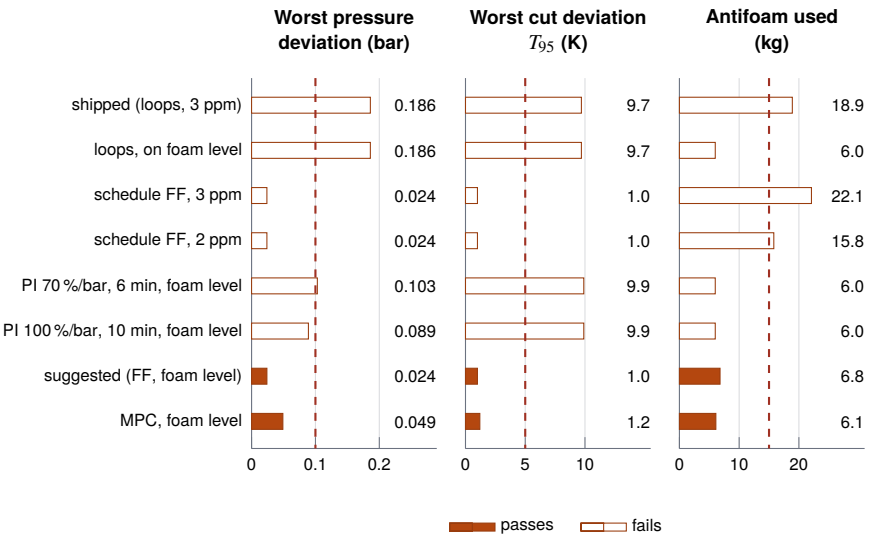


Figure 11.4 Three columns of Table 11.5 against the challenge's limits (dashed): the worst pressure deviation against 0.1 bar, the worst deviation of the true HCGO 95 % point against 5 K and the antifoam against 15 kg. Solid bars pass the challenge, white bars fail it. Notice that the pressure PI at 100 % per bar reaches the pressure band but not the cut, and that the feedforward on a fixed ratio fails only on antifoam.

the band over 95 % of the run, passed it, and only grading the cut at every instant shows what the feedforward adds. Retuning the HCGO PI does not rescue the loops. With the pressure PI at 100 % per bar, gains of 0.02 to 0.5 % per K and integral times of 5 to 30 min leave a worst deviation of 9.7 K at best, and the settings from 0.2 % per K, or with a 5 min integral time from 0.04 % per K, swing 35 to 116 K off on reports that arrive 15 to 25 min late. The PI alone also moves the compressor 13 041 % of its range over the two days against 2930 % for the feedforward, 4.5 times as far. That second advantage belongs to the transmitter's noise. With the noise switched off the PI at 100 % per bar travels 48 % and the feedforward 1149 %, so the high-gain loop's travel is noise passed to the compressor, while the feedforward's is the schedule it follows.

11.6 Technology in Practice

Figure 11.5 draws the loops of Section 11.4 with the hardware a refinery installs.

What Measures It

The feedforward rests on the drum's skin thermocouples. In refiners' practice they sit around the skirt and the lowest circumferential welds, the operators read them beside the quench water flow and the drum pressure, and strain gauges in the axial and hoop directions next to them show the stress in the wall, which is largest during the water quench [18]. Skin temperatures alone do not give what the gauges give, but they help ensure an adequate warm-up [18]. The 2 min lag and 1 K of noise of Section 11.3 stand for such a thermocouple. The heavy gas oil cut comes from an online distillation analyzer that reports every 10 min behind a 5 min sample line, and the outage comes from detectors that see through the drum's wall.

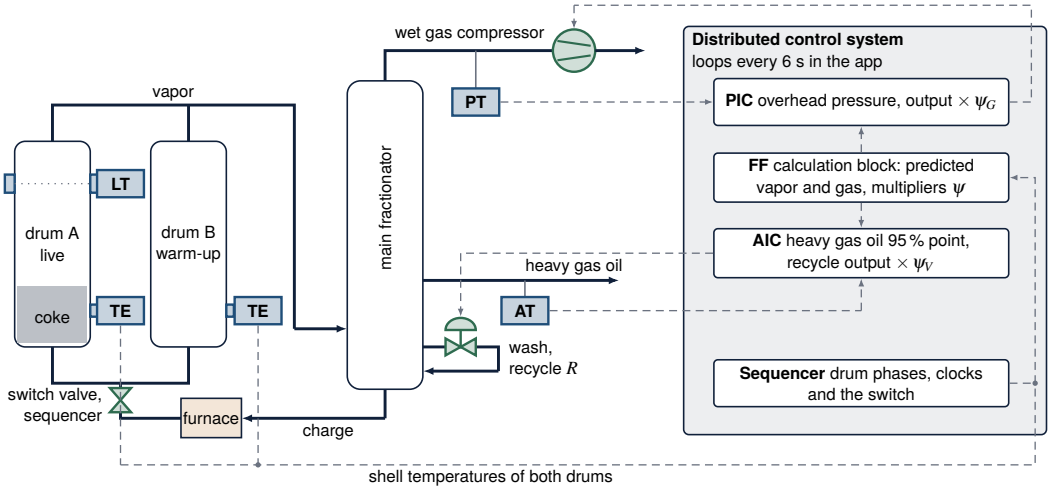


Figure 11.5 The fractionator's loops and the schedule feedforward with their hardware. Lettered boxes are sensing devices: the overhead pressure transmitter (PT), the online distillation analyzer on the heavy gas oil (AT), the drum shell thermocouples (TE) and the radiation level detectors across the live drum (LT). Valve symbols and the compressor are final elements, and dashed arrows are signals. Notice that the feedforward reads the sequencer and the shells, not the fractionator, and writes only multipliers on the outputs of the loops beside it. The coil outlet loops, the overhead temperature loop and the antifoam are left out. Not to scale, and not a piping diagram.

TECHNOLOGY **Sensor: Radiation Level Detectors on the Drum**

Gamma rays pass through a vessel's wall and are weakened by the material inside, so a source on one side and a detector on the other give an approximate level. The detector counts pulses in a gas-filled tube. A cobalt source has a half-life of 5.3 years and a cesium source 30 years, so cesium drifts less. Accuracy varies with the density of the contents, coating on the wall, source decay and detector drift, and such gauges are used where no other instrument is practical, delayed coking vessels among them, since the coke is cut out by a high-pressure water jet [19]. In a pilot coker a scanned gamma densitometer told gas, foam, liquid and coke apart by density [14]. The chapter's detectors read each drum's foam top with a 1 min lag and 0.15 m of noise.

What Moves It

The wet gas compressor is the handle of the pressure loop and of the feedforward. The coil outlets are held by the firing of each pass, the recycle by the wash and the foam by an antifoam injection pump. Foam is the reason the pump exists: in a pilot coker, runs without antifoam foamed over, plugged the overhead line and ended the run [14], the event the latched FOAM-OVER of Section 11.3 stands for.

TECHNOLOGY Final Element: Wet Gas Compressor Speed

A compressor's gas flow can be set by its speed instead of by a valve. Varying the speed uses less energy than throttling the discharge or recycling gas to the suction, but a variable-speed machine often cannot change flow as fast as a control valve and has limited options for a fail-safe state. A centrifugal compressor surges, violently, when the pressure across it is too high and its flow too low, and a recycle valve from discharge to suction prevents it [19]. One refiner answers a pressure drop after a drum switch by raising the compressor's suction pressure setpoint, which lowers its speed [20]. The chapter's compressor is the fraction u of its capacity in Equation (11.6), with no surge term.

Where the Controller Runs

A refinery runs such loops in a distributed control system, in which a supervisory level oversees loops that are sometimes nested, with cycle times from milliseconds to minutes [21]. The app executes its loops every 6 s and its MPC every 2 min, inside that range. The five loops are PI blocks. The sequencer is sequence logic that steps each drum through its phases and moves the switch valve, and the feedforward is a calculation block that multiplies the outputs of the compressor, reflux and recycle loops, as Equation (11.8) writes it. The sequencer's phases and clocks are known exactly because the control system runs them, and that is what lets the feedforward act at the switch rather than after it.

IN PRACTICE Warm the Drum, Then Switch

Refiners describe the switch as this chapter models it. In a 2011 panel, one refiner explained that a switch onto a relatively cold drum lowers the system pressure until the new drum reaches its normal overhead temperature [20]. Its practice was to watch the drum's skin temperatures to avoid a switch to a cold drum, to start the antifoam about four hours before the switch, and, if the pressure still fell, to raise the compressor's suction pressure setpoint. In a 2010 panel a refiner set the minimum preheat before a switch from strain studies, and an engineering contractor described how limits on the skin temperature's rate of change can enter the warm-up and switch controls of the distributed control system [18].

What Is Modernizing It

Three developments touch this unit. The first is drum life. Strain gauges at high temperature and laser inspection of bulges join the internal inspection that is made only at major turnarounds [22], and one refiner recommends strain gauges on new drums and fits them to old ones [18]. The second is a view into the drum. A scanned gamma densitometer in a pilot coker detected foam forming and collapsing, by itself and under antifoam, and classified the coke morphology [14], where the chapter's detectors read one foam top. The third is the cut between analyzer reports. Predicting a quality that is sampled slowly from fast measurements is the original and still dominant use of soft sensors [23], and model predictive control, with more than 4600 reported applications, is most common in refining [15]. The chapter's MPC moves its analyzer bias only when a result arrives.

11.7 Limitations

Every size, yield, rate and limit is a teaching value fitted to no operating coker. The yield correlation has the published form on the CCR, with teaching slopes for the COT and the drum pressure, and the pseudo boiling curve is the app's own. A drum is lumped, so the bed's structure, the liquid inside it and the drum's thermal fatigue, the real reason a cycle is not shortened at will [10], are absent, and the new drum's deficit is written as cracking deferred rather than condensation on the walls. The foam has no ceiling and stands 16 m to 18 m tall on case 5's residue while the outage loop doses little, a height a real unit would not tolerate because the foam's top would move with every change of vapor or pressure. The fractionator has no trays, pumparounds or wash-bed coking, and its pressure model condenses all condensable vapor at once. A shell that warms at exactly one declared rate and a switch on the stroke of the schedule would not occur in a real unit, but each isolates the one mechanism its case is about.

The control layers carry simplifications of their own. There is no pass flow balancing, so the run length is set by pass B. The feedforward's model is exact as shipped, and only its UA_d can be given an error. A high-gain pressure PI alone meets the challenge's pressure band, so the grade that calls for the feedforward is the HCGO cut, and it is taken on the simulator's true 95 % point, which no instrument in a real unit reports without the analyzer's delay. The MPC has no run-length term and spends the coils for yield up to its soft skin limit. The detector outage differs from the truth by up to 0.68 m, mostly the detectors' noise.

OPEN CHALLENGE **Drums Not Built for Shorter Cycles**

The model's drum has no wall stress, yet the wall limits the cycle. Coke drums were designed to ASME Section VIII, Division 1, and not at first for cyclic loads. As some plants cut the cycle from 24 to 12 hours or less, through-wall cracks came more often and earlier [22]. Hot and cold spots in the wall raise bulges, and the internal inspection that finds them happens only at major turnarounds, about every four years. Most drum failures come from high cyclic thermal stresses, and they grow more frequent as throughput rises and cycles shorten [10, 22]. Progress would be a sequencer that plans the warm-up, the switch and the quench against the measured wall stress.

11.8 Conclusion

A delayed coker hosts a batch step inside a continuous plant, and this chapter put the drum cycle, the furnace coils and the fractionator's loops on one clock. A switch onto a drum left cold holds back 16.9 t/h of vapor, 3.1 times a warmed one, and dips the fractionator pressure 0.178 bar on the loops alone. A feedforward built from the schedule and the shell thermocouples keeps the dip to 0.024 bar and the HCGO 95 % point within 1 K, degrades gracefully with a model error of 20 %, and cannot make the vapor the steel takes, which only the warm-up saves.

Ten degrees of COT trade 1.9 points of coke for coils that coke 2.2 times as fast, antifoam on the drum's outage uses 6.2 kg where a safe fixed ratio uses 20.6 kg, and a heavier residue shortens the cycle until the decoking floor of 12 h leaves only the feed rate.

The drum and foam models are lumped and declared rather than fitted, and a high-gain PI also holds the graded pressure band, so only the HCGO cut, graded at every instant, separates it from the feedforward. A fuller model would add a foam ceiling, wall condensation and thermal stress in the drum, and pass flow balancing. The coker is a plant on which to practice feedforward from a known schedule.

KEY POINTS

- A delayed coker cracks residue in one of two drums that take turns, and each switch sends the charge into a drum colder than its vapor.
- A switch onto a cold drum holds back 16.9 t/h of vapor, 3.1 times a warmed one, and dips the pressure 0.178 bar on the loops alone.
- A schedule feedforward keeps the dip to 0.024 bar but cannot make the vapor the steel takes, which only the warm-up saves.
- A high-gain pressure PI also holds the pressure band, so only the heavy gas oil cut, graded at every instant, separates it from the feedforward.

Where else this applies. The backwash sequencer of Volume 3, Chapter 6 and the diversion sequence of Volume 8, Chapter 13 are batch steps on a known schedule beside a continuous plant.

References

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

[2] M. Baldea and I. Harjunkoski, “Integrated production scheduling and process control: A systematic review,” *Computers & Chemical Engineering*, vol. 71, pp. 377–390, 2014.

[3] A. N. Sawarkar, A. B. Pandit, S. D. Samant, and J. B. Joshi, “Petroleum residue upgrading via delayed coking: A review,” *The Canadian Journal of Chemical Engineering*, vol. 85, no. 1, pp. 1–24, 2007.

[4] K. Wisecarver, “Delayed coking,” in *Springer Handbook of Petroleum Technology*, ser. Springer Handbooks. Cham: Springer International Publishing, 2017, pp. 903–913.

[5] R. Dittmar and G. Reinig, “Model predictive FCCU main fractionator control – robustness and adaptation aspects,” *IFAC Proceedings Volumes*, vol. 30, no. 6, pp. 1201–1206, 1997.

[6] J. Yang, G. Shen, P. Li, W. Wei, C. Zhang, and J. Xiao, “Kinetic model for the deep-severity thermal reaction in the coke drum of delayed coking,” *Energy & Fuels*, vol. 31, no. 3, pp. 2681–2686, 2017.

[7] X.-L. Zhou, X. Di, G.-X. Yu, R.-X. Lu, and C.-L. Li, “Simulation of delayed coking reaction in coke drum,” *Petroleum Science and Technology*, vol. 28, no. 3, pp. 277–285, 2010.

[8] A. C. Silva, C. McGreavy, and M. F. Sugaya, “Coke bed structure in a delayed coker,” *Carbon*, vol. 38, no. 15, pp. 2061–2068, 2000.

[9] P. E. Eaton and J. Williams, “Delayed coker furnace fouling control – laboratory correlation to field experience,” in *CORROSION 2008*. New Orleans, LA: NACE International, 2008, pp. 1–7.

[10] K. Kirkpatrick, L. Antalffy, B. Millet, and G. Miller, “Extending drum life while increasing throughput in a delayed coking unit,” in *ASME 2017 Pressure Vessels and Piping Conference, Volume 3A: Design and Analysis*. Waikoloa, Hawaii, USA: American Society of Mechanical Engineers, 2017, paper V03AT03A006.

[11] Occupational Safety and Health Administration, “OSHA technical manual, section IV, chapter 2: Petroleum refining process,” U.S. Department of Labor, Washington, DC, osha.gov/otm/section-4-safety-hazards/chapter-2, read November 2026.

[12] J. H. Gary and G. E. Handwerk, *Petroleum Refining*. CRC Press, 2001.

[13] J. G. Speight, *The Chemistry and Technology of Petroleum*. CRC Press, 2014.

[14] M. Volk, Jr., K. D. Wisecarver, and C. M. Sheppard, “Fundamentals of delayed coking joint industry project,” Office of Scientific and Technical Information (OSTI), Tech. Rep., 2003.

[15] 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.

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

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

[18] American Fuel & Petrochemical Manufacturers, “Question 53: During the coking cycle, how are drum skin temperatures used to monitor drum wall condition?” AFPM Q&A Database, panel answers by two refiners and an engineering contractor, 2010, afpm.org, question and answer search.

[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] American Fuel & Petrochemical Manufacturers, “Question 51: What are some steps that can be taken to minimize re-foaming while steaming to the fractionator?” AFPM Q&A Database, panel answers by two refiners and an engineering contractor, 2011, afpm.org, question and answer search.

[21] 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.

[22] J. Chen, “Fatigue behaviour of coke drum materials and its application to extend the service lives of coke drums,” Ph.D. dissertation, University of Alberta, 2014.

[23] 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.

11.9 Exercises

11.1 Concept check · 2 min

At the scheduled switch of case 2 the fractionator pressure falls from 1.00 to 0.941 bar g, although the charge, the coil outlet temperature and the compressor's capacity have not changed. What takes the gas away?

(a) The switch valve shuts for a moment, and the charge reaches neither drum

(b) The new drum stands at a lower pressure, and the overhead gas flows back into it

(c) The new drum's cooler shell takes heat from its contents, so less of the charge cracks

(d) The warm-up condensate drains to blowdown and carries part of the wet gas with it

11.2 Concept check · 2 min

On the cold switch of case 8 the schedule feedforward cuts the pressure dip from 0.178 to 0.024 bar. Which quantity does it leave unchanged?

(a) The time at which the wet gas compressor starts to move

(b) The swing of the true HCGO 95 % point after the switch

(c) The top reflux that the overhead loop and its multiplier send

(d) The vapor that the cold drum holds back over the next four hours

11.3 Concept check · 1 min

In case 5 a residue that foams three times as much runs on a fixed antifoam ratio of 3 ppm, and it foams over 14.2 h into a 16 h fill. Why near the end of the fill and not at its start?

(a) The foam stays about 9 m thick while the bed rises under it, so the outage shrinks

(b) The vapor rate rises as the drum fills, so the steady foam grows through the fill

(c) The antifoam is used up over the fill, so the dose falls as the switch comes nearer

(d) The drum pressure rises as the drum fills, and the pressure factor swells the foam

11.4 Theory and calculation · 6 min

The design residue has a Conradson carbon of 17.5 wt%. The fresh feed is 120 t/h, the recycle 12.0 t/h, the coil outlet 495 °C and the drum pressure 1.70 bar g. A drum is 6.8 m across and 26 m to the vapor line, and the coke bed has a density of 0.90 t/m^3.

(a) Use Equation (11.1) to find the coke the charge leaves in the drum, in t/h and as a share of the fresh feed. *Checkpoint:* 36.6 t/h, 30.5 % of the fresh feed.

(b) Find the coke per meter of bed and the rate at which the bed rises. *Checkpoint:* 32.7 t/m and 1.12 m/h.

(c) A fill starts with an empty drum and the foam rides on the bed at 3.0 m. Find the outage at the end of a 16 h fill, and compare it with the FOAM HIGH line at 4 m. *Checkpoint:* 5.1 m, 1.1 m above the FOAM HIGH line.

11.5 Theory and calculation · 7 min

Take the drum constants of Equation (11.2): $C_m = 60\,\mathrm{MJ/K}$, $UA_d = 18\,\mathrm{kW/K}$ and $h_{eff} = 1200\,\mathrm{kJ/kg}$, with the drum vapor at $T_{d0} = \mathrm{COT} - 45\,\mathrm{K}$ and the coil outlet at 495 °C.

(a) Find the vapor held back for each kelvin of shell below the vapor, in t/h per K, and the shell's time constant on line. *Checkpoint:* 0.054 t/h per K and 55.6 min.

(b) In case 8 the warmed drum goes on line with its shell at 350.4 °C and the cold drum at 137.1 °C. Find the vapor each holds back at the switch, and their ratio. *Checkpoint:* 5.38 t/h and 16.9 t/h, a ratio of 3.14.

(c) Let the shell approach the vapor with the time constant of (a) and ignore the pitch that cracks back. Find the total vapor each drum holds back and the part of it held in the first four hours. Show that the total equals $C_m(T_{d0} - T_m)/h_{eff}$. *Checkpoint:* 15.6 t and 5.0 t in total, 15.4 t and 4.9 t in four hours.

11.6 Theory and calculation · 6 min

The overhead pressure follows the balance of Equation (11.6), $C_P\,dP/dt = G - u\,Q_c(P + p_{atm})/(P_0 + p_{atm})$, with $C_P = 0.645\,\text{t/bar}$, $Q_c = 17.3\,\text{t/h}$, $P_0 = 1.0\,\text{bar g}$ and $p_{atm} = 1.013\,\text{bar}$. At design the compressor stands at 75.2 % and the fractionator receives 95.4 t/h of vapor.

(a) Find the wet gas at design and the time constant of the overhead pressure. *Checkpoint:* 13.0 t/h and 6.0 min.

(b) The warmed switch of case 4 holds back 5.37 t/h of vapor with the gas share of the design vapor. Find the fall in wet gas, and the compressor position, from 75.0 %, that holds the pressure at P_0, which is what the multiplier ψ_G of Equation (11.8) does. *Checkpoint:* 0.73 t/h and 70.8 %.

(c) If nothing moved the compressor, how far would the pressure settle below P_0 for that fall in gas? Compare with the 0.907 bar g that case 1 reaches at the same kind of switch with no loop. *Checkpoint:* 0.11 bar, to about 0.89 bar g.

11.7 App challenge · 8 min

Load scenario 4, Pressure at the switch: loops or schedule feedforward. The sequencer switches the drums at 120 min, onto a drum at the end of its warm-up. Use the default seed and the default speed of 60 min per second, and export the CSV file after each run.

(a) Run as shipped (Loops) to 240 min. Find the lowest `pressure_bar_g` and its time, and the lowest `comp_pct` and its time. *Checkpoint:* 0.940 bar g at 129.2 min and 73.4 % at 137.4 min, within 0.002 bar, 0.1 % and 0.2 min.

(b) Choose Loops + schedule FF, press Reset and run to 240 min again. Find the range of `pressure_bar_g`, and the lowest `comp_pct` and its time. *Checkpoint:* 0.988 to 1.014 bar g, and 70.8 % at 122.0 min.

(c) In each run, which moved first, the pressure or the compressor? Say what the feedforward knew at 120 min that the pressure loop did not.

11.8 App challenge · 9 min

Load scenario 8, A skipped warm-up. Drum B goes on line at 2 h after its warm-up, and drum A, whose warm-up is skipped, goes on line cold at 18 h. Use the default seed and a speed of 300 min per second, run to 1260 min and export the CSV file.

(a) Read each new drum's shell temperature from the note the switch posts, and the peak of `truth_deficit_tph` after each switch. *Checkpoint:* 350 °C and 137 °C. 5.38 t/h and 16.90 t/h.

(b) Find the lowest `pressure_bar_g` in the three hours after each switch, and its dip below the setpoint of 1.00 bar g. *Checkpoint:* 0.940 bar g and 0.822 bar g, dips of 0.060 and 0.178 bar, within 0.002 bar.

(c) Choose Loops + schedule FF, press Reset and run to 1260 min again. Find the lowest pressure after the cold switch and the peak of `truth_deficit_tph` after it. *Checkpoint:* 0.976 bar g and 16.90 t/h.

Explain the ratio of the two peaks in (a), and why (c) changes the dip but not the peak.

11.9 App challenge · 9 min

Load scenario 9, Challenge, two days on the coker. Use the default seed and a speed of 600 min per second, so that each 2880 min run takes about five seconds. Read the So far tile when the run stops at 48 h.

(a) Run as shipped. Read the worst HCGO deviation, the worst pressure deviation and the antifoam used, and the verdict. *Checkpoint:* 9.7 °C, 0.186 bar, 18.9 kg. The target is missed.

(b) Choose Loops + schedule FF and set Antifoam to On foam level, press Reset and run the two days again. *Checkpoint:* 1.0 °C, 0.024 bar, 6.8 kg. The target is met.

(c) Use Table 11.5 to say which of the two changes fixes which failure of (a), and name the event at which the pressure and the HCGO left their bands in (a).

11.10 Application area · 4 min

In case 7 pass B reaches the 650 °C skin limit on coil day 40.0 at a coil outlet of 495 °C, and on day 63.7 at 490 °C. At the start of the run the cooler outlet makes 31.55 against 30.61 % coke on 120 t/h of fresh feed.

(a) Treat each run as lasting from a clean coil to the limit. How many decokes a year does each outlet need, and how much more coke a year does the cooler outlet make? *Checkpoint:* 9.1 against 5.7 decokes, 3.4 fewer. 1.13 t/h of coke, about 9900 t a year.

(b) Name what a planner must know to choose between the two, and two things the simulator leaves out of this trade.

11.11 Application area · 3 min

In case 6 a residue at 22 wt% CCR fills the drum faster. On the 16 h cycle the outage reaches 2.97 m at the switch. A 14 h cycle keeps 4.26 m, but the next drum goes on line at 268 °C instead of 339 °C, and a 12 h cycle at 151 °C.

(a) Why does a shorter cycle make a colder switch? Use the 0.054 t/h per kelvin of Exercise 11.5 to find the vapor held back at the switch for the 16, 14 and 12 h cycles. *Checkpoint:* 6.0, 9.8 and 16.1 t/h.

(b) At 26 wt% and 135 t/h no cycle keeps the outage above 4 m. Which handle is left, and what sets the shortest cycle?

11.12 Review the agent · 5 min

An AI assistant reads Table 11.5 and writes: “The schedule feedforward is unnecessary. A pressure PI at 100 % per bar with a 10 min integral time keeps the pressure within 0.089 bar over the two days, inside the 0.1 bar band at all four seeds, and it needs no model of the drum. Use this tuning on the loops and keep antifoam on foam level.” Find the flaw in the proposal and say how you would check it in the app.