INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 6

Fluid Catalytic Cracking Simulator for Internal Heat Balance Feedback and Selective Control

Pages 127–160

Edited by John D. Hedengren

How to cite this chapter. "Fluid Catalytic Cracking Simulator for Internal Heat Balance Feedback and Selective Control," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 127–160.

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

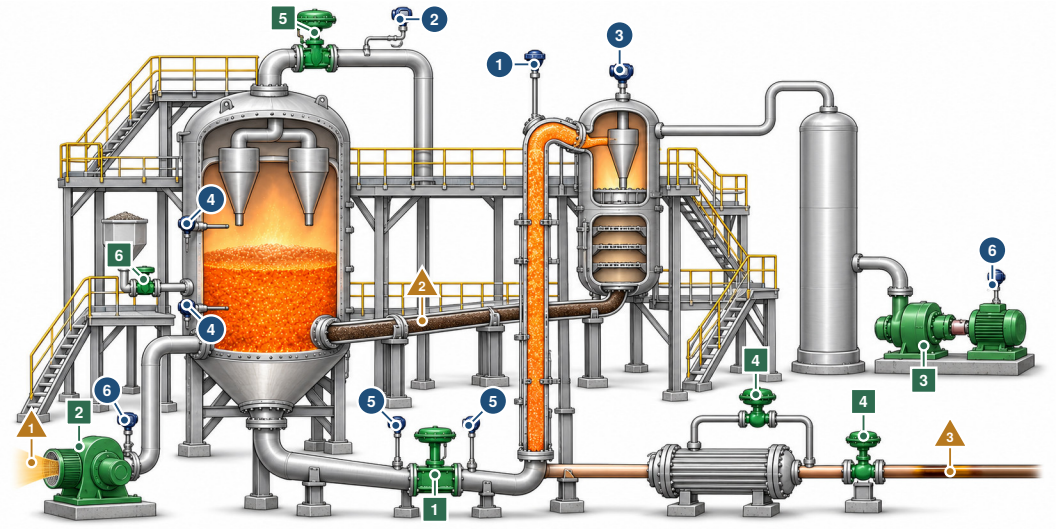
The app for this chapter
https://apmonitor.com/apps/fcc-unit/

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

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

CHAPTER 6

Fluid Catalytic Cracking Simulator for Internal Heat Balance Feedback and Selective Control



The controller moves 1 catalyst slide valve 2 air blower 3 wet gas compressor 4 feed rate and feed preheat 5 flue gas valve 6 fresh catalyst addition

The controller reads 1 riser outlet temperature 2 oxygen in the flue gas 3 reactor pressure 4 bed and plenum temperatures 5 pressure drop across the slide valve 6 air flow and compressor speed

What disturbs it 1 hot afternoon air 2 catalyst activity 3 a heavier feed

The orange bed of the regenerator, where the coke burns off, is the only source of heat for the riser at the center, which the hot catalyst reaches through the slide valve at the bottom. Illustration, not to scale.

IN PLAIN TERMS **A Cracker That Is Its Own Furnace**

The main idea. A cat cracker breaks heavy oil into gasoline on a hot powder called catalyst. The powder picks up a carbon crust called coke, and burning the coke off is the only heat the plant has.

The challenge. Hotter cracking makes more coke and more coke makes more heat, so one small valve move about doubles in effect over ninety minutes. The plant also runs against the limits of its machines, and which one gives out first changes with the weather and the oil.

Why it matters. Running by one fixed number leaves fuel unmade on good days and overheats the plant on bad ones.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Sections 6.2 and 6.5, the Did you know? boxes in Sections 6.1 and 6.6, the Try it boxes in Sections 6.3 and 6.4, and Key points in Section 6.10. No equations.
Study	65 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	15 min	This page, Section 6.6 with Figure 6.5, Section 6.7, Section 6.8 and the Open challenge box in Section 6.9.

The control question A cat cracker has no furnace: its only heat is the coke it burns. How much does that internal heat balance amplify and slow every move, and how should one feed rate be shared among equipment limits that move with the weather, the feed and the severity?

Learning objectives

- Measure the heat balance loop gain from one slide valve step, and explain why the first two minutes understate the process gain.
- Explain from the coke balance why more air lowers the carbon on the catalyst but not the settled burn rate, and why afterburn peaks between partial and full burn.
- Compute the blower capacity, the catalyst circulation and the slide valve differential from the model, and the feed that a moving limit costs.
- Explain how a low selector with external reset shares the feed among four constraint controllers, and why the regenerator bed needs an override on the feed preheat instead.
- Compare the selector structure with a predictive controller on time above limits, riser outlet error, feed and margin.

Prerequisites Sections 1.4.4 (FCC heat balance and loop gain), 1.4.7 (selectors and external reset), 1.6.2 (PI loops), 1.6.3 (predictive control) and 1.6.4 (cases), and the throughput override of Chapter 3.

The app https://apmonitor.com/apps/fcc-unit/

See also Volume 4, Chapter 3 for a selector whose integrators track the applied command, Volume 4, Chapter 4 and Volume 4, Chapter 7 for a feed rate pushed against limits that change. Volume 6, Chapter 7 for amplification by a recycle.

6.1 Introduction

Opening one valve by three percent on the unit described here moves the riser outlet temperature 2.11 °C in the first two minutes and 7.17 °C in the end, and the end arrives ninety minutes later. Nothing else was touched. The regenerator dense bed, which no controller and no operator asked to move, settles 5.25 °C higher than it started. An engineer who reads the first two minutes of that trend and fits a model to it has measured a plant 3.4 times weaker and very much faster than the one underneath, and a controller tuned on that reading meets the real gain an hour later.

The cause is that the unit is its own furnace. Vaporized gas oil meets hot catalyst at the bottom of a riser and cracks for three seconds on the way up, the reaction is endothermic, and the catalyst is the only source of heat. The reaction leaves coke on the catalyst as it goes. The coked catalyst is stripped and dropped into a regenerator where the coke is burned with air, and that combustion is the only thing that reheats the catalyst. Extra catalyst therefore makes the riser hotter, a hotter riser cracks harder, harder cracking makes more coke, and the coke is the fuel. The fraction of any disturbance that the plant returns to itself is 0.52 on this model, so by Equation (1.8) the amplification is $1/(1-0.52) = 2.09$ and the settling time belongs to the regenerator rather than to the riser.

DID YOU KNOW? Most of the World's Gasoline

Fluid catalytic cracking produces the majority of the world's gasoline, and at the start of 2014 it ran in over 300 of the world's 646 refineries [1]. About 2300 tonnes of its catalyst are made every day [1]. The unit simulated in this chapter charges 250 t/h of gas oil.

The class of plant is larger than refining. Any unit that supplies its own reaction heat from a stream it produced closes the same path, and the family includes autothermal reactors, regenerative burners, feed and effluent exchangers around exothermic beds such as the one of Chapter 8, and every hot recycle, whose material counterpart is the snowball of Volume 6, Chapter 7. What such a plant shares is a loop gain below unity that inflates the steady-state gain, stretches the settling time and climbs with severity, and the same arithmetic produces multiple steady states once the gain approaches one, which has been documented for this unit specifically [2, 3]. A second general class sits on top

of the first. Severity and throughput both earn money on a cat cracker, the economic optimum lies outside the feasible region, and which of several machines runs out first changes with the weather, the feed and the catalyst. A low selector over one shared throughput handle is the classical answer [4] (Section 1.4.7), deciding which loop should own that handle inside each operating region has a documented design procedure [5], and raising a plant's rate until its slowest machine says stop is what a coordinating layer above the units is built to do [6].

However, a transfer function fitted to a step test carries neither behavior. The internal feedback is invisible because it is already inside the identified gain, and the identified gain is the one the test was long enough to see. The constraint set is invisible because a linear model has no limits in it, and the limit that matters today is not the one that mattered yesterday. This chapter works with a browser-based riser and regenerator described completely enough for a reader to rebuild it, and takes five measurements on it. They are the heat balance loop gain and the two very different process gains one loop shows at two minutes and at steady state, the gap between what more air does to the burn rate at a frozen carbon inventory and what it does once the unit settles, the afterburn peak that sits between partial burn and full burn so that the road passes through it in both directions, and what a low selector with external reset over four constraint controllers, with an override on the fifth constraint the shared handle cannot reach, is worth against a fixed throughput, and what a predictive controller on the same four handles costs against it. The measured results are that the heat balance amplification is 2.09 with the steady main loop gain 2.6 °C per percent of valve against 0.74 °C at two minutes, that the settled coke burn rate moves 0.3 % for a 9.3 % increase in air while the carbon on the regenerated catalyst falls by a third, and that over the 420 min graded shift the selective layer earns $95,391 with no time above the cyclone limit against $84,532 and 296 min above it for a fixed line-up. The plant has fifteen states, a closed-form coke burn, a closed-form slide valve and a damped Newton steady-state solver that has to be kept off a second, cold solution branch. The four-lump kinetics, the well-mixed regenerator and the selector arithmetic are all standard.

Section 6.2 traces how a cracking unit that burns its own coke came to be controlled the way it is. Section 6.3 builds the riser, stripper, regenerator, pressure and slide valve balances, puts the five limits on one scale and lists every constant. Section 6.4 measures what the heat balance does: the loop gain, the burn rate against the air, the afterburn peak and the one limit the feed rate cannot reach. Section 6.5 describes the three regulatory loops, the low selector with external reset, the override on the feed preheat and the predictive controller. Section 6.6 runs the fourteen scenarios through severity, constraint handover, the two failures that hold every setpoint, the graded challenge and the comparison with the predictive controller. Section 6.7 sets the instruments and final elements of a real unit beside the model. Sections 6.8 to 6.10 give the limitations, the economics of the constraint layer and a summary, and the exercises follow.

6.2 How the Cat Cracker Came To Be Controlled

Catalytic cracking began as a fixed-bed process in which the catalyst was regenerated in place, and it moved to a circulating solid in the 1940s, a step the History box below dates, precisely so that the coke could be burned off continuously in a second vessel while the first vessel kept cracking [1, 7]. The circulating solid also carried the combustion heat back, and that is where the control problem was created rather than solved. A design that recovers reaction heat from the product of the reaction has no independent heat input left, so every operating variable acts through the same energy balance and the reactor temperature became a quantity that is arrived at rather than set. Zeolite catalyst in the 1960s made the reaction fast enough that essentially all of the cracking happens in the transfer line, which shortened the reaction side to seconds and left the regenerator holding the whole of the plant's inertia [1, 7].

> ⧖ HISTORY **1942: The Catalyst Starts to Circulate**
>
> The first catalytic crackers held their catalyst still. The Houdry process, started up in 1936 in a New Jersey plant that processed 15 000 barrels a day, used a fixed bed reactor [1]. In 1938 a consortium of oil and engineering companies set out to develop a new cracking process. It built on a suggestion of Lewis and Gilliland, that a slow flow of gas through a powder might lift it enough to make it flow like a liquid. A pilot plant of 100 barrels a day at Baton Rouge, Louisiana, moved powdered catalyst through a reactor and a regenerator, and the first commercial unit started up in the middle of 1942. By the end of the war 34 units were running in the United States. The second Baton Rouge unit, started in June 1943, was still in operation when the review cited here appeared in 2015 [1].

The consequences of the internal energy balance were mapped much later than the technology. The observation that a cat cracker has more than one steady state at the same set of external conditions, a hot well-regenerated branch and a cold coked-up branch, was published as an operating fact with the underlying heat balance arithmetic behind it [2], and the full picture of where the branches lie, which are stable and how the boundary moves with feed quality was worked out in a four-part study of the current generation of units [3, 8]. That study also introduced the argument that a unit with far more measurements than manipulated variables should be controlled partially, holding a small number of dominant variables and letting the rest go where the balances put them [9]. The kinetic side developed in parallel and converged on lumping. Cracking selectivity was reduced to a small number of pseudo-components with second-order gas oil cracking and first-order overcracking of the gasoline lump [10], and the four-lump form that adds coke as a separate product and gives the gasoline yield a maximum against severity is the version a heat balance needs [11]. Combustion in the regenerator acquired its own literature once units began running lean, because carbon monoxide that leaves the dense bed finds oxygen in the freeboard where there is no catalyst to absorb the heat, and the freeboard was analyzed as an ignition problem in its own right [12].

Control practice moved in three steps. The first was regulatory, and it is still underneath everything else. Riser outlet temperature is held on the regenerated catalyst slide valve, flue gas oxygen on the air blower, and reactor pressure on the wet gas compressor, with a fourth loop holding the regenerator pressure on the flue gas slide valve. Which variables to pair, and which of the many available measurements are worth closing at all, was posed as a control structure selection problem and answered for the FCC specifically [9, 13]. The second step was a public dynamic model. A full simulator of a Model IV unit, with the catalyst circulation set by pressure balance rather than by a valve, was published so that control methods could be compared on the same plant [14], and it became the reference case for a generation of FCC control studies. The third step was multivariable constraint control. An industrial reactor and regenerator controller reported in 1993 pushed severity and throughput against the machine limits rather than tracking setpoints, and reported the payout in exactly those terms [15]. The regulatory structures that preceded it did not disappear, because a low selector and an override need one function block and no model, and they remain the layer that acts when the multivariable controller is off [4, 5]. The question of which fence a unit is allowed to stand against, and how much throughput the answer is worth, is the same question a coordinating layer asks [6].

6.3 Process Model

Regenerator R-101 and reactor R-102 exchange catalyst around one loop. 250 t/h of gas oil preheated to 220 °C meets 1700 t/h of catalyst at the bottom of the riser, a catalyst-to-oil ratio of 6.80, and leaves the riser at 525 °C with 75.00 % of the feed converted by weight into 46.00 % gasoline, 23.58 % light gas and LPG and 5.42 % coke. The catalyst delivers 95.0 MW to the riser, of which 72.7 MW heats and vaporizes the feed, 18.4 MW cracks it, 2.4 MW goes to the steam and 1.5 MW leaves the shell. The stripper misses a further 0.142 kg/s of hydrocarbon, so 3.905 kg/s of coke at 7.25 % hydrogen reaches the regenerator and releases 130.1 MW burning with 204 t/h of air, which holds the dense

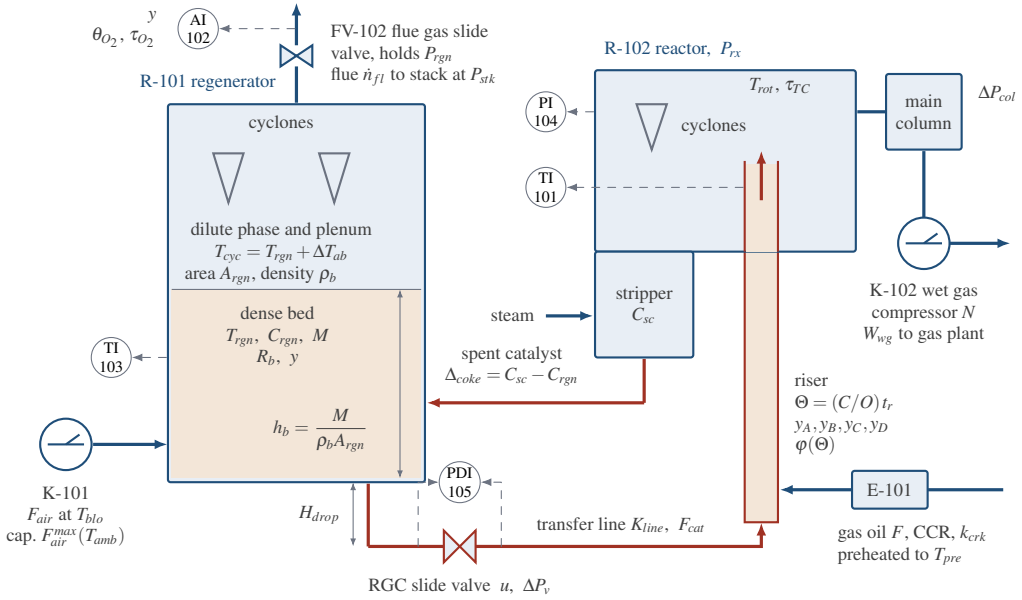


Figure 6.1 Regenerator R-101 and reactor R-102 with one loop of catalyst between them. Gas oil F preheated to T_{pre} in E-101 meets regenerated catalyst F_{cat} at the foot of the riser and cracks along catalyst contact Θ into the four lumps y_A to y_D at activity φ, leaving at T_{rot}. Coked catalyst at C_{sc} is stripped and drops into the dense bed, where air F_{air} from blower K-101 burns it at rate R_b leaving oxygen mole fraction y, so the bed holds M kilograms of catalyst at T_{rgn} with C_{rgn} carbon left on it. Carbon monoxide that escapes the bed burns above it and raises the plenum to T_{cyc}. Circulation returns down a standpipe of elevation H_{drop} through the regenerated catalyst slide valve at opening u and differential ΔP_v, and the transfer line takes K_{line} of differential as the square of the flow. FV-102 holds the regenerator pressure and K-102 holds the reactor pressure P_{rx}, against which ΔP_v is measured.

bed at 700 °C with 2.00 % oxygen by volume left in the flue gas. Carbon on the regenerated catalyst is 0.050 % by weight, delta coke is 0.827 %, the slide valve is 45.0 % open with 57.0 kPa across it, and the gross margin is $15,348 per hour. The vessels, the catalyst loop, the instrument tags and every symbol used below are drawn in Figure 6.1.

Table 6.1 lists everything the control system touches. The clock is in minutes, flows are in kg/s and displayed in t/h with the factor of 3.6 written out, duties are in kW, temperatures in °C, pressures in kPa absolute, inventories in kg and coke as a mass fraction of catalyst. Every derivative below therefore carries a factor of 60 where a per-second quantity enters a per-minute rate. Fifteen quantities are integrated. Ten of them are the plant and five are first-order actuator positions.

6.3.1 The Riser, Which Puts Itself Out

Cracking is marched along the riser against catalyst contact Θ, the product of the catalyst-to-oil ratio and the vapor residence time, because that product and not the length is what an operator changes. Four lumps are carried, gas oil y_A, gasoline y_B, light gas and LPG y_C, and coke y_D, with second-order

Table 6.1 Classification of Variables. Nominal Values Are Those of the Solved Design Point at Which the Basic Scenarios Open.

Symbol	Role	Nominal	Range
T_{rot} riser outlet	CV, measured	525.0	495–555 °C SP
y flue gas oxygen	CV, measured	2.00	0.6–5.0 vol % SP
P_{rx} reactor	CV, measured	310	250–390 kPa a SP
U nearest fence	CV, computed	26.1	60–100 % SP
U_{rgn} bed approach	CV, computed	−0.1	60–100 % SP
u RGC slide valve	MV	45.0	0–100 %
F_{air} air blower	MV	204	100–250 t/h
N compressor speed	MV	85.0	55–100 %
F feed rate	MV, selector	250	150–340 t/h
T_{pre} feed preheat	MV, override	220	175–300 °C
catalyst addition	MV, operator	2.0	0.5–8 kg/min
CCR carbon residue	DV, measured	0.60	0–1.8 wt %
k_{crk} crackability	DV, unmeasured	1.00	0.75–1.15
T_{amb} ambient	DV, measured	15	−5–42 °C
γ CO promoter	DV, unmeasured	1.00	0.02–1.50
stripper condition	DV, unmeasured	1.00	0.40–1.20
a_{er} standpipe aeration	DV, unmeasured	1.00	0.60–1.10
feed metals	DV, unmeasured	1.00	0.2–6.0
catalyst losses	DV, measured	2.0	0–30 kg/min
T_{rot} (°C)	state	525.0	300–700
T_{rgn} (°C)	state	700.0	400–1000
C_{rgn} (wt fraction)	state	0.000500	10^{-7}–0.06
T_{cyc} (°C)	state	703.6	400–1100
y duct oxygen (vol %)	state	2.00	0–21
P_{rx}, P_{rgn} (kPa a)	states	310, 350	130–600, 130–650
M inventory (kg)	state	165000	40000–320000
u_{fl} flue gas valve (%)	state	45.6	5–95
a catalyst activity	state	1.000	0.10–1.60
five actuator positions	states	as above	as above

cracking of the gas oil and first-order overcracking of the gasoline, as shown in Equation (6.1).

$$\begin{aligned}
\frac{dy_A}{d\Theta} &= -\varphi\,(a_1 + a_2 + a_3)\,y_A^2 \\
\frac{dy_B}{d\Theta} &= \varphi\,\left[a_1 y_A^2 - (a_4 + a_5)\,y_B\right] \\
\frac{dy_C}{d\Theta} &= \varphi\,\left(a_2 y_A^2 + a_4 y_B\right) \\
\frac{dy_D}{d\Theta} &= \varphi\,\left(a_3 y_A^2 + a_5 y_B\right)
\end{aligned} \tag{6.1}$$

where $a_i = k_i\,\kappa_i \exp\left[-(E_i/R)\left(T_{rot}^{-1} - T_{rot,0}^{-1}\right)\right]$ are the five rate coefficients evaluated at the riser outlet temperature in kelvin against the kinetic reference $T_{rot,0}$, κ_i is the feed crackability multiplier and applies to a_1 and a_2 only, and φ is the catalyst activity. All four lumps are mass fractions of the feed and sum to unity, so both sides of (6.1) are dimensionless per unit of contact.

The activity term is the part worth reading. Activity at any point in the riser is set by the coke that is already on the catalyst at that point, which is what came around from the regenerator plus what

this same reaction has laid down since the feed nozzles, as shown in Equation (6.2).

$$\varphi(\Theta) = a \exp\left[-\alpha \cdot 100 \left(C_{rgn} + \frac{y_D(\Theta)}{C/O}\right)\right] \qquad C/O = F_{cat}/F \tag{6.2}$$

where a is the equilibrium catalyst activity, α is the activity lost per weight percent of coke on catalyst and the factor of 100 converts the mass fraction to weight percent. The riser therefore extinguishes itself as it travels, with φ falling from 0.928 at the feed nozzles to 0.281 at the outlet at the design point, and carbon on the regenerated catalyst is charged twice, once directly in (6.2) and again through every kilogram of coke that the lost activity failed to prevent.

Contact is set by the circulation and by a vapor residence time that shortens as the vapor rate rises and lengthens with pressure, as shown in Equation (6.3).

$$\Theta = \frac{F_{cat}}{F} t_r, \qquad t_r = \operatorname{sat}\left(t_{r,0}\sqrt{\frac{F_0}{F} \cdot \frac{P_{rx}}{P_{rx,0}}},\ 1.0,\ 9.0\right) \tag{6.3}$$

where $t_{r,0}$ is the design residence time in seconds, F_0 is the design feed rate and the saturation bounds are in seconds. The exponent of one half rather than one is deliberate. Pure ideal gas scaling gives one, but the moles expand as the feed cracks and the catalyst slip velocity rises with gas density, and a full exponent puts a positive feedback path through the reactor pressure that walks the unit off its heat balance in manual operation.

Carbon residue in the feed goes to coke essentially quantitatively and is taken out of the unconverted bottoms rather than out of nothing, which is the standard additive coke treatment, as shown in Equation (6.4).

$$c_{add} = \min(0.9\, y_A,\ k_{CCR}\,\mathrm{CCR}) \qquad y_A \leftarrow y_A - c_{add}, \qquad y_D \leftarrow y_D + c_{add} \tag{6.4}$$

where CCR is the feed carbon residue in weight percent and k_{CCR} converts it to a coke mass fraction. Equation (6.1) is marched with the explicit midpoint rule in $n_r = 14$ equal steps of Θ, which was chosen because explicit Euler needed forty steps before the coke yield stopped moving in the third figure and the operating map evaluates roughly a thousand of these per second.

6.3.2 The Riser Energy Balance

The riser and reactor hold one lumped thermal capacity, and the catalyst is the only heat source in it, as shown in Equations (6.5) and (6.6).

$$Q_{dem} = F\left[c_{p,oil}(T_{rot} - T_{pre}) + \Delta H_{vap}\right] + F x \Delta H_{rx} + F_{st} c_{p,st}(T_{rot} - T_{st}) + Q_{loss,rx} \tag{6.5}$$

$$C_{rx}\frac{\mathrm{d}T_{rot}}{\mathrm{d}t} = 60\left[F_{cat} c_{p,cat}(T_{rgn} - T_{rot}) - Q_{dem}\right] \tag{6.6}$$

where $x = 1 - y_A$ is the conversion, $F_{st} = f_{st}F + F_{strip}$ is the dispersion plus stripping steam, $C_{rx} = M_{strip}c_{p,cat} + M_{shell,rx}c_{p,steel}$ is the combined catalyst and steel capacity in kJ/°C, and the factor of 60 turns kW into kJ/min. There is no other term in (6.6). Lowering the preheat grows the first bracket, and with the riser outlet held the circulation has to grow with it, which is the whole of the preheat handle.

6.3.3 The Stripper, and Why Throughput Reaches the Regenerator

Stripping steam is a fixed rate, so the steam-to-catalyst ratio falls as the circulation rises and entrained hydrocarbon rides into the regenerator on the catalyst, as shown in Equation (6.7).

$$C_{st} = c_{st,0}\exp\left[k_{st}\left(\frac{1}{r_{st}} - 1\right)\right], \qquad r_{st} = \frac{F_{strip}\,\eta_{st}/F_{cat}}{F_{strip}/(6.8F_0)} \tag{6.7}$$

where η_{st} is the stripper condition multiplier and r_{st} is the steam-to-catalyst ratio as a fraction of its design value, which is four grams of steam per kilogram of catalyst. The carryover rate is $R_{st} = F_{cat}\min(C_{st}, 0.004)$ in kg/s. That stream is hydrogen rich, so per kilogram it releases 30 % more heat than coke and needs 12 % more oxygen, and both arrive exactly when the unit is being pushed. Without (6.7) the regenerator temperature would not move with throughput at all.

Catalytic coke and stripper carryover burn in the same bed and are blended into one pseudo-coke, as shown in Equation (6.8).

$$\begin{aligned} R_c &= F y_D + R_{st} \\ \Delta H_c^{mix} &= \frac{F y_D \Delta H_c + R_{st} H_{st}}{R_c} \\ w_C^{mix} &= \frac{F y_D w_C + R_{st} w_{C,st}}{R_c} \end{aligned} \tag{6.8}$$

where R_c is the total coke reaching the regenerator in kg/s and hydrogen on coke, which with delta coke is the pair a unit engineer reads the regenerator by [16], is $100(1 - w_C^{mix})$, or 7.25 % at the design point.

6.3.4 Coke Burn and Excess Oxygen, Solved Together

The dense bed is treated as well mixed, so the oxygen that the combustion sees is the oxygen in the flue gas. The burn rate is first order in carbon on catalyst and first order in oxygen mole fraction, and the oxygen balance closes on the same mole fraction, which makes the pair solvable exactly as shown in Equation (6.9).

$$\begin{aligned} R_b = \beta y, \qquad F_g y = \dot{n}_{O_2,in} - \nu R_b \implies y &= \frac{\dot{n}_{O_2,in}}{F_g + \nu\beta} \\ \beta = k_b M C_{rgn} \frac{P_{rgn}}{P_{rgn,0}} \exp\left[-\frac{E_b}{R}\left(T_{rgn}^{-1} - T_{rgn,0}^{-1} \right) \right] & \end{aligned} \tag{6.9}$$

where $F_g = F_{air}/M_{air}$ is the air in kmol/s, $\dot{n}_{O_2,in} = 0.2095 F_g$, $\nu = w_C^{mix}/12.011 + (1 - w_C^{mix})/(4 \times 1.008)$ is the oxygen demanded per kilogram of coke burned completely, and k_b is the combustion rate constant. Read at a frozen carbon inventory, (6.9) says that more air really is more burn, because y climbs toward 20.95 % and R_b approaches 0.2095β. Read at steady state it says the opposite, because the coke burned has to equal the coke the riser sends and the riser does not know how much air there is. That single line is the whole of the result in Section 6.4.2.

6.3.5 Afterburn Above the Bed

Below about one percent oxygen a growing fraction of the carbon leaves the dense bed as carbon monoxide, and above the bed there is no catalyst to absorb heat, so whatever monoxide meets whatever oxygen up there burns in the cyclones, as shown in Equation (6.10).

$$\begin{aligned} f_{CO} &= \frac{1}{1 + \gamma (y/y_{CO})^{n_{CO}}}, \qquad \iota = \frac{1}{1 + \mathrm{e}^{-(T_{rgn} - T_{ign})/w_{ign}}} \\ \dot{n}_{CO} &= f_{CO} \frac{R_b w_C^{mix}}{12.011}, \qquad \dot{n}_{ab} = \iota \min(\dot{n}_{CO},\, 2 y F_g) \\ \Delta T_{ab} &= \frac{\dot{n}_{ab} \Delta H_{CO}}{(F_{air} + R_b)\, c_{p,fl}} \end{aligned} \tag{6.10}$$

where γ is the carbon monoxide promoter multiplier, y_{CO} is the bed oxygen at which half the carbon still leaves as monoxide, n_{CO} sets the sharpness of the transition, ι is an ignition sigmoid of width w_{ign} about T_{ign}, and ΔH_{CO} is the heat of monoxide combustion in kJ/kmol. The min in (6.10) is what makes the afterburn a peak rather than a slope, because it needs both reactants and each runs out at one end of the sweep. Heat released in the bed is the full combustion heat less whatever walked out as monoxide, $Q_{dense} = R_b \Delta H_c^{mix} - \dot{n}_{CO} \Delta H_{CO}$, so a bed that afterburns is a bed that is cooling.

6.3.6 The Regenerator Energy and Carbon Balances

The dense bed carries the plant's inertia, and the air is a cooling load as well as an oxidant, as shown in Equations (6.11) and (6.12).

$$\begin{aligned} C^{th}_{rgn}\frac{\mathrm{d}T_{rgn}}{\mathrm{d}t} &= 60\,[Q_{dense} - F_{cat}c_{p,cat}(T_{rgn}-T_{rot}) \\ &\quad -F_{air}c_{p,air}(T_{rgn}-T_{blo}) \\ &\quad -R_b c_{p,fl}(T_{rgn}-T_{rot}) - Q_{loss,rgn}] \end{aligned} \tag{6.11}$$

$$\begin{aligned} \frac{\mathrm{d}M}{\mathrm{d}t} &= \dot{m}_{add} - \dot{m}_{loss} \\ \frac{\mathrm{d}C_{rgn}}{\mathrm{d}t} &= \frac{60(R_c - R_b) - C_{rgn}\,\mathrm{d}M/\mathrm{d}t}{M} \end{aligned} \tag{6.12}$$

where $C^{th}_{rgn} = Mc_{p,cat} + M_{shell,rgn}c_{p,steel}$ and $\dot{m}_{add}$ and $\dot{m}_{loss}$ are the fresh catalyst addition and the cyclone losses in kg/min. At the design point the 130.1 MW released in the bed splits into 95.0 MW carried to the riser, 31.8 MW spent heating air from the blower discharge to bed temperature, 0.8 MW into the flue gas and 2.5 MW through the shell. The air term in (6.11) is why raising throughput does not raise the bed temperature. More feed brings proportionally more coke and proportionally more air, and the second of those is a cooling load.

Two first-order states follow the bed, the cyclone and plenum temperature and the flue gas duct, as shown in Equation (6.13).

$$\begin{aligned} \tau_{cyc}\frac{\mathrm{d}T_{cyc}}{\mathrm{d}t} &= T_{rgn} + \Delta T_{ab} - T_{cyc} \\ \tau_{duct}\frac{\mathrm{d}y_{duct}}{\mathrm{d}t} &= 100y - y_{duct} \end{aligned} \tag{6.13}$$

where y_{duct} is the flue gas oxygen reported in volume percent and $\tau_{duct} = 0.25$ min is the duct holdup, which is a different quantity from the analyzer lag τ_{O_2} of Figure 6.1. The lagged T_{cyc} is the cyclone temperature that the application displays, logs and scores, and the afterburn fence constrains it through the cyclone thermocouple. Raising τ_{cyc} from 0.4 to 10 min in scenario 6 leaves the plenum at 704.2 °C one minute after the promoter loss at 240 min, against 729.0 °C at the default.

6.3.7 Pressures and the Wet Gas Compressor

Reactor pressure is one gas holdup between the light gas the reaction makes and the volume the compressor swallows, as shown in Equation (6.14).

$$\begin{aligned} \frac{\mathrm{d}P_{rx}}{\mathrm{d}t} &= \frac{60R(T_{sys}+273.15)}{V_{sys}M_{wg}}\,(W_{wg} - W_{out}) \\ W_{wg} &= Fy_C, \qquad W_{out} = \frac{N}{100}Q_{wgc}\frac{(P_{rx}-\Delta P_{col})M_{wg}}{R(T_{suc}+273.15)} \end{aligned} \tag{6.14}$$

where N is the compressor speed in percent, Q_{wgc} is the suction volumetric capacity at full speed, ΔP_{col} is the main column pressure drop to the compressor suction, and R is the gas constant. The machine is a fixed volumetric capacity, so once N reaches 100 % the reactor pressure is set by the reaction rather than by a controller.

Regenerator pressure has its own holdup and its own slide valve, held on the regenerator pressure itself rather than on a differential, as shown in Equation (6.15).

$$\begin{aligned} \frac{\mathrm{d}P_{rgn}}{\mathrm{d}t} &= \frac{60R(T_{rgn}+273.15)}{V_{rgn}} \\ &\quad \times \left(\dot{n}_{fl} - K_{fl}u_{fl}\sqrt{P_{rgn}-P_{stk}}\right) \\ \frac{\mathrm{d}u_{fl}}{\mathrm{d}t} &= K_{flv}\left(P_{rgn} - P^{SP}_{rgn}\right) \end{aligned} \tag{6.15}$$

where $\dot{n}_{fl} = F_g + R_b\left[-(1 - w_C^{mix})/(4 \times 1.008)\right] + \frac{1}{2}\dot{n}_{CO}$ is the flue gas molar rate after the change in moles that combustion produces, and u_{fl} is the flue gas slide valve position in percent. Holding P_{rgn} and not a differential is what lets the compressor limit reach the catalyst circulation, because a reactor pressure that rises against a fixed regenerator pressure squeezes the differential the catalyst slide valve has to work with.

6.3.8 The Slide Valve, Which Is Not a Flow Source

The regenerated catalyst slide valve is a restriction in a circuit, and the transfer line downstream of it takes differential pressure as the square of the flow. Writing $F_{cat} = K_s u\sqrt{\Delta P_v}$ with $\Delta P_v = \Delta P_{av} - K_{line}(F_{cat}/F_{cat,0})^2$ gives an implicit equation that solves in closed form, as shown in Equation (6.16).

$$\begin{aligned} F_{cat} &= \frac{K_s u\sqrt{\Delta P_{av}}}{\sqrt{1 + K_s^2 u^2 K_{line}/F_{cat,0}^2}} \\ \Delta P_{av} &= P_{rgn} - P_{rx} + \frac{580\, a_{er}\, g\,(H_{drop} + h_b)}{1000} \\ h_b &= \frac{M}{\rho_b A_{rgn}} \end{aligned} \tag{6.16}$$

where K_s is the valve coefficient, u is the opening in percent, a_{er} is the standpipe aeration multiplier, $F_{cat,0} = 6.8F_0$ is the design circulation, and the constant 580 is the aerated catalyst density in kg/m^3 at design aeration. The head term is in kPa. The closed form matters, because marching (6.16) by substitution oscillates and then diverges above about seventy percent open, which is exactly the region where the low-differential limit lives and therefore the only region a student examines closely. At the design point the standpipe and vessel differential supply 119.0 kPa, the transfer line takes 62.0 kPa of it, and 57.0 kPa is left across the valve against a 25 kPa minimum and a 14 kPa interlock.

6.3.9 Catalyst Activity, Actuators and Economics

Metals in the feed land on the catalyst and never leave, and the only answer is fresh catalyst, as shown in Equation (6.17).

$$\tau_a \frac{\mathrm{d}a}{\mathrm{d}t} = \mathrm{sat}\left(\frac{a_{max}}{1 + k_{me}\, m/\dot{m}_{add}},\ 0.1,\ 1.6\right) - a \tag{6.17}$$

where m is the feed metals multiplier, a_{max} is the fresh catalyst activity and τ_a is the turnover time, which is set to 90 min here against weeks on a real unit so that the effect can be watched in one session. Five actuators are first-order lags on commands the controller has already clamped, as shown in Equation (6.18).

$$\begin{aligned} \tau_s \frac{\mathrm{d}u}{\mathrm{d}t} &= u_{cmd} - u, & \tau_b \frac{\mathrm{d}F_{air}}{\mathrm{d}t} &= F_{air,cmd} - F_{air} \\ \tau_w \frac{\mathrm{d}N}{\mathrm{d}t} &= N_{cmd} - N, & \tau_f \frac{\mathrm{d}F}{\mathrm{d}t} &= F_{cmd} - F \\ \tau_p \frac{\mathrm{d}T_{pre}}{\mathrm{d}t} &= T_{pre,cmd} - T_{pre} \end{aligned} \tag{6.18}$$

No actuator can therefore walk past its stop, and the slide valve state is also clipped after each step.

Gross margin is products less feed, power and fresh catalyst, as shown in Equation (6.19).

$$\begin{aligned} J &= 3.6F\left(y_B p_{gaso} + y_C p_{light} + y_A p_{btms} - p_{feed}\right) \\ &\quad - \left(W_{blo} + W_{wgc}\right) p_{pow} - 60\,\dot{m}_{add}\, p_{cat} \\ W_{blo} &= \frac{1.005\,F_{air}\,(288.15)}{\eta_{blo}}\left[\left(\frac{P_{rgn}+30}{101.3}\right)^{0.286} - 1\right] \\ W_{wgc} &= \frac{2.0\,W_{wg}\,(T_{suc}+273.15)}{\eta_{wgc}} \\ &\quad \times \left[\left(\frac{P_{dis}}{P_{rx}-\Delta P_{col}}\right)^{0.200} - 1\right] \end{aligned} \tag{6.19}$$

where J is the gross margin in dollars per hour, p_{gaso}, p_{light}, p_{btms} and p_{feed} are in dollars per tonne, p_{pow} is in dollars per kilowatt hour and p_{cat} is in dollars per kilogram. Both machines are priced from an adiabatic head with a fixed efficiency, so 1.005 and 2.0 are the specific heats of air and of wet gas in kJ/(kg K), 288.15 and $T_{suc}+273.15$ are the two suction temperatures in kelvin, and the exponents 0.286 and 0.200 are $(\gamma-1)/\gamma$ for each gas. Each pressure ratio is floored at 1.05 and the compressor suction pressure at 20 kPa, so neither power turns negative when a pressure collapses. Both powers are in kW and the prices appear in Table 6.2.

6.3.10 The Five Fences on One Scale

Each limit is expressed as the fraction of its own approach span that has been used up, so a comparison across four different units means something, as shown in Equation (6.20).

$$U_i = 100\left(1 - \frac{\text{limit}_i - \text{PV}_i}{\text{span}_i}\right) \tag{6.20}$$

The five are the air blower against its capacity, the wet gas compressor against 100 % speed, the cyclone temperature against its metal limit, the slide valve differential against its minimum, and the regenerator dense bed against its metal limit. Blower capacity is not a constant, because a centrifugal machine takes in a volume rather than a mass, as shown in Equation (6.21).

$$F_{air}^{max} = F_{air,0}^{max}\,\frac{288.15}{273.15 + T_{amb}} \tag{6.21}$$

so the fence itself moves with the weather, from 230.4 t/h at 15 °C to 213.4 t/h at 38 °C. The spans were chosen so that each fence the feed rate can reach is worth a similar amount of feed. Measured as local slopes at the design point with the riser outlet and the oxygen held, one span is worth 48.5 t/h of feed on the blower, 49.0 on the compressor and 62.9 on the slide valve differential, which is what lets one gain and one integral time serve four constraint controllers. The cyclone fence barely moves with feed while the oxygen is held, and it is reached through the oxygen falling once the blower runs out. The regenerator span says something different, because thirty degrees is the whole authority of the feed preheat handle.

> **TRY IT IN THE APP** **11 · The fence moves**
>
> Load the case, set the speed to 20× and press Run. Nothing else has to be set. At 90 min the afternoon warms from 15 to 38 °C. Watch the blower capacity fall from 230.4 to 213.4 t/h and the feed rate follow it down from 284.5 to 261.5 t/h. In the evening the feed walks back up to 287.8 t/h.

6.3.11 Parameters, Initialization and Numerics

The 105 constants are listed in Table 6.2. Every one is shown and adjustable in the application's model panel except the 31 marked hidden, which are neither displayed nor adjustable and are the quantities

a student has to infer from a response. Two further constants are fixed in code and never exposed, the gas constant $R = 8.314\,\mathrm{kJ/(kmol\,K)}$ and the molecular weight of air $M_{air} = 28.85\,\mathrm{kg/kmol}$.

Each scenario opens at a solved steady state rather than at a guess. A manual scenario is solved with the operator's four handles given and both temperatures free. An automatic scenario is solved with the riser outlet temperature held, the oxygen on setpoint and the reactor pressure on setpoint, and the manual stations are then moved to match, so that manual operation is bumpless from the first second and no scenario spends its opening ten minutes walking to its own setpoint. The riser loop's integral is seeded net of the feedforward bias of (6.24), so a case that opens away from the design feed rate starts with the slide valve where the solution put it. A constraint scenario asks for more feed than exists on purpose, so the requested throughput is bisected downward until the solved point is inside every fence with room to spare, which is taken as a slide valve below 88 % open and a nearest fence below 88 % of its span.

Three balances are solved simultaneously in either case, the riser energy balance, the regenerator energy balance and the coke balance, with a damped Newton on a numerical Jacobian. Successive substitution around the same circle would contract only by the loop gain of about one half per pass, while the Newton finishes in six or seven iterations at every line-up used here. Two details of that solver are not decoration. The Jacobian column for the carbon on the regenerated catalyst is perturbed relatively rather than absolutely, because five ten-thousandths perturbed like a temperature falls below the noise floor of the riser march. Trust-region step limits and, in valve mode, an eight-step continuation from the design line-up keep the solve on the hot branch, because this plant also has a cold, coked-up solution that satisfies all three balances equally well [2, 3]. With the oxygen held at 2.0 % and the reactor pressure on setpoint, a Newton iteration started from a cold guess converges to riser outlet temperatures of 442 to 445 °C with 0.50 to 0.66 % carbon on the regenerated catalyst for slide valve positions from 42 to 55 %.

Integration is classical fourth-order Runge-Kutta at a fixed 0.02 min (1.2 s) step, the controller executes every 0.1 min (6 s), and the historian records ten samples per simulated minute, which a run exports as a CSV file of 129 columns. The engine, its seed (20260913 here) and its speed settings are those of Section 1.6.1. Speeds up to 20×, twenty simulated minutes per wall-clock second, are what the ninety-minute regenerator tail needs.

6.3.12 Assumptions

Model detail is concentrated on the heat balance and the constraint set. Each simplification below is stated with the conditions under which it fails.

Four lumps and one activity term. The scheme reproduces conversion, the gasoline maximum against severity and the coke yield, which is what the heat balance needs. It says nothing about gasoline octane, LPG olefinicity or cycle oil quality, none of which is a function of conversion alone, and real yield models carry ten to twenty lumps for exactly that reason [10, 11].

Carbon residue goes to coke quantitatively and comes out of the bottoms. That is the standard additive coke treatment. Real resid coke also poisons the catalyst and changes the metals loading, and neither is coupled here.

Catalyst deactivation is compressed by a factor of several hundred. The activity turnover in (6.17) is 90 min against weeks on a real unit. The direction and the trade are right and the clock is not, so no interval reported for the metals scenario predicts a real catalyst addition schedule.

One well-mixed dense phase at one temperature and one oxygen concentration. Real beds have a bubble phase that bypasses the catalyst, radial gradients of twenty degrees or more, and an air distributor whose condition is most of the difference between a unit that afterburns and one that does not [8].

Afterburn is one lumped dilute-phase reaction with an ignition sigmoid. The shape is right and the peak is in the right place, but (6.10) is a correlation rather than a freeboard model [12] and it cannot represent the local hot spots that actually destroy cyclones.

Symbol	Meaning	Value	Unit	Vis.
Feed and steam				
F_0	design feed rate	69.444	kg/s	H
$c_{p,oil}$	feed sensible heat capacity	2.45	kJ/kg °C	A
ΔH_{vap}	feed heat of vaporization	300	kJ/kg	A
ΔH_{rx}	heat of cracking per kg converted	354	kJ/kg	A
f_{st}	dispersion steam, fraction of feed	0.03	—	A
F_{strip}	stripping steam	2.0	kg/s	A
$c_{p,st}$	steam heat capacity	2.10	kJ/kg °C	H
T_{st}	steam supply temperature	250	°C	H
Riser kinetics				
$T_{rot,0}$	kinetic reference temperature	525	°C	H
$t_{r,0}$	vapor residence time at design	3.0	s	A
k_1	gas oil to gasoline rate constant	0.233302	—	A
k_2	gas oil to light gas rate constant	0.091886	—	A
k_3	gas oil to coke rate constant	0.018851	—	A
k_4	gasoline overcracking rate constant	0.01200	—	A
k_5	gasoline to coke rate constant	0.00260	—	A
E_1	activation energy, gasoline	68	kJ/mol	A
E_2	activation energy, light gas	89	kJ/mol	A
E_3	activation energy, coke	48	kJ/mol	A
E_4	activation energy, overcracking	105	kJ/mol	A
E_5	activation energy, gasoline to coke	48	kJ/mol	A
α	activity lost per wt % coke on catalyst	1.50	per wt %	A
k_{CCR}	additive coke from feed carbon residue	0.0090	per wt %	A
n_r	riser integration steps	14	—	H
Stripper				
$c_{st,0}$	hydrocarbon left on spent catalyst	0.00030	—	A
k_{st}	stripping sensitivity to steam ratio	1.20	—	A
H_{st}	heat of combustion, unstripped oil	43000	kJ/kg	A
$w_{C,st}$	carbon fraction, unstripped oil	0.860	—	A
Regenerator				
$T_{rgn,0}$	combustion reference temperature	700	°C	H
k_b	coke combustion rate constant	2.36699	—	A
E_b	coke combustion activation energy	150	kJ/mol	A
w_C	carbon fraction of catalytic coke	0.930	—	A
ΔH_c	heat of complete coke combustion	33014	kJ/kg	A
—	oxygen mole fraction in air	0.2095	—	H
ΔH_{CO}	heat of CO combustion	283000	kJ/kmol	H
y_{CO}	bed oxygen at half CO selectivity	0.0030	—	A
n_{CO}	sharpness of the CO transition	3.0	—	A
T_{ign}	dilute-phase CO ignition temperature	665	°C	A
w_{ign}	ignition transition width	18	°C	H
$c_{p,air}$	air heat capacity, blower to bed	1.10	kJ/kg °C	H
$c_{p,fl}$	flue gas heat capacity	1.15	kJ/kg °C	H
$c_{p,cat}$	catalyst heat capacity	1.15	kJ/kg °C	A
T_{blo}	blower discharge temperature	190	°C	A
$Q_{loss,rgn}$	regenerator shell heat loss	2500	kW	A
$Q_{loss,rx}$	reactor and riser heat loss	1500	kW	A
A_{rgn}	regenerator cross-section	78.5	m^2	A
ρ_b	dense bed density	480	kg/m^3	A
Thermal inertia				
M_{strip}	catalyst in reactor and stripper	25000	kg	A
$M_{shell,rx}$	reactor vessel steel	60000	kg	A
$M_{shell,rgn}$	regenerator steel and refractory	250000	kg	A

Symbol	Meaning	Value	Unit	Vis.
Pressure and wet gas				
$P_{rx,0}$	design reactor pressure	310	kPa a	H
$P_{rgn,0}$	design regenerator pressure	350	kPa a	H
ΔP_{col}	column drop to compressor suction	100	kPa	A
M_{wg}	wet gas molecular weight	30	kg/kmol	A
T_{suc}	compressor suction temperature	40	°C	A
T_{sys}	reactor and column vapor temperature	127	°C	H
V_{sys}	reactor, column and receiver volume	270	m^3	A
Q_{wgc}	compressor suction capacity at 100 %	7.9613	m^3/s	A
V_{rgn}	regenerator freeboard volume	400	m^3	H
P_{stk}	flue gas stack pressure	105	kPa a	H
K_{fl}	flue gas slide valve coefficient	0.002651475	—	H
P^{SP}_{rgn}	regenerator pressure setpoint	350	kPa a	A
K_{flv}	flue gas valve controller gain	0.15	%/min kPa	H
Catalyst loop				
H_{drop}	bed floor to slide valve elevation	9.5	m	A
K_{line}	transfer line loss at design flow	62	kPa	A
K_s	RGC slide valve coefficient	1.39032	—	A
ΔP_{min}	differential at loss of metering	25	kPa	A
ΔP_{trip}	low-differential interlock	14	kPa	A
M_0	design catalyst inventory	165000	kg	A
τ_a	catalyst activity turnover time	90	min	A
k_{me}	activity lost per metals ratio	0.60	—	A
a_{max}	fresh catalyst activity	1.30	—	H
Actuators				
τ_s	slide valve travel	0.12	min	A
τ_b	air blower response	0.35	min	A
r	blower rate limit	2.5	kg/s min	A
τ_w	compressor speed response	0.20	min	A
τ_f	charge pump and feed flow loop	0.60	min	A
τ_p	feed preheat train response	6.0	min	A
Instruments				
τ_{TC}	riser outlet thermocouple	0.25	min	A
τ_{O_2}	oxygen analyzer response	0.50	min	A
θ_{O_2}	oxygen sample line delay	0.60	min	A
τ_{Pm}	pressure transmitter	0.05	min	H
τ_{bed}	regenerator bed thermocouple	0.40	min	H
t_{opt}	feasible optimum search interval	5.0	min	H
Constraints				
$F^{max}_{air,0}$	blower capacity at 15 °C	64.0	kg/s	A
T^{max}_{rgn}	dense bed metallurgical limit	730	°C	A
T^{max}_{cyc}	cyclone and plenum limit	760	°C	A
—	regenerator high-temperature trip	750	°C	A
—	cyclone high-temperature trip	790	°C	A
s_{air}	blower approach span	10.0	kg/s	H
s_{wgc}	compressor approach span	14	%	H
s_{rgn}	dense bed approach span	30	°C	H
s_{cyc}	cyclone approach span	45	°C	H
$s_{\Delta P}$	slide valve approach span	32	kPa	H
Economics				
p_{gaso}	gasoline value	780	$/t	A
p_{light}	light gas and LPG value	520	$/t	A
p_{btms}	cycle oil and slurry value	430	$/t	A
p_{feed}	gas oil feed cost	520	$/t	A

Table 6.3 The Slide Valve Step of Scenario 1, 45.0 to 48.0 % in Manual With Air, Feed and Preheat Held. Percentages Are of the Settled Change.

Time after the step (min)	2	22	41	83
ΔT_{rot} (°C)	2.11	4.53	5.73	6.79
of settled 7.17 °C (%)	29	63	80	95
ΔT_{rgn} (°C)	−0.41	1.70	3.31	4.74
of settled 5.25 °C (%)	−8	32	63	90

The stripper is one exponential in the steam-to-catalyst ratio. That gives hydrogen on coke the right sensitivity to circulation, which is what makes the regenerator temperature move with throughput at all, but a real stripper is a tray or packing efficiency and it fails in ways (6.7) cannot show.

There is no standpipe fluidization model. The differential in (6.16) falls to the interlock and stops there. The flow reversal and the pressure oscillation that a real low-differential event produces are absent, and they are the reason such an event is dangerous rather than merely expensive.

There is no fractionator. Reactor pressure is one gas holdup swallowed by one compressor characteristic and the main column is a fixed pressure drop, so nothing here knows about gasoline endpoint, slurry circulation, or the wet gas molecular weight changing with severity.

No catalyst cooler, no flue gas expander, no CO boiler, no feed recycle and no torch oil. Each is a handle on a real unit, and the catalyst cooler in particular is what a resid unit installs to answer exactly the hot regenerator problem that the heavy feed scenario poses [16].

Perfect repeatability is a teaching convenience. A fixed step and a fixed scan produce one trajectory on every machine, which no distributed control system does.

6.4 What the Heat Balance Does

6.4.1 Measuring the Loop Gain

The experiment is one slide valve move in manual operation with the air rate, the feed rate and the preheat all held. Opening the regenerated catalyst valve from 45.0 to 48.0 % raises the circulation from 1700 to 1751.5 t/h. Solving the riser energy balance of (6.5) and (6.6) alone for that circulation, with the regenerator frozen where it was, gives a riser outlet temperature of 528.38 °C against 524.96 °C, a rise of 3.42 °C. What the unit actually settles at is 532.12 °C, a rise of 7.17 °C, and the dense bed rises 5.25 °C from 699.9 to 705.2 °C without anyone asking it to.

The ratio of those two numbers is the amplification, and the loop gain follows from it, as shown in Equation (6.22).

$$A = \frac{\Delta T_{rot}}{\Delta T_{rot}^{open}} = \frac{7.17}{3.42} = 2.094, \qquad L = 1 - \frac{1}{A} = 0.523 \tag{6.22}$$

where L is the fraction of any disturbance that the plant returns to itself through the coke it makes. A loop gain of 0.523 does two things at once. It leaves the plant stable, because the sum converges, and it makes the plant slow, because each pass around the loop has to wait for the regenerator. The recorded trajectory in Table 6.3 shows both. Push severity up and L climbs, from 0.534 at 525 °C to 0.550 at 535 °C for a one percent valve step at line-ups that hold 2.0 % oxygen, which is why a unit that has always been steady can turn twitchy the week somebody raises the riser outlet setpoint, and why the multiplicity that appears once the gain reaches unity is a property of the same arithmetic [2, 3].

Two features of Table 6.3 are worth separating. The dense bed initially falls, by 0.46 °C at three minutes, because the extra catalyst carries heat out of the regenerator before the extra coke arrives to pay for it, and it crosses back through zero at 8.8 min. The carbon on the regenerated catalyst is a third integrator inside the same loop and it overshoots, rising from 0.0494 to 0.0577 % by weight 13 min after the step before settling at 0.0545. Three coupled integrators with one of them slow is why nobody tunes an FCC main loop aggressively.

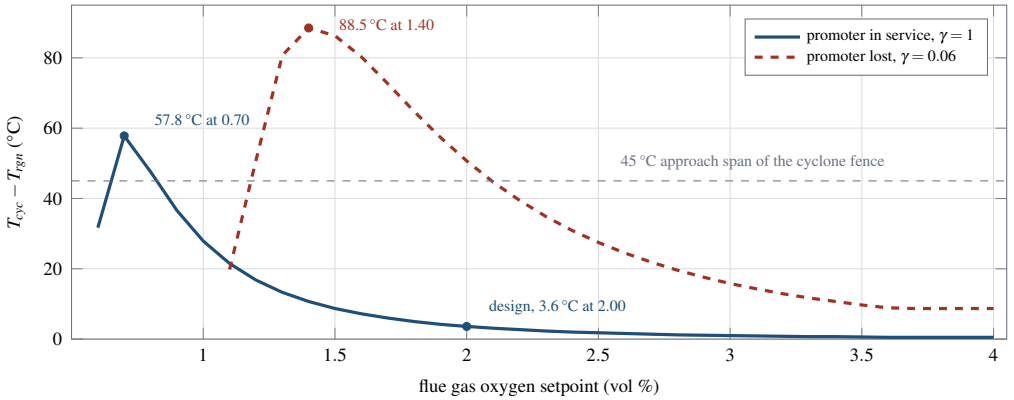


Figure 6.2 Afterburn against the excess oxygen setpoint at settled operating points, swept on the shipped model at 250 t/h and 525 °C. The peak is between partial burn and full burn rather than at either end, because the reaction needs carbon monoxide and oxygen together. The solid line is the promoter in service and the dashed line the promoter lost. Losing the promoter moves the peak right and up and puts 50.7 °C of afterburn at the design oxygen. Both curves flatten above about 3.6 % because the blower runs out of capacity and the oxygen stops at 3.71 and 3.62 %, and the promoter-lost curve is drawn only above 1.1 %, below which the steady-state solve leaves the hot branch.

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.4.2 More Air Is Not More Burn

The two halves of this claim are not the same statement, and (6.9) contains both. At a frozen carbon inventory of 0.050 % by weight and a frozen bed temperature, raising the air from 56.7 to 66.3 kg/s, an increase of 16.9 %, raises the instantaneous burn rate by 15 %, because the rate is first order in oxygen mole fraction. At steady state the same move is worth almost nothing. Scenario 2 takes the air from 204 to 223 t/h, and 75 min later the burn rate has gone from 3.906 to 3.920 kg/s, a change of 0.3 %. It cannot do anything else, because the regenerator burns exactly the coke the riser sends it and the riser does not know how much air there is.

What the extra air buys is a lower carbon inventory. The excess oxygen rises from 2.03 to 3.59 % by volume and the carbon on the regenerated catalyst falls by a third, from 0.0494 to 0.0333 % by weight. What it costs is heat, because the air term in (6.11) is a cooling load, so the bed settles at 691 °C instead of 700 and the riser follows it down from 525 to 519 °C. Taking the air the other way, down to 162 t/h, takes the unit apart. By $t = 300$ min carbon has built to 0.46 % by weight on the regenerated catalyst, the activity in (6.2) has gone with it, conversion has fallen from 75.0 to 55.5 % and the margin from $15,347 to $7,019 per hour.

6.4.3 The Afterburn Peak

Sweeping the oxygen setpoint across settled operating points rather than frozen ones is the only honest way to draw the afterburn characteristic, because the carbon inventory moves with the oxygen and it is the pair together that decides how much monoxide leaves the bed. The result is in Figure 6.2.

The peak with the promoter working is 57.8 °C above the bed at 0.70 % oxygen by volume, against 3.6 °C at the design 2.00 % and 1.0 °C at 3.00 %. The road between partial burn and full burn therefore goes through the worst of it in both directions, which is the answer to the suggestion that an operator can simply add air. Losing the promoter moves the peak to 88.5 °C at 1.40 % and leaves 50.7 °C at the design oxygen, and the only remaining answer is a great deal more air. Scenario 6 asks for 4.0 % and does not get it, because the blower saturates at 230.4 t/h, the oxygen stops at 3.62 % and 8.7 °C of afterburn is what is left. The dense bed meanwhile gets cooler as the oxygen falls, from

700 to 686.6 °C at 0.85 % and 672.9 °C at the peak, because the heat in the escaping monoxide leaves the bed with it. An operator watching the bed thermocouple alone sees a unit that is behaving well, right up until the cyclones fail.

> **TRY IT IN THE APP** **6 · Afterburn, and why you cannot just change the air**
>
> Load the case and press Run. The oxygen setpoint is taken down to 0.85 % and the dense bed cools as the oxygen falls. Watch the gap between the plenum and the bed, not the bed alone. The promoter is lost at 240 min. It then asks for 4.0 % oxygen, the blower saturates at 230.4 t/h, the oxygen stops at 3.62 % and 8.7 °C of afterburn remains.

6.4.4 The Fence the Feed Rate Cannot Reach

Four of the five limits respond to throughput. The regenerator dense bed does not, and the arithmetic says why. Bed temperature follows delta coke, which is coke per kilogram of circulating catalyst, and at a held riser outlet temperature delta coke barely moves with throughput because more feed brings proportionally more catalyst and proportionally more air. Measured on this model, fifty tonnes an hour of extra feed at 525 °C moves the dense bed by 4 °C and moves it the wrong way, downward. The handle that does reach the bed is the feed preheat, whose whole travel from 175 to 300 °C is worth 30.8 °C of bed temperature, which is where the thirty degree approach span in (6.20) comes from.

Raising the preheat by sixty degrees at a held riser outlet takes the circulation from 1700 down to 1373 t/h, the catalyst-to-oil ratio from 6.80 to 5.49, and the conversion from 75.00 to 69.95 %, and it makes the regenerator hotter by 15.5 °C. Almost everybody predicts cooler, because less coke is burned in absolute terms. Less coke is correct, but it is burned onto a much smaller stream of catalyst, and the bed temperature follows the ratio and not the rate. Cutting the preheat is not free either. At 180 °C of preheat the unit makes 77.6 % conversion with the slide valve 61 % open and 39.6 kPa across it against a 25 kPa minimum, so defending one fence has spent two others.

Pause. Try Exercise 6.2 before reading on.

6.5 Controller Design

6.5.1 The Implemented PI Algorithm

All eight controllers are the velocity-form PID of Section 1.6.2, with the derivative taken on the measurement, and all eight ship with $\tau_D = 0$, so the derivative term is present in the implementation and inactive in every default tuning. One execution is Equation (6.23), which is Equations (1.16) to (1.19) with an output slew limit added.

$$
\begin{aligned}
e_k &= \sigma\,(\mathrm{SP}_k - \mathrm{PV}_k) \\
P_k &= K_c e_k, \qquad D_k = -K_c \tau_D \,\sigma \left.\frac{\mathrm{d}\mathrm{PV}^f}{\mathrm{d}t}\right|_k \\
u_k^{unsat} &= u_{bias} + \Sigma_k^{PD} + I_{k-1} \\
\Sigma_k^{PD} &= \Sigma_{k-1}^{PD} + \Delta P_k + \Delta D_k \\
u_k &= \mathrm{slew}\left(\mathrm{sat}(u_k^{unsat}, u_{min}, u_{max}),\, u_{k-1},\, r\right) \\
I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - u_k^{unsat}}{T_t}\right] \Delta t_c
\end{aligned}
\tag{6.23}
$$

where $\sigma = \pm 1$ is the action, PV^f is the measurement after an optional derivative filter that no loop here uses, Σ^{PD} accumulates the proportional and derivative increments without limit so that it telescopes to exactly $P + D$ for the gains in force, I_{k-1} is the integral accumulated through the previous scan, r is the output rate limit, which only the oxygen loop uses, and T_t is the back-calculation tracking time,

Table 6.4 The Implemented Loops as Coded, With Default Tunings. Action $\sigma = +1$ Is Reverse Acting and $\sigma = -1$ Is Direct Acting. The Air Ceiling Is the Blower Capacity of Equation (6.21).

CV	MV	K_c	τ_I (min)	σ	Output limits
T_{rot}	slide valve	0.55	3.5	+1	0–100 %
y	air blower	1.6	6.0	+1	28–F_{air}^{max} kg/s
P_{rx}	compressor	0.80	0.8	−1	55–100 %
U_{air}	feed rate	0.32	9.0	+1	41.7–94.4 kg/s
U_{wgc}	feed rate	0.32	9.0	+1	41.7–94.4 kg/s
U_{cyc}	feed rate	0.32	9.0	+1	41.7–94.4 kg/s
$U_{\Delta P}$	feed rate	0.32	9.0	+1	41.7–94.4 kg/s
U_{rgn}	feed preheat	1.00	10	+1	175–300 °C

Table 6.5 Measured Loop Responses in Scenario 4. The Open-Loop Column Is a 1 % Slide Valve Step at a Frozen Air Rate.

Time after the step (min)	1	2	20	settled
ΔT_{rot} (°C per %)	0.47	0.74	1.55	2.57
Closed-loop setpoint steps				
	t_{63} (min)		t_{95} (min)	
T_{rot}, 525 to 532 °C	4.8		13	
P_{rx}, 310 to 330 kPa	0.4		2.1	

which no loop sets and which therefore equals τ_I throughout. Anti-windup is back-calculation on the integral alone, and it runs only when there is reset action to unwind. While the output is clipped or rate limited, $(u_k - u_k^{unsat})/T_t$ drives I toward the value that reproduces the applied output, removing $\Delta t_c/\tau_I$ of the excess per scan, or 1/35 on the riser outlet loop. The clamp acts on the emitted output alone and never on Σ^{PD}, for the reasons given in Section 1.6.2. Manual operation solves the same expression for I, so transfer is bumpless in both directions. The controller period is $\Delta t_c = 0.1$ min throughout, five times the integration step.

6.5.2 Three Regulatory Loops on Three Different Plants

The regulatory layer is three loops, and the striking thing about them is that no two are the same kind of plant. Their pairings, directions and default tunings are in Table 6.4, and their measured responses are in Table 6.5.

The riser outlet loop on the slide valve has a steady gain, for a one percent step, of 2.6 °C per percent of valve at a frozen air rate and 2.4 °C with the oxygen loop closed, and only 0.74 °C of it, 29 %, has arrived at two minutes. A loop tuned on the first two minutes is tuned for a plant 3.5 times weaker and very much faster than the one underneath it, and the error is in the direction that produces an aggressive controller. Closed at the default tuning, a 7 °C setpoint step is 63 % complete in 4.8 min and 95 % complete in 13 min, with two thirds of a degree of overshoot.

The oxygen loop on the air blower is awkward for the opposite reason. Its steady gain is 0.31 % by volume per kg/s of air, and the immediate response is a thirteenth of that, because three kilograms a second delivers 0.07 of its eventual 0.92 % in the first minute, and the rest arrives as the carbon inventory in (6.12) burns down to its new level, 90 % of it by 7.7 min and 95 % by 12.5 min. On top of that the analyzer sits behind 0.60 min of sample line and 0.50 min of its own response, modeled as a transport delay followed by a first-order lag rather than as one lumped filter, because a delay inside a filter is not the same loop and does not detune the same way. The blower is also the one actuator with a rate limit, 2.5 kg/(s min), because its ceiling in (6.21) moves with the weather and an unrestrained controller takes a four kilogram per second jump in the ceiling all at once.

The reactor pressure loop on the compressor is the only ordinary one, at −2.9 kPa per percent of speed and settled in two minutes, and it is in service in every mode including manual. That is not a shortcut. The catalyst slide valve differential in (6.16) is measured against the reactor pressure, so every kilopascal the reactor gains comes off the differential the circulation is metered with. The 20 kPa setpoint step in Table 6.5 takes the compressor from 85.0 to 78.0 % and the slide valve differential from 57 to 38 kPa.

A feed rate feedforward completes the layer. A feed rate change is a load on the riser rather than a disturbance the temperature loop should have to discover, and the constraint layer above knows exactly how much feed it is about to add, as shown in Equation (6.24).

$$u_{bias} = K_{ff}\,(F_{cmd} - F_0) \tag{6.24}$$

where K_{ff} defaults to 1.50 % of valve per kg/s of feed, which is the travel that holds the catalyst-to-oil ratio where it was. It enters (6.23) as the controller bias, so the feedback loop only ever corrects what the feedforward got wrong and the manual tracker accounts for it automatically. Without it the evening feed increase in the graded challenge, which climbs at the 3.6 t/h per minute the selector allows and adds 16 t/h in the five minutes after the ambient falls at 310 min, pulls the riser outlet 4.4 °C off setpoint. With it, the same increase costs 1.3 °C.

6.5.3 The Low Selector and Its External Reset

Four constraint controllers share the feed rate through a low selector that also carries the operator's own throughput target, as shown in Equation (6.25) and drawn in Figure 6.3.

$$F_{SP} = \min(F_{target},\ u_{air},\ u_{wgc},\ u_{cyc},\ u_{\Delta P}) \tag{6.25}$$

Every one of the four holds its own fence at the same constraint target U_{SP}, which defaults to 97 % of the approach span, and every one of the four carries the same gain and the same integral time, which is possible only because (6.20) put them on a common scale. The same structure, a selector whose losing integrators track the command applied, guards the fuel valve of the turbofan in Volume 4, Chapter 3.

External reset is the detail that makes the structure work, and it costs one line. Each constraint controller is given the output the selector actually applied last scan, plus a slack of 0.10 kg/s, as its own upper output limit, as shown in Equation (6.26).

$$u_i^{max} = \mathrm{sat}(F_{sel} + \delta,\ F_{min},\ F_{max}), \qquad \delta = 0.10\,\mathrm{kg/s} \tag{6.26}$$

The back-calculation term in (6.23) then parks each loser's integral at that limit. With $T_t = \tau_I$ the fixed point of the integral recursion is exactly $I = u_i^{max}$ whatever the error, so a controller a hundred percent away from its own fence sits on the limit, 0.10 kg/s above the selected output, with its integral parked rather than wound to the ceiling and its unclamped demand above the limit by its proportional term alone. The instant its own error goes negative, its proportional term takes it below the others and it wins the actuator in one scan. The application verifies this on itself by running the push scenario for 150 min and requiring that at least two idle controllers sit strictly above the selected output and that the worst of them is within 2.0 kg/s of it. The same slack is also a feed ramp limit, because one slack per six-second scan is 3.6 t/h per minute, about as fast as a charge pump, a feed heater and a riser outlet temperature controller will follow together.

🕮 **Pause.** Try Exercise 6.3 before reading on.

6.5.4 The Override on the Second Actuator

The regenerator dense bed cannot be reached by the feed rate, so a fifth controller holds it with the feed preheat, low-selected against the operator's preheat target, as shown in Equation (6.27).

$$\begin{aligned} T_{pre}^{SP} &= \min\!\left(T_{pre}^{target},\ \mathrm{PID}_{rgn}(U_{SP}, U_{rgn})\right) \\ u_{rgn}^{max} &= T_{pre,k-1}^{SP} + \delta_p \end{aligned} \tag{6.27}$$

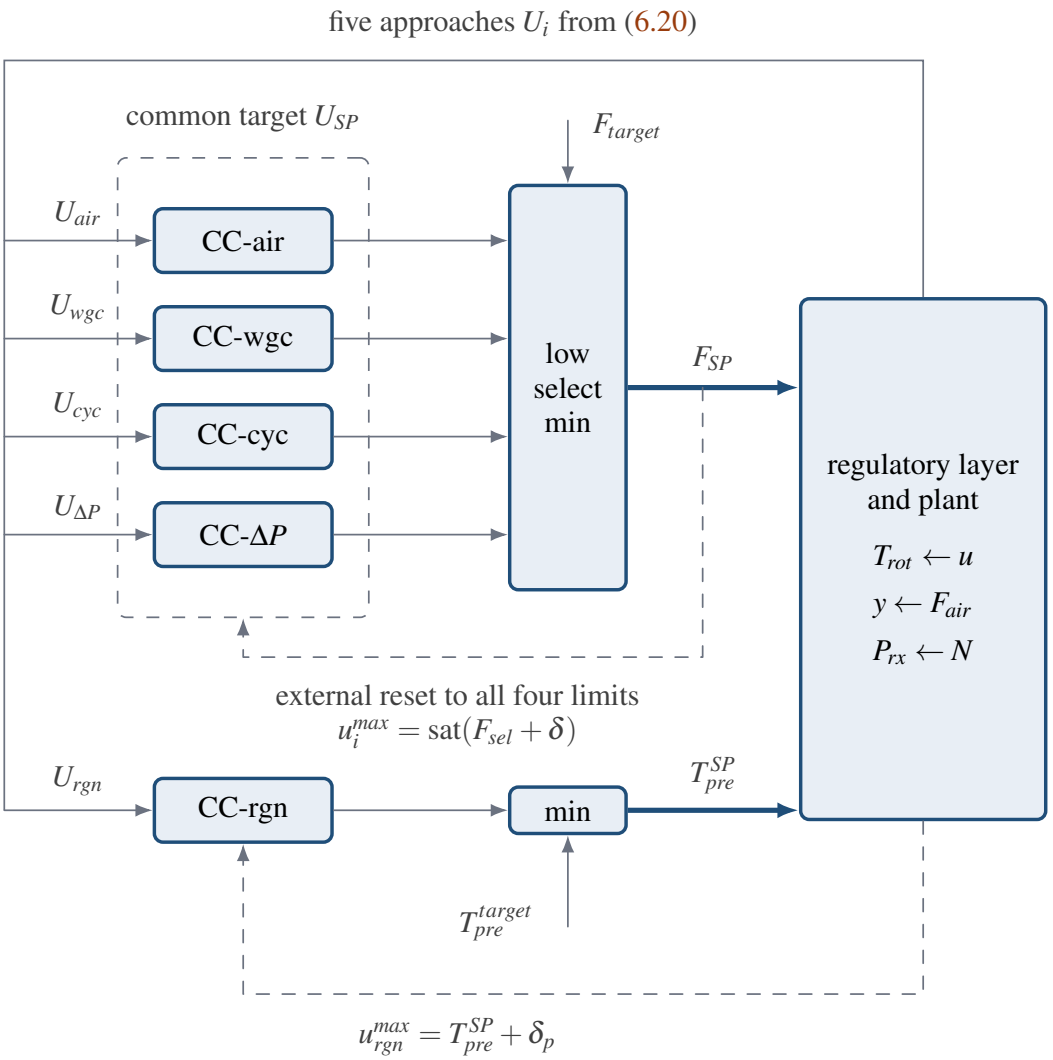


Figure 6.3 The constraint layer. Four controllers hold four fences at one common target U_{SP} and share the feed rate through a low selector that also carries the operator's throughput target, so the selector can only ever take feed away. Each loser is given the selected output plus a slack δ as its own upper output limit, which is external reset and is what parks its integral instead of letting it wind to the ceiling. The regenerator bed does not respond to feed rate, so its controller acts on the feed preheat through a second low select against the operator's preheat target.

where $T_{pre,k-1}^{SP}$ is the preheat setpoint applied on the previous scan, $\delta_p = 3.0\,^{\circ}\text{C}$ is the preheat slack, and the same external reset argument applies. The override is on a second actuator rather than on the shared one for a physical reason and not a structural preference, which is that cutting preheat sends more catalyst around for the same riser outlet temperature and spreads the same coke over more of it. Nothing else available moves delta coke. The cost is that the extra circulation walks straight into the slide valve differential and the blower, so defending the regenerator spends two of the other four fences.

Order inside one control execution never varies. The measurements are formed first, each through its own filter and, for the oxygen, its delay line, then the five fence approaches of (6.20) are evaluated,

then the tuning constants and the weather-dependent blower ceiling are applied, then the feasible optimum marker is refreshed if its own five-minute clock has expired, then the feedforward, then the three regulatory loops or their manual trackers, then the low selector of (6.25), and last the preheat override. The three regulatory loops therefore act on instrument signals, the riser outlet thermocouple, the delayed and lagged analyzer and the pressure transmitter. The constraint controllers read instruments too: the afterburn and bed approaches use the cyclone and bed thermocouples (0.40 min lags), the other three the metered air rate, compressor speed and slide valve differential, with no lag.

6.5.5 Model Predictive Control

A fourth Control option, the model predictive controller (MPC) of Section 1.6.3, replaces the selector and the override with one quadratic program on the four handles constraint control moves, u, F_{air}, F and T_{pre}, leaving the pressure loop in service and the riser outlet setpoint, and so the severity, to the operator [17, 18]. With only u and F_{air} free, a diagnostic copy holds the differential in scenario 10 only by leaving the riser outlet 24 °C under setpoint all run, and leaves scenario 7's bed above 730 °C for 353 of 400 min. Every $T_s = 0.5$ min it plans $N_p = 60$ steps, whose moves are held over eight blocks (1, 1, 2, 2, 4, 6, 10, 34), and applies the first move. Its model is the balances of Section 6.3 at the declared parameters and nominal disturbances, never the run's, linearized once at the design point with the pressure loop closed, with 21 states (the plant's less M and a, the five actuators, the pressure loop and six instrument lags). Its RMS residual over two hours is 1.1 to 16.4 % of the plant's response to slide valve, feed and preheat steps but 44 to 108 % on the riser outlet and bed for air steps, because the oxygen balance curves near 2 %, so ±2 kg/s of air both cool the riser outlet (1.92 and 0.73 °C, the model's ∓1.00). A steady-state Kalman filter at every scan estimates five input disturbances from the three thermocouples, the analyzer, the compressor speed and the applied commands, which makes the plan offset-free [19]. The plan minimizes Equation (6.28).

$$\mathscr{J} = \sum_{k=1}^{N_p} \left[q_T \left(\frac{e_{T,k}}{2} \right)^2 + q_O \left(\frac{e_{y,k}}{0.3} \right)^2 - c_F \frac{F_k}{\bar{F}} - c_P \frac{T_{pre,k}}{\bar{T}_{pre}} \right] + S \sum_{j=1}^{8} \sum_{m=1}^{4} \left(\frac{\Delta v_{m,j}}{s_m} \right)^2 + \mathscr{J}_s \tag{6.28}$$

where e_T and e_y are the predicted riser outlet and oxygen errors, $\bar{F}$ and $\bar{T}_{pre}$ the feed and preheat ranges, $\Delta v_{m,j}$ block j's move of handle m on scales s_m of 10 %, 2.5 kg/s, 20 t/h and 100 °C, $q_T = 1$, $q_O = 0.3$, $c_F = 5$, $c_P = 2$ and $S = 300$, and $\mathscr{J}_s$ prices each soft slack σ, in spans, at $1000(\sigma^2 + \sigma)$, the oxygen floor's at a tenth and quadratically. Hard limits are the shipped clamps and air slew, the selector's 3.6 t/h per minute feed rise, 5 % of valve and 5 °C of preheat per decision, and a valve ceiling set each decision from the measured ΔP_v by the declared law (6.16) so that ΔP_v stays above 25.96 kPa, its target at $U_{SP} = 97$ %. Soft limits are $U_i \le U_{SP}$ on the bed, cyclones, compressor and air, and $y \ge y^{SP} - 0.3$ %, a floor the MPC alone carries, all from the fourth step (2 min) so that an unavoidable early excess buys no sudden move, which hides nothing, since no case crosses a limit in its first 2 min. A dual active-set solver [20] takes 1.32 to 1.37 ms a plan in Chrome (load average 0.7 to 4.1), 0.27 % of a core at 1×. $S = 200$, between the tuning grid's 100 and 300, also meets its rule (IAE 113.4). Lowering S to 10 cuts scenario 13's riser outlet IAE from 127.4 to 55.6 °C min, but below $S = 20$ the challenge rings in a 3 min limit cycle. A rejected plan hands the handles to constraint control, whose seven controllers track the applied commands exactly with no margin, so each moves by one execution of its own law (checked to 10^{-9}, noise off).

Table 6.6 The Fourteen Scenarios and the Mechanism Each Is Built Around

#	Scenario	What it shows
1	The heat balance by hand	loop gain from one valve step
2	Air up a sixth, then down a fifth	the coke balance sets the burn rate
3	Preheat 220 to 280 to 180 °C	the only handle on the bed
4	The three loops, holding still	three plants, one unit
5	Severity to 552 °C	overcracking, then two fences in order
6	Oxygen down to 0.85 %, promoter lost	the afterburn peak
7	A heavier feed slate	the fence feed rate cannot reach
8	Metals to four times normal	perfect control hiding a failure
9	Standpipe aeration fails	the fences are not independent
10	Ask for 330 t/h	the selector and external reset
11	Ambient 15 to 38 to 12 °C	a constraint that moves
12	Severity, then aeration, under the selector	three binding fences, one structure
13	Overrides against a QP	the fences as one program
14	Graded challenge, seven hours	architecture before tuning

Table 6.7 Severity Sweep at 250 t/h, Scenario 5. Yields Are Weight Percent of Feed and Every Figure Is Read at the Settled Point. At the 552 °C Setpoint the Valve Is Fully Open and the Riser Outlet Settles at 551.8 °C.

T_{rot} (°C)	525	535	545	552
Conversion	75.00	78.17	81.01	82.72
Gasoline	46.00	46.27	45.88	45.20
Light gas and LPG	23.58	26.33	29.41	31.69

> ⧖ HISTORY **1978: Predictive Control Starts on a Cat Cracker**
>
> Model predictive control was founded in the late 1970s by two groups working independently, and both had a cat cracker in hand [21]. Richalet and coworkers introduced their method in 1978 with an explicit process model, a receding horizon and constraints on inputs and outputs. One of its three applications was a distillation column of a catalytic cracker, with two manipulated and three controlled variables and a sample time of 3 min. Cutler and Ramaker, at an oil company, used a piecewise linear model to predict the future behavior of a catalytic cracking unit, and reported the method in use there since 1974 [21]. The predictive controller of Section 6.5.5 plans every 0.5 min.

6.6 Results

Fourteen scenarios are arranged from the heat balance outward, listed in Table 6.6. Each fixes the opening mode, the control architecture, the ten tuning constants, the disturbance settings and any timed events, and is run, credited and exported as Section 1.6.4 describes.

6.6.1 Severity, and the Machine That Stops It

Scenario 5 walks the riser outlet setpoint up in three steps with the basic loops in service and the feed rate fixed at 250 t/h. The yields are in Table 6.7.

The gasoline yield peaks at about 535 °C and falls above it, because gasoline is an intermediate lump and a_4 in (6.1) is cracking it to light gas faster than a_1 is making it. Severity is therefore not a number to maximize, which is the practical content of the four-lump scheme [11]. What stops the sweep is the wet gas compressor, since every kilogram of the light gas leaves through one machine and the machine reaches 100 % speed at 540 °C on this feed rate. Past that the reactor pressure climbs because the compressor cannot take the gas away, and the rising reactor pressure eats the slide valve differential in (6.16). The 552 °C setpoint drives the valve fully open, the riser outlet stops at

Table 6.8 Constraint Handover Under the Low Selector at $U_{SP} = 97\,\%$ and a Throughput Target of 330 t/h. Every Figure Is the Settled Value After the Named Change.

	base 15 °C	hot 38 °C	evening 12 °C	severity 540 °C
Binding fence	air	air	air	ΔP
Feed rate (t/h)	284.5	261.5	287.8	245.1
Blower capacity (t/h)	230.4	213.4	232.8	230.4
Slide valve ΔP (kPa)	38.4	51.0	36.5	26.0
Conversion (%)	74.06	74.69	73.97	79.72
Gross margin ($/h)	17235	15986	17409	15746

551.8 °C, and the differential settles at 13.97 kPa, under the 14 kPa interlock, which the application flags at $t = 367$ min. Two fences arrive in order from one setpoint. On settled operating points at this throughput the gross margin peaks at $16,082 per hour at 542.4 °C, 2.5 °C past the 539.9 °C at which the compressor reaches full speed, and falls above it. The compressor limit therefore sits close to the profitable operating point, which is how a unit with its machines sized to that point behaves.

> **+ DID YOU KNOW? Cracking Harder Makes Less Gasoline**
>
> From the simulator (Table 6.7). Raising the riser outlet setpoint from 525 to 552 °C lifts conversion from 75.00 to 82.72 %, yet gasoline peaks at 46.27 % at 535 °C and falls to 45.20 %. The gain goes to light gas, which rises from 23.58 to 31.69 % and has to leave through the compressor.

6.6.2 Constraint Handover

Scenario 10 asks for 330 t/h, which the unit cannot make, and scenarios 11 and 12 then move the binding fence four times without touching a controller. Table 6.8 carries all four settled points.

The first column is the selector doing its job. The feed rate walks up and stops at 284.5 t/h with the air blower held at 97.0 % of its approach span, worth $17,235 per hour against the $15,348 of the design line-up, and the other three constraint controllers sit a fraction above the selected output with their integrals parked.

The second column is the argument for constraint control stated in one number. Warming the afternoon from 15 to 38 °C costs the blower 7.4 % of its mass capacity by (6.21), from 230.4 to 213.4 t/h, and the unit gives up 23 t/h of feed for it without anybody changing a setpoint. In the evening at 12 °C the capacity goes to 232.8 t/h and the feed rate walks back up to 287.8 t/h. Without the layer, somebody has to decide in advance how much feed the unit can take, and because the fence moves with the weather that number has to be safe on the worst afternoon of the year.

The fourth column is the fence changing hands three times. Fifteen degrees of extra severity puts the light gas yield from 23.06 to 27.84 % and takes the feed rate off the blower and onto the slide valve differential at 72 min while the reactor pressure climbs, then onto the wet gas compressor at 96 min, which holds the feed near 248 t/h once the pressures settle. From about 150 min the blower stands at only 47 % of its span, but its controller sits on the shared external reset limit, and on many scans the compressor controller's output reaches the same limit. The selector breaks such an exact tie toward the fence nearest its target, so the label names the compressor on every step from 150 to 200 min, and over the five minutes before the aeration event the feed stays between 247.6 and 248.5 t/h. Then the standpipe aeration deteriorates and the differential takes it for good at 245.1 t/h with 26.0 kPa across the valve. Three machines, three answers, one selector, and no logic changed. The structure does not know or care which constraint is active, which is why the answer to the question of what this unit is limited by is always today, at this severity, on this feed, in this weather. A drill that meets a different binding limit in each soil, in Volume 4, Chapter 7, makes the same point on one machine.

6.6.3 The Override, and Two Failures That Hold Every Setpoint

Scenario 7 switches to a heavier gas oil with the constraint layer in service, raising the carbon residue from 0.60 to 1.25 % by weight and dropping the crackability five percent. Additive coke by (6.4) raises delta coke from 0.827 to 0.99 %, the dense bed peaks at 733.3 °C against a 730 °C limit and spends 36 min above it with the cyclones reaching 742.9 °C, and the override cuts the feed preheat from 220 to 189.8 °C. The bed settles at 729.1 °C, which is exactly what the layer was asked for, the 730 °C limit less three percent of its thirty degree span. Conversion ends at 72.07 % and the margin at $13,008 per hour, and the feed rate never moves at all.

Scenarios 8 and 9 are the two failures a controlled variable cannot show. Putting four times the normal metals in the feed leaves the riser outlet temperature within 1.1 °C, the oxygen within 0.14 % and the reactor pressure within 0.12 kPa of setpoint for the first 220 min, while the slide valve opens from 45 to 72.3 %, the circulation rises from 1700 to 2023 t/h, the differential falls from 57 to 31 kPa, the bed cools from 700 to 670 °C and conversion falls three points to 72.05 % at a cost of $871 per hour. Nothing alarms. The evidence is in the actuator position and the yields, which is an argument for putting conversion, catalyst-to-oil ratio and valve position on the operator's screen as first-class numbers.

Scenario 9 is the mechanical version and it is the clearest demonstration that the fences are not independent. It runs on the basic loops at 285 t/h, and losing 26 % of the standpipe aeration takes the head out of (6.16), so the valve opens from 62.8 to 95.3 % to hold the temperature and the differential falls from 38.2 to 16.7 kPa, below the 25 kPa minimum. Opening the valve makes it worse, because the transfer line takes its share as the square of the flow. Somebody then raises the reactor pressure 25 kPa to give the compressor room, which is a sensible thing to do about a different problem, and the differential is measured against exactly that pressure. The valve pegs at 100 %, the differential settles at 11.5 kPa and stays under the 14 kPa interlock, the riser outlet collapses to 496.4 °C, conversion to 63.09 % and the margin to $12,387 per hour. Less severity cannot recover it, because the riser outlet setpoint at its 495 °C minimum still leaves 12.2 kPa. More preheat can, because 280 °C cuts the circulation from 1699 to 1550 t/h and restores 21.9 kPa at 69.2 % conversion.

6.6.4 What the Structure Is Worth

The graded challenge is seven hours with weak tuning, the feed rate fixed at 215 t/h, nobody watching any of the five limits, and five timed events, a warm afternoon at 40 min, a heavier slate at 110, the promoter falling away at 190, the standpipe deteriorating at 250 and the evening cooling at 310. The application judges the attempt at $t = 420$ min, on the case as given, and credits it when the margin is at least 8 % above the $84,532 the unit makes left alone, the riser outlet has stayed within ±2.5 °C of setpoint over the last third of the shift, the bed has spent under an hour above 730 °C and the cyclones under 5 min above 760 °C, and the slide valve differential has never fallen below 25 kPa. Table 6.9 compares leaving the unit alone against the suggested line-up over that 420 min shift.

The suggested line-up earns 12.8 % more money, 4.8 points clear of the threshold, and 88 more tonnes of gasoline, while the shipped line-up fails the rule on four counts: no margin above the reference, a worst riser outlet error of 2.52 °C, and 296 min each for the bed and the cyclones. The more important rows of Table 6.9 are the ones underneath. Leaving the unit alone puts the plenum 38 °C past the 790 °C trip at its worst and holds it above the 760 °C metal limit for 296 min of the 420 min shift, which on a real unit is a cyclone replacement. Every part of the line-up earns its place. Constraint control means the feed rate follows whichever fence is nearest instead of a number guessed in the morning. A throughput target well above what the unit can make means the selector is always the thing deciding, because a selector never adds feed and a target at or below the unit's capability leaves it nothing to take away. A constraint target of 93 % rather than 100 leaves something for the next disturbance, and there are four of them coming. 2.3 % oxygen rather than 2.0 buys margin against the afterburn of Figure 6.2 at the cost of a little air, which is the cheapest thing in the scenario against the most expensive.

Table 6.9 The Graded Challenge, Both Runs Read at $t = 420$ min. The Suggested Line-Up Is the Constraint Architecture at $U_{SP} = 93\,\%$ With a 330 t/h Target, 530 °C Severity, 2.3 % Oxygen and the Tuned Loops.

	left alone	suggested
Gross margin over the shift ($)	84532	95391
Average margin ($/h)	12076	13627
Gasoline made (t)	678	766
Feed processed (t)	1505	1692
Time above the 730 °C bed limit (min)	296	32
Time above the 760 °C cyclone limit (min)	296	0
Peak plenum temperature (°C)	828.5	757.2
Time below the 25 kPa differential (min)	0	0
Worst riser outlet error, last third (°C)	2.52	1.26
Riser outlet at the end (°C)	532.2	530.0

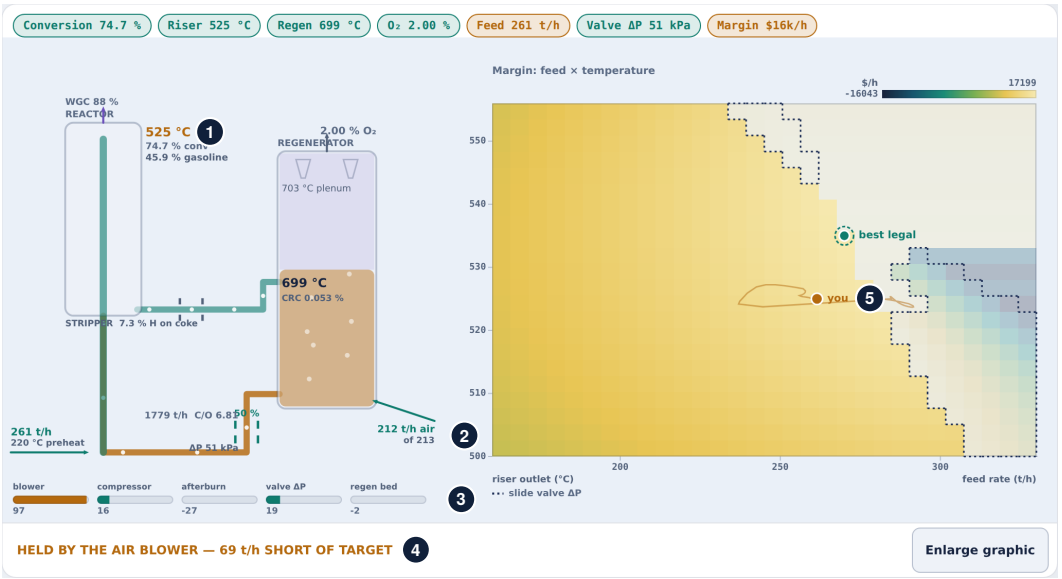


Figure 6.4 Constraint handover in scenario 11 at $t = 250$ min, 160 min after the afternoon warmed from 15 to 38 °C at $t = 90$ min and 10 min before the evening arrives, cropped to the plant view with numbered callouts. The riser outlet (1) reads 525.0 °C and the excess oxygen 2.00 %, both back on setpoint after a brief excursion when the capacity fell, while the blower capacity (2) has stepped down from 230.4 to 213.4 t/h and the feed rate has followed it from 284.5 to 261.5 t/h against an unchanged 330 t/h operator target. The air blower approach (3) sits on its 97 % target with the three other feed fences below it at 16, −27 and 19 % and the regenerator bed at −2, and the banner (4) names the blower and the 69 t/h it costs. The map puts the current point (5) below and left of the ringed best legal one, and the badge reads $16k per hour, the gross margin having fallen from $17,235 to $15,986 over the afternoon. In the full window, outside this crop, the trends cover the last three hours.

6.6.5 Comparison With the Predictive Controller

Table 6.10 runs seven cases under the shipped structure and under the MPC with only Control changed, with the same seeds, events, instruments and durations, no preview, the suggested line-up on the challenge, and 600 to 880 solves per MPC run, none rejected.

Table 6.10 Basic Loops (BL) or Constraint Control (CC) Against the MPC. Floor Is Time Below the MPC's Own Oxygen Floor.

Case (min)		IAE_T °C min	Bed min	ΔP min	IAE_O % min	Floor min	Feed t	Margin $
3 (300)	BL	317.6	0	0	58.95	51.8	1250	74988
	MPC	57.2	0	0	6.45	0.0	1227	74209
6 (440)	BL	272.5	0	0	101.00	128.6	1833	113371
	MPC	38.9	0	0	37.33	8.8	1813	112300
7 (400)	CC	97.5	35.6	0	22.29	17.2	1667	87464
	MPC	55.6	0	0	5.92	6.6	1665	87294
10 (300)	CC	18.4	0	0	3.72	0.0	1424	86198
	MPC	15.9	0	0	20.12	0.0	1430	86461
11 (400)	CC	219.5	0	0	45.49	40.2	1840	111647
	MPC	31.7	0	0	27.79	9.0	1837	111530
12, 13 (420)	CC	471.1	0	37.2	32.55	6.4	1782	113282
	MPC	127.4	0	0	15.79	0.0	1764	112271
14 (420)	CC	140.3	32.2	0	34.93	29.7	1692	95391
	MPC	127.9	0	0	49.41	2.4	1699	95754

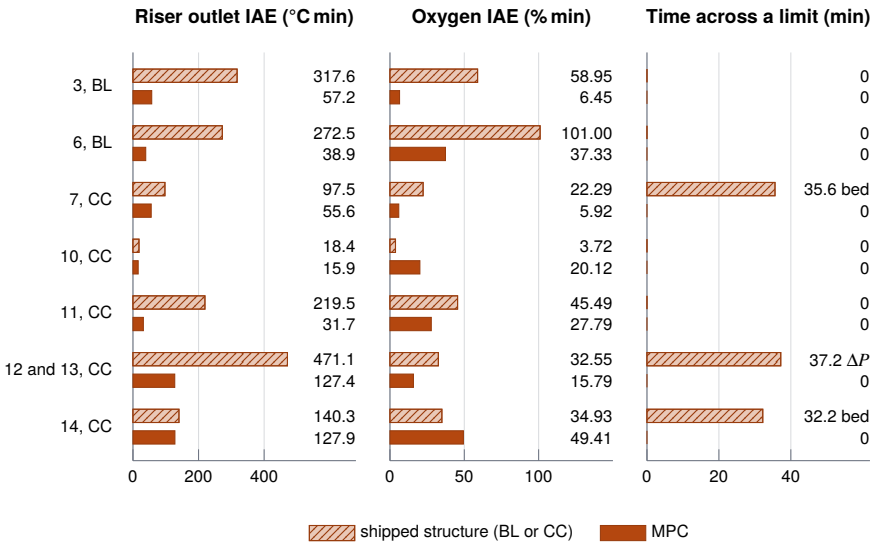


Figure 6.5 Three columns of Table 6.10 as bars, the shipped structure hatched and the MPC solid, for the seven cases in the table's order. The third panel draws the time above the bed limit or below the differential minimum, whichever the case has. Notice that the MPC shortens every bar of the first panel and has none in the third, while its oxygen bar is the longer one in cases 10 and 14. Feed and margin are not drawn.

The MPC lowers the riser outlet IAE in every case and never crosses the bed limit or the differential minimum. On scenario 7 the override cuts the preheat after the bed passes its target (peak 733.3 °C at 52.5 min) and the plan before (peak 729.1 °C at 214.7 min). Oxygen goes both ways. Below the MPC's own floor it spends 8.8 against 128.6 min on scenario 6, at an oxygen IAE of 37.33 against 101.00, but on that scorecard IAE it loses scenario 10 (20.12 against 3.72 % min) and the challenge (49.41 against 34.93), holding the oxygen under setpoint inside its floor to keep the air off its fence and the feed up. Figure 6.5 draws these three comparisons as bars.

It also earns less on scenarios 12 and 13, 1764 against 1782 t and $112,271 against $113,282, with 84 percent of the gap between the severity step and the aeration loss (581.6 against 596.4 t from 60 to 200 min, Figure 6.6). There the selector cuts the feed later, 263 against 248 t/h at 90 min, and its differential is under 25 kPa for 23.7 of its 37.2 min, while its compressor and blower reach their full approach span at 69.0 and 73.4 min without passing it, as the MPC's compressor does at 70.0 min.

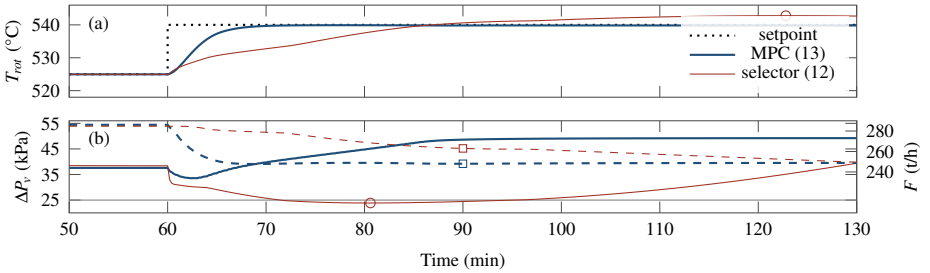


Figure 6.6 Scenarios 12 (selector, thin lines) and 13 (MPC, thick lines). (a) True riser outlet against the dotted setpoint, the selector's peak of 542.8 °C at 122.8 min circled. (b) Slide valve differential (solid, left) with its 25 kPa minimum and the selector's 23.9 kPa at 80.6 min circled, and feed rate (dashed, right), 263 and 248 t/h at the squares.

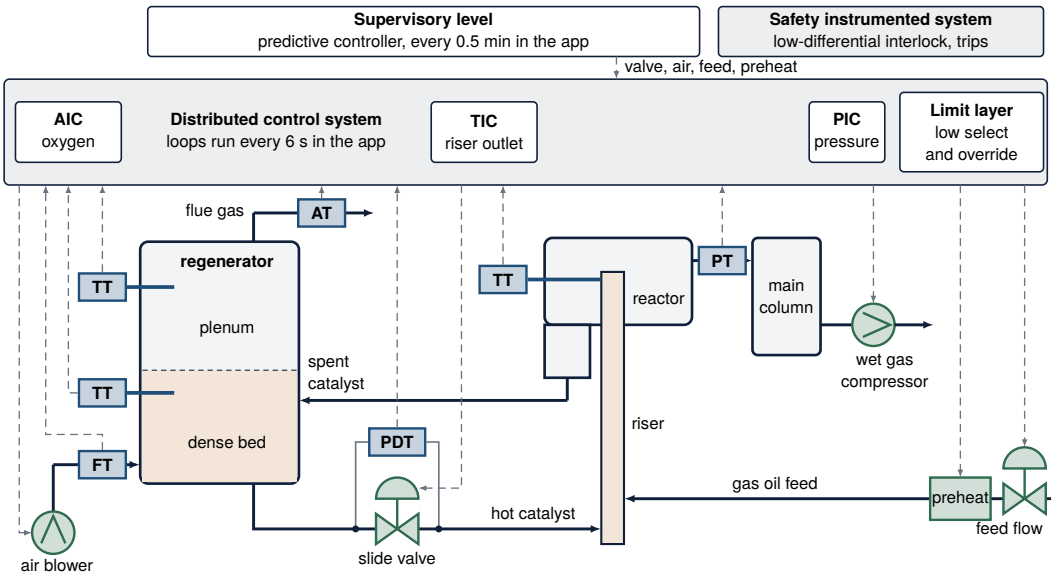


Figure 6.7 The loops of the cat cracker with their hardware. Lettered boxes are sensing devices, namely thermocouples with temperature transmitters (TT), the oxygen analyzer (AT) behind its sample line, the pressure transmitter (PT), the differential pressure transmitter (PDT) across the slide valve and the air flow transmitter (FT). Valve and machine symbols are final elements. Dashed arrows are signals. Notice that the limit layer reads the air flow, the two regenerator thermocouples, the slide valve differential and the compressor speed, and writes only the feed rate and the preheat, so one feed rate answers to whichever machine is nearest its limit. The low-differential interlock sits outside the control system. Not to scale, and not a piping or safety diagram.

On scenarios 3 and 6 it gives up 23 and 20 t by cutting a feed that basic loops never move.

6.7 Technology in Practice

Figure 6.7 draws the loops of Section 6.5 with the hardware a refinery installs.

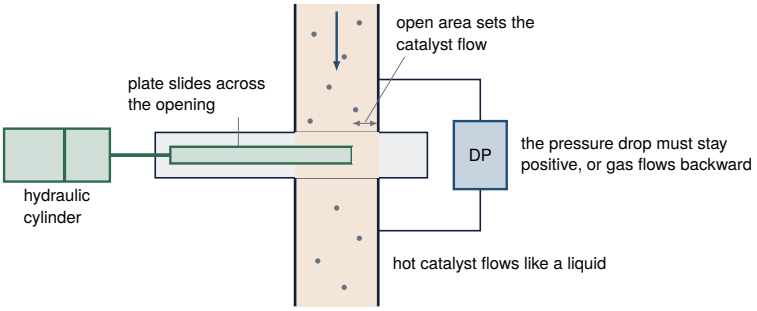


Figure 6.8 Principle of a slide valve in a catalyst standpipe. Notice that the plate only narrows the opening. The flow through it is driven by the pressure difference that the transmitter reads, which is why Equation (6.16) treats the valve as a restriction and not as a flow source.

What Measures It

The riser outlet, the dense bed and the plenum are read by thermocouples. The flue gas oxygen is read by an analyzer. One kind is a zirconium oxide cell, a solid electrolyte between platinum electrodes with flue gas on one side and reference air on the other, whose voltage is converted to an oxygen concentration [22]. An analyzer away from the process needs a sample system, and analyzer problems quite often start there. A plugged line leaves the analyzer reporting stale sample with no way of knowing it [22]. The delay θ_{O_2} and the lag τ_{O_2} of Table 6.2 stand for such a line and its analyzer. Refineries read catalyst levels and densities by differential pressure, because fluidized catalyst exerts a hydraulic pressure like a liquid, and catalyst that is not fluidized exerts none and behaves like fine sand [23]. That is the physical content of the aeration multiplier a_{er} in Equation (6.16).

> **TECHNOLOGY Sensor: Thermocouple in a Thermowell**
>
> Two wires of different metals joined at one end produce a voltage roughly proportional to temperature, and that junction is a thermocouple [22]. Thermocouples are less accurate than resistance thermometers but more rugged and wider in range. The Chromel and Alumel type is listed for about 32 to 2300 °F, an upper end near 1260 °C. The sensor sits in a thermowell, a pressure-tight sheath that lets it be replaced under pressure. The mass of the well adds a first-order lag, and an air gap at its tip adds a second [22]. The lags of 0.25 min at the riser outlet and 0.40 min in the bed stand for the well. An open circuit, the most common failure, can read as a wildly varying temperature, so the transmitter is set to drive its reading high or low, whichever is safer [22].

What Moves It

The regenerated catalyst slide valve controls the flow of hot catalyst to the riser [24]. The model carries it as the restriction of Equation (6.16) behind the travel lag τ_s, and Figure 6.8 shows its principle. The other handles are the air blower, the wet gas compressor, the charge pump and the preheat train, each a first-order lag in Equation (6.18).

> **TECHNOLOGY Final Element: Catalyst Slide Valve**
>
> Slide valves are a feature of most cat cracker designs. They are gate-type valves that open and close by sliding a disc across an orifice [23]. Their parts must move freely at high temperature, so the design keeps clearances, typically 0.04 to 0.06 inches wide, and a slide valve leaks a little even when new and fully closed [23]. The service is severely erosive, because the flowing catalyst is abrasive and wears the internals [24]. The differential pressure across the valve should always be positive. If it falls too low to sustain circulation, catalyst and gas can reverse and mix air with hydrocarbon [23]. At the refinery of that report the control system began to close both slide valves at 2 psi, which is close to the 14 kPa interlock of this chapter.

Where the Controller Runs

A refinery runs such loops in a distributed control system, in which a supervisor sends set points to distributed controllers that act through actuators from sensor feedback, with cycle times from milliseconds to minutes [25]. The app executes its controllers every 6 s and its predictive controller every 0.5 min, inside that range. The three regulatory loops, the four constraint controllers, the low selector and the preheat override would sit in the distributed controllers as function blocks. A controller with integral action that a selector has de-selected tends to wind up. One remedy feeds the selected output back to the integral term of each block, which pneumatic practice called external reset [22]. That is Equation (6.26). The low-differential interlock and the temperature trips belong to a safety instrumented system of sensors, logic solvers and final elements, designed to IEC 61511 and often independent of the control system [25].

> **IN PRACTICE** **A Slide Valve Is Not a Block Valve**
>
> What the model's slide valve never does is wear. Two federal investigations of cat cracker explosions, in 2015 and 2018, each found a spent catalyst slide valve eroded after five to six years of service, closed and still passing gas [23, 24]. In the first case the valve had been tested by closing it part way, which proved the mechanism and not the seal, and its position said nothing about the catalyst level above it. The stated lesson was to measure a safeguard directly, here by differential pressure, with a minimum that shuts the unit down [24]. The second report asked that API RP 553, the recommended practice for refinery control valves, cover the body of the slide valve [23].

What Is Modernizing It

Three developments bear on this chapter. The first is predictive control. Its early applications in the petrochemical industry had sampling times of several hours, and today these have largely come down to minutes and seconds. It often runs as a supervisory layer above the PID loops [21], where the app's predictive controller writes its four handles directly. The second is the safety system. Historically it stood alone, and some modern ones are designed to communicate with the control system [25]. The third is guidance. After the explosions of the In practice box the investigators asked in 2022 for a public technical publication on the safe operation of cat crackers, with recommended monitoring during startup, shutdown and standby [23], modes this chapter's simulator does not run.

6.8 Limitations

The application is not a yield model and not a regenerator design tool. With four lumps, one well-mixed bed and a correlated afterburn it cannot support a claim about a real unit's octane, its cyclone temperatures or how much feed it could take tomorrow. The kinetic constants were fitted so that this unit sits at the design point in Section 6.3 and reproduces a gasoline maximum, not against any real unit's yield data. Delta coke and hydrogen on coke behave correctly in direction and roughly in magnitude, which is what the heat balance needs, and nothing beyond that should be read from them. The steady-state solver finds the hot branch by continuation, and it declines to converge on the cold branch below about 1.1 % oxygen with the promoter lost, which is why the dashed curve in Figure 6.2 stops there rather than continuing to zero. That non-convergence is honest about the model and unhelpful as a picture of the plant, since a real unit does have an operating regime down there.

Two omissions matter more than the rest. There is no standpipe fluidization model, so a low-differential event ends at the interlock rather than in the flow reversal that makes it dangerous, and there is no catalyst cooler, which is the piece of equipment a real resid unit installs to answer precisely the problem scenario 7 poses. The individual assumptions and where each one breaks are set out in Section 6.3.12. One linearization at 250 t/h leaves the MPC without the afterburn's turn below 1 % oxygen, the oxygen balance's curvature and the gains that move with throughput.

6.9 Discussion

The economics of the constraint layer can be read directly, because the application searches the throughput and severity plane every five minutes for the best point that breaks no limit and carries the answer beside the achieved margin. At the design conditions that best legal point is 290 t/h at 530 °C, worth $17,982 per hour against the $15,348 the design line-up collects, so the fixed line-up leaves $2,634 per hour on the table. The selective layer of Figure 6.3 recovers $1,887 per hour of that gap in the base case of Table 6.8 without any optimizer in the loop. What remains is the severity dimension, which no controller in this application moves, and that is the natural boundary between a regulatory constraint layer and a real-time optimizer.

The value of a point of conversion sets the scale for everything else. Raising severity from 525 to 528.2 °C at fixed throughput buys 1.06 points of conversion and $248 per hour, about $234 per point, which is close to two million dollars a year on a unit of this size. That is why a cat cracker carries advanced control at all, and it is why the reactor and regenerator were among the first units in a refinery to receive a multivariable constraint controller [15].

The model also supports methods beyond the implemented structures. The fifteen-state plant with five explicit limits and a written objective is the natural form for model predictive control. The predictive controller of Section 6.5.5 already absorbs the selector and the override into one optimization with the constraints written in, but it leaves severity to the operator, and a layer that moves severity and throughput together rather than one at a time is the next step [6]. The choice of which variables to hold and which to let float is the partial control question that this unit posed [9], and the application makes it concrete because holding all five limits with two actuators is not possible and the student can watch which one gives. Moving horizon estimation on the same equations would recover the catalyst activity and the standpipe aeration from the slide valve position and the yields, turning the two invisible failures of scenarios 8 and 9 into soft sensors.

> **OPEN CHALLENGE** **A Feed That Is Too Easy to Crack**
>
> Light feeds strain the heat balance of this chapter from the other side. A 2015 review names the growing use of tight oils as a challenge for cat crackers [1]. These oils are low in carbon residue and aromatics, which makes the coke deposit, and with it the heat balance of the unit, harder to control. In the review's words the feed is in some respects too easy to crack. Its high naphtha and gas yields can also disturb the columns and gas plant downstream and limit throughput. Process changes and added heavier feed answer this only in part [1]. Progress would be a unit that holds its heat balance from such a feed to the heavy slate of scenario 7.

6.10 Conclusion

This chapter studied a fluid catalytic cracking riser and regenerator, a plant that supplies its own reaction heat and then has to be operated against limits it moves itself. The simulator integrates fifteen states, with a four-lump riser marched against catalyst contact, a coke burn and a flue gas oxygen balance solved together in closed form, and a slide valve whose implicit equation is solved rather than iterated. A three percent slide valve move in manual operation raises the riser outlet 7.17 °C against the 3.42 °C the extra catalyst alone accounts for, a heat balance loop gain of 0.523 and an amplification of 2.09, and only 29 % of that answer has arrived at two minutes. The settled coke burn rate moves 0.3 % for a 9.3 % increase in air while the carbon on the regenerated catalyst falls by a third, and the afterburn peaks at 57.8 °C above the bed at 0.70 % oxygen rather than at either end of the sweep. A low selector over four constraint controllers with external reset holds 284.5 t/h against a 330 t/h request and gives back 23 t/h when the afternoon warms 23 degrees, while a fifth controller on the feed preheat parks the dense bed at 729.1 °C on a heavy feed the shared handle cannot help with. Over the 420 min graded shift with five timed events the structure earns $95,391 with no time above the cyclone limit against $84,532 and 296 min above it. A model predictive controller on the same handles crosses no limit in seven cases, at 1.0 % less feed on scenario 13 and a worse oxygen

IAE on two. A fuller model would add a catalyst cooler as a third actuator on the heat balance, a standpipe fluidization model, a fractionator so that the wet gas rate stops being one yield lump, and an extension of the shipped MPC to optimize severity as well as throughput.

KEY POINTS

- A cat cracker has no furnace. The coke it burns is its only heat, so every move comes back to it amplified, 2.09 times here, and slowly.
- Only 29 % of the riser outlet response has arrived at two minutes, so a loop tuned on that reading comes out too aggressive.
- More air lowers the carbon on the catalyst but not the settled burn rate, which the riser sets.
- A low selector with external reset lets four constraint controllers share one feed rate without windup.
- The regenerator bed does not respond to feed rate, so its override acts on the feed preheat.

Where else this applies. A selector whose losing integrators track the applied command guards the fuel valve in Volume 4, Chapter 3, a feed rate pushed against limits that change returns in Volume 4, Chapter 4 and Volume 4, Chapter 7, and amplification by a recycle in Volume 6, Chapter 7.

References

[1] E. T. C. Vogt and B. M. Weckhuysen, "Fluid catalytic cracking: recent developments on the grand old lady of zeolite catalysis," *Chemical Society Reviews*, vol. 44, no. 20, pp. 7342–7370, 2015.

[2] W. M. Edwards and H. N. Kim, "Multiple steady states in FCC unit operations," *Chemical Engineering Science*, vol. 43, no. 8, pp. 1825–1830, 1988.

[3] A. Arbel, I. H. Rinard, R. Shinnar, and A. V. Sapre, "Dynamics and control of fluidized catalytic crackers. 2. multiple steady states and instabilities," *Industrial & Engineering Chemistry Research*, vol. 34, no. 9, pp. 3014–3026, 1995.

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

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

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

[7] A. A. Avidan and R. Shinnar, "Development of catalytic cracking technology. a lesson in chemical reactor design," *Industrial & Engineering Chemistry Research*, vol. 29, no. 6, pp. 931–942, 1990.

[8] A. Arbel, Z. Huang, I. H. Rinard, R. Shinnar, and A. V. Sapre, "Dynamic and control of fluidized catalytic crackers. 1. modeling of the current generation of FCC's," *Industrial & Engineering Chemistry Research*, vol. 34, no. 4, pp. 1228–1243, 1995.

[9] A. Arbel, I. H. Rinard, and R. Shinnar, "Dynamics and control of fluidized catalytic crackers. 3. designing the control system: Choice of manipulated and measured variables for partial control," *Industrial & Engineering Chemistry Research*, vol. 35, no. 7, pp. 2215–2233, 1996.

[10] V. W. Weekman and D. M. Nace, "Kinetics of catalytic cracking selectivity in fixed, moving, and fluid bed reactors," *AIChE Journal*, vol. 16, no. 3, pp. 397–404, 1970.

[11] L.-S. Lee, Y.-W. Chen, T.-N. Huang, and W.-Y. Pan, "Four-lump kinetic model for fluid catalytic cracking process," *The Canadian Journal of Chemical Engineering*, vol. 67, no. 4, pp. 615–619, 1989.

[12] H. de Lasa and A. Errazu, "Ignition of a fluidized bed catalytic cracking regenerator: Freeboard region influence," in *Fluidization*. Boston, MA: Springer US, 1980, pp. 563–570.

[13] M. Hovd and S. Skogestad, "Procedure for regulatory control structure selection with application to the FCC process," *AIChE Journal*, vol. 39, no. 12, pp. 1938–1953, 1993.

[14] R. C. McFarlane, R. C. Reineman, J. F. Bartee, and C. Georgakis, "Dynamic simulator for a model IV fluid catalytic cracking unit," *Computers & Chemical Engineering*, vol. 17, no. 3, pp. 275–300, 1993.

[15] P. Grosdidier, A. Mason, A. Aitolahti, P. Heinonen, and V. Vanhamäki, "FCC unit reactor-regenerator control," *Computers & Chemical Engineering*, vol. 17, no. 2, pp. 165–179, 1993.

[16] R. Sadeghbeigi, *Fluid Catalytic Cracking Handbook*, 3rd ed. Oxford, UK: Elsevier, 2012.

[17] 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.

[18] 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.

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

[20] 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.

[21] M. Schwenzer, M. Ay, T. Bergs, and D. Abel, “Review on model predictive control: an engineering perspective,” *The International Journal of Advanced Manufacturing Technology*, vol. 117, no. 5–6, pp. 1327–1349, 2021.

[22] 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.

[23] U.S. Chemical Safety and Hazard Investigation Board, “FCC unit explosion and asphalt fire at Husky Superior refinery, Superior, WI, incident date April 26, 2018,” U.S. Chemical Safety and Hazard Investigation Board, Washington, DC, Investigation Report 2018-02-I-WI, 2022.

[24] ——, “ExxonMobil Torrance refinery electrostatic precipitator explosion, Torrance, California, incident date February 18, 2015,” U.S. Chemical Safety and Hazard Investigation Board, Washington, DC, Investigation Report 2015-02-I-CA, 2017.

[25] 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.

6.11 Exercises

6.1 Concept check · 2 min

In scenario 1 the regenerated catalyst slide valve opens from 45 to 48 % in Manual with the air, the feed and the preheat held. The regenerator dense bed ends 5.25 °C hotter, but for the first few minutes it falls, by 0.46 °C at three minutes. Why does it fall first?

(a) The carbon on the regenerated catalyst drops at once, so the bed burns less coke for a few minutes

(b) The bed thermocouple lags the bed by a few minutes and shows the old temperature drifting

(c) The extra catalyst carries heat out of the bed before the extra coke it brings back has arrived to burn

(d) The higher circulation lowers the bed level, so the thermocouple reads the cooler dilute phase

6.2 Concept check · 2 min

The heavier feed of scenario 7 takes the regenerator dense bed over its 730 °C limit while the riser outlet is held at 525 °C. Which handle brings the bed back, and why?

(a) The feed rate cut, because less feed sends less coke to the regenerator to burn

(b) The fresh catalyst addition raised, because a more active catalyst lays down less coke per tonne

(c) The slide valve opened, because more circulation carries more of the bed’s heat to the riser

(d) The feed preheat lowered, because more catalyst must circulate and the coke is spread over more of it

6.3 Concept check · 1 min

In scenario 10 the air blower controller holds the feed rate at 284.5 t/h. What are the other three constraint controllers doing?

(a) Each sits on its upper limit, the selected feed plus 0.10 kg/s, with its integral parked there

(b) Each has wound its integral up to the top of the feed range and will unwind when its fence comes near

(c) Each has been put in manual by the selector and restarts from the selected feed when its fence is reached

(d) Each tracks the selected feed exactly, so all four outputs are equal and the selector picks the first

6.4 Theory and calculation · 7 min

In scenario 1 the slide valve opens from 45.0 to 48.0 % in Manual. Solving the riser energy balance alone, with the regenerator frozen, gives a riser outlet rise of 3.42 °C. The unit settles 7.17 °C higher, and 2.11 °C of that has arrived two minutes after the move.

(a) Find the amplification A and the loop gain L, and the factor by which Equation (1.8) stretches a time constant inside the loop. *Checkpoint:* $A = 2.10$, $L = 0.523$, a factor of 2.10.

(b) For a one percent valve step the loop gain is 0.534 at 525 °C and 0.550 at 535 °C. Find the amplification at each, and the percentage by which the settled gain grows between them. *Checkpoint:* 2.15 and 2.22, a growth of 3.6 %.

(c) An engineer fits a process gain to the first two minutes. Find that gain in °C per percent of valve, the settled gain, the fraction of the answer seen and how many times too weak the fitted plant is. *Checkpoint:* 0.70 against 2.39 °C per percent; 29 %, 3.4 times too weak.

6.5 Theory and calculation · 6 min

The blower capacity follows Equation (6.21) with $F_{air,0}^{max} = 64.0\,\mathrm{kg/s}$ at 15 °C. Its approach span is 10.0 kg/s, and one span of the blower fence is worth 48.5 t/h of feed with the riser outlet and the oxygen held.

(a) Find the capacity in t/h at 15, 38 and 12 °C, and the fraction of mass capacity lost from 15 to 38 °C. *Checkpoint:* 230.4, 213.4 and 232.8 t/h; 7.4 %.

(b) With the blower holding the feed at a fixed approach, estimate the feed given up in the afternoon and gained in the evening, and compare with the measured 284.5, 261.5 and 287.8 t/h of Table 6.8. *Checkpoint:* 22.9 t/h lost and 3.3 gained, against 23.0 and 3.3 measured.

6.6 Theory and calculation · 8 min

Use Equation (6.16) with $K_s = 1.39032$, $K_{line} = 62\,\mathrm{kPa}$, $F_{cat,0} = 6.8F_0$ and $F_0 = 69.444\,\mathrm{kg/s}$, $P_{rgn} - P_{rx} = 40\,\mathrm{kPa}$, $H_{drop} = 9.5\,\mathrm{m}$, $M = 165\,000\,\mathrm{kg}$, $\rho_b = 480\,\mathrm{kg/m^3}$, $A_{rgn} = 78.5\,\mathrm{m^2}$ and $g = 9.81\,\mathrm{m/s^2}$.

(a) Find h_b, ΔP_{av}, the circulation at $u = 45\,\%$ in t/h and the differential ΔP_v across the valve. *Checkpoint:* 4.38 m, 119.0 kPa, 1700 t/h and 57.0 kPa.

(b) The standpipe aeration falls to $a_{er} = 0.74$. Find ΔP_{av}, and the opening and the differential that hold the design circulation. *Checkpoint:* 98.4 kPa; 56.3 % open with 36.4 kPa.

(c) At $a_{er} = 0.74$, find the circulation and the differential at 70 and 100 % open. *Checkpoint:* 1824 t/h with 27.1 kPa; 1967 t/h with 15.4 kPa.

6.7 App challenge · 8 min

Open scenario 1, “The heat balance by hand”, which starts in Manual, with the default seed. Set the speed to 20× and press Run. At about 10 min move Regenerated catalyst slide valve from 45 to 48 %, run on to at least 135 min, and export the CSV. Let t_s be the time of the last row in which `slide_cmd_pct` still reads 45.0. Read the change of each column from the row at t_s.

(a) Read `rot_C` at $t_s + 2$ and `trgn_C` at $t_s + 3$ minutes. *Checkpoint:* $+2.11 \pm 0.10$ and $-0.46\,°\mathrm{C} \pm 0.05$.

(b) Read both at $t_s + 120$ minutes, and the peak of `crc_wtpct` with its time after t_s. *Checkpoint:* $+7.05 \pm 0.05$ and $+5.09\,^{\circ}\mathrm{C} \pm 0.05$; 0.0577 % at 13 min.

(c) Take the settled rise of 7.17 °C from Section 6.4. What fraction of it has arrived at 2 and at 120 minutes, and which state sets the slow part? *Checkpoint:* 29 % and 98 %; the regenerator bed.

6.8 App challenge · 8 min

Open scenario 7, "A heavier feed, and the fence feed rate cannot reach", with the default seed. It runs on Constraint control, and the heavier feed arrives at 20 min. Set the speed to 20×.

(a) Run to 120 min and export the CSV. Read the peak of `trgn_C` and its time, the lowest `tpre_C` and its time, `trgn_C` and `tpre_C` at 120 min, and the minutes above the bed limit in the last row of the column `metric_over_rgn_min`. *Checkpoint:* 733.3 °C at 52.5 min; 186.5 °C at 81.7 min; 729.0 and 189.2 °C; 35.6 min.

(b) At any time between 120 and 150 min, switch the mode to Manual and set Feed rate to 200 t/h. Run on 60 min and read `trgn_C` and `rot_C` 15 and 60 minutes after the cut. *Checkpoint:* Bed 733.7 to 733.8 °C and riser 555.0 to 555.1 °C at 15 minutes; bed 742.1 to 742.2 °C and riser 561.0 to 561.1 °C at 60.

6.9 App challenge · 8 min

Open scenario 11, "The fence moves", with the default seed. It runs on Constraint control with a 330 t/h target, the afternoon warms to 38 °C at 90 min and the evening cools to 12 °C at 260 min. Run 400 min at 20× and export the CSV.

(a) Read `feed_tph`, `air_max_tph` and `margin_usd_h` at 85, 250 and 400 min, and the `fence` column at each. *Checkpoint:* 284.5, 230.4, 17235; 261.5, 213.4, 15986; 287.8, 232.8, 17408; air blower each time.

(b) Find the lowest `feed_tph` between 90 and 130 min and its time, and the highest `util_air_pct` in the same window. *Checkpoint:* 236.7 t/h at 92.0 min; 143.7 % at 90.1 min.

(c) Compute the feed given up per tonne of lost air capacity in the afternoon, and explain why the feed falls below its new level first. *Checkpoint:* 1.35 tonnes of feed per tonne of air capacity.

6.10 Application area · 3 min

An operator on a unit with the promoter in service reduces the excess oxygen setpoint from 2.00 to 0.85 % to save blower power, and reports that the regenerator dense bed has cooled from 700 to 686.6 °C.

(a) Is the cooler bed good news? Use Figure 6.2. *Checkpoint:* No: the monoxide leaving the bed burns above it.

(b) The promoter feeder then plugs. Why is adding air no longer a complete answer on this unit? *Checkpoint:* The blower saturates and 8.7 °C of afterburn remains.

6.11 Application area · 4 min

Section 6.9 gives the economics of the unit at its design conditions.

(a) One point of conversion is worth $234 per hour at fixed throughput. What is it worth in a year of 8760 hours on stream? *Checkpoint:* About $2.05 million.

(b) The best legal point is worth $2,634 per hour more than the design line-up, and the selector recovers $1,887 per hour of it. What fraction is recovered, what remains per year, and which layer would have to recover it? *Checkpoint:* 72 %; about $6.5 million a year; an optimizer that moves severity.

6.12 Review the agent · 5 min

An AI assistant is asked to tune the riser outlet loop. It runs scenario 4, “The three basic loops, holding still”, switches to Manual, steps the slide valve from 45.0 to 46.0 % and reads the riser outlet: up 0.47 °C at one minute and 0.74 °C at two, when the trend looks flat. It reports a first-order plant with $K_p = 0.74\,^{\circ}\mathrm{C}$ per percent and $\tau = 1$ min, and proposes lambda tuning, $K_c = \tau/(K_p\lambda)$ with $\lambda = 1$ min and $\tau_I = \tau$, which gives $K_c = 1.35$ percent per °C, “faster than the shipped 0.55”. Find the flaw and say how you would check it.