INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 4

A Crude Atmospheric Tower Simulator for Degrees of Freedom and Triangular Interaction

Pages 81–103

Edited by John D. Hedengren

How to cite this chapter. "A Crude Atmospheric Tower Simulator for Degrees of Freedom and Triangular Interaction," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 4, pp. 81–103.

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

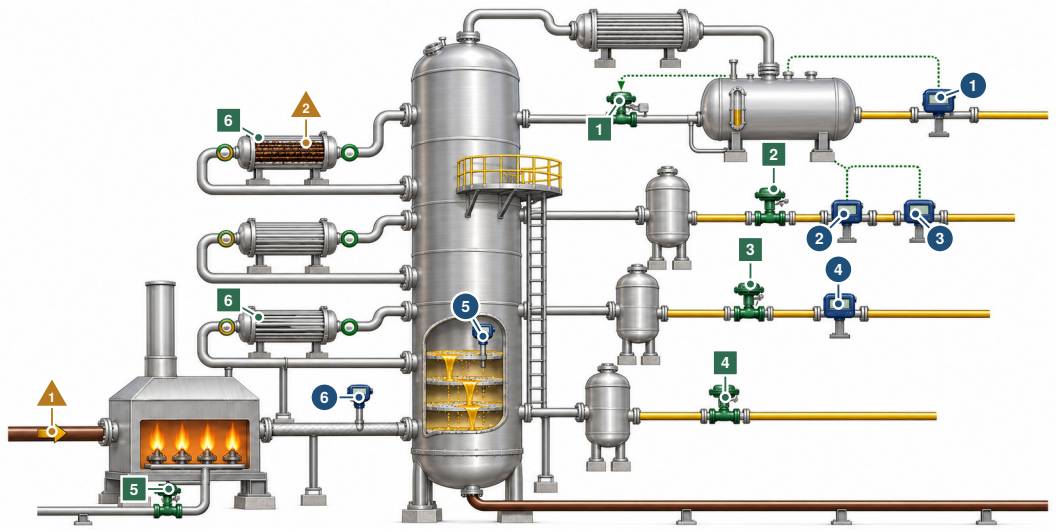
The app for this chapter
https://apmonitor.com/apps/crude-tower/

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

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

CHAPTER 4

A Crude Atmospheric Tower Simulator for Degrees of Freedom and Triangular Interaction



The controller moves 1 reflux valve 2 kerosene draw valve 3 diesel draw valve 4 gas oil draw valve 5 heater firing 6 heat taken out by the three pumparounds

The controller reads 1 naphtha end point 2 kerosene flash point 3 kerosene freeze point 4 diesel cloud point 5 overflash, the liquid that wets the wash zone 6 heater outlet temperature

What disturbs it ▲1 a heavier crude ▲2 a fouled top pumparound ▲ more charge

The dotted signal lines from the naphtha analyzer at the upper right and from the kerosene analyzers below it lead back toward the overhead drum and its one reflux valve, so only one of the readings can be held. Illustration, not to scale.

IN PLAIN TERMS Four Fuels, Four Flows, Eight Promises

The main idea. A crude tower boils crude oil and splits it into four fuels, from a light one near the top to a heavy one lower down. Operators decide where one fuel ends and the next begins by setting four flows.

The challenge. Each fuel must pass two tests, so there are eight promises and only four flows to keep them. It is like slicing a loaf from one end. A thicker slice moves every cut after it and none before it.

Why it matters. Only half of the promises can be held by control, and a plant that tries to hold all eight gets controllers that fight each other.

THREE WAYS TO READ THIS CHAPTER

Explore	10 min	This page, the Did you know? boxes in Sections 4.1 and 4.5.7, the History boxes in Sections 4.2 and 4.4, the Try it box in Section 4.5, and Key points in Section 4.9. 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 4.5 with Figure 4.3, Section 4.6, Section 4.7 and the Open challenge box in Section 4.8.

The control question A crude tower cuts one boiling curve into four products, each with two end specifications, using four handles. Which four of the eight numbers can be held, which way do the draws interact, and in what order should the loops be brought into service?

Learning objectives

- Count the degrees of freedom of a tower that makes products in series, and show that the naphtha end point and the kerosene flash point are one cut point.
- Compute cut points, tails and the overflash from the assay and the material and heat balances.
- Derive the gains and the inverse response of the top cut from the reflux power law and the tail.
- Explain why the gas oil draw is paired with the overflash and how the override protects the wash zone.
- Choose a commissioning order from the direction of interaction, and compare four loops with one predictive controller on measured runs.

Prerequisites Sections 1.4.3 (cut points and degrees of freedom), 1.4.2 (heater balance), 1.6.2 (PI loops), 1.6.3 (predictive control) and 1.6.4 (cases).

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

See also Volume 6, Chapter 2 for pairing on a triangular gain matrix and Volume 6, Chapter 7 for choosing which flow a structure holds. Volume 4, Chapter 2 for a protective override and Volume 4, Chapter 4 for throughput pushed against machine limits.

4.1 Introduction

A crude atmospheric tower is the first unit in a refinery and the last one a control course reaches, and the reason is that it does not fit the shape a course teaches with. Four products leave the tower as consecutive slices of one boiling curve. Each carries two specifications, an end point and a property such as a flash point or a freeze point, so there are eight numbers to hold. There are four places the curve can be cut. A student who counts the loops before counting the degrees of freedom will pair a controller with the naphtha end point and another with the kerosene flash point and then discover that the two are the same variable. Measured on the tower modeled here, taking the top reflux from 135 t/h to 170 t/h moves the naphtha end point from 191.4 °C to 176.0 °C and the kerosene flash point from 58.2 °C to 48.8 °C. Not identical, and not independent either. Two loops on one valve is the outcome, and on a panel it looks like a loop that has stopped working.

The misconception is that a specification is a controlled variable. On a unit that makes products in series from a single feed, half the specifications are consequences of the other half, and which half is which is an operating decision that changes with the season and the market. Hold the naphtha end point when the reformer is the constraint. Hold the kerosene flash point when jet fuel is worth more than gasoline. What cannot be done is to hold both, and no amount of tuning changes that. Counting independent handles before pairing loops is the cheapest engineering available on any unit, and it is the step that a two-input two-output example does not motivate.

DID YOU KNOW? Thirty to Fifty Trays in a Steel Cylinder

A crude tower is a steel cylinder about 120 feet high [1] that holds 30 to 50 fractionation trays [2]. All the crude oil a refinery runs passes through it first. On 1 January 2026 the 130 operable refineries of the United States could charge 19.2 million barrels of crude oil per stream day [3].

The second structural fact is that the interaction here runs one way. A draw moves its own boundary and every boundary below it and nothing above it, which makes the plant lower triangular. Measured, twenty more tonnes an hour of diesel leaves the kerosene cut point exactly where it was, while twenty more tonnes an hour of kerosene moves the diesel cut point by 17.6 °C. Towers with side strippers and side rectifiers have this structure generally [4], and so do many other units that separate a feed into consecutive fractions. Section 1.4.3 derives both facts from the material balance of the running example.

This chapter works with an interactive crude tower whose assay, tail sharpening and heat balance are hidden, so that the degrees of freedom must be counted from experiment, the triangular structure discovered rather than asserted, and a commissioning order chosen with a measurement to check it against. The model is an assay-sliced tower with its pumparound heat recovery, its furnace and its overflash override, written out in the detail needed to rebuild it. The scenario sequence contains a commissioning result that contradicts a standard heuristic and is measured rather than argued. The application runs in a browser without installation and exports its complete time history. The assay model, pumparound heat recovery and override control are established practice. Section 4.2 gives the domain history. Section 4.3 states the model: the cut points on the assay, the tails and the internal reflux, the heat balance and the overflash. Section 4.4 describes the four loops and their directions, the top loop selector, the overflash override and the optional predictive controller. Section 4.5 reports the measured results, from the degrees of freedom and the triangular interaction through the commissioning order, the pumparounds, the constraint and the graded challenge to the comparison with predictive control. Section 4.6 sets the loops beside the instruments, flow loops and control system of a real unit, Section 4.7 states the limitations, Section 4.8 the three results that generalize and Section 4.9 the conclusion. The exercises at the end of the chapter work through the same scenarios.

4.2 Background

Crude oil is a continuum of hydrocarbons and a refinery begins by cutting it into fractions by boiling range. The atmospheric tower does that first cut, and everything downstream, the reformer, the hydrotreaters, the cracking units, is fed by one of its draws. The measurement that describes the feed is the true boiling point assay, a laboratory distillation curve that reports the temperature at which each cumulative volume fraction has boiled. Characterizing a crude is still an active measurement problem, with spectroscopic methods used to predict the curve without running the distillation [5], and the product specifications are written on a different and cheaper laboratory distillation, ASTM D86, whose relationship to the true curve has its own literature [6].

HISTORY 1908 to 1922: The Pipe Still Takes Over From the Bench of Stills

Early continuous refineries passed crude oil through a bench of stills, as many as twelve or more. A book of 1922 records that tubular stills had largely replaced the bench in recent years [7]. One of the first plants to top crude by other means than the conventional stills was erected in 1908. In the Trumble system that the book describes, the crude was heated against hot distillates and residue, passed through a coil of pipes in a furnace and entered a single evaporator, a vessel 6 feet in diameter and 25 feet high [7]. The vapors left together and were condensed in stages into seven fractions, whose quality could be varied to some extent through the steam and cooling water of each stage [7]. The furnace, the single tower and the preheat against hot products are those of this chapter, but the side cuts were still condensed outside the tower.

Table 4.1 Classification of Variables at the Design Operating Point

Symbol	Role	Nominal	Range
EP naphtha	CV or consequence	184.1 °C	150–215 °C
FP kerosene	CV or consequence	53.9 °C	38–75 °C
FR kerosene freeze	CV	−47.0 °C	−58 to −34 °C
CP diesel cloud	CV	−1.8 °C	−14 to 8 °C
Φ overflash	CV, constraint	3.28 vol%	> 2 vol%
L top reflux	MV	150 t/h	60–320 t/h
D_K	MV	74 t/h	20–160 t/h
D_D	MV	102 t/h	30–200 t/h
D_G	MV	51 t/h	0–130 t/h
COT	MV, economic	377 °C	340–400 °C
$Q_{TPA,KPA,DPA}$	MV, economic	18/14/10 MW	≤35/30/25 MW
charge	DV, measured	662 m^3/h	400–850 m^3/h
crude heaviness	DV, unmeasured	0	−0.4 to 0.8
Q_F	constraint	76.5 MW	≤ 85 MW

Ranges are the spans the application accepts, as setpoints for the controlled variables and as valve or slider positions for the others.

Control of the tower developed from tray temperatures. Before continuous analysis, each draw was set from a tray temperature and adjusted when the laboratory reported a product outside its specification, with the cut point inferred rather than measured. That arrangement is workable because the temperature at a draw tray is close to a cut point, and it fails whenever the crude changes, because a different crude has a different relationship between temperature and composition at the same tray. The modern answer is either an inferential built from several temperatures and flows or a direct analyzer, and in both cases the interesting work is the layer above (Section 1.4.6).

Multivariable control arrived on crude units early and for a specific reason. The tower has more measured qualities than handles, several hard constraints including the furnace and the wash zone, and an economic objective that changes with the crude slate, which is exactly the problem a model predictive controller with an inequality-constrained objective is built for [8]. Dynamic simulation and nonlinear model-based quality control of a crude tower were demonstrated in the mid 1990s [9], and subsequent work has concentrated on tuning those controllers [10], on integrating them with real time optimization [11], and on the reduced models a planning layer needs above them [12, 13]. Surrogate models for design optimization of the same unit are a current subject [14]. The regulatory layer underneath all of it is still four loops and an override, and it is that layer this application is about. One such multivariable controller is also offered as an option above that layer, and Section 4.5.7 measures it against the four loops.

Pumparound heat recovery arrived with energy prices. A pumparound withdraws hot liquid from a tray, exchanges it against incoming crude and returns it cooler, which both recovers heat and condenses vapor. Measured on this tower with all four loops in automatic, taking the top pumparound from 12 MW to 26 MW raises the crude preheat from 238.7 °C to 264.2 °C and cuts the furnace firing from 81.0 MW to 70.5 MW, so 10.5 MW of fuel follow from one setpoint. It is the largest single energy lever on the unit and it is not free, for reasons Section 4.7 states. The preheat train that receives this heat, and the fouling that erodes it, are the subject of Chapter 3.

4.3 Process Model

Tower T-101 charges 662 m^3/h of a light sweet crude through fired heater F-101 to a flash zone, and makes overhead naphtha, kerosene, diesel and atmospheric gas oil with a residue from the bottom. Three pumparounds recover heat into the crude preheat train. Figure 4.1 shows the process arrangement and principal manipulated variables.

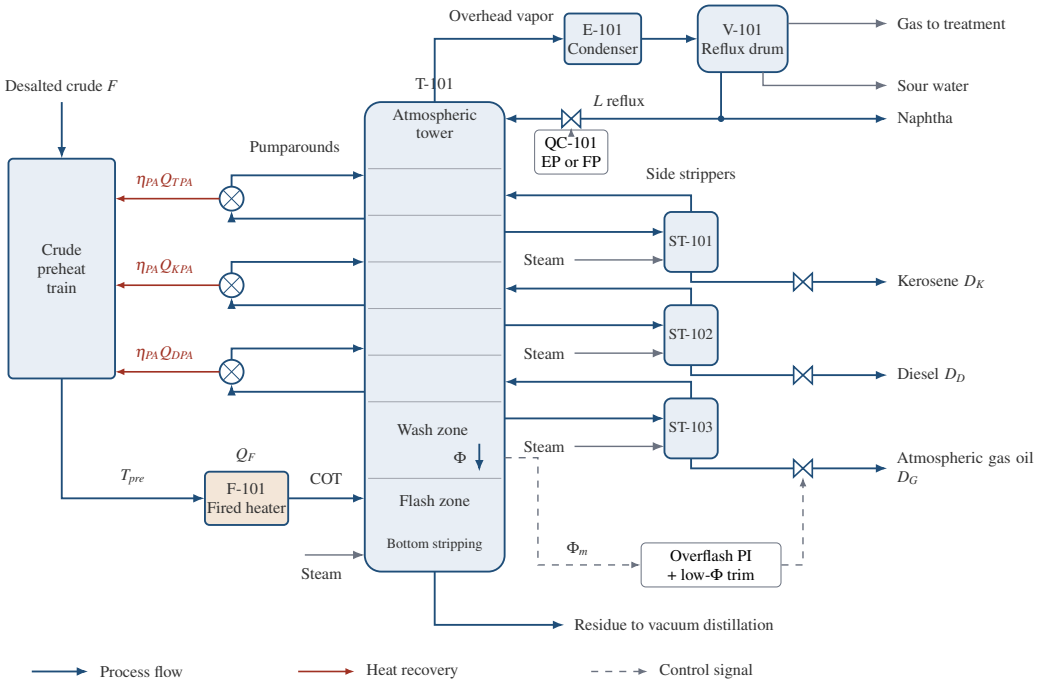


Figure 4.1 Simplified process and control schematic of atmospheric tower T-101. Desalted charge F receives a recovered fraction η_{PA} of the heat removed by three separate pumparound circuits (Q_{TPA}, Q_{KPA} and Q_{DPA}), each returning cooled liquid above its withdrawal tray. Preheated crude at T_{pre} passes through F-101, with firing duty Q_F, and enters the flash zone after reaching coil outlet temperature COT. Condenser E-101 and drum V-101 separate overhead gas, sour water and hydrocarbon liquid, which splits into naphtha and reflux L. QC-101 adjusts L using either naphtha end point EP or kerosene flash point FP. Side strippers receive steam, return vapor above their respective draw trays, and discharge kerosene D_K, diesel D_D and atmospheric gas oil D_G. Overflash Φ is the internal liquid descending through the wash zone. Its measured estimate Φ_m feeds the gas oil PI controller and a downstream protective trim that reduces D_G at low overflash in both Manual and Auto. The kerosene and diesel quality loops, pumps, level and pressure controls, and downstream product cooling are omitted for clarity. Steam and overhead separation are shown as process context, not as independently simulated units.

4.3.1 Cut Points and the Assay

The crude is described by a true boiling point curve $T_{bp}(v)$ relating cumulative volume fraction to boiling temperature, interpolated linearly between anchor points from 20 °C at $v = 0$ to 750 °C at $v = 1$. A heaviness h raises the curve by $h\,(40 - 25v)$ °C, more at the light end than at the heavy end. Four cut fractions $v_1 \dots v_4$ divide it, and each product is the slice between two of them. The naphtha cut is placed by the top reflux, Equation (4.1), which is a power law fitted so that more reflux pulls the naphtha cut down.

$$v_1 = v_{1,ref}\left(\frac{L_{ref}}{L}\right)^{p}(1 + 0.35\,h) \tag{4.1}$$

The remaining three are placed by the material balance. The cut below draw i lies at the naphtha fraction plus the volume of every draw above it, and each cut point T_i follows its place on the assay with the first-order lag τ_i, Equation (4.2).

$$v_{i+1} = v_1 + \sum_{j \le i}\frac{D_j}{\rho\,F}, \qquad \tau_i\,\frac{dT_i}{dt} = T_{bp}(v_i) - T_i \tag{4.2}$$

where D_1, D_2 and D_3 are the kerosene, diesel and gas oil draws (D_K, D_D, D_G), F is the volumetric charge and $\rho = 860\,\mathrm{kg/m^3}$ is one density for the crude and all of its products. That equation is the triangular structure. A draw enters the cut points below it and none above it, as in Equation (1.4), which writes the same balance with the naphtha counted as the first draw.

Pause. Try Exercise 4.2 before reading on.

4.3.2 Tails and Internal Reflux

A cut is not sharp. Each product carries a tail of material from the fraction on either side, and the tail width narrows as the internal reflux in that section rises, Equation (4.3).

$$\sigma_i = \frac{\beta}{\max(0.2, R_i)^n}, \qquad R_i = \frac{L_i}{D_i} \tag{4.3}$$

where σ_i is the tail spread in degrees Celsius, D_i is the product drawn at the top of section i (the overhead naphtha for the top section), $n = 1$, and L_i is the liquid running down through section i. The liquid starts as the top reflux and, section by section, gains the vapor condensed by the pumparound above it, Q_{PA}/λ with $\lambda = 250\,\mathrm{kJ/kg}$ (14.4 t/h per megawatt), and loses the draw taken off at its top. Product properties are linear correlations on the cut points and tails. The naphtha end point is $T_1 + \sigma_1$, the kerosene flash point is $T_1 - \sigma_1 - 112\,^\circ\mathrm{C}$, and the kerosene freeze point and the diesel cloud point are linear in the 95 percent points $T_2 + \sigma_2$ and $T_3 + \sigma_3$.

This is where a draw does its second job. Liquid drawn off at a tray no longer runs down through the trays below it, so a larger draw thins the section underneath and widens its own tail. Measured on the diesel draw, thirty more tonnes an hour moves the diesel and gas oil boundary from 337.6 °C to 363.9 °C, a shift of 26.4, while the diesel 95 percent point moves from 340.4 °C to 367.8 °C, which is 27.4. The extra degree is the tail, which widened from 2.82 °C to 3.91 °C. The second effect is about four percent of the first, and it means the loop gain is not constant.

4.3.3 Heat Balance and the Overflash

The furnace heats the preheated crude to the flash zone and vaporizes part of it, Equation (4.4), and the pumparounds set the preheat, Equation (4.5).

$$Q_F = \dot{m}\left[c_p\,(T_{fz} - T_{pre}) + \lambda\,\phi_v\right], \qquad T_{fz} = \mathrm{COT} - 2\,^\circ\mathrm{C} \tag{4.4}$$

$$T_{pre} = T_{cold} + \frac{\eta_{PA}\,(Q_{TPA} + Q_{KPA} + Q_{DPA})}{\dot{m}\,c_p} \tag{4.5}$$

where $\dot{m} = \rho F$ is the crude mass flow, T_{fz} is the flash zone temperature after a fixed transfer line drop, $\phi_v = T_{bp}^{-1}(T_{fz} + \Delta T_{fl})$ is the fraction vaporized, read on the assay $\Delta T_{fl} = 25\,^\circ\mathrm{C}$ above the flash zone temperature to stand for the stripping steam, and the preheat follows Equation (4.5) with the lag τ_{pre}. Equation (4.4) is the heater balance of Equation (1.1) with the vaporization term kept. The overflash is the liquid that overflows the wash zone above the flash zone, and it is what wets the packing there. It is the difference between what the flash zone vaporizes and what the draws take, $\Phi = 100\,(\phi_v - v_4)$ in volume percent, so it falls when the coil outlet falls or when the draws rise. With every draw held, the overflash moves 0.185 vol% per degree of coil outlet temperature, and eight degrees off the coil outlet takes it from 3.28 vol% to 1.80 vol%, below its 2 vol% minimum, while the kerosene freeze point does not move at all.

4.3.4 Assumptions

The tower is represented by cut points on an assay rather than by trays. There is no stage-by-stage calculation, no vapor-liquid equilibrium and no tray hydraulics, so flooding, weeping and the capacity limit that ends a real push are absent.

Table 4.2 Model Parameters

Symbol	Meaning	Value	Class
c_p	crude specific heat	2.60 kJ/kg/K	adjustable
η_{PA}	pumparound heat recovered	0.75	adjustable
T_{cold}	crude before pumparounds	173 °C	adjustable
ΔT_{fl}	flash curve below TBP	25 °C	adjustable
λ	latent heat of traffic	250 kJ/kg	fixed
β	tail spread at unit R	12.0 °C	adjustable
n	tail sharpening exponent	1.0	hidden
$v_{1,ref}$, L_{ref}	reference naphtha cut	0.200, 150 t/h	hidden
p	reflux power on the top cut	0.45	adjustable
τ_1–τ_4	cut point lags	4, 8, 12, 15 min	adjustable
τ_v	draw and reflux flow lag	1.2 min	adjustable
τ_{COT}	furnace thermal lag	5 min	adjustable
τ_{PA}	pumparound duty lag	6 min	adjustable
τ_{pre}	preheat train lag	14 min	adjustable
θ	product analyzer cycle	12 min	adjustable
τ_{an}	analyzer averaging lag	2.0 min	adjustable
Φ_{min}	wash zone dry below	2.0 vol%	visible
$Q_{F,max}$	furnace firing limit	85 MW	visible
COT_{max}	coil outlet limit	395 °C	visible

The tail model is empirical. A real separation sharpness depends on tray count, efficiency and stripping steam as well as on internal reflux, and the relationship between a tail and a D86 specification is a correlation rather than a calculation.

Stripping steam is folded into a fixed offset between the true boiling point curve and the flash curve, so steam rates cannot be studied.

The pumparound sections have no vapor traffic model above the return, which is the largest single limitation and is discussed in Section 4.7.

The crude is a single assay with a heaviness parameter that shifts the curve and tilts it linearly. A real crude switch changes the shape of the curve and not just its position and slope, and it changes sulfur, acid number and metals, none of which appear here.

Product properties are correlations fitted to the sliced assay rather than laboratory methods. The D86 relationship to the true curve is itself an approximation [6].

4.4 Controller Design

4.4.1 Velocity Form PI

Every loop uses the velocity-form PID of Section 1.6.2, Equations (1.16) to (1.19), executed every $\Delta t_c = 0.2$ min. No loop in this application uses derivative action ($\tau_D = 0$), so each is a PI controller, and the tracking time of the back-calculation is $T_t = \tau_I$. The controller direction is the $a = \pm 1$ of Equation (1.16). In manual the integral follows the operator's valve position, which makes the transfer to automatic bumpless, and in automatic the operator's station follows the controller output. The quality loops act on analyzer readings that are sampled every $\theta = 12$ min and averaged with a 2 min lag, while the overflash loop acts on a continuous measurement with a 0.8 min lag. Integration is fixed-step fourth-order Runge-Kutta at 0.05 min (Section 1.6.1).

4.4.2 The Four Loops and Their Directions

The plant gains at the design point, taken as local slopes of the steady-state model, are –0.448 °C/t/h from top reflux to naphtha end point (–0.425 °C/t/h averaged over a 10 t/h increase, the figure the application quotes), 0.692 °C/t/h from kerosene draw to kerosene freeze point, 0.411 °C/t/h from diesel draw to diesel cloud point, and –0.176 vol%/t/h from gas oil draw to overflash. The section

time constants are 4, 8, 12 and 15 min. The tunings the automatic scenarios load set each controller gain near the reciprocal of its plant gain, the internal-model-control gain for a closed-loop time constant equal to the open-loop one [15], giving K_c of 2.5, 1.4, 2.4 and 5.5. The integral times of 6, 12, 18 and 22 min are one and a half times the section time constants rather than equal to them.

Two of the four loops are direct acting and two are reverse acting, and getting the signs right requires reading the assay rather than the flow sheet. More reflux makes a lighter naphtha, so a rising end point calls for more reflux and the end point loop is direct acting ($a = -1$). The flash point controller on the same valve is direct acting for the same reason. More kerosene draw takes heavier material into the kerosene, so its freeze point rises with the draw and that loop is reverse acting ($a = +1$), as is the diesel loop. More gas oil draw takes liquid that would otherwise have been overflash, so a falling overflash calls for less draw and the overflash loop is direct acting.

> ⧖ HISTORY **1973: Which Variables, Which Valves, Which Links**
>
> In 1966 the relative gain array gave engineers a measure of interaction for deciding which valve to pair with which measurement [16]. The wider question was put in a 1973 critique of chemical process control theory: which variables should be measured, which inputs manipulated, and what links made between the two sets. Academic work, the critique held, had largely neglected these structural decisions [16]. A series of papers in 1980 answered in part with a rule, to always control the active constraints, and the pairing of the gas oil draw with the wash zone in this chapter is in that spirit. The review that gives these dates still makes the practical case for single loops. They can be tuned one at a time without a model of the interactions, and the gain from a multivariable controller has to be weighed against the cost of obtaining and maintaining its model [16].

4.4.3 The Top Loop Has Two Objectives and One Valve

The top loop selector chooses whether the reflux holds the naphtha end point or the kerosene flash point. It cannot hold both, and the application makes the choice explicit rather than offering two loops that would fight. Each objective has its own controller, and the one not selected tracks the reflux so that a switch between them is bumpless. Measured in scenario 2, the end point loop takes the end point from 184.0 °C to 176.0 °C by raising the reflux from 150 t/h to 170 t/h, and the kerosene flash point follows from 53.9 °C to 48.8 °C, a ratio of 0.63.

🕮 **Pause.** Try Exercise 4.1 before reading on.

4.4.4 The Overflash Override

The gas oil draw is paired with the overflash rather than with a gas oil specification. A slightly heavy gas oil is giveaway. A dry wash zone is coke on the packing and a shortened run. A one-directional override acts on the gas oil draw downstream of its loop. Whenever the measured overflash Φ_m is below its minimum, the requested draw is multiplied by $1 - \min(0.6, (\Phi_{min} - \Phi_m)/1.5)$, so the cut grows by 6.7 percent of the draw per 0.1 vol% of shortfall up to 60 percent, without waiting for the loop to integrate. The override is not tied to the loop mode and also acts with the gas oil draw in manual. The PI controller neither sees the cut nor back-calculates against it, so its integral keeps acting on the overflash error alone, and because the cut vanishes continuously as the measured overflash returns to the minimum, the hand-back to the loop is bumpless. The override is proportional, so it limits the time below the minimum without removing it. Pairing a manipulated variable with a constraint rather than with a product is the right choice whenever the constraint is more expensive to violate than the specification is. Chapter 5 protects its wash bed with an override of the same kind, and Volume 4, Chapter 2 treats a protective override on a machine.

4.4.5 Model Predictive Control

A second structure is offered beside the four loops, never in their place. Section 1.6.3 describes the predictive control option shared by the series, and this section gives what is specific to the tower. With Control set to MPC and the tower in automatic, one plan moves the top reflux and the three

Table 4.3 Identified Coefficients of the Controller's Model

Handle	Quality	a	b	Gain
L	EP (°C)	−0.0881	−0.3603	−0.4484
L	FP (°C)	+0.0882	−0.3602	−0.2720
L	FR (°C)	−0.0059	−0.2252	−0.2311
L	CP (°C)	−0.0029	−0.1351	−0.1380
L	Φ (vol%)	+0.0600	—	+0.0600
D_K	FR (°C)	+0.0328	+0.6587	+0.6915
D_K	CP (°C)	+0.0029	+0.3952	+0.3981
D_K	Φ (vol%)	−0.1756	—	−0.1756
D_D	CP (°C)	+0.0153	+0.3952	+0.4105
D_D	Φ (vol%)	−0.1756	—	−0.1756
D_G	Φ (vol%)	−0.1756	—	−0.1756

Coefficients are per t/h of the handle, a on the flow lag and b on the section lag of Equation (4.6). The gain is $a+b$.

draws together. Every $T_s = 1$ min, five controller scans, a quadratic program chooses all four handles over $N = 120$ decisions, two hours ahead, and only its first move is applied, for one period [8, 17]. The loop switches then say only which loops take over if a plan fails, and the wash zone override still trims the gas oil draw downstream.

The controller's model maps $u = (L, D_K, D_D, D_G)$ to the five qualities that carry both a specification and an instrument the loops read, the end point, the flash, freeze and cloud points and the overflash. A first design with eight outputs, adding a gas oil end point and two tray temperatures, was not built, because the application has no gas oil end point analyzer or specification and its one tower temperature is read by no loop. Step tests at the design point, in manual with the analyzer noise off (5 t/h of reflux and 3 t/h of draw, up and down, 120 min of the true qualities), gave second-order-plus-dead-time fits that all placed a lag near 1 min beneath one of 4, 8 or 12 min. That is the structure of the plant equations, so the model is written on it, Equation (4.6), with coefficients fitted by linear least squares.

$$\begin{aligned} \tau_v \dot{x}_{v,j} &= u_j - u_{j,0} - x_{v,j}, \qquad \tau_s \dot{x}_{s,j} = x_{v,j} - x_{s,j} \\ \hat{y}_q &= y_{q,0} + \sum_j \left(a_{qj} x_{v,j} + b_{qj} x_{s(q),j} \right) + \beta_q \end{aligned} \tag{4.6}$$

Here $u_{j,0}$ is handle j at $t = 0$, $x_{v,j}$ its flow lag state with $\tau_v = 1.2$ min, and $x_{s(q),j}$ the lag of the section whose cut point quality q reads, τ_1 for the end point and flash point, τ_2 for the freeze point and τ_3 for the cloud point, while the overflash reads the flow lags alone. $y_{q,0}$ is the reading at the start and β_q the output bias. The lags are the declared values of Table 4.2, not the run's, so an edited plant parameter is a model mismatch. The ten-state model's coefficients are in Table 4.3, with residuals of at most 0.003 percent of each response, and a repeat with 1 t/h steps agrees within 0.08 percent of each gain from 1 to 120 min. The flash point's apparent 1.5 min dead time is an inverse response that the opposite signs of its a and b reproduce.

The plan predicts the true qualities, and the analyzers enter through the estimator. Each prediction passes through a copy of its instrument, sampled in the scan in which the real analyzer publishes and averaged with the declared 2 min lag, or for the overflash the declared 0.8 min lag, and every scan $\beta_q \leftarrow \beta_q + \lambda_q (z_q - y_{q,0} - \tilde{y}_q - \beta_q)$. Here z_q is the reading the loops see, $\tilde{y}_q$ the copy's output, $\lambda_q = \Delta t_c / \tau_b$ with $\tau_b = 6$ min for the analyzers and $\lambda_q = 1$ for the overflash, which makes the prediction offset free [18]. Charge, crude, coil outlet and pumparounds reach the plan only through the biases, and there is no preview. The plan minimizes Equation (4.7) over twelve move blocks of

1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15, 75 decisions.

$$J = \sum_{k=1}^{N} \left[q\,(\hat{y}_{t,k} - r_t)^2 - w\left(D_{K,k} + \tfrac{p_D}{p_K} D_{D,k} + \tfrac{p_G}{p_K} D_{G,k} \right) \right] + S \sum_{m=1}^{12} \|\Delta u_m\|^2 + \sum_i \rho_i \left(\sigma_i^2 + \sigma_i \right) \tag{4.7}$$

$\hat{y}_{t,k}$ is the top quality the selector holds, k decisions ahead, with target r_t, $q = 1$ per °C^2, $S = 0.1$ per (t/h)^2 with Δu_1 measured from the commands applied now, and $w = 0.05$ per t/h scales indicative values of $p_K = 260$, $p_D = 230$ and $p_G = 120$ dollars per tonne above residue. The constraints are Equation (4.8).

$$\begin{aligned} & u_{min} \le u_m \le u_{max}, && \sigma_i \ge 0 \\ & \widehat{FR}_k \le e_k^{FR} + \sigma_{FR}, && \widehat{CP}_k \le e_k^{CP} + \sigma_{CP} \\ & \hat{\Phi}_k \ge e_k^{\Phi} - \sigma_{\Phi}, && \hat{\Phi}_k \ge e_k^{W} - \sigma_W \end{aligned} \tag{4.8}$$

The handle ranges of Table 4.1 are hard. The soft one-sided bands sit on the loops' own targets (−47 and −1.8 °C, 3.3 vol%) rather than on the specifications, so that both structures work to the same numbers, with a second overflash band at $\Phi_{min} + 0.6$ vol% and a band on the unheld top quality at its specification (38 or 185 °C). Each slack is priced at $\rho_i = 10^4$, or 10^5 for the wash zone band. Each edge is a funnel, $e_k = \bar{y} + \max(0,\, h_k - \bar{y})\, f_k$ for an upper limit $\bar{y}$ and mirrored for a lower one, where h_k is the prediction with the present commands held, $f_k = 1$ for the first minute and $\exp(-(t_k - 1)/\tau_q)$ after it, and τ_q is the declared lag of the quality's section, or 1.2 min for the wash zone band. The top quality is a setpoint because only one of the end point and flash point can be held, and lined out on scenario 3 with the noise off the value holds the true end point 0.0190 °C below target (the plan's own horizon-end prediction is 0.0202 °C). A dual active-set solver [19] handles each plan, and S and τ_b were chosen once on scenario 10 for the least integrated error among settings that rejected no plan and moved the handles no more than the loops.

Under the plan each loop tracks the applied command with no preload, its output equal to the command at every scan of scenarios 6, 8 and 10 with the noise off. Any plan with a solver status other than optimal, or with a first move that is infinite, undefined or beyond a handle range, is discarded with a historian mark, and the loops take over. Switching to Loops at five instants moved each handle by exactly one loop execution, at most 0.0165 t/h of reflux and 0.1377 t/h of gas oil, and a switch to manual leaves every station where the plan put it. The plant runs on the run's parameters and the instruments sample, average and add noise, so neither coincides with this model. In Chrome a plan took 4046 and 4126 µs on average in two measurements on scenario 10, 0.41 percent of a core at real time and 48.6 to 49.5 percent at the 120 times top speed, where the application kept 58.5 to 59.7 fps against 59.5 to 59.9 fps under the loops.

4.5 Results

Eleven scenarios run from slicing the assay by hand to a graded challenge. Each scenario is credited separately (Section 1.6.4). A case is complete when the learner performs the actions its rule lists and the simulation reaches the listed state. The challenge is judged on the case as given, and changing a model parameter or a disturbance setting voids the attempt. Credit accumulates toward certificates of progress and of completion for the module.

4.5.1 Degrees of Freedom

Table 4.4 reports the top reflux swept across its working range with the top loop in manual and every draw held at its design value, read once the run has settled.

The measured ratio of the two movements is about 0.63 over 150 t/h to 170 t/h. It depends on the assay and on how the tail narrows as the traffic rises. The tower has four cut points and eight end

Table 4.4 One Valve Moving Two Specifications

Quantity	$L = 135\,\mathrm{t/h}$	$L = 170\,\mathrm{t/h}$
Naphtha end point (°C)	191.4	176.0
Kerosene flash point (°C)	58.2	48.8

specifications, so four of the eight are choices and four are consequences, and the only question is which four to choose.

4.5.2 Triangular Interaction

With every draw held, twenty more tonnes an hour of diesel leaves the kerosene cut point exactly where it was, and twenty more tonnes an hour of kerosene moves the diesel cut point by 17.6 °C. The closed loop shows the same direction in scenario 3. Lowering the freeze point target from −47 °C to −50 °C with all four loops in automatic takes the kerosene draw from 74.0 t/h to 69.8 t/h and the diesel draw from 102.1 t/h to 106.2 t/h, while the reflux and the gas oil draw (50.9 t/h) stay where they were. Lowering the cloud point target by 3 °C instead moves the diesel draw to 94.7 t/h and the gas oil draw to 58.1 t/h, and the kerosene draw stays at 74.1 t/h. The interaction runs one way, down the tower, which is what makes a tower with four interacting loops tunable one loop at a time. Volume 6, Chapter 2 finds the same triangular structure in the gain matrix of a reverse osmosis train.

4.5.3 Commissioning Order

Scenario 7 brings the four loops into automatic twenty minutes apart from the bottom of the tower upward, the gas oil loop at 10 min, the diesel loop at 30, the kerosene loop at 50 and the top loop at 70, and a heavier crude (heaviness 0.35) arrives at 120 min. Summed over the reflux and the three draws, the valve travel over the first 120 min is 3.8 t/h. Bringing the top loop in at 10 min and the kerosene loop at 30, alongside the scripted gas oil and diesel switches, raises it to 6.4 t/h, and most of the difference (2.1 t/h of 2.6 t/h) is in the reflux. A strictly top-down order (top at 10 min down to gas oil at 70, with the scripted switches removed) takes 5.7 t/h. After all four are in service the schemes behave identically, with peak overflash deviations of 4.99 vol% and 5.01 vol% and peak freeze point deviations of 12.02 °C and 12.05 °C on the crude change, within 0.3 percent of each other, so the entire difference is in the commissioning transient.

The reason is the triangular structure. Closing the top loop moves the naphtha boundary and therefore every boundary below it. Downstream loops closed first are already in service to absorb that, while downstream loops closed last let the drift accumulate in manual until an operator picks it up. The familiar heuristic, commission the inner or fastest loop first, is about cascades, where the interaction points from the inner loop outward. The statement that covers both is to bring up the loop that receives the disturbance before the loop that generates it.

4.5.4 Pumparound Heat Recovery

Table 4.5 reports scenario 5, in which the top pumparound steps from 18 MW to 26 MW at 30 min and to 12 MW at 180 min with all four loops in automatic, read at 30, 180 and 360 min. Each megawatt of duty raises the preheat by 1.82 °C and saves 0.75 MW of firing, which is η_{PA}. The kerosene loop takes 1.7 t/h more draw at 26 than at 12 MW, because the sharper split puts more kerosene into the cut at the same freeze point.

Ten and a half megawatts of fuel and a sharper split follow from one setpoint, which raises the obvious question of why not take all of it. In this model nothing stops it, and that is a limitation stated in Section 4.7.

Table 4.5 Top Pumparound Duty Against Preheat, Firing and Sharpness

Quantity	12 MW	18 MW	26 MW
Crude preheat (°C)	238.7	249.6	264.2
Furnace firing (MW)	81.0	76.5	70.5
Kerosene and diesel tail (°C)	3.51	2.65	2.00
Kerosene draw (t/h)	73.1	74.1	74.8

TRY IT IN THE APP **5 · Pumparounds: heat against sharpness**

Load the scenario and press Run. Nothing has to be set. At 30 min the top pumparound steps from 18 to 26 MW, and at 180 min it drops to 12 MW. Watch the furnace firing beside the tower. It settles at 70.5 MW with the duty at 26 MW and at 81.0 MW with it at 12 MW, because each megawatt recovered saves 0.75 MW of firing.

4.5.5 The Constraint Binds First

With every draw held, eight degrees off the coil outlet takes the overflash from 3.28 vol% to 1.80 vol%, below its minimum, while the kerosene freeze point does not move at all. The sensitivity is 0.185 vol% per degree and nothing else. With the four loops in automatic the gas oil loop absorbs an eight-degree loss without the override, and the overflash bottoms at 2.70 vol%, so scenario 6 takes 22 °C off the coil outlet at 30 min. The overflash crosses its minimum at 34.2 min, the override engages at 35.1 min and removes at most 2.95 t/h beyond the loop's request, and the overflash bottoms at 1.89 vol% at 42.2 min before the override releases at 58.4 min, as shown in Figure 4.2. The wash zone is dry for 23 min in all. The gas oil draw settles at 27.5 t/h before the coil outlet is restored at 200 min and is back at 50.4 t/h by 360 min, and the kerosene freeze point stays within 0.09 °C of its target throughout. The constraint binds long before any product notices, which is the argument for pairing the gas oil draw with it.

Pause. Try Exercise 4.3 before reading on.

4.5.6 The Graded Challenge

The challenge runs 700 min with analyzer noise of 0.8 °C and three disturbances, a heavier crude (heaviness 0.35) at 60 min, a charge increase from 662 m^3/h to 740 m^3/h at 260 min, and a fouled top pumparound (18 MW to 12 MW) at 460 min. It starts with the top loop alone in automatic and cautious tuning. It is met when, with the targets left as given, the overflash stays within 0.65 vol% of target and the kerosene freeze point within 1.5 °C over the last third of the run, and the wash zone is dry for no more than 15 min in all. With the shipped configuration (top loop alone in automatic, cautious tuning, coil outlet 377 °C) the worst deviations over the last third are 4.01 vol% and 1.92 °C and the wash zone is dry for 174.2 min, so the run fails on all three counts. With the suggested configuration (all four loops in automatic with K_c and τ_I of 1.5 and 10 min on the top loop, 0.8 and 24 min on the kerosene loop, 2.4 and 18 min on the diesel loop and 14 and 22 min on the gas oil loop, coil outlet 381 °C, kerosene and diesel pumparounds at 24 and 18 MW) they are 0.28 vol% and 0.65 °C with 1.0 min dry, inside the limits by 0.37 vol%, 0.85 °C and 14.0 min. Both runs use the analyzer noise seed the application reuses on every run. The suggested configuration also passes with each of seeds 1 to 400, at worst 0.46 vol%, 1.24 °C and 2.0 min dry. The suggested tuning differs from that of the other automatic scenarios. The top and kerosene loops act on the noisy analyzers and both move the kerosene cut, so they run softer (K_c 1.5 and 0.8 against 2.5 and 1.4), while the gas oil loop acts on the noise-free overflash and runs at $K_c = 14$ rather than 5.5. The wash zone is then dry only for the minute after the instantaneous crude switch. The heat moves are not scored. They keep the firing at 81.7 MW after the pumparound fouls, where the same tuning without them fires 93.0 MW.

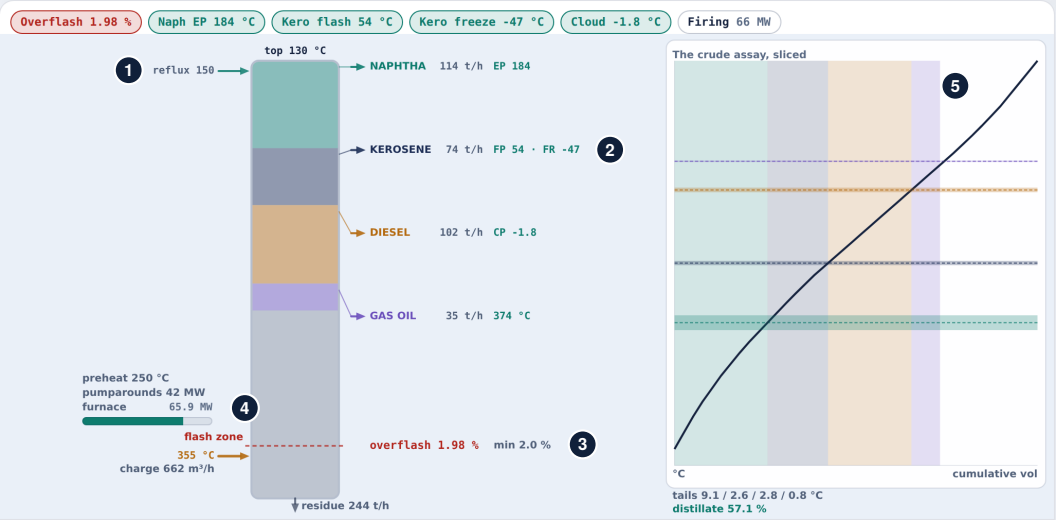


Figure 4.2 The crude tower application in scenario 6 at 55 min, 25 min after the coil outlet temperature stepped from 377 °C to 355 °C with all four loops in automatic, cropped to the plant view with numbered callouts. (1) The top reflux, the handle of the top loop. (2) The kerosene draw with its flash and freeze points. (3) The overflash at 1.98 vol%, just below its 2 vol% minimum, where the override holds the gas oil draw at 35.0 t/h against a loop request of 35.7 t/h. (4) The furnace firing of 65.9 MW. (5) The inset, which slices the assay at the four cut points. The naphtha end point, kerosene freeze point and diesel cloud point stay on their targets (184.0, −47.0 and −1.8 °C), and in the full window, outside this crop, the trends cover the last 30 min.

4.5.7 Comparison With Model Predictive Control

Table 4.6 sets the four loops against the plan of Section 4.4.5 on scenarios 10, 5, 6, 9 and 8 and on the challenge. Each row runs from $t = 0$ with the case's seed, instruments and noise, the loops at shipped tuning (suggested on the challenge), the plan at scenario 10's settings, with no preview and the same targets. IAE and travel are the scorecard's tiles on the true qualities. The off-specification minutes are not tiles but the CSV accumulators `metric_off_naph_min`, `metric_off_kero_min` (flash below 38 °C or freeze above −40 °C) and `metric_off_diesel_min`, integrated every 0.05 min step. Figure 4.3 draws three of the table's columns for the six cases that both structures run.

In scenario 10 only the reflux is asked to move, but a lighter naphtha moves every cut beneath it. Under the loops the freeze point falls to −50.69 °C at 39.4 min, because the reflux has moved the naphtha and kerosene cut and the kerosene draw has not, and swings back to −45.95 °C at 62.4 min. The plan's freeze point never falls below −48.10 °C, as shown in Figure 4.4, because its model carries the triangle and raises the kerosene draw with the reflux, to 74.98 t/h by 31 min while the loops' draw is still at 74.06. On the step back the loops carry the end point to 187.59 °C at 158.2 min, off specification for 8.35 min, and after 150 min the plan's end point never exceeds 183.98 °C. The plan wins all four integrated errors, the freeze point by 18.1 against 123.9 °C · min, with 98.8 against 270.2 t/h of travel and the same 223.0 k$ of draw value.

On scenario 9 the plan cuts the end point error from 304.7 to 35.1 °C · min but loses the cloud point, 318.7 against 179.6, because the higher charge lightens every cut and the plan's cloud point, never above its −1.8 °C target, gives away inside a one-sided band that only the value's pull closes, more slowly than the diesel loop's integral. On scenario 6 the wash zone is dry for 8.3 against 23.45 min, and on scenario 5 the plan loses only the freeze point, 30.8 against 26.8, with 10.3 t/h of travel against 33.8. The draw value stays within 0.3 percent of the loops' in every shared case.

Table 4.6 Four Loops Against Model Predictive Control

Case (window)	Structure	IAE, true quality from $t = 0$				Travel	Off specification (min)			Wash zone	Draw value	Plans,
		EP	FR	CP	Φ	(t/h)	naphtha	kerosene	diesel	dry (min)	(k$)	rejected
10, end point steps (270 min)	loops	148.2	123.9	77.4	39.4	270.2	8.35	0	0	0	223.0	—
	predictive	66.0	18.1	10.0	4.1	98.8	0	0	0	0	223.0	270, 0
5, pumparound traffic (360 min)	loops	12.6	26.8	20.8	10.0	33.8	0	0	0	0	292.9	—
	predictive	7.4	30.8	15.4	7.2	10.3	0	0	0	0	293.0	360, 0
6, coil outlet (360 min)	loops	12.6	6.5	3.6	182.7	63.0	0	0	0	23.45	285.0	—
	predictive	7.3	6.6	7.6	86.7	54.8	0	0	0	8.3	285.4	360, 0
9, charge (600 min)	loops	304.7	223.0	179.6	99.0	519.4	45.8	0	0	0	572.7	—
	predictive	35.1	200.7	318.7	52.4	89.4	13.4	0	0	0	571.4	600, 0
8, heavier crude (480 min)	loops	371.7	294.5	193.1	98.2	542.5	16.5	8.9	21.25	19.2	385.9	—
	predictive	769.3	443.3	250.0	81.6	212.2	79.8	26.0	47.45	3.5	385.0	480, 0
	bias every scan[a]	511.1	287.0	167.5	75.7	308.2	57.45	20.2	34.75	3.5	—	480, 0
	reflux gains ×0.72[a]	567.5	332.5	182.3	71.7	201.9	50.7	20.55	34.2	3.5	—	480, 0
	both[a]	361.2	206.8	119.6	63.0	265.9	34.0	11.55	26.85	3.5	—	480, 0
11, challenge, suggested (700 min)	loops	766.1	569.8	399.4	95.9	924.9	91.6	16.05	17.65	1.0	611.0	—
	predictive	820.9	638.7	543.9	125.3	639.4	71.55	21.35	37.2	2.95	610.7	700, 0
11, challenge, shipped[b] (700 min)	predictive	821.2	638.5	539.5	119.5	636.8	71.65	21.35	38.25	3.0	603.2	700, 0

IAE is in °C · min, and in vol% · min for the overflash Φ. Travel sums the four handles' travel tiles. Off-specification minutes are the CSV accumulators named in the text, at their 0.05 min resolution. The draw value prices the kerosene, diesel and gas oil made at 260, 230 and 120 dollars per tonne above residue. [a] Diagnosis copies of the predictive controller, not in the application: the analyzer biases updated every scan instead of through the 6 min filter, the model's reflux coefficients on the four analyzer qualities scaled by 0.72, or both. [b] Control set to MPC and every other setting as shipped. The scored target is met (worst deviations 0.55 vol% and 0.94 °C over the last third, 3.0 min dry), while the furnace fires 93.0 MW at 700 min, above its 85 MW limit for the last 439.4 min, exactly as the loops fire with the same settings. Solve times are given in Section 4.4.5.

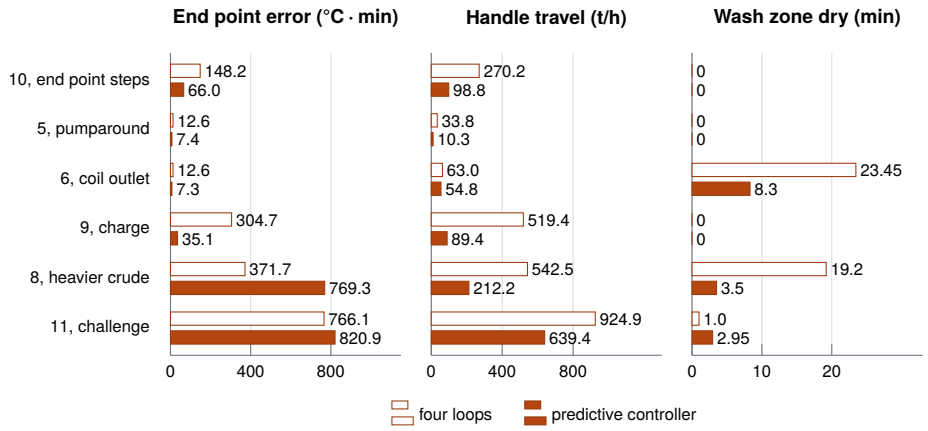


Figure 4.3 Three columns of Table 4.6 for the six cases that both structures run, the challenge with the suggested settings. White bars are the four loops and solid bars the predictive controller. Notice that the plan travels less in every case, wins the end point error everywhere but on the heavier crude and the challenge, and shortens the dry time of the wash zone in scenarios 6 and 8.

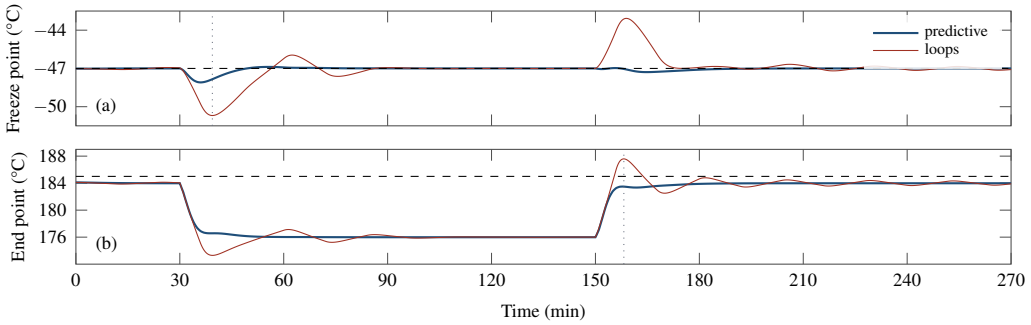


Figure 4.4 Scenario 10 under the predictive controller and under the four loops, with the end point target stepped from 184 to 176 °C at 30 min and back at 150 min. (a) True kerosene freeze point against its −47 °C target. (b) True naphtha end point against the 185 °C specification. The dotted verticals mark 39.4 min, the loops' lowest freeze point, and 158.2 min, the loops' highest end point.

The plan loses scenario 8 on every quality, an end point error of 769.3 against 371.7 °C · min and 79.8 against 16.5 min of naphtha off specification, while keeping the wash zone dry for 3.5 against 19.2 min. The first cause is the model's gain, though not through the crude itself. At the design reflux of 150 t/h the heavier crude makes the end point 1.14 times as sensitive to the reflux as the model says. The climb of the reflux to about 260 t/h that the crude forces flattens the power law of Equation (4.1), to 0.47 of the model's gain there and 0.72 over the climb, so each analyzer result is answered with too little reflux. The second cause is the 6 min bias filter, which adds its lag to the analyzer's. In diagnosis copies that are not in the application the end point error falls to 511.1 with the bias updated every scan, to 567.5 with the reflux gains scaled by 0.72, and to 361.2 with both, against the loops' 371.7, yet the naphtha is still off specification for 34.0 against 16.5 min.

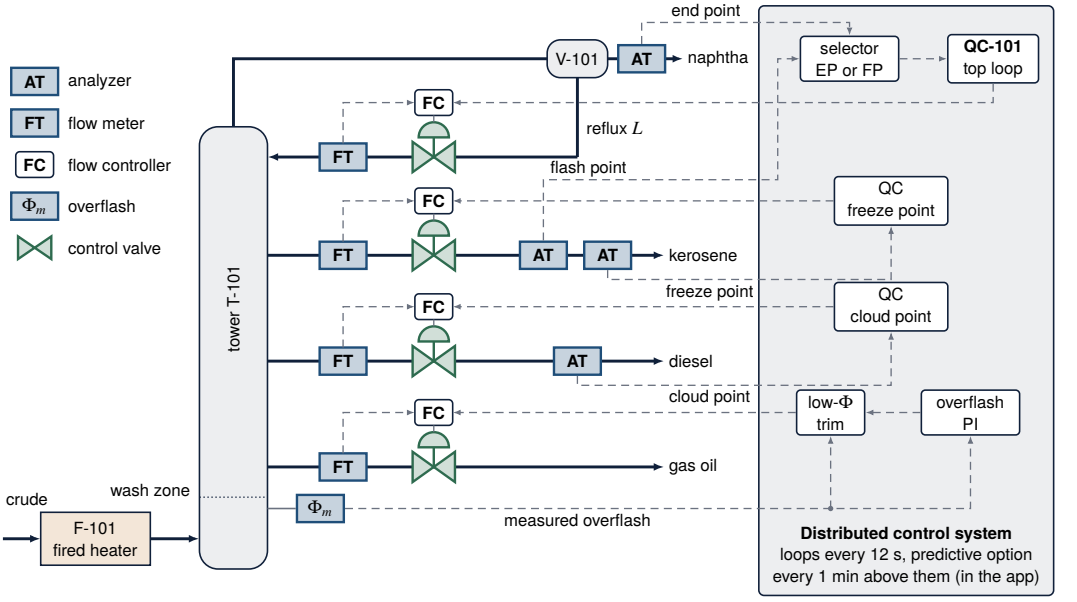


Figure 4.5 The four loops with their hardware. Lettered boxes are sensing devices, valve symbols are final elements and dashed arrows are signals. Notice that every quality controller writes the set point of a flow controller (FC, drawn at its valve) and no valve directly, that the top loop reads two analyzers through a selector, and that the wash zone trim sits between the overflash controller and its flow loop. Not to scale, and not a piping or safety diagram.

> ✦ **DID YOU KNOW? The Controller That Won the Step Lost the Crude**
>
> From the simulator (Table 4.6). On the end point steps of scenario 10 the predictive controller integrates 66.0 °C · min of end point error against the loops' 148.2. On the heavier crude of scenario 8 the same controller integrates 769.3 against 371.7, yet it keeps the wash zone dry for 3.5 against 19.2 min.

For information, the challenge with Control on MPC and everything else as shipped meets its scored target, a worst overflash deviation of 0.55 vol% at 564.0 min, a worst freeze point deviation of 0.94 °C at 533.0 min and 3.0 min dry. In the same run the furnace fires 93.0 MW, over its 85 MW limit for the last 439.4 min, and the naphtha is off specification for 71.65 min, neither of which the target scores. The firing is identical under the loops, because the duty of Equation (4.4) depends only on the coil outlet, the preheat and the crude, never on a handle of either structure. With the suggested settings both fire 81.7 MW and both meet the target, the plan with 0.54, 0.94 and 2.95 against the loops' 0.28, 0.65 and 1.0. The challenge is unchanged. Halving the integration step leaves the winner of every tile unchanged in all six cases, and playback at 1, 20 and 120 times real time is byte identical to the headless run.

4.6 Technology in Practice

Figure 4.5 draws the four loops of Section 4.4 with the hardware a refinery installs in place of the model's ideal instruments.

What Measures It

Four of the five controlled variables are product qualities, read by analyzers that take a sample and report once a cycle. The cycle θ and the lag τ_{an} of Table 4.2 stand for the sample system and the

analysis time of such an instrument. The fifth, the overflash, is read in the model without a cycle, through a 0.8 min lag. A variable that can be determined only at a low sampling rate or by off-line analysis is the original and still dominant use of the soft sensor, a model that predicts it from faster measurements [20].

> **TECHNOLOGY Sensor: An Analyzer That Reports Once a Cycle**
>
> Many industrial analyzers sit away from the process, and tubes, heaters, filters and pumps bring them a steady stream of sample. Problems with process analyzers quite often stem from that sample system [21]. If the line plugs, the analyzer measures stale liquid and has no way of knowing it, so its reading no longer follows the process [21]. The process gas chromatograph shows why such readings come in steps. It injects a small fixed volume of sample, vaporized first if it is a liquid, and separates the species by the time each needs to pass a column. Cycle times of minutes are not uncommon, and some dead time is inevitable [21]. The model's analyzers publish every 12 min through a 2 min averaging lag (Table 4.2).

What Moves It

The top reflux and the three draws are flows, and the model gives each the flow lag τ_v. The coil outlet temperature and the pumparound duties are operator settings in the application, with lags of 5 and 6 min. A guide for refinery energy managers puts the crude that enters the tower at 340 to 390 °C [2], a range that holds the 377 °C of Table 4.1.

> **TECHNOLOGY Final Element: A Flow Loop Under Every Handle**
>
> The handles of this chapter are flows in tonnes per hour, so each is a loop of its own: a meter, a flow controller and a control valve. The quality loop above sends the flow set point, and the flow controller counters a change in supply pressure before the outer loop has to react [21]. Such a cascade requires an inner variable faster than the outer one, and liquid flow answers a valve at once [21]. The most common pressure-based flow element is the orifice plate, whose differential pressure rises with the square of the flow. Such meters are commonly understood not to hold good accuracy much below a third of their full range [21]. The model gives each flow loop one lag, 1.2 min, and an exact meter.

Where the Controller Runs

A refinery runs such loops in a distributed control system, in which a supervisory level sends set points to distributed controllers that act through the actuators from the feedback of the sensors [22]. A survey of industrial practice puts basic control at every second and constraint control at every minute, and builds the conventional constraint layer from PID algorithms, lead-lag blocks and high and low select logic [8]. The four loops of this chapter execute every 0.2 min in the application, and the selector on the top loop and the trim on the gas oil draw belong to that conventional layer. The predictive option plans every 1 min, the period the survey names.

> **IN PRACTICE On a Working Unit, a Temperature Loop Stands Under the Quality**
>
> An account of one industrial crude unit shows what stands between a quality and a valve [11]. A temperature controller at the tower top sits under the predictive controller, which holds the flash point of the diesel product, and keeps the top reflux inside its bounds, by moving that controller's set point. The 95 % point of the diesel, an ASTM D86 specification, is held mainly through the heavy diesel reflux and the furnace outlet temperature. Ten outputs are kept by eight inputs once a minute, most of them inside zones. The main disturbance is the composition of the crude, and the optimizer above acts once a crude analyzer has detected a change and the unit has reached steady state [11].

What Is Modernizing It

The first of three developments is predictive control with an optimizer above it. A survey taken in mid-1999 counted 4542 linear predictive applications, 1985 of them in refining, the largest single

block [8]. On the unit of the In practice box a rigorous steady-state model sets targets for the crude feed rate, the kerosene draw, the furnace outlet temperature and five more variables [11]. The second is the soft sensor. Its weak point is upkeep: after launch its performance usually deteriorates as the process changes, and most soft sensors have no automated maintenance, which leaves the work to people and makes it a significant cost [20]. The third is the faster analyzer. Programming the column temperature of a process chromatograph during the cycle gives the separation of a long column with the dead time of a much shorter one [21].

4.7 Limitations

The largest limitation is the pumparound. The application models the heat recovered and the liquid condensed below the return, and it does not model the vapor traffic lost above it. On a real tower, taking more pumparound duty makes the section above it softer, so the trend in Table 4.5 would reach a point of diminishing return that this model does not show. A real tower also has a hydraulic capacity, and a pumparound circulating rate is a large flow through trays that can flood.

The tower is cut points on an assay rather than trays. Nothing hydraulic can be studied, tray efficiency is absent, and stripping steam is a fixed offset rather than a manipulated variable. The furnace firing is the duty the coil outlet requires rather than a limited actuator, so crossing the 85 MW limit raises an alarm without capping the firing or pulling the coil outlet down.

Product properties are correlations on the sliced assay. A student would not learn here that a kerosene can fail on sulfur or on smoke point, or that a diesel cetane number depends on composition rather than on boiling range alone.

The crude heaviness disturbance shifts and tilts the assay rather than reshaping it, so a real crude switch is more complicated than the one modeled, and none of the corrosion, fouling or catalyst poisoning consequences of a crude change appear.

The regulatory layer implemented here is what a real unit carries underneath a multivariable controller. The application includes one such controller only as an option, whose bands sit on the loops' own targets and whose product values are indicative, so the economic trade between throughput and quality that a real unit makes continuously [11] is still set by the student through those targets rather than by the software.

The predictive controller's model is linear, fitted at one design point and blind to every disturbance, which reaches it only through the analyzer biases, and on the heavier crude of scenario 8 the plant's end point gain at 260 t/h is 0.47 of the model's. The furnace firing is in neither structure, so neither can act on the 85 MW limit.

The analyzer and overflash readings start at the design targets, not at the tower's state, so a setting chosen before Run reaches the readings only over their averaging lag, and the first minutes of such a run show the change late.

4.8 Discussion

Three results generalize past crude towers.

The first is that specifications and controlled variables are different things. On any unit making consecutive fractions from one feed there are more specifications than handles, the surplus ones are consequences, and which are which is an operating decision rather than a control design. The arithmetic is simple and the consequence of skipping it is two loops on one valve, which presents on a panel as a loop that has stopped working.

The second is that the direction of interaction determines the commissioning order, and the familiar rule is stated for the wrong direction. Measured here, bringing the top two loops in first costs 1.7 times as much valve travel as commissioning from the bottom up (6.4 t/h against 3.8 t/h). The general form covers cascades and triangular plants both: bring up the loop that receives the disturbance before the loop that generates it.

The third is that a manipulated variable is often better paired with a constraint than with a product. The overflash moves 0.185 vol% per degree of coil outlet while the kerosene freeze point moves by nothing, so the constraint is the fast and sensitive measurement on that handle as well as the expensive one to violate. Where those two coincide the pairing is obvious, and where they do not the question of what the valve is for has to be answered explicitly.

OPEN CHALLENGE **Optimizing a Tower That Is Seldom at Rest**

The conventional optimizer above a process unit updates a steady-state model from data taken at steady state and then computes new set points [23]. A 2022 review reports the wait for steady state as a fundamental limiting factor of the method, which it finds less used in practice than one would expect [23]. On a crude tower the main disturbance is the crude itself. On one industrial unit, model uncertainty between the layers left some variables oscillating short of the optimum until the middle layer was retuned [11]. Scenario 8 shows the same weakness in small. Progress would be an optimizer that learns from transient data, a family of methods the review surveys [23], proven through a crude switch on a running unit.

4.9 Conclusion

A crude tower is where counting loops stops being enough. The simulated tower of this chapter makes four products from one assay with four handles and eight specifications, and moving the top reflux from 135 t/h to 170 t/h takes the naphtha end point from 191.4 °C to 176.0 °C and the kerosene flash point from 58.2 °C to 48.8 °C together, so the two cannot be held independently. The interaction between the four draws runs one way down the tower, and commissioning along that direction costs 3.8 t/h of valve travel against 6.4 t/h with the top two loops brought in first, and once every loop is in service the peak deviations on a crude change agree within 0.3 percent. The gas oil draw is paired with the overflash rather than with a product because eight degrees off the coil outlet takes the overflash from 3.28 vol% to 1.80 vol% with the draws held while the kerosene freeze point does not move, and because a dry wash zone costs a run rather than a giveaway. Ten and a half megawatts of furnace fuel follow from fourteen megawatts of top pumparound duty, and the same duty narrows the kerosene and diesel tail from 3.51 °C to 2.00 °C. An optional predictive controller that carries the triangle in its model integrates 18.1 against 123.9 °C · min of freeze point error on an end point step with 37 percent of the loops' travel, and loses the heavier crude. A fuller model would add a vapor traffic model above each pumparound return so that the heat recovery sweep reaches its own limit, a hydraulic capacity so that flooding can end a push, and a second crude whose assay changes shape rather than position. As it stands, the simulated tower is a place to count degrees of freedom before pairing loops.

KEY POINTS

- A tower with four handles and eight specifications can hold four of them. The other four are consequences.
- The naphtha end point and the kerosene flash point are one cut point and cannot be held independently.
- The interaction runs one way, down the tower. Bring up the loop that receives the disturbance before the loop that generates it.
- The gas oil draw is paired with the overflash: the constraint binds long before any product notices.
- A predictive controller with the triangle in its model wins the end point step and loses the heavier crude.

Where else this applies. A triangular gain matrix returns in Volume 6, Chapter 2, a wash bed protected by an override in Chapter 5, and the choice of which flow a structure holds in Volume 6, Chapter 7.

References

[1] 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.

[2] 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.

[3] U.S. Energy Information Administration, "Refinery capacity report 2026, table 1: Number and capacity of operable petroleum refineries by PAD district and state as of January 1, 2026," U.S. Department of Energy, Washington, DC, Tech. Rep., 2026.

[4] W. L. Luyben, "Control of a distillation column with side stripper and side rectifier," *Chemical Engineering Research and Design*, vol. 161, pp. 38–44, 2020.

[5] C. Pasquini and A. F. Bueno, "Characterization of petroleum using near-infrared spectroscopy: Quantitative modeling for the true boiling point curve and specific gravity," *Fuel*, vol. 86, no. 12-13, pp. 1927–1934, 2007.

[6] R. N. Santos, E. R. Lima, and M. L. Paredes, "Astm d86 distillation curve: Experimental analysis and premises for literature modeling," *Fuel*, vol. 284, p. 118958, 2021.

[7] J. Kewley, "Fractional distillation as applied in the petroleum industry," in *Distillation Principles and Processes*, S. Young, Ed. London: Macmillan, 1922, pp. 319–360, chapters XXXI and XXXII, pp. 334–351.

[8] 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.

[9] C.-B. Chung and J. B. Riggs, "Dynamic simulation and nonlinear-model-based product quality control of a crude tower," *AIChE Journal*, vol. 41, no. 1, pp. 122–134, 1995.

[10] A. S. Yamashita, A. C. Zanin, and D. Odloak, "Tuning the model predictive control of a crude distillation unit," *ISA Transactions*, vol. 60, pp. 178–190, 2016.

[11] P. A. Martin, A. C. Zanin, and D. Odloak, "Integrating real time optimization and model predictive control of a crude distillation unit," *Brazilian Journal of Chemical Engineering*, vol. 36, no. 3, pp. 1205–1222, 2019.

[12] F. Li, F. Qian, M. Yang, W. Du, and V. Mahalec, "Product tri-section based crude distillation unit model for refinery production planning and refinery optimization," *AIChE Journal*, vol. 67, no. 2, 2020.

[13] Y. Yang and P. I. Barton, "Refinery optimization integrated with a nonlinear crude distillation unit model," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 205–210, 2015.

[14] Y. Xiong, X. Shi, Y. Ma, and Y. Chen, "Optimization design of crude oil distillation unit using bi-level surrogate model," *Frontiers in Control Engineering*, vol. 4, 2023.

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

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

[17] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

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

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem – a review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

4.10 Exercises

4.1 Concept check · 2 min

An engineer adds a second PI loop on the top reflux of this tower, so that the reflux is asked to hold the naphtha end point at 184 °C and the kerosene flash point at 50 °C at the same time. What should the engineer expect?

(a) Both targets are met, because the two loops read different analyzers on separate cycles

(b) Both targets are met after retuning, because a decoupler separates two outputs of one handle

(c) The flash point loop wins, because the flash point moves more than the end point per tonne of reflux

(d) At most one target is met, because both are read from the one cut point that the reflux places

4.2 Concept check · 1 min

With every other draw and the reflux held at the design line-up, the diesel draw is raised by 20 t/h. Which cut points have moved once the tower settles?

(a) The diesel and gas oil cut and the gas oil and residue cut, and none above them

(b) All four cuts, the naphtha and kerosene cut least because it is farthest away

(c) Only the diesel and gas oil cut, because the gas oil draw itself has not changed

(d) The kerosene and diesel cut and the diesel and gas oil cut, on either side of the diesel

4.3 Concept check · 2 min

In scenario 6 the coil outlet falls 22 °C with all four loops in automatic, and the measured overflash drops below its 2 vol% minimum. Which statement describes the override on the gas oil draw?

(a) It replaces the gas oil loop and holds the overflash at its minimum until the coil outlet returns

(b) It scales down the loop's request more as the shortfall grows, which shortens the time below the minimum without removing it

(c) It writes the gas oil loop's integral, so that the loop does not wind up while the draw is cut

(d) It acts only with the gas oil loop in automatic, so a draw held in manual has no protection

4.4 Theory and calculation · 9 min

The top cut follows Equation (4.1) with $v_{1,ref} = 0.200$, $L_{ref} = 150\,\mathrm{t/h}$, $p = 0.45$ and $h = 0$, and the tail follows Equation (4.3) with $\beta = 12.0\,^{\circ}\mathrm{C}$ and $n = 1$, where $R_1 = L/D_N$ and the naphtha made is $D_N = v_1 \rho F$ with $\rho F = 569.3\,\mathrm{t/h}$. Between $v = 0.15$ and $v = 0.25$ the assay is the straight line through 145 °C at 0.15, 175 °C at 0.20 and 205 °C at 0.25. Take steady state, so $T_1 = T_{bp}(v_1)$.

(a) Find v_1, T_1, σ_1, the naphtha end point and the kerosene flash point at $L = 150\,\mathrm{t/h}$. *Checkpoint:* 0.200, 175.0 °C, 9.11 °C, 184.1 °C and 53.9 °C.

(b) Repeat at $L = 170\,\mathrm{t/h}$, and find the ratio of the flash point's change to the end point's change. *Checkpoint:* 0.1890, 168.4 °C, 7.60 °C, 176.0 °C, 48.8 °C and 0.63.

(c) Differentiate at $L = 150\,\mathrm{t/h}$ to find dT_1/dL and $d\sigma_1/dL$. Use them to find the gains of the end point and the flash point, and compare with a and b of Table 4.3. *Checkpoint:* -0.360 and $-0.0881\,^{\circ}\mathrm{C/t/h}$; gains -0.448 and $-0.272\,^{\circ}\mathrm{C/t/h}$.

4.5 Theory and calculation · 8 min

Use the design line-up: charge $F = 662\,\mathrm{m^3/h}$ with $\rho F = 569.3\,\mathrm{t/h}$, $v_1 = 0.200$, draws of 74, 102 and 51 t/h, coil outlet 377 °C, pumparound duties summing to 42 MW and the parameters of Table 4.2. The assay passes through 333 °C at 0.50, 358 °C at 0.55, 383 °C at 0.60 and 410 °C at 0.65, linearly between them.

(a) Find the diesel and gas oil cut point, and how far it moves when the kerosene draw rises by 20 t/h. *Checkpoint:* 337.6 °C, moving 17.6 °C.

(b) Find the fraction vaporized, the overflash, the overflash's sensitivity to the coil outlet and the coil outlet drop, with every draw held, that brings the overflash to 2.0 vol%. *Checkpoint:* 0.6315, 3.28 vol%, 0.185 vol% per °C and 6.9 °C.

(c) Find the rise in preheat and the firing saved per megawatt of pumparound duty, and apply them to the scenario 5 step from 12 to 26 MW. *Checkpoint:* 1.82 °C and 0.75 MW per megawatt; 25.5 °C and 10.5 MW.

4.6 Theory and calculation · 5 min

The loops' plant gains at the design point are −0.448 °C/t/h (reflux to end point), 0.692 °C/t/h (kerosene draw to freeze point), 0.411 °C/t/h (diesel draw to cloud point) and −0.176 vol%/t/h (gas oil draw to overflash), with section time constants of 4, 8, 12 and 15 min.

(a) Find the controller gain $1/|K|$ for each loop and compare with the gains the automatic scenarios load. *Checkpoint:* 2.23, 1.45, 2.43 and 5.68, against 2.5, 1.4, 2.4 and 5.5.

(b) The measured overflash reads 1.9 vol% while the gas oil loop requests 40 t/h. Find the draw the override sends, and the measured overflash below which the cut stops growing. *Checkpoint:* 37.3 t/h; 1.1 vol%.

4.7 App challenge · 8 min

Load scenario 2, One valve, two specifications, and change no setting. The naphtha end point target steps from 184 to 176 °C at 30 min, and at 150 min the selector switches the top loop to the kerosene flash point with a 54 °C target. Run to 360 min at 20 min per wall second (about 18 s) and export the CSV file. Read the columns `naph_ep_C`, `kero_flash_C` and `reflux_tph`.

(a) Read the three at 30 and at 150 min, and find the ratio of the flash point's change to the end point's change. *Checkpoint:* 184.0, 53.9, 150.3 t/h; 176.0, 48.8, 170.1 t/h; ratio 0.63.

(b) After the switch, find the highest end point and its time, and say how long the end point stays above its 185 °C specification. *Checkpoint:* 188.0 °C at 170.4 min, above the specification from 165.6 to 184.8 min.

(c) Read the three at 360 min. *Checkpoint:* 184.3 °C, 54.0 °C and 149.6 t/h.

Accept temperatures within 0.1 °C, the reflux within 0.1 t/h and times within 0.2 min.

4.8 App challenge · 8 min

Load scenario 6, The loop that is not on a product, and change no setting. The coil outlet steps from 377 to 355 °C at 30 min and back at 200 min. Run to 360 min at 20 min per wall second and export the CSV file.

(a) From the columns `overflash_pct` and `ago_override`, find the first row below 2 vol%, the first row with the override on and the last row with it on. *Checkpoint:* 34.2, 35.1 and 58.4 min.

(b) Find the lowest overflash and its time, and read the Wash zone dry tile. *Checkpoint:* 1.887 vol% at 42.2 min; 23 min.

(c) Read `ago_tph` at 200 and at 360 min, and find the largest distance of `kero_freeze_C` from its −47 °C target. *Checkpoint:* 27.5 t/h, 50.4 t/h and 0.09 °C.

Accept times within 0.1 min and other values within one unit in the last digit given.

4.9 App challenge · 9 min

Load scenario 7, Commissioning order, and why the usual rule is backwards here, and keep Control on Loops. The script switches the gas oil, diesel, kerosene and top loops to automatic at 10, 30, 50 and 70 min, and a heavier crude arrives at 120 min.

(a) Run as shipped, at 1 min per wall second from 115 min, and pause between 119.5 and 120.2 min. Add the four MV travel tiles. *Checkpoint:* 3.79 to 3.82 t/h, 3.81 t/h at 120 min.

(b) Reload, run at 1 min per wall second, and switch the Top loop to Auto as the clock reads 10 min and the Kerosene loop as it reads 30 min. Pause as in (a) and add the tiles again. *Checkpoint:* 6.39 t/h with the top loop switched at 10.0 min, about 5.97 to 6.57 t/h for a switch between 9.5 and 10.5 min.

(c) Run both cases on to 180 min, export each file, and find after 120 min the largest distance of `overflash_pct` from 3.3 vol% and of `kero_freeze_C` from –47 °C. *Checkpoint:* 4.99 and 12.02 °C bottom-up, 5.01 and 12.05 °C with the top two loops early.

4.10 Application area · 3 min

In scenario 5 more top pumparound duty only helps: the preheat rises, the firing falls and the kerosene and diesel split sharpens.

(a) Name the effects that bound pumparound duty on a real tower and that this model leaves out. *Checkpoint:* The softer section above the return, flooding and the exchanger network.

(b) The step from 12 to 26 MW saves 10.5 MW of firing. Find the fired duty saved in a year of continuous operation. *Checkpoint:* About 92 000 MWh.

4.11 Application area · 3 min

A planner proposes to use the application to decide whether a new crude cargo will make jet fuel and diesel on specification.

(a) Give two reasons from the model why the application cannot answer this. *Checkpoint:* The heaviness only shifts and tilts the assay, and the properties are boiling-range correlations.

(b) Name two specifications a real kerosene or diesel can fail that the application does not compute. *Checkpoint:* Sulfur or smoke point for kerosene, cetane number for diesel.

4.12 Review the agent · 5 min

An AI assistant is asked how to pass the graded challenge with the least effort and replies as follows. “Set Control to MPC and change nothing else. The scored target is met, with a worst overflash deviation of 0.55 vol%, a worst freeze point deviation of 0.94 °C and 3.0 min dry. The predictive controller moves all four handles together, so it also keeps the furnace inside its 85 MW limit, and the heat moves of the suggested configuration are no longer needed.” Find the flaw and say how to check it.