INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 12 Control in Petrochemicals and Polymers

CHAPTER 2

An Ethane Cracking Furnace Simulator for Balancing Schemes and Constrained Severity

Pages 24–47

Edited by John D. Hedengren

How to cite this chapter. "An Ethane Cracking Furnace Simulator for Balancing Schemes and Constrained Severity," in *Control in Petrochemicals and Polymers*, John D. Hedengren, Ed., vol. 12 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 24–47.

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

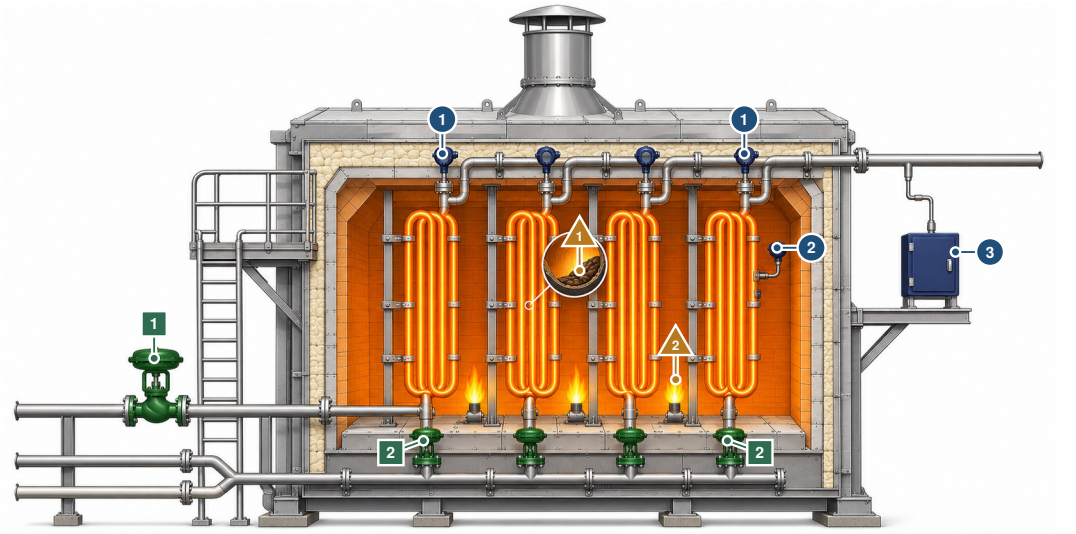
The app for this chapter
https://apmonitor.com/apps/ethylene-furnace/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v12-petrochemicals-polymers/#ethylene-furnace

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

CHAPTER 2

An Ethane Cracking Furnace Simulator for Balancing Schemes and Constrained Severity



The controller moves 1 fuel valve to the burners, which sets the firebox temperature
2 four pass flow valves, which move feed between the passes

The controller reads 1 four coil outlet temperatures 2 tube metal temperature (one drawn, every pass has one)
3 conversion analyzer on the effluent

What disturbs it ▲1 coke growing inside the tubes ▲2 a degraded burner that makes one pass coke faster
▲ changes in feed rate and dilution steam

Notice that the outlet heads read the gas on one side of the coke, while the limit that ends the run is on the tube metal on the other side. Illustration, not to scale.

IN PLAIN TERMS **Four Hot Tubes and One Hidden Limit**

The main idea. A cracking furnace heats ethane gas in four steel tubes until it breaks apart into ethylene. Each tube slowly grows a crust of carbon, called coke, between the gas and the metal, and the furnace must stop now and then to burn it out.

The challenge. The usual control evens out how hot the gas leaves each tube, yet a crusted tube reads cool there while its metal runs hot. Evening out the wrong reading can end the run sooner than doing nothing.

Why it matters. The whole furnace must stop when its one hottest tube reaches the limit of its metal.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	This page, the History boxes in Sections 2.2 and 2.5.5, the Did you know? boxes in Sections 2.3 and 2.6, the Try it boxes in Sections 2.4 and 2.6, and Key points in Section 2.10. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 2.6 with Figure 2.2, Section 2.7, Section 2.8 with its Open challenge box, and the second Open challenge in Section 2.9.

The control question A cracking furnace comes off line when one tube reaches its metal limit, yet most pass balancing schemes equalize the coil outlet temperatures on the other side of the coke. What does balancing the measured variable do to the run length, and which measurement can find the severity that makes the most ethylene per elapsed hour?

Learning objectives

- Explain why a coked pass reads cold at its outlet and hot at its skin, and compute the metal spread that equal outlets leave.
- Derive the run length from the self-accelerating coking law and compare it with a straight line through the first day.
- Compare balancing on coil outlet, on tube metal and not at all on a furnace with one degraded burner, from measured run lengths.
- Locate the severity optimum from the on-stream rate, and explain why neither the conversion nor the ethylene rate can find it.
- Judge where a predictive controller's tube metal ceiling belongs against the override and the decoke trigger.

Prerequisites Sections 1.4.2 (coking and run length), 1.4.8 (back-off from a limit), 1.6.2 (velocity form) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ethylene-furnace/

See also Volume 6, Chapter 4 for fouling that walks one effect toward a hard temperature limit under an override. Volume 9, Chapter 4 for an estimator that carries an unmeasured bias, and Volume 9, Chapter 8 for an optimum that sits on a constraint that moves with the weather. Volume 4, Chapter 2 for a protective override that acts only near its limit.

2.1 Introduction

A balancing scheme equalizes what it measures. On a cracking furnace that is usually the coil outlet temperature, because outlet thermocouples are cheap and reliable and every furnace has them. The constraint that ends the run is on the tube metal, on the other side of a coke layer that is growing at a different rate in every pass. Measured on the furnace modeled here, with four passes starting at coke resistances of 1.0, 2.0, 3.0 and 4.5×10^{-4} m^2 K/W and the flows trimmed until every outlet matches to within half a degree, the tube metal spread is still 16.1 °C. Twenty days later, with the outlets still matched, it is 34.0 °C. That spread is run length. The box comes off when the hottest tube reaches 1100 °C, and three of the four passes arrive there with room they never used.

The misconception the simulator corrects is that balancing a measured variable balances the process. Coke sits between the metal and the gas, so it makes the metal hotter for a given flux and the gas colder. A coked pass therefore reads low on the instrument the scheme uses and high on the instrument the constraint is written on, and the two disagree about which pass needs attention. Measured on a furnace whose fourth pass cokes 1.8 times as fast as the others because of a degraded burner, balancing on coil outlet temperature gives a run of 841 h against 888 h for no balancing at all. The scheme is working perfectly, the instrument is accurate, the tuning is correct, and the result is a shorter run than doing nothing, because the scheme is moving feed away from the pass that needed more of it.

The trap is general. Parallel heat exchangers balanced on outlet temperature rather than on approach, parallel reactors balanced on conversion rather than on bed temperature, and any set of parallel paths where the limit lives on one side of a resistance and the instrument on the other, all have it. The cracking furnace is the clearest example because the resistance grows during the run and the growth is self-accelerating: the coking rate rises by an e-fold every 30 °C of tube metal temperature, so a pass that starts slightly ahead pulls further ahead for the rest of the campaign. Coke formation and its effect on run length prediction is an active experimental and modeling subject [1, 2].

This chapter works with an interactive cracking furnace whose coking rate, radiant coefficient and wall resistance are hidden, so that the divergence must be observed, the balancing variable chosen against a measured run length, and a severity optimum found that neither the production trend nor the conversion trend can locate. The cracking kinetics, the coking correlation and override control are standard. Section 2.2 gives the domain history. Section 2.3 writes out the four pass radiant box, its coking dynamics and its instrument chain in the detail needed to rebuild it, and Section 2.4 derives the metal spread a balancing scheme cannot remove and the positive feedback that sets the run length. Section 2.5 describes the firing and severity loops, the two balancing schemes, the override, the two decoke policies and a model predictive controller that the application offers beside its loops (Section 2.5.5). Section 2.6 measures the scenario sequence, which contains a balancing result that is worse than no control and a severity optimum invisible to every continuous measurement on the unit, and compares the predictive controller with the loops on three scenarios (Section 2.6.6). Section 2.7 names the hardware and the practice behind the model. Sections 2.8, 2.9 and 2.10 give the limitations, the results that carry beyond cracking furnaces and a summary. The application exports its complete time history, as every simulator of this volume does (Section 1.6.4).

2.2 Background

Steam cracking is the largest chemical process in the world by energy consumption, and ethylene is its product. Ethane and dilution steam are heated in a fired tubular coil to about 845 °C in well under a second, and the cracking reactions that make ethylene also make coke, which deposits on the inside of the tube. Every furnace therefore runs a campaign and then comes off line for a decoke, and the campaign length is one of the numbers an ethylene plant is managed by.

> ⧖ HISTORY **1918 to 1921: Ethane Cracked in a Fired Tube**
>
> In the 1910s the chemist George O. Curme Jr. worked at an industrial research institute in Pittsburgh. During the First World War the government, looking for another source of ethylene, asked him about making it from gas oil. Curme suggested ethane instead, because it requires less energy to crack. A small plant was begun at Buffalo, New York, in the summer of 1918. It made its first ethylene in January 1920, and a fire destroyed it later that year. The work moved to a site with plenty of natural gas near Charleston, West Virginia [3]. There a small works-scale plant started in the summer of 1921, and its new tubular cracking process was demonstrated in three-inch tubes inside a gas-fired refractory furnace [4]. Cracking furnaces today still heat the feed in tubes fired from outside by gas or oil burners [5].

Three features of the unit set its control problem and each has its own literature. The first is the coking itself. Coke thickness is not measured, the rate depends on tube metal temperature, on severity and on dilution steam, and run length prediction from coke properties is still being refined against experiment [1]. The second is severity. Higher coil outlet temperature gives higher conversion and lower selectivity, so the ethylene rate and the ethylene yield move in opposite directions, and the effect of coil outlet temperature on furnace performance has been studied directly [6]. The third is the decoke itself, which takes the furnace out of service for about thirty hours and has been optimized as a process in its own right [2]. Scheduling the decokes of a bank of furnaces against a plant demand is a cyclic scheduling problem of the same family as heat exchanger cleaning scheduling [7].

Control developed from the firing. The earliest arrangement is a firebox temperature controlled by fuel gas, with the coil outlet read and corrected by an operator. Coil outlet temperature control on the firing came next and is universal now. Pass balancing arrived with individual pass flow control, and it is almost always configured on coil outlet temperature for the instrument reason given above. Tube skin thermocouples exist on most furnaces and are treated as protection rather than as control measurements, which is the habit this application is built to question. Above the regulatory layer, severity control on an effluent analyzer and constraint handling against the tube metal limit are what

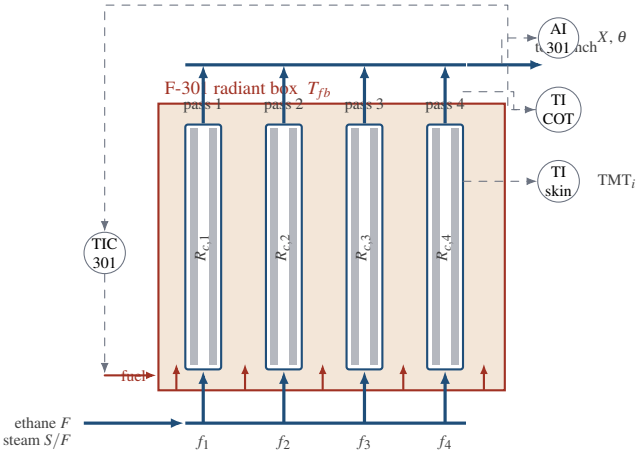


Figure 2.1 Ethane cracking furnace F-301. Ethane F with dilution steam ratio S/F is split into four pass flows f_i entering a radiant box at firebox temperature T_{fb}. Each coil carries a growing internal coke resistance $R_{c,i}$ between the gas and the tube metal, so each pass has its own coil outlet temperature and its own tube metal temperature TMT_i. AI-301 reports ethane conversion X on an analyzer cycle θ. TIC-301 manipulates the firing to hold the average coil outlet temperature, and the four pass flow trims balance the box on whichever variable is selected.

a multivariable controller provides on a modern plant [8], and the same layer now carries emissions and energy objectives as well [9].

The mathematics of what happens when a fired tube runs away is older than any of that. Parametric sensitivity in fixed bed reactors, the property that a small change in an operating parameter produces a disproportionate change in the peak temperature, was given its classical treatment by van Welsenaere and Froment in 1970 [10] and generalized later [11], and Section 1.4.3 applies it to an exothermic tube. The ethylene oxide reactor of Chapter 6 is such a tube. A cracking coil is not a fixed bed and its cracking reactions are endothermic, so it does not run away thermally. What it does instead is the subject of Section 2.6: it degrades on a timescale of weeks through a positive feedback between metal temperature and coking rate.

2.3 Process Model

Furnace F-301 cracks 40 t/h of ethane in four parallel passes of 10 t/h each, with a dilution steam ratio of 0.33 kg/kg, in a radiant box at 1150.9 °C. The coil outlet is 845.0 °C, giving 65.0 % conversion at 80.0 % selectivity for 20.80 t/h of ethylene, with 42.64 MW of absorbed duty, 106.6 kW/m^2 of flux and a tube metal temperature of 1004.9 °C against an 1100 °C limit. Figure 2.1 shows the arrangement.

> ✦ **DID YOU KNOW? The Largest-Volume Organic Chemical**
>
> Ethylene is the organic material that industry makes in the largest volume. In 2013 the plants of Europe could make about 25 million tonnes of it a year, about a quarter of the world's capacity [12]. The furnace of this chapter makes 20.80 t/h at its design point.

2.3.1 Flux, Metal Temperature and the Coke Layer

Each pass absorbs radiant flux from the box through the tube wall and the coke layer. The metal temperature is the gas temperature plus the flux times the resistance on the outside of the gas, Equation (2.1), and the flux is what the firebox can drive through the total resistance, Equation (2.2).

Table 2.1 Classification of Variables at the Start of a Clean Run

Symbol	Role	Nominal	Range
COT	CV, continuous	845.0 °C	820–880 °C SP
X (AI-301)	CV, analyzer	0.650	0.55–0.78 SP
TMT_{max}	constraint	1004.9 °C	$\leq$ 1100 °C
T_{fb}	MV	1150.9 °C	900–1400 °C
pass trims	MV	0 %	±35 %
F	DV, measured	40 t/h	25–52 t/h
S/F	DV, measured	0.33	0.20–0.60
pass skew	DV, unmeasured	1×	0.8–3×
coke factor	DV, unmeasured	1×	0.5–3×
$R_{c,i}$	states	0	growing

The first is the balance of Equation (1.1) in Chapter 1, written for one pass.

$$\mathrm{TMT}_i = \mathrm{COT}_i + q_i\,(R_w + R_{c,i}) \tag{2.1}$$

$$q_i = \frac{h_{fb}\,(T_{fb} - \mathrm{COT}_i)}{1 + h_{fb}\,(R_w + R_{c,i})} \tag{2.2}$$

where $h_{fb} = 0.73\,\mathrm{kW/(m^2\,K)}$ is a linearized radiant coefficient and $R_w = 1.50 \times 10^{-3}\,\mathrm{m^2\,K/W}$ is the clean wall and film resistance. Equation (2.2) is the closed form that results from eliminating the metal temperature. The outlet temperature of each pass then follows from its energy balance, Equation (2.3), in which the absorbed flux pays for heating the hydrocarbon and steam from the 600 °C crossover and for the endothermic cracking.

$$q_i \frac{A}{4} = \dot{m}_i\, c_p\,(\mathrm{COT}_i - T_{in}) + f_i\, X_i\, \Delta H_c \tag{2.3}$$

where $\dot{m}_i$ is the hydrocarbon plus steam flow of pass i, f_i its hydrocarbon flow, $c_p = 2.60\,\mathrm{kJ/(kg\,K)}$ and $\Delta H_c = 4600\,\mathrm{kJ/kg}$ of ethane converted. At the design point each pass absorbs 10.66 MW, 8.31 MW of it for the cracking. Both sides are monotone in COT_i, so the outlet is found by bisection, which converges whatever the coke resistance.

Equations (2.1) and (2.2) carry the lesson of the application. Coke raises $R_{c,i}$, which lowers q_i and therefore lowers the gas temperature the pass reaches, while raising the metal temperature for whatever flux does get through. A coked pass reads cold on the outlet and hot on the skin.

🕮 **Pause.** Try Exercise 2.1 before reading on.

2.3.2 Coking, and Why It Accelerates

The coke resistance grows at a rate that depends on the metal temperature, on severity and on dilution steam, Equation (2.4).

$$\frac{\mathrm{d}R_{c,i}}{\mathrm{d}t} = k_c\, \psi_i\, e^{(\mathrm{TMT}_i - T_{c,ref})/\theta_c} \left(\frac{X}{X_0}\right)^n \left(\frac{S/F}{(S/F)_0}\right)^{-m} \tag{2.4}$$

with $\theta_c = 30\,^\circ\mathrm{C}$ per e-fold, $T_{c,ref} = 1004\,^\circ\mathrm{C}$, $n = 1.8$, $m = 0.55$, $(S/F)_0 = 0.33$ and X the conversion of the pass. The factor ψ_i is the furnace coke factor times, for pass 4 only, the pass skew from a degraded burner. The exponential in metal temperature closes a positive feedback: a hotter tube cokes faster, which makes it hotter still. Nothing in a standard control scheme opposes it.

Table 2.2 Model Parameters

Symbol	Meaning	Value	Class
A	radiant surface, all passes	$400\,\mathrm{m}^2$	adjustable
h_{fb}	linearized radiant coefficient	$0.73\,\mathrm{kW/(m^2\,K)}$	adjustable
R_w	clean wall and film	$1.50 \times 10^{-3}\,\mathrm{m^2\,K/W}$	adjustable
k_c	coking rate at reference	$2.44 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$	adjustable
θ_c	degrees per e-fold in coking	$30\,^\circ\mathrm{C}$	adjustable
n	coking against severity	1.8	adjustable
m	steam suppression of coking	0.55	adjustable
X_0	conversion at 845 °C	0.650	adjustable
σ_0	selectivity at X_0	0.800	adjustable
θ	severity analyzer cycle	0.35 h	adjustable
TMT_{max}	tube metal limit	1100 °C	visible
decoke time	out of service	30 h	adjustable

2.3.3 Cracking and Severity

Conversion rises with coil outlet temperature and selectivity falls with conversion, Equations (2.5) and (2.6).

$$X = X_0 + \frac{\mathrm{d}X}{\mathrm{d}T}\left(\mathrm{COT} - \mathrm{COT}_0\right) \tag{2.5}$$

$$\sigma = \sigma_0 + \frac{\mathrm{d}\sigma}{\mathrm{d}X}\left(X - X_0\right) \tag{2.6}$$

with $\mathrm{d}X/\mathrm{d}T = 0.00420/^\circ\mathrm{C}$ and $\mathrm{d}\sigma/\mathrm{d}X = -0.52$. The ethylene rate is $FX\sigma$, and because X and σ move in opposite directions and the coking rate goes as $X^{1.8}$, the on-stream rate has an interior optimum that neither X nor the ethylene rate can locate.

2.3.4 Assumptions

The coil is a lumped pass rather than a distributed reactor. There is no axial profile, so the coke thickness is uniform along the tube and the peak metal temperature is at the outlet, which on a real coil it usually is but not always.

Cracking is a conversion and selectivity correlation rather than a free radical mechanism, so the effect of feedstock composition on the product slate cannot be studied and the model is calibrated for ethane only.

The radiant transfer is linearized into a single coefficient. A real firebox has a temperature distribution, and the pass skew parameter is the only representation of it here.

The coke layer grows and is removed completely by the decoke. Real decoking is a controlled burn whose duration depends on how much coke is present, and it does not always return the tube to its clean condition [2].

Tube metal thermocouples are treated as accurate. On a real furnace they are notoriously difficult, they read the tube surface rather than the mid wall, and a substantial part of the argument about balancing on tube metal in industry is about whether the measurement can be trusted.

There is no downstream quench or transfer line exchanger, so the coking that occurs there and often limits a real campaign is absent.

2.4 Dynamic Analysis

2.4.1 Why Equal Outlets Do Not Give Equal Metal

Combining Equations (2.1) and (2.2) and eliminating the flux gives the metal temperature of a pass directly in terms of the firebox and its own resistance, Equation (2.7).

$$\mathrm{TMT}_i = \mathrm{COT}_i + \frac{h_{fb} R_i \left(T_{fb} - \mathrm{COT}_i\right)}{1 + h_{fb} R_i}, \qquad R_i = R_w + R_{c,i} \tag{2.7}$$

Setting every COT_i equal, which is what an outlet balancing scheme achieves, leaves the second term varying with R_i alone. Differentiating at fixed outlet temperature gives the spread a balancing scheme cannot remove, Equation (2.8).

$$\left.\frac{\partial \,\mathrm{TMT}}{\partial R}\right|_{\mathrm{COT}} = \frac{h_{fb}\left(T_{fb} - \mathrm{COT}\right)}{\left(1 + h_{fb} R\right)^2} \tag{2.8}$$

At the moment the outlets of scenario 3 first agree to within half a degree, the firebox is at 1178.9 °C and the mean resistance is 1.77×10^{-3} m^2 K/W, so the numerator is $0.73 \times 334 = 244\,\mathrm{kW/m^2}$, the denominator is $(1 + 1.29)^2 = 5.2$, and the sensitivity is about 4.6×10^4 °C per m^2 K/W. The coke spread of 3.5×10^{-4} m^2 K/W between passes 1 and 4 therefore predicts 16.4 °C of metal spread, and the running model reports 16.1 °C. The spread is not a control failure. It is the resistance between the instrument and the constraint, and no amount of trimming on the near side of it can reach.

> **TRY IT IN THE APP** **3 · Equal outlets, unequal metal**
>
> Load the case and press Run. Touch nothing. Watch the hottest and coldest pass traces on the coil outlet strip come together, within half a degree by 8.2 h. Then watch the Spread across the four tubes trace under the tube metal strip: 16.1 °C at that moment and 34.0 °C by day twenty.

2.4.2 The Positive Feedback That Sets the Run Length

Differentiating Equation (2.4) along a pass with the outlet, and therefore the conversion, held constant gives Equation (2.9).

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\mathrm{d}R_c}{\mathrm{d}t}\right) = \frac{\mathscr{S}}{\theta_c}\left(\frac{\mathrm{d}R_c}{\mathrm{d}t}\right)^2 \tag{2.9}$$

where $\mathscr{S} = \mathrm{d}\,\mathrm{TMT}/\mathrm{d}R_c$ is the metal sensitivity along the path the furnace takes. For one pass trimmed against the others at a fixed firebox it is Equation (2.8). For four passes coking together with the outlet held by the firing, Equation (2.3) fixes the flux, so $\mathrm{TMT} = \mathrm{COT} + qR$ and $\mathscr{S} = q$ exactly. With $\mathscr{S}$ constant the coking rate is $\dot{R}_0/(1 - t/t^*)$ with $t^* = \theta_c/(q\dot{R}_0)$, and the metal rises as $-\theta_c \ln(1 - t/t^*)$, which is Equation (1.2) of Section 1.4.2. At the design point $q = 106.6\,\mathrm{kW/m^2}$ and $\dot{R}_0 = 2.51 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$, so $t^* = 1120\,\mathrm{h}$ and the hottest tube reaches 1100 °C at 1073 h, which is the run measured in scenario 2 in Section 2.6. The coking rate grows as the square of itself, which is a finite-time blow-up rather than an exponential and is why the last week of a campaign is worse than the first month. It is also why a pass that starts a campaign slightly ahead finishes far ahead, and why a degraded burner is visible on a tube metal trend days before it is visible anywhere else. A furnace does not run away thermally in the sense of van Welsenaere and Froment [10], because the cracking reactions are endothermic. It runs away on a timescale of weeks through the coke.

Pause. Try Exercise 2.3 before reading on.

2.5 Controller Design

2.5.1 Velocity Form PID

Both PID loops use the shared velocity-form PID of Section 1.6.2, with the derivative on the measurement, executed every $\Delta t_c = 0.02\,\mathrm{h}$ (72 s). The error and the proportional and derivative terms

e_k, P_k and D_k are those of Equation (1.14), and the rest of the algorithm, in the symbols used below, is Equation (2.10).

$$\begin{aligned} s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias} + s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left(\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{\tau_t} \right) \Delta t_c \end{aligned} \tag{2.10}$$

Here s_k is the running sum S_k of Equation (1.15), c_k the demanded output v_k of Equation (1.16), u_k the output applied and τ_t the tracking time T_t of Equation (1.17). Neither loop uses derivative action ($\tau_D = 0$), so both are PI. The bias is zero and I is seeded at reset with the initial output. Anti-windup acts on the integral alone, by back-calculation with $\tau_t = \tau_I$ [13], and s_k is never truncated, so the tube metal override can act on the severity loop through I. Manual mode sets I so that c_k equals the operator's value, and in automatic mode the manual firebox value follows the controller output after every execution. Integration is fixed-step fourth-order Runge-Kutta at 0.005 h, and the historian logs a row every 67 steps, 0.335 h.

2.5.2 Firing, Severity and Balancing

TIC-301 holds the flow-weighted average coil outlet temperature on the firebox temperature with $K_c = 3.0$ °C of firebox per °C, $\tau_I = 0.10$ h and output limits of 900 and 1400 °C. Above it, a severity loop holds the analyzer conversion by moving the coil outlet setpoint, with $K_c = 120$ °C per unit of conversion, $\tau_I = 1.5$ h and output limits of 820 and 880 °C. Both loops have $a = +1$, because a cold outlet calls for more firing and a low conversion for a hotter outlet. These are the tuned values most scenarios load. Scenario 1 and the challenge start from cautious values, $K_c = 0.8$ and $\tau_I = 0.6$ h on the firing and $K_c = 30$ and $\tau_I = 6$ h on the severity.

The balancing scheme moves the four pass flow trims to equalize either the four coil outlet temperatures or the four tube metal temperatures. It is implemented as an integral-only trim at rate $K_B = 0.0035$ per °C per hour on the deviation of each pass from the mean of the four, with each trim limited to ±35 % of the pass flow and the trims shifted to sum to zero so that the total feed is unchanged. Which variable it equalizes is the subject of Section 2.6.

2.5.3 The Tube Metal Override

A one-directional override reduces the coil outlet setpoint as the hottest tube approaches its limit, starting 12 °C below it with a gain of 2.0 °C per degree. It can only pull down, and because it is proportional on the hottest tube it switches on and off at the controller scan while the tube sits at the edge of its band. A tube run above its metal limit is replaced at the next turnaround, and a percent of conversion is worth less than that. The severity loop is not badly tuned when the override acts. It is optimizing an objective that has temporarily stopped being the objective, and a single PID has no way to know that. A protective override of the same one-directional kind guards a compressor's surge line in Volume 4, Chapter 2.

The override acts after the severity controller, so the limits of Equation (2.10) never see it. Left alone, the severity integral would see conversion short, wind up toward 880 °C, and leave the proportional pull-down fighting the top of the range. While the override acts, the severity integral therefore receives external reset. Each scan adds $(u_k^o - u_k)\Delta t_c / \tau_I$ to I_k, where u_k^o is the setpoint the override applies, which is the back-calculation of Equation (2.10) with the override's output in place of the clamp. From a decoke until the restart ramp of the next subsection ends, the severity loop tracks the outlet target it held before the decoke, then returns bumplessly.

2.5.4 Decoke Policy

The decoke is started by the operator, after a fixed number of hours (1000 h by default) or when the hottest tube reaches a threshold (1094 °C by default, six degrees under the limit). Each takes the

furnace out for 30 h. The feed is withdrawn no faster than the firebox comes down, so the tube metal does not rise as the coils empty. During the burn each pass carries 0.2 t/h, the firing holds their outlet at the 880 °C coil outlet limit (a firebox near 889 °C, below the firing loop's 900 °C floor), and the coke burns off with a time constant of 0.8 h. The feed and the outlet setpoint then return over 4 h from 880 °C to target, and the outlet peaks at 881.4 °C after a decoke at 700 h in scenario 2. The same choice between a fixed interval and a trigger on a measured margin, against a fouled jacket and a wash, is weighed in Chapter 11.

2.5.5 Model Predictive Control

A selector placed after the existing switches keeps the loops above as the default and offers a model predictive controller (MPC) as the alternative supervisory layer [8, 14]. Section 1.6.3 describes the option in general, and this subsection gives what is specific to the furnace. Under MPC, in automatic, on line and after the restart ramp, one quadratic program takes the place of both the severity loop and the balancing trim. TIC-301 remains in service beneath it, the tube metal override remains armed above it and acts on the planned setpoint exactly as it acts on the severity loop's, and the decoke policy is untouched. Every $T_s = 0.1$ h, which is five controller scans, the plan chooses the coil outlet setpoint COT_{sp} and the trims δ_1, δ_2 and δ_3 of passes 1 to 3, as fractions of the pass flow, and sets $\delta_4 = -(\delta_1 + \delta_2 + \delta_3)$, so feed moves between passes and the total never changes. It predicts $N_h = 10$ periods ahead (1 h, adjustable from 0.5 to 3 h) with move blocks of 1, 2 and 7 periods, decides moves away from the commands in force, applies the first block and holds it until the next plan. Planning begins on the call that primes the run, so a run started under MPC is planned from $t = 0$.

> ⧖ HISTORY **1986: A Pyrolysis Furnace Under Constrained Predictive Control**
>
> The first industrial predictive controllers handled constraints in a somewhat ad hoc way. The next step posed dynamic matrix control as a quadratic program, with the limits on inputs and outputs written into the problem explicitly. The method was first described in a 1983 conference paper, and a fuller description followed in 1986. It closed with results from a pyrolysis furnace, where the controller adjusted the fuel gas pressure of three burners to control the stream temperature at three places in the furnace. The tests showed the input constraints enforced as the furnace moved and the three temperatures decoupled. A later survey counts the algorithm as the second generation of the technology [8]. The plan of Section 2.5.5 solves the same kind of quadratic program every 0.1 h.

The plan's model is the pass model of Equations (2.1) to (2.3) and (2.5), evaluated with the declared parameter values of Table 2.2 and the nominal disturbances of Table 2.1 and linearized by central differences about clean tubes, four passes of 10 t/h and an outlet of 845 °C. Appended to it are the declared firebox lag of 0.035 h, pass flow lag of 0.010 h and outlet transmitter lag of 0.012 h, and TIC-301 itself, executed as in Equation (2.10) at the gain and integral time on its faceplate. Composing five scans with the input held gives Equation (2.11), the furnace's form of Equation (1.18).

$$\tilde{\xi}_{k+1} = \Phi\,\tilde{\xi}_k + \Gamma\,\tilde{v}_k, \qquad \hat{y}_k = \Omega\,\tilde{\xi}_k + \hat{b} \tag{2.11}$$

Tildes mark departures from the design point. The seven states ξ are the firebox temperature, the four pass flows, the filtered outlet and the integral of TIC-301, the inputs are $v = (\mathrm{COT}_{sp}, \delta_1, \delta_2, \delta_3)$, and $\hat{y}$ lists the predicted conversion $\hat{X}$, the flow-weighted outlet, the four tube metal temperatures $\widehat{\mathrm{TMT}}_i$ and the four departures $\hat{z}_i - \bar{\hat{z}}$ of each pass from the mean of the four, where z is the outlet or the metal temperature as the balancing switch selects. The offsets $\hat{b}$ are described below. About the design point a pass outlet gains 0.355 °C and its metal 0.692 °C per degree of firebox, and they lose 10.86 °C and 5.19 °C per t/h of the pass's own flow. Shifting one percent of the feed onto a pass therefore cools its metal by 0.52 °C, while each degree of outlet setpoint warms every tube by 1.95 °C and raises conversion by 0.0042. Stepped by ±1 °C of setpoint and ±1 % of trim on clean tubes, under the tuned and the cautious firing tuning alike, the simulated furnace and the model differ

by no more than 0.24 % of the response from 0.1 to 2 h. Step-test fits were not used, since the plan has to predict the furnace with TIC-301 closed at whatever tuning its faceplate carries and the pass model yields every slope in closed form.

Each pass outlet and tube metal temperature is read on every scan, the same readings the balancing trim uses, so the offset of each from the model is measured rather than estimated. Each entry of $\hat{b}$ is the reading less the model's present value, held constant over the horizon, and it is through these offsets that the coke enters the plan. The firebox and flow states are the model's own response to the commands already sent, and the filtered outlet and the integral are read from TIC-301. Conversion is the only late measurement. The plan computes it from the four outlets with Equation (2.5) and adds a bias $\hat{b}_X$ that each scan moves by $0.04\,(X_{AI} - X_{cp} - \hat{b}_X)$, where X_{AI} is the analyzer reading and X_{cp} a copy of the analyzer, with the same 0.35 h cycle and 0.08 h averaging on the same clock, fed with the computed conversion. The bias is frozen through a decoke and the restart ramp and removes any steady offset in the conversion law [15].

Each plan minimizes Equation (2.12), which adds a linear price on each slack to the terms of Equation (1.19).

$$\begin{aligned} J = & \sum_{k=1}^{N_h} w_X \left(\hat{X}_k - X_{sp}\right)^2 \\ & + \lambda_m \sum_{j=1}^{3} \left[\Delta\mathrm{COT}_{sp,j}^2 + \sum_{i=1}^{3} \left(100\,\Delta\delta_{i,j}\right)^2 \right] \\ & + \sum_{i=1}^{4} \left(\rho_T\,\varepsilon_i^2 + \mu_T\,\varepsilon_i + \rho_z\,\zeta_i^2 + p_z\,\zeta_i\right) \end{aligned} \tag{2.12}$$

Here $w_X = 1/0.006^2$ and X_{sp} is the conversion target. With the severity loop in manual the first term instead holds the flow-weighted outlet on the operator's setpoint at a weight of one per square degree. The move weight is $\lambda_m = 10$, Δ is the change of block j in degrees or fractions, measured from the command in force for $j = 1$, and the factor 100 puts the trims in percent. The slacks ε_i of the tube rows cost $\rho_T = 10^3$ per square degree and $\mu_T = 10^4$ per degree, an exact penalty, so the plan never predicts a crossing that its handles could avoid. The slacks ζ_i of the spread rows cost $p_z = 3$ per degree and $\rho_z = 0.01$ per square degree, and because each row has one slack for the whole horizon, p_z prices the largest excess of each pass over the mean. The rows are Equation (2.13).

$$\begin{aligned} & 820 \le \mathrm{COT}_{sp,k} \le 880, \qquad |\delta_{i,k}| \le 0.35 \\ & \widehat{\mathrm{TMT}}_{i,k} \le \overline{T} + \left[\mathrm{TMT}_{i,0} - \overline{T}\right]^+ e^{-kT_s/\tau_f} + \varepsilon_i \\ & \hat{z}_{i,k} - \bar{\hat{z}}_k \le \left[z_{i,0} - \bar{z}_0\right]^+ e^{-kT_s/\tau_f} + \zeta_i, \qquad \varepsilon_i,\ \zeta_i \ge 0 \end{aligned} \tag{2.13}$$

The first line is hard and repeats the clamps of the severity loop and the trim, with the upper setpoint bound lowered if the coil outlet limit is set below 880 °C and pass 4 entered as two rows per block. Zero trims with any setpoint in range satisfy every hard row and the other rows are soft, so a plan always exists. The ceiling is $\overline{T} = \mathrm{TMT}_{max} - \Delta_o - 1\,^\circ\mathrm{C}$, a back-off of one degree under the override's start with $\Delta_o = 12\,^\circ\mathrm{C}$, read as the override reads it, which is 1087 °C as shipped. With $[\cdot]^+$ the positive part, a row already violated starts at the present reading and closes onto its limit with $\tau_f = 0.1$ h, because a single slack per row would otherwise excuse any later violation up to the present one. With balancing off the trims stay where the operator set them and the spread rows are dropped. Replacing the closing rows by flat ones moves the severity IAE of scenario 10 from 2.176 to 2.178 and the hours over the limit in scenario 8 from 228.51 to 228.45, so none of the results below depends on them.

Five design choices need a reason. The controller is a selector of its own rather than a further value of an existing switch, since it takes the setpoint and the trims together and either switch would misname it. It moves the outlet setpoint rather than the firing, so TIC-301 stays in service, the setpoint keeps the severity loop's range and the override acts on it unchanged. Every tube has its own soft row

Table 2.3 Scenario Sequence

#	Scenario	Teaching point
1	Four tubes that do not agree	the two measurements disagree
2	The coil outlet loop	the constraint is elsewhere
3	Equal outlets, unequal metal	balancing the wrong variable
4	Balancing on the constrained variable	what it costs
5	A burner problem on one corner	worse than no control
6	The optimum neither trend can see	severity economics
7	Dilution steam	a second optimization
8	The tube metal override	objective changes
9	When to decoke	schedules are open loop
10	Balance the passes before the override	one plan, two handles
11	Challenge: sixty days	graded run

instead of one row on the hottest, because a maximum is not linear in the inputs. It decides every 0.1 h rather than every minute, because a minute is not a whole number of 72 s scans and five scans per decision cut the number of plans fivefold. Its estimator is the measured offsets with the analyzer bias rather than a Kalman filter, because every pass is read without delay, the furnace answers in minutes while the coke builds over days, and an input-disturbance Kalman filter, measured once on a separate copy of the code, moved the severity IAE of scenario 10 by at most 0.016 and the hours over the limit in scenario 8 by at most 3.4 h, changing no ranking. Where a disturbance reaches no reading directly, the filter has to carry it as a state, as the fuel bias of Volume 9, Chapter 4 is carried.

The solver's answer is refused when it is suboptimal, holds a nonfinite number or leaves a bound. The severity loop, which has tracked the planned setpoint, and the balancing trim, which has followed the planned trims, then take both handles with one execution each and no bump, the trend is marked and a note is posted, and control returns to the plan at its next acceptable solution. With the severity loop in manual the setpoint stays where the plan left it. The plan reads only what the shipped loops read, which is the pass outlet and metal readings, the outlet transmitter, the analyzer, the integral and tuning of TIC-301, the setpoints and switches, the commands applied, the tube metal limit and margin, and the coil outlet limit. Apart from those two limits, read where the override and the severity loop read them, it never reads a run parameter, the coke or a disturbance, so a plant parameter a learner changes is model error to the plan. The quadratic program is solved by a dual active-set method [16] in a mean of 165 µs per plan in Chrome on the development computer, between 163 and 167 µs over four measurements, which is 0.17 percent of one core at 1 h per second.

2.6 Results

Eleven scenarios run from watching four tubes diverge to a graded sixty-day challenge. Table 2.3 lists them, and Section 1.6.4 describes how a scenario is loaded, credited and exported. Section 2.11 sets runs from several of them.

2.6.1 The Constraint Is Not on the Manipulated Variable

With the coil outlet loop closed in scenario 2, the outlet stays within 0.01 °C of 845 °C for the first 1000 h while the firebox walks up from 1150.9 to 1217.8 °C and the hottest tube from 1004.9 to 1071.8 °C. With the override switched off, the run at this severity is 1073 h, which is 44.7 days, and it ends because of a number nobody is controlling. With the override armed as shipped, the outlet setpoint starts to come down at 1050.2 h, when the hottest tube reaches 1088 °C, but coke grows at any severity, and with nobody decoking a tube crosses 1100 °C at 1134.0 h and spends 16.8 h above it by 1200 h. The firing still has room, because the firebox peaks at 1242 °C against its 1400 °C limit.

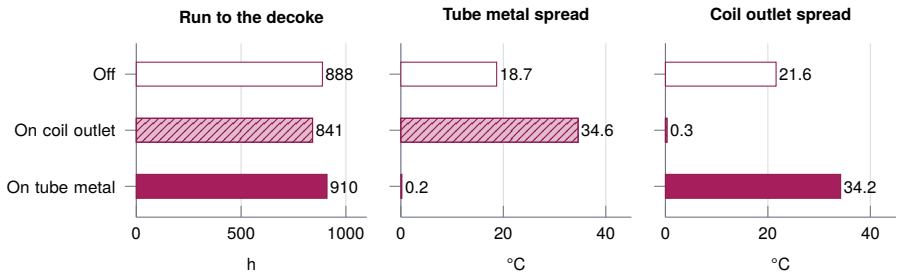


Figure 2.2 The three balancing schemes of Table 2.4 on the furnace with one degraded burner: white bars with balancing off, hatched bars balancing on coil outlet, solid bars balancing on tube metal, with the table's value beside each bar. Notice that the scheme that brings the outlets together drives the tube metal apart and gives the shortest run of the three.

Table 2.4 Balancing Schemes on a Furnace With One Degraded Burner

Balancing	Run (h)	TMT spread (°C)	COT spread (°C)
Off	888	18.7	21.6
On coil outlet	841	34.6	0.3
On tube metal	910	0.2	34.2

It is the tube metal that runs out, and tube metal is a consequence of firing rather than a thing anybody sets.

2.6.2 Balancing on the Wrong Variable

Scenario 3 starts four passes at coke resistances of 1.0, 2.0, 3.0 and 4.5×10^{-4} m^2 K/W with balancing on coil outlet temperature. The outlet spread falls from 12.5 to under 0.5 °C at 8.2 h and stays there, while the tube metal spread rises from 10.8 to 16.1 °C at the same moment, 20.9 °C by day ten and 34.0 °C by day twenty. The override first acts at 336.7 h, and pass 4 first crosses its limit at 455.3 h and spends 4.1 h above it by day twenty.

Balancing on tube metal, in scenario 4 from the same starting coke with the decoke at 1094 °C, converges the metal temperatures and separates the outlet temperatures, which is the same trade in the other direction. Just before the decoke the metal spread is 0.15 °C against 28.8 °C for outlet balancing, and the outlet spread is 16.0 °C against 0.14 °C. The run length barely moves, 396.6 h against 392.4 h, and with balancing off it is 408.8 h, the longest of the three. A one-off difference in coke is partly self-limiting, because with equal flows the coked pass runs a colder outlet, cracks less and therefore cokes more slowly, and its metal spread grows only from 10.8 to 13.2 °C over the run. Balancing on the constraint pays when something keeps pushing one pass ahead, which is the case of the next subsection. The cost is a product cost. Four passes cracking at four different outlet temperatures make four different severities, so the effluent is a blend. On an ethane cracker that is tolerable, and on a naphtha cracker, where the yield structure is far more temperature sensitive, it is the reason this is argued about.

2.6.3 Worse Than No Control

Table 2.4 reports a furnace whose fourth pass cokes 1.8 times as fast as the others because of a degraded burner (scenario 5). Each run starts from clean tubes with the override armed and ends with a decoke when the hottest tube reaches 1094 °C, and the spreads are read on line just before the decoke. Figure 2.2 draws the three rows side by side.

The middle row is the result worth reading twice. Balancing on coil outlet temperature gives a shorter run than not balancing at all, because coke makes pass 4 read cold and the scheme therefore

Table 2.5 Severity Against Instantaneous Rate and On-Stream Rate

COT (°C)	X	σ	C_2H_4 (t/h)	Run (h)	Per elapsed hour (t/h)
835	0.608	0.822	19.99	2516	19.75
845	0.650	0.800	20.80	1073	20.23
850	0.671	0.789	21.18	700	20.31
855	0.692	0.778	21.54	454	20.20
865	0.734	0.756	22.21	186	19.12

takes feed away from exactly the pass that needed more of it. Nothing is broken. The instrument is accurate, the tuning is correct and the loop is doing what it was asked.

> **TRY IT IN THE APP** **5 · A burner problem on one corner**
>
> Run the case to its decoke with Pass balancing on coil outlet and read the Run tile, 841 h. Reset, set Pass balancing to tube metal before pressing Run, and repeat, 910 h. A third run with balancing off decokes at 888 h, later than the run balanced on the outlets.

The first sign of the degraded burner is a single tube metal temperature pulling away from the other three. With outlet balancing, pass 4 is 0.8 °C hotter than pass 1 after three days, 3.0 °C after ten and 19.1 °C after thirty, while the four outlets agree to within 0.2 °C, as shown in Figure 2.3. Neither balancing scheme fixes the burner and neither is supposed to. One of them buys time to schedule the repair and the other spends it.

2.6.4 A Severity Optimum No Continuous Measurement Can Find

Table 2.5 reports five coil outlet setpoints, each held from clean tubes until the hottest tube reaches its 1100 °C limit, at a dilution steam ratio of 0.33. They are measured in scenario 2 with the override off and the decoke triggered on tube metal at 1100 °C, and the last column is the on-stream rate when the furnace comes back on line after its 30 h decoke.

The instantaneous ethylene rate rises monotonically across the whole range and so does the conversion. Neither can find the optimum, because both are still improving where the optimum has already gone past. The only number that finds it is ethylene divided by elapsed time, Equation (1.3) of Chapter 1, and that number cannot be measured until after the decoke.

Pause. Try Exercise 2.2 before reading on.

The top is also very flat, with 20.23, 20.31 and 20.20 t/h at 845, 850 and 855 °C, while the run falls from 1073 h to 454 h across the same ten degrees. Fifteen degrees above the peak the on-stream rate is 5.9 % lower and fifteen degrees below it 2.7 % lower. Optima with flat tops and steep sides should be approached from the flat side. Scenario 6 poses the same question through the severity loop, with the override armed and the decoke at 1094 °C. Conversion targets of 0.62, 0.65, 0.68 and 0.70 decoke at 2026, 1119, 622 and 420 h and give 19.91, 20.23, 20.31 and 20.18 t/h per elapsed hour once the furnace is back on line, so the optimum a learner finds there is the same flat top.

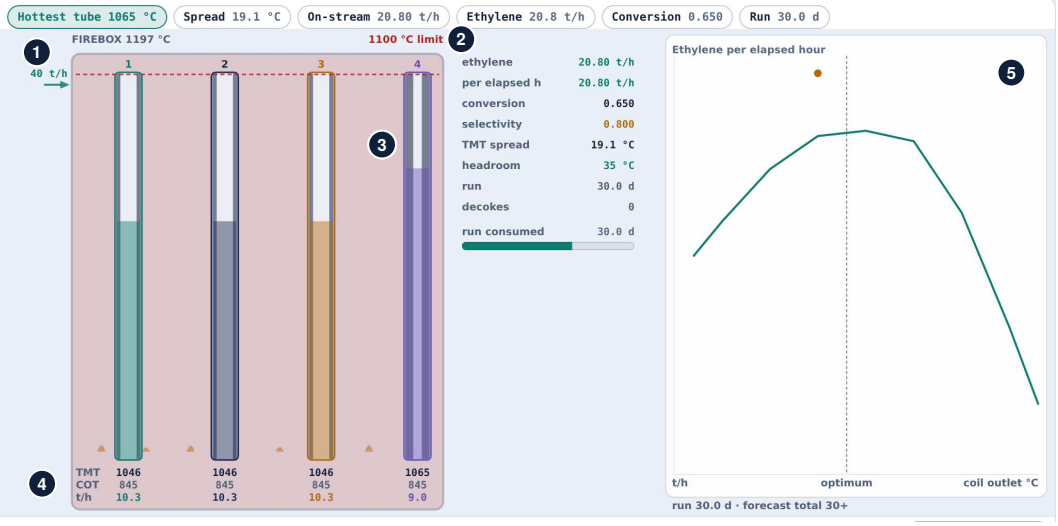


Figure 2.3 The cracking furnace application in scenario 5, the degraded burner case, with balancing on coil outlet temperature at 720 h (30.0 days), cropped to the plant view. The firebox at 1197 °C is drawn with its four coils, each showing its coke layer and its tube metal temperature as a bar against the 1100 °C limit. (1) The status chip reads the hottest tube. (2) The dashed line is the 1100 °C limit. (3) Pass 4, coking 1.8 times as fast, is the hottest tube at 1065.3 °C against 1046.2 °C for the other three. (4) The rows under the coils read each pass's tube metal, outlet and flow: all four outlets read 845 °C (spread 0.16 °C) because the balancing has cut the flow of pass 4 to 9.02 t/h against 10.33 t/h. (5) Beside the furnace readouts is the measured on-stream rate against coil outlet temperature.

> ✦ **DID YOU KNOW?** **Fastest While Running, Slowest Overall**
>
> From the simulator (Table 2.5). At a coil outlet of 865 °C the furnace makes ethylene faster than at any other setting in the table, 22.21 t/h. Counted per elapsed hour it makes the least, 19.12 t/h, because its run lasts only 186 h before the decoke.

2.6.5 The Graded Challenge

The challenge runs 1400 h, about 58 days, on a furnace whose fourth pass already cokes 1.5 times as fast as the others at a coke factor of 1.2, with measurement noise of 0.6 °C on the outlet and 0.25 % on the analyzer. Dilution steam falls to 0.28 at 200 h, the feed rises to 45 t/h at 600 h and the pass 4 skew rises to 2.2 at 900 h. It starts with the cautious tuning, the outlet setpoint at 852 °C, the severity loop in manual, no balancing, the override armed and no automatic decoke. It is met when, judged at 1400 h, the coil outlet stays within 5 °C of its setpoint on every logged sample from 924 h (the last third, restart ramps included), a tube spends less than 0.1 h above its metal limit, and the furnace finishes above 18.5 t/h of ethylene per hour of elapsed time, with the model parameters and disturbances as given.

As shipped, nobody decokes the furnace. A tube crosses 1100 °C at 538 h, the firebox reaches its 1400 °C limit at 634 h, and the coke then chokes the flux until the outlet falls to the 600 °C crossover temperature, so the run ends with a worst outlet error of 220 °C, 858.2 h over the limit and 10.67 t/h. The suggested configuration (the tuned loops, the severity loop in automatic at 0.650, balancing on tube metal, and a decoke at 1090 °C) decokes at 642.4, 987.0 and 1305.0 h and passes with a worst outlet error of 3.53 °C, an hour into the last restart and 1.47 °C inside the band, no time over the limit

Table 2.6 The Shipped Loops Against the Predictive Controller

#	Structure	Severity IAE (h)	Over limit (h)	First over (h)	Run (h)	C_2H_4 (t)	Per elapsed hour (t/h)	Peak TMT (°C)	Plans, mean solve (µs)
8	loops	129.296	260.46	219.53	–	6163	–	1399.0	–
	MPC	113.709	228.51	251.48	–	6691	–	1398.8	4800, 133
5	loops	0.585	0	–	841.3	17489	20.071	1094.0	–
	MPC	13.772	0	–	1040.4	21290	19.889	1094.0	10405, 108
10	loops	2.986	0	–	–	2439	–	1095.4	–
	MPC	2.176	0	–	–	2427	–	1087.3	1200, 111
	loops, on TMT	2.015	0	–	–	2420	–	1095.3	–

Scenario 8 runs 480 h of coking, scenario 5 (one degraded burner) runs to its decoke and scenario 10 runs 120 h. Severity IAE is the integral of the absolute conversion error on the logged rows. Per elapsed hour divides the ethylene of the run by the run plus its 30 h decoke. Peak TMT is taken over every integration step. No plan was rejected. The mean solve time is wall clock in Node.js on the development computer and varies from run to run.

and 20.67 t/h, 2.17 t/h above the target. The app gives this one run whether the settings are made before Run or carried through a Reset. On four other noise seeds the worst error is 3.52 to 3.93 °C. The decoke trigger with the tuned outlet loop alone also passes (3.95 °C, 19.70 t/h), but neither passes without the other.

2.6.6 The Predictive Controller Against the Shipped Loops

Three scenarios were run under both structures with the same seed, events, instruments and duration, and with measurement noise off. The loops keep each case's own settings and shipped tuning. The plan uses the settings scenario 10 ships, a horizon of 1 h, $\lambda_m = 10$ and $p_z = 3$, with no retuning for any case, and neither structure has a preview. Table 2.6 reports the severity IAE of the scorecard, the time with a tube above 1100 °C, the run, the ethylene made, the peak tube metal and the plans.

The coking case, scenario 8, runs 480 h with nobody decoking. The plan holds the balanced tubes on its 1087 °C ceiling and surrenders only the severity the ceiling requires, where under the loops the override chatters in its band. The first tube crosses 1100 °C at 251.5 h instead of 219.5 h, the time over the limit falls from 260.5 to 228.5 h, the severity IAE from 129.3 to 113.7, and the furnace makes 6691 t of ethylene against 6163 t, 8.6 % more. Once the setpoint reaches its 820 °C floor both structures run away, as the case is meant to show.

Scenario 5 run to its first decoke with balancing on coil outlet temperature is a loss. The plan equalizes the outlets much as the shipped trim does until pass 4 reaches the ceiling late in the run, then moves feed onto pass 4 and cuts severity. Its run lasts 1040.4 h against 841.3 h and makes 21 290 t against 17 489 t, but at a severity IAE of 13.772 against 0.585, so the ethylene per hour of run and decoke is 19.889 t/h against 20.071 t/h, 0.9 % less. The ceiling accounts for all of it. At 1087 °C it sits 7 °C under the 1094 °C decoke trigger, so the plan reaches the trigger only after its trims are at their limit and its setpoint is on its floor, while the override lets the tube into its band at full severity. In a diagnostic run that no scenario ships, with the trigger lowered to 1086 °C under the ceiling, the plan decokes at 791.1 h and the loops at 791.8 h, and both earn 20.040 t/h. The plan carries no price on run length or on the decoke, and the comparison does not retune it for this case.

Scenario 10 was built for this comparison. It starts from the furnace of scenario 5 about 780 h into its run, with coke resistances of 4.0×10^{-4} on passes 1 to 3 and 10.6×10^{-4} m^2 K/W on pass 4, which still cokes 1.8 times as fast, the severity loop in automatic at 0.650, balancing on coil outlet temperature and the override armed. Pass 4's tube starts at 1077 °C, 11 °C under the override's start. Figure 2.4 follows the 120 h under both structures. The plan first equalizes the outlets, to within 0.5 °C at 2.6 h against 10.6 h for the shipped trim. Doing so takes feed off pass 4, down to 8.5 t/h, and pass 4's tube reaches the ceiling at 13.7 h. From then on the plan shifts feed back onto pass 4, reaching the trim limit of 13.5 t/h at 57.2 h, and holds the tube on the ceiling. The override never

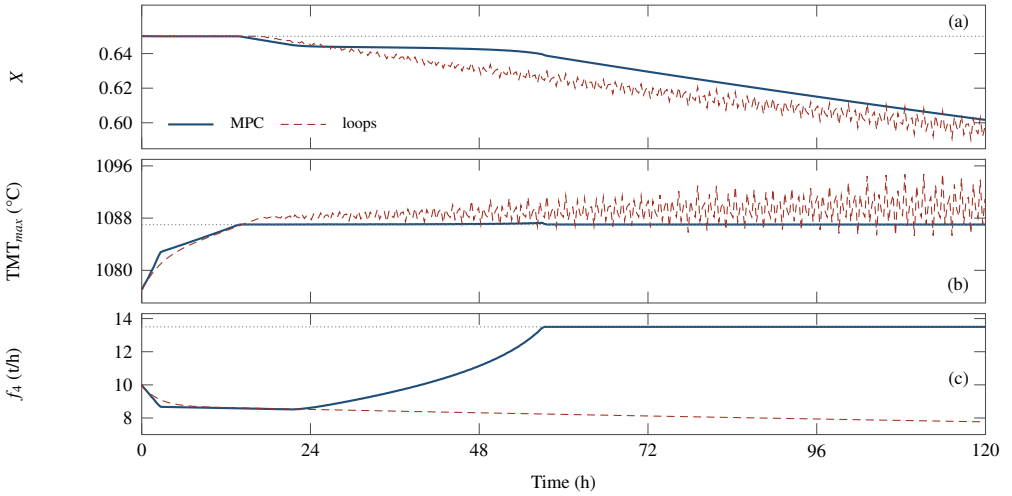


Figure 2.4 Scenario 10 over its 120 h under the predictive controller (thick solid) and the shipped loops (thin dashed), on every logged row. (a) Ethane conversion against its 0.650 target (dotted). (b) Hottest tube metal temperature and the plan's 1087 °C ceiling (dotted), one degree under the override's start. (c) Pass 4 flow and the 35 % trim limit at 13.5 t/h (dotted).

acts and the hottest tube peaks at 1087.3 °C. Under the loops the override first acts at 16.8 h and is acting on 153 of the 359 logged rows, pass 4 falls to 7.8 t/h, and the hottest tube reaches 1095.4 °C (1094.8 °C on the logged rows). Conversion first drops 0.006 below target at 24.7 h under the loops and 25.7 h under the plan. After that the loops' conversion swings with the override, from 0.6275 to 0.6344 over hours 47 to 49 and from 0.6129 to 0.6236 over hours 71 to 73, where the plan's stays within 0.6424 to 0.6427 and 0.6289 to 0.6302. The severity IAE is 2.176 against 2.986 and the peak tube 1087.3 °C against 1095.4 °C, for 2427 t of ethylene against 2439 t, with a mean outlet spread of 44.8 °C against 0.7 °C.

The plan is not the best structure on this case. The loops balancing on tube metal score 2.015, and most of that gap comes from where the plan's ceiling sits rather than from its 1 h view. Balancing on tube metal itself, the plan scores 2.140, so the balancing variable explains 0.036 of the 0.161 difference. The rest comes from the loops letting the hot tube into the override's band, to 1095.3 °C with a mean hottest tube of 1087.85 °C after the first day, while the plan holds 1087.0 °C with a mean of 1086.16 °C. With the override margin cut to 6 °C, which lifts the plan's ceiling to 1093 °C, the plan on tube metal scores 1.858 with no tube above 1093.04 °C, while the loops on tube metal at the same margin score 1.760 but spend 0.87 h above 1100 °C and peak at 1101.4 °C. Over these five days, leaving the flows alone scores best of all, 1.933 under the loops and 2.093 under the plan. Balancing on tube metal leads on days 2 and 3 and trails on days 4 and 5, because feed moved onto pass 4 brings passes 1 to 3 up to the ceiling, 1087.0 °C against 1062.7 °C at 120 h under the plan, and their coke then grows faster, to 8.94×10^{-4} against 6.23×10^{-4} m^2 K/W.

On the challenge, reported for information, choosing the plan and nothing else fails. Nothing decokes in the shipped settings, so every tube runs away once the setpoint is on its floor, with a worst outlet error of 220 °C, 798.5 h over the limit and 11.00 t/h. With the suggested settings and the plan the challenge is met, with a worst outlet error of 3.49 °C, no time over the limit and 21.00 t/h, against 3.53 °C and 20.67 t/h under the suggested loops. The unscored measures show the cost. The severity IAE is 4.978 against 3.015, because the suggested decoke trigger of 1090 °C again sits above the plan's ceiling, and the hottest tube peaks at 1090.1 °C, so no limit the score leaves out is broken.

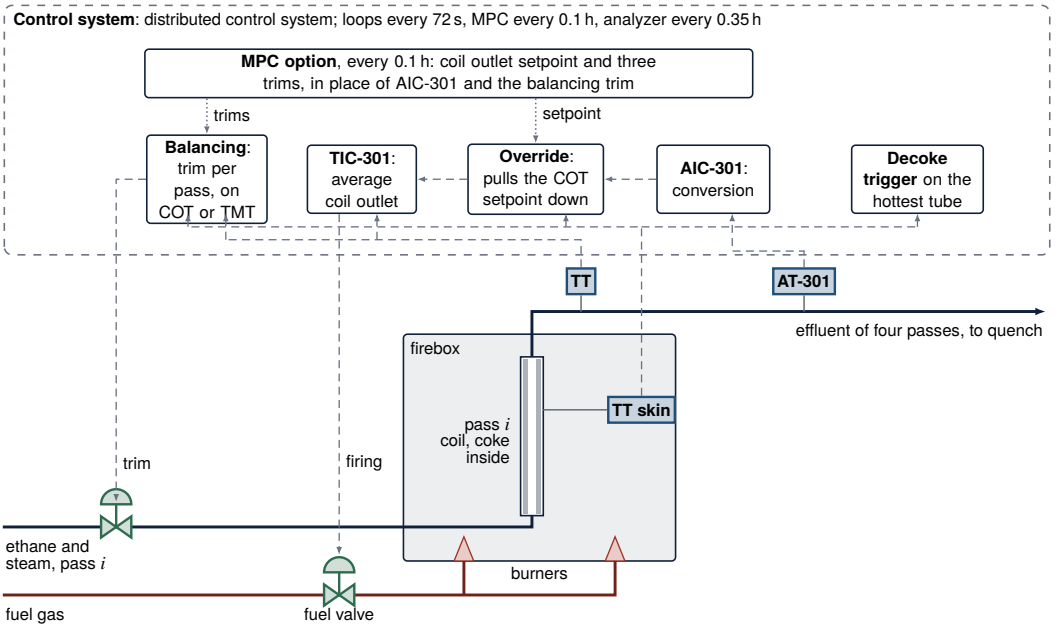


Figure 2.5 The loops of furnace F-301 with their hardware, one pass of four drawn. Blue boxes are sensing devices: TT on the coil outlet of each pass, TT skin on its tube metal and AT-301, the effluent analyzer. Green symbols are final elements: the fuel valve to the burners and the pass flow valve. Dashed arrows are signals, the dotted arrows the outlet setpoint and trims that the MPC option writes in place of the severity loop and the balancing trim. Notice that the balancing block can read either the outlet or the tube metal of each pass, and that only the override and the decoke logic read the tube metal by default. Not to scale, and not a piping diagram.

2.7 Technology in Practice

Figure 2.5 draws the firing loop, the severity loop, the pass balancing and the tube metal override of Section 2.5 with the hardware of a cracking furnace in place of the schematic's symbols.

What Measures It

Each pass carries a thermocouple at its coil outlet, and the firing loop holds the flow-weighted average of the four. Each pass also carries a tube metal reading, which the override, the decoke trigger and the balancing on tube metal use. The severity loop reads an online analyzer on the effluent that reports conversion every 0.35 h. The chapter gives the outlet transmitter a lag of 0.012 h and takes the tube metal readings as exact, which is the assumption the Limitations question.

> **TECHNOLOGY** **Sensor: Tube Metal Temperature**
>
> In furnaces of this kind tube walls commonly run between 700 and 1200 °C. A thermocouple welded to a tube or set into its wall reads continuously, but in the firebox it does not last long, reads the weld rather than the true surface, and is prone to erroneous readings. Operators also survey the tubes with infrared pyrometers and cameras, which receive the energy the wall emits and reflects [17]. Thermocouples are rugged but less accurate than resistance thermometers. An open circuit, their most common failure, can show as a wildly varying reading, so the transmitter is set to go high on it where excess temperature is the danger. A pyrometer sees only the surface, and its main weakness is inaccuracy, since the tube's emittance must be known [18]. The chapter takes its tube readings as exact.

What Moves It

Two kinds of final element carry the handles. The fuel valve to the burners sets the firing, and in a cracking furnace the tubes are heated from outside by gas or oil burners [5]. The chapter models the result as a firebox temperature behind a lag of 0.035 h, within 900 and 1400 °C. The four pass flow valves carry the trims that move feed between passes while the total stays fixed.

> **TECHNOLOGY** **Final Element: The Pass Flow Valve**
>
> Each pass has its own flow control valve, and the balancing scheme trims all four. A pneumatic actuator sets the stem in proportion to its air pressure, against a spring. Friction in the packing makes the stem move in stair steps, called stiction, and leaves it unresponsive to small reversals of the signal. Valves with poor maintenance have shown hysteresis beyond 10 %, and a positioner corrects the gap between signal and stem. The fail position is chosen from the hazard rather than from habit, the way a cooling valve that would starve an engine is made to fail open [18]. The chapter's trims move without friction, behind a lag of 0.010 h, within ± 35 % of the pass flow.

Where the Controller Runs

A distributed control system has a supervisory level that sends set points to the field controllers, its loops are sometimes nested or cascaded, and their cycle times range from milliseconds to minutes [19]. The app's loops execute every 0.02 h (72 s), its MPC every 0.1 h, and its analyzer reports every 0.35 h. The structure maps onto standard blocks: a PI temperature controller on the firing, cascaded under a PI analyzer controller with external reset, an integral trim per pass with its sum held at zero, a proportional override that can only pull the outlet setpoint down, and a decoke trigger in the sequence logic.

> **IN PRACTICE** **Burning the Coke Out**
>
> Almost all plants decoke with steam and air. The air burns the coke off the inside of the coils while steam passes with it to prevent local overheating, at coil outlets of up to 800 °C. A decoke typically lasts 18 to 48 h, and in a European survey decoking took on average about 3 % of the run time. Cyclones or wet scrubbers catch the dust in the gas leaving. Every plant that answered dopes its feed with sulfur compounds against coke, and the burn is pushed hard only near its end, since high temperatures could damage the coils. The European reference document under Directive 2010/75/EU sets out these practices [12]. The chapter's decoke takes 30 h.

What Is Modernizing It

Three developments touch this chapter. The first is the layer above the loops. Crackers use multivariable control with online analyzers, performance and constraint controls, often with online optimization, and for a gas feed the severity they manage is the conversion [12]. The predictive controller of Section 2.5.5 is a small example of that layer, and its history reaches back to a pyrolysis furnace [8]. The second is the coil itself. Coils with less nickel or with an inert inner coating retard coke, mainly in its first layer, and recoating may be needed after a few decokes [12]. Sensors that detect fouling early and high-temperature coatings have been studied as well, and modern cracking furnaces reach fuel efficiencies of typically 92 to 94 % [5]. The third is the heat source. Electric cracking would replace the burners, and with them the handle that TIC-301 moves, by resistance or other electric heating [20].

2.8 Limitations

The coil is lumped rather than distributed, so there is no axial temperature profile and no way to study where along the tube the coke actually forms.

Cracking is a conversion and selectivity correlation calibrated for ethane. A naphtha furnace has a far richer product slate and a much stronger yield dependence on temperature, which is exactly where the balancing trade in scenario 4 becomes expensive, and the application cannot show it.

The tube metal thermocouple is treated as accurate. In industry the argument about balancing on tube metal is largely an argument about whether the measurement can be trusted, and a simulator that assumes it can understates the difficulty of adopting the scheme it recommends.

The transfer line exchanger downstream is absent, along with the coking that occurs in it and often ends a real campaign before the radiant coil does.

The decoke restores the tube exactly and takes a fixed thirty hours. Neither is true, and real decoke duration depends on how much coke is present [2]. The restart is also simplified. The decoke holds the outlet at its 880 °C limit and the setpoint ramps down to target as the feed returns, where a real furnace would bring the outlet up from a cooler steam standby. Dilution steam is tied to the hydrocarbon flow here, so the nearly empty coils cannot be held cool with the firing at its 900 °C floor.

The predictive controller predicts with clean tube slopes, and coke weakens them. With 7.5×10^{-4} m^2 K/W of coke on every pass, the outlet's response to the firebox and the metal's response to the pass flow are both 14.4 % weaker. At the start of scenario 10 the hot pass's metal answers its own flow 26.7 % more weakly than the model does, at −3.80 °C per t/h against the model's −5.19 °C. The plan therefore expects more cooling per tonne it moves than it gets, and only the next measured offset corrects it. The model also leaves out a parameter changed by the learner, which it treats as error, and the decoke and restart ramp, during which it stands aside. The plan has no price on run length or on the decoke, which is why it loses scenario 5. Its back-off of one degree under the override is fixed rather than chosen against the decoke trigger or the thermocouple's error.

The plan's cost in computing is visible at the top speed. At 600 h per second in Chrome the plans take 41.9 % of the wall clock, the page draws 15.3 frames per second against 50.5 under the loops, and the run advances about 456 h per second rather than 600. The logged run is identical at every speed, and at 200 h per second the page keeps 59.6 frames per second.

Finally, the application models one furnace. An ethylene plant runs a bank of them, and the scheduling problem of staggering their decokes against a plant demand is a large part of the real economics [7] and is entirely outside this model.

> ⚠ **OPEN CHALLENGE** **How Fast Will This Furnace Coke?**
>
> Coke formation is still one of the most important problems of steam cracking. Its properties had hardly been studied for a decade, and reported values of some of them differed by orders of magnitude. A 2020 study measured them on industrial samples and showed what the gap costs: across the values in use, the predicted run length of an industrial furnace on ethane ranged from 20 to 227 days. Progress means coke properties measured well enough for run length models to predict [1]. The coking law of Equation (2.4) is exact in the simulator and uncertain on every real furnace.

2.9 Discussion

Three results generalize past cracking furnaces.

The first is that a balancing scheme equalizes what it measures, and if the constraint lives on the other side of a resistance the scheme can make things worse. The measured case here is 841 h against 888 h for no control, and it arises from a sign: the resistance makes the constrained variable rise and the measured variable fall, so the scheme moves the wrong way. Before configuring any balancing loop it is worth writing down what is between the instrument and the limit.

The second is that an optimum can be invisible to every continuous measurement on a unit. Conversion and ethylene rate both rise monotonically through the optimum in Table 2.5, and only production per hour of elapsed time finds it. Where that happens, an optimizer needs a model rather than a measurement, and the model has to include the thing that ends the run [8].

The third is the shape of the optimum rather than its location. Five degrees above the peak costs 0.5 % of the on-stream rate and fifteen degrees costs 5.9 %, while fifteen degrees below it costs 2.7 %. An optimum with a flat top and a steep side should be approached from the flat side, which is a statement about where to put the setpoint rather than about how to find the peak.

OPEN CHALLENGE **A Cracking Furnace Without Flames**

Gas-fired cracking furnaces emit about 366 million tonnes of carbon dioxide a year, and improvements to the present designs cannot reach the cuts now required. Electric heating avoids the heat lost with the flue gas, and resistance heating could be retrofitted to existing crackers. It still has to be tested and demonstrated, and the furnace's heat is woven into the rest of the plant: the steam it raises drives the turbines of the compressors, which would become electric motors. Progress means a demonstrated electric cracker and a plant whose heat integration is redesigned around it [20].

2.10 Conclusion

A cracking furnace ends its campaign when one tube reaches its metal limit, and the instrument that scheme designers reach for is on the wrong side of the coke. In the simulated furnace of this chapter, four passes whose coil outlet temperatures are matched to within half a degree still carry 16.1 °C of tube metal spread, and on a furnace with one degraded burner, balancing on coil outlet temperature gives a run of 841 h against 888 h for no balancing at all and 910 h for balancing on tube metal. The severity optimum is equally hidden. Conversion and ethylene rate both rise monotonically from 835 to 865 °C while ethylene per hour of elapsed time peaks near 850 °C at 20.31 t/h and falls to 19.12 t/h by 865, and the run length falls from 2516 h to 186 h across the same range.

A model predictive controller offered beside the loops shortens the time above the tube limit in the coking case from 260.5 to 228.5 h, but its ceiling 7 °C under the decoke trigger costs 0.9 % of on-stream rate on the degraded burner furnace. A fuller model would add an axial profile so that the location of the coke can be studied, a naphtha feedstock so that the yield cost of unequal pass severities becomes visible, a transfer line exchanger, and a bank of furnaces so that decoke staggering can be scheduled. The furnace as it stands shows what a balancing scheme actually balances.

KEY POINTS

- A balancing scheme evens out what it measures, and on a cracking furnace the limit is on the tube metal, on the other side of the coke from the outlet thermocouple.
- Four passes with outlets matched to within half a degree still carry 16.1 °C of tube metal spread.
- With one degraded burner, balancing on coil outlet temperature gives a run of 841 h against 888 h for no balancing and 910 h for balancing on tube metal.
- Ethylene per elapsed hour peaks near 850 °C while conversion and the ethylene rate keep rising, so no continuous measurement finds the optimum.

Where else this applies. Fouling that walks toward a hard temperature limit under an override returns in Volume 6, Chapter 4, and a protective override that acts only near its limit in Volume 4, Chapter 2.

References

[1] M. Geerts, S. H. Symoens, P. A. Reyniers, G. B. Marin, M.-F. Reyniers, and K. M. Van Geem, "Steam cracking coke properties and their influence on furnace run length predictions: Experimental and modeling study," *Industrial & Engineering Chemistry Research*, vol. 59, no. 52, pp. 22 460–22 472, 2020.

[2] Y. Zhang, W. Cai, and F. Xin, "Simulation and optimization of coil decoking in ethane pyrolysis furnace," *Chemical Engineering Communications*, vol. 196, no. 8, pp. 950–968, 2009.

[3] M. Rosenzweig, "A petrochemical centennial approaches," Chemical Processing, August 7, 2019, chemicalprocessing.com/articles/2019/a-petrochemical-centennial-approaches, read November 2026.

[4] R. D. Stief, *A History of Union Carbide Corporation: From the 1890s to the 1990s*. Danbury, CT: Carbide Retiree Corps, 1998, chapter 2, read in the copy at carbidecamps.net, November 2026.

[5] M. Neelis, E. Worrell, and E. Masanet, "Energy efficiency improvement and cost saving opportunities for the petrochemical industry: An ENERGY STAR guide for energy and plant managers," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-964E, 2008.

[6] K. Barazandeh, O. Dehghani, M. Hamidi, E. Aryafard, and M. R. Rahimpour, "Investigation of coil outlet temperature effect on the performance of naphtha cracking furnace," *Chemical Engineering Research and Design*, vol. 94, pp. 307–316, 2015.

[7] J. Tian, Y. Wang, and X. Feng, "Simultaneous optimization of flow velocity and cleaning schedule for mitigating fouling in refinery heat exchanger networks," *Energy*, vol. 109, pp. 1118–1129, 2016.

[8] S. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[9] W. Shen, Z. Tian, L. Zhao, and F. Qian, "Life cycle assessment and multiobjective optimization for steam cracking process in ethylene plant," *ACS Omega*, vol. 7, no. 18, pp. 15 507–15 517, 2022.

[10] R. van Welsenaere and G. Froment, "Parametric sensitivity and runaway in fixed bed catalytic reactors," *Chemical Engineering Science*, vol. 25, no. 10, pp. 1503–1516, 1970.

[11] M. Morbidelli and A. Varma, "Parametric sensitivity and runaway in fixed-bed catalytic reactors," *Chemical Engineering Science*, vol. 41, no. 4, pp. 1063–1071, 1986.

[12] H. Falcke, S. Holbrook, I. Clenahan, A. López Carretero, T. Sanalan, T. Brinkmann, J. Roth, B. Zerger, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of large volume organic chemicals," Publications Office of the European Union, Luxembourg, EUR 28882 EN, 2017.

[13] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[14] 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.

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

[16] 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.

[17] S. James, "Reformer and heater tubes temperature accuracy: Comparing practical infrared techniques," Conference paper, 2015, irinfo.org/images/articles/pdf/02-01-2015-james.pdf, read November 2026.

[18] 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.

[19] 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.

[20] M. E. H. Tijani, H. Zondag, and Y. Van Delft, "Review of electric cracking of hydrocarbons," *ACS Sustainable Chemistry & Engineering*, vol. 10, no. 49, pp. 16 070–16 089, 2022.

2.11 Exercises

2.1 ⌕ Concept check · 2 min

Pass 4 of a furnace balanced on coil outlet temperature carries more coke than the other three passes. Compared with the others, what do its outlet and skin thermocouples read, and which way does the balancing scheme move its feed?

(a) Its outlet and its skin both read hot, so the scheme adds feed to pass 4

(b) Its outlet and its skin both read cold, so the scheme takes feed from pass 4

(c) Its outlet reads cold and its skin reads hot, so the scheme takes feed from pass 4

(d) Its outlet reads hot and its skin reads cold, so the scheme adds feed to pass 4

2.2 ⌕ Concept check · 1 min

Table 2.5 lists five coil outlet setpoints. Which quantity, read on the furnace, locates the best of them?

(a) The analyzer conversion, which levels off at the best setpoint

(b) The instantaneous ethylene rate, which peaks where selectivity losses overtake conversion gains

(c) The selectivity, which is highest at the setpoint with the longest run

(d) The ethylene per hour of elapsed time with the decoke counted, known only after the decoke

2.3 Concept check · 2 min

All four passes coke together while the firing holds the coil outlet. How does the coking rate develop over the run?

(a) It grows as the square of itself, so it would become unbounded at a finite time t^*

(b) It grows exponentially in time, with a fixed e-folding time set by θ_c

(c) It stays at its first-day value, because the held outlet fixes the conversion

(d) It falls slowly, because the growing coke lowers the flux through each tube

2.4 Theory and calculation · 6 min

When the outlets of scenario 3 first agree to within half a degree, the firebox is at 1178.9 °C, the outlets are at 845 °C and the mean resistance $R = R_w + R_c$ is 1.77×10^{-3} m^2 K/W. Take $h_{fb} = 0.73\,\mathrm{kW/(m^2\,K)}$.

(a) Evaluate the sensitivity of Equation (2.8). *Checkpoint:* 4.64×10^4 °C per m^2 K/W.

(b) Pass 4 carries 3.5×10^{-4} m^2 K/W more coke than pass 1. Estimate the metal spread and compare it with the 16.1 °C the model reports. *Checkpoint:* 16.2 °C.

(c) Evaluate the same sensitivity on clean tubes at the design point ($T_{fb} = 1150.9\,^\circ\mathrm{C}$, COT = 845.0 °C, $R = R_w = 1.50 \times 10^{-3}\,\mathrm{m^2\,K/W}$), and say why it is larger. *Checkpoint:* 5.09×10^4 °C per m^2 K/W.

2.5 Theory and calculation · 7 min

The clean furnace runs at its design point: tube metal 1004.9 °C, flux 106.6 kW/m^2, coil outlet held at 845 °C, conversion and steam ratio at their reference values and coke factor one. Take $k_c = 2.44 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$, $T_{c,ref} = 1004\,^\circ\mathrm{C}$, $\theta_c = 30\,^\circ\mathrm{C}$ and a metal limit of 1100 °C.

(a) Find the coking rate $\dot{R}_0$ from Equation (2.4) and the rate at which the metal starts to rise. *Checkpoint:* 2.51×10^{-7} m^2 K/(W h) and 0.0268 °C/h.

(b) Find t^* and the run length to the limit. *Checkpoint:* 1120 h and 1073 h, within 2 h.

(c) Find the run length that a straight line through the first day's metal slope predicts, and its ratio to the run of (b). *Checkpoint:* 3550 h, 3.3 times too long.

2.6 Theory and calculation · 6 min

A planner notes that a coil outlet of 865 °C makes 6.8 % more ethylene than 845 °C and proposes the change. Use Equations (2.5) and (2.6) with $F = 40\,\mathrm{t/h}$, $X_0 = 0.650$, $\sigma_0 = 0.800$, $\mathrm{d}X/\mathrm{d}T = 0.004\,20/^\circ\mathrm{C}$ and $\mathrm{d}\sigma/\mathrm{d}X = -0.52$, and a decoke of 30 h.

(a) Find the conversion, the selectivity and the ethylene rate at 850 °C. *Checkpoint:* 0.671, 0.789 and 21.18 t/h.

(b) The run at 850 °C lasts 700 h. Find the on-stream rate with Equation (1.3). *Checkpoint:* 20.31 t/h.

(c) The rows of Table 2.5 at 835 and 865 °C make 19.99 and 22.21 t/h for 2516 and 186 h. Find their on-stream rates, how far each lies below the result of (b), and answer the planner. *Checkpoint:* 19.75 t/h, 2.7 to 2.8 % below, and 19.12 t/h, 5.8 to 5.9 % below.

2.7 App challenge · 8 min

Load scenario 3, Equal outlets, unequal metal, and change no setting (Pass balancing on coil outlet, default seed). Run at 60 h per second until the clock passes 480 h, pause and export the CSV file.

(a) Find the first row on which `cot_spread_C` is below 0.5 °C. Give its time and the tube metal spread `spread_C` on that row. *Checkpoint:* 8.4 h (the row at 8.375 h) and 16.1 °C, within 0.1 °C.

(b) Read `spread_C` on the last row at or before 240 h and on the last row at or before 480 h, and the largest `cot_spread_C` after the row of (a). *Checkpoint:* 20.9 °C and 34.0 °C, within 0.1 °C. The outlet spread stays under 0.5 °C.

(c) Compare (a) with Exercise 2.4 and explain why the outlet scheme cannot remove the spread, and why it grows.

2.8 App challenge · 8 min

Load scenario 5, A burner problem on one corner, as shipped (Pass balancing on coil outlet, decoke on tube metal at 1094 °C, default seed). Run at 60 h per second until the clock passes 950 h, pause and export the CSV file. Press Reset, set Pass balancing to On tube metal before pressing Run, and repeat. Press Reset again, set Pass balancing to Off, and repeat.

(a) In each file, read `run_hours` on the first row with `decoking` equal to 1. *Checkpoint:* 841.3 h, 910.4 h and 887.9 h, within 0.2 h.

(b) On the row at 719.9 h of each file, read `flow0_tph` and `flow3_tph`. *Checkpoint:* 10.32 and 9.03 t/h, 9.66 and 11.01 t/h, 10.00 and 10.00 t/h, within 0.02 t/h.

(c) Explain from (b) why balancing on coil outlet gives a shorter run than no balancing at all.

2.9 App challenge · 9 min

Load scenario 10, Balance the passes before the override, as shipped (Supervisory control on MPC, default seed). Run at 10 h per second to about 115 h, switch to 1 h per second, and pause when the clock reads between 120 and 121 h. Read the IAE tile of the Conversion SP / PV loop on the scorecard and export the CSV file. Press Reset, set Supervisory control to Loops before pressing Run, and repeat.

(a) Give the IAE under each structure. *Checkpoint:* MPC 2.176 to 2.225, Loops 2.986 to 3.040, as the tile prints them.

(b) In each file, over the 359 rows up to 120 h, count the rows on which `tmt_override` is 1, and find the largest `tmt_max_C`. *Checkpoint:* MPC 0 rows and 1087.3 °C. Loops 153 rows and 1094.8 °C.

(c) Explain why the override never acts under the MPC, and what the plan moves instead.

2.10 Application area · 4 min

A plant installs a decoke system that takes 20 h instead of 30 h. Each run starts from clean tubes, so the run lengths of Table 2.5 do not change.

(a) With Equation (1.3) and the table's ethylene rates and runs, find the on-stream rate at 845, 850, 855 and 865 °C. *Checkpoint:* 20.42, 20.59, 20.63 and 20.05 t/h.

(b) Say where the optimum moves and why, and what setpoint you would choose.

2.11 Application area · 3 min

The predictive controller holds each tube under a ceiling one degree below the override's start, and the override starts Δ_o below the 1100 °C limit. Skin thermocouples read the tube surface, and suppose the one on pass 4 reads 5 °C below the true metal.

(a) Find the ceiling for $\Delta_o = 12\,^{\circ}\mathrm{C}$ and for $\Delta_o = 6\,^{\circ}\mathrm{C}$, and the true metal of pass 4 when its reading sits on the ceiling. *Checkpoint:* 1087 °C and 1093 °C, with true metal 1092 °C and 1098 °C.

(b) Use Equation (1.13) to say what the 13 °C between the limit and the shipped ceiling has to cover, and which margin you would ship.

2.12 Review the agent · 5 min

An AI assistant reads Table 2.6 and writes: “In scenario 5 the predictive controller makes 0.9 % less ethylene per elapsed hour than the loops. The cause is its one-hour prediction horizon, which cannot see the coke that builds over weeks. Set the Prediction horizon to its 3 h maximum and the plan will beat the loops on scenario 5.” Find the flaw, and say how you would check the claim in the app.