INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 12 Control in Petrochemicals and Polymers

CHAPTER 1

Introduction: Quality and Safety Limits From Cracker to Pellet

Pages 1–23

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Quality and Safety Limits From Cracker to Pellet," 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. 1, pp. 1–23.

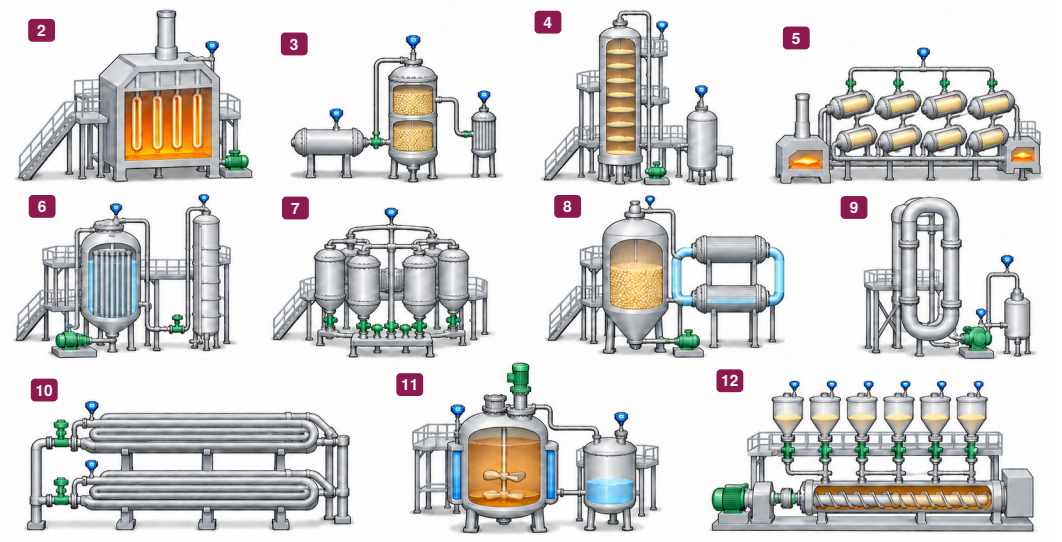
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 1

Introduction: Quality and Safety Limits From Cracker to Pellet



The chapters 2 cracking furnace 3 acetylene converter 4 propylene splitter 5 dehydrogenation beds 6 ethylene oxide reactor 7 paraxylene moving bed 8 gas-phase polyethylene 9 polypropylene loop 10 tubular LDPE reactor 11 butyl rubber reactor 12 additive feeders

Notice that the top row makes and purifies the olefins, the middle row converts or separates them and holds two polymer reactors, and the bottom row holds two more and the feeders that dose additives into the resin. Illustration, not to scale.

IN PLAIN TERMS **From cracked gas to plastic pellets**

The main idea. A petrochemical plant heats gas and light oil until the molecules break into small pieces such as ethylene. Other units clean the pieces and join most of them into long chains that leave as plastic pellets and rubber.

The challenge. Each unit sells a quality, like purity or how easily a plastic melts, that is checked only every few minutes or hours. Each unit also runs near a limit it must never cross, like a tube that gets too hot.

Why it matters. Good control decides how much product can be sold and keeps the plant safe.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.4 and 1.4.6, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example P-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.8 with Table 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A petrochemical complex cracks ethane, propane or naphtha into olefins, separates them to polymer grade, converts part of them into intermediates and polymerizes the rest into resins that leave the site as pellets. Each step sells a property that is measured rarely, a purity, a selectivity, a melt index, a Mooney viscosity or an additive content. Each step also runs close to a limit that ends the campaign or damages the plant if it is crossed, a tube metal temperature, a runaway, a decomposition, a cavitating

pump or a refrigeration floor. This volume answers one question: how does each unit hold a quality it seldom sees against a safety or equipment limit it must never cross, and which control techniques recur along the chain from the cracker to the pellet?

Seven control problems define the industry, and most chapters meet more than one of them.

- *Severity is bought with run length.* A hotter cracking furnace makes more product per hour and cokes its tubes faster, and the coke ends the run on a limit no loop controls.
- *Selectivity and runaway are one event.* The wanted and the unwanted reaction share a catalyst, the unwanted one releases more heat and speeds up faster with temperature, so a bed that loses selectivity also loses its margin to runaway.
- *Heat removal sets the capacity of a polymer reactor.* A polymerization releases several megajoules for each kilogram of polymer, and the reactor makes only as much as its recycle gas, its jacket or its injection pattern can carry away.
- *Grades change on a schedule.* A polymer plant makes one grade after another, and the cost of each change is the tonnage made between specifications.
- *Quality is inferred.* Analyzers report every few minutes, laboratories every few hours, and product quality is carried between them by temperatures, gas ratios, balances and models.
- *Some units run in cycles.* Regenerative beds, simulated moving beds, furnace decokes and reactor washes repeat, and they are corrected once a cycle from the outcome of the last.
- *The limit is measured in the wrong place.* The tube metal sits behind coke, a hot spot runs ahead of its thermocouple, and an impeller sees a lower pressure than the transmitter reads.

Eleven chapters answer the question on eleven plants, each with an app that the reader runs in a browser. Chapters 2 to 5 make and purify the olefins: an ethane cracking furnace, an acetylene converter, a propylene splitter and a bank of propane dehydrogenation beds. Chapters 6 and 7 are an oxidation reactor and an aromatics separation. Chapters 8 to 11 are four polymer reactors, and Chapter 12 doses the additives into the resin.

A reader who has worked through the volume can do the following on a unit that is new to the reader.

- Predict the run length of a fouling or coking surface from its first days and find the severity that maximizes production per elapsed hour.
- Compute how much heat the unwanted reaction makes and how close a cooled bed sits to the end of its cool operating point.
- Say how a polymer reactor removes its heat, what limits it and what spends the margin.
- Estimate the off-grade tonnage of a grade change and the preparations that reduce it.
- Correct an inferred quality from a late laboratory result, and tune a loop that acts once a cycle.
- Name, for each unit, the quality held, how often it is seen, the limit that binds and the structure that guards it.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, ratio, cascade and feedforward from Volume 3, selectors and overrides from Volume 4, dead time and estimation from Volume 5, run-to-run control from Volume 8 and predictive control from Volume 9. The refinery that supplies naphtha and aromatics is Volume 11.

1.2 A Short History

The milestones of Figure 1.1 follow the chain of Figure 1.2, from the cracker and its first products to the polymer reactors, and then to the control that came to both.

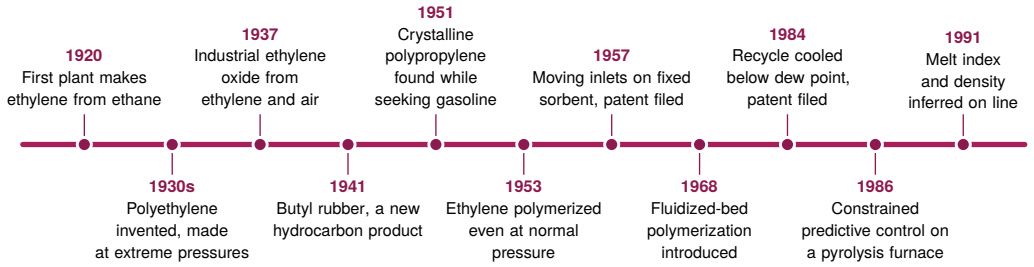


Figure 1.1 Milestones of petrochemicals, polymers and their control. The dates are from [1] (1920), [2] (1930s and 1951), [3] (1937), [4] (1941), [5] (1953), [6] (1957), [7] (1968), [8] (1984), [9] (1986) and [10] (1991). A patent is dated by its filing where the card says so. Notice that the products came first and the control of their quality came last. Events are evenly spaced, not to a time scale.

> ⧖ HISTORY **1920 to 1958: Ethylene and Its First Oxide**
>
> The first plant to make ethylene from ethane started up in January 1920 [1]. The small works-scale plant that followed in 1921 cracked the gas in three-inch tubes inside a gas-fired refractory furnace [11], a furnace of the kind Chapter 2 models. Ethylene oxide was once made by a costly two-stage route through a chlorohydrin, with a large load of effluent. Industrial production by direct oxidation of ethylene with air started in 1937, and processes on oxygen followed in 1958. The direct route replaced the old one because of the cost of chlorine and the effluent, and in 2013 the oxygen process was the only route in use in the European Union [3]. Chapter 6 runs that process in its recycle loop.

A 1941 paper announced butyl rubber as a new hydrocarbon product. Its authors had found that long chains of simple olefins such as the butylenes carry most of the properties of rubber, but that the saturated polymers cannot be vulcanized in the usual way [4]. A little diolefin such as isoprene, copolymerized with isobutylene at low temperature, gave a polymer of very low unsaturation that could still be cured with sulfur [12]. That is the recipe of Chapter 11. In 1957 a patent was filed for a separation in a fixed bed of sorbent whose inlets and outlets are moved in steps, so that the bed behaves as if the solid flowed against the liquid [6]. Chapter 7 moves its ports the same way.

> ⧖ HISTORY **1951 to 1954: Polyolefins at Low Pressure**
>
> The polyethylene of the 1930s needed pressures of 20 to 30 thousand pounds per square inch and gave a branched, low-density polymer. In 1951 J. Paul Hogan and Robert L. Banks, trying to turn propylene into gasoline over a chromium catalyst, made crystalline polypropylene instead. In less than a year they had a process for a stiffer, high-density polyethylene that needed only a few hundred pounds per square inch [2]. Near the end of 1953 Karl Ziegler's group found that ethylene polymerizes even at normal pressure on catalysts that are easy to prepare [5], and early in 1954 Giulio Natta's group polymerized propylene with them [13]. The reactors of Chapter 8 and Chapter 9 run at 21.7 and 38 bar, while the tube of Chapter 10 still runs at 2500 bar.

Polymerization in a fluidized bed of growing powder was introduced in 1968 and now accounts for about 40 percent of the world's polyethylene capacity [7]. A patent filed in 1984 cooled the recycle gas below its dew point and returned the two-phase mixture to the bed, because the polymer a reactor of a given size can make is set by the heat it can withdraw [8]. That is the condensed mode of Section 1.4.4 and of Chapter 8. Control reached these plants through models. A 1986 description of constrained predictive control closed with results from a pyrolysis furnace, where it held three stream temperatures through the fuel gas pressure of three burners [9]. In 1991 the melt index and density in an industrial fluidized-bed reactor were predicted on line from temperatures and gas compositions,

with the model corrected from infrequent laboratory results [10]. Chapter 2 and Chapter 8 build on both ideas.

1.3 Notation and Conventions

Temperatures are in °C, differences of temperature in kelvin, and absolute temperatures in a rate law in kelvin. Mass flows are in t/h and duties in MW. The coil outlet temperature of a furnace is COT and its tube metal temperature TMT. A conversion is X and a selectivity S, the share of the converted feed that becomes the wanted product. An activation energy is written as an activation temperature E/R in kelvin. A resistance to heat transfer is R_f in $\mathrm{m^2\,K/W}$, a bed or polymer inventory is M_b and its residence time τ_s. Melt index is MI, and $q = \ln \mathrm{MI}$. A cycle-to-cycle loop counts time in reports k. The framework's h is its integration step and its a the controller action, so neither is used here for anything else.

The chapters keep their own symbols, and several collide. Chapter 2 writes selectivity as σ, coke resistance as R_c and the analyzer cycle as θ. Chapter 3 uses R for the hydrogen ratio, L for the ethylene lost and a for catalyst activity, as Chapter 5 does. Chapter 8 uses L for bed level and R_H for the hydrogen ratio, while Chapter 9 writes the ratio r_H and the melt flow rate Q. The m_j of Chapter 7 are flow-rate ratios and the m of Chapter 11 is the cooling margin. Mooney viscosity is ML in Chapter 11, and the additive concentrations of Chapter 12 are in ppm by mass.

1.4 From Cracker to Pellet on One Chain

1.4.1 The Olefin, Aromatics and Polymer Chain

Figure 1.2 draws the chain and marks where each chapter's plant sits in it [14]. Ethane, propane or naphtha is cracked with steam in fired furnaces, the cracked gas is quenched and compressed, and a cold train separates it into cuts. The two-carbon cut passes an acetylene converter before ethylene is split from ethane, and the three-carbon cut goes to a propylene splitter, which also takes the propylene of a dehydrogenation plant [15]. Ethylene feeds an oxidation reactor and two polyethylene processes, propylene a polypropylene reactor, and isobutylene with a little isoprene a butyl rubber reactor. Paraxylene comes from the aromatics of a reformer and is separated by adsorption. The resins are compounded with additives into pellets.

The chapters take the plants one at a time, and the techniques they share are developed below on one running example, a small chain called P-1. P-1 cracks ethane in four identical furnaces, one of which is followed here. Part of its ethylene goes to a cooled oxidation tube and 20 t/h to a gas-phase polyethylene reactor, and a bank of regenerative beds makes propylene from propane. Every number of P-1 is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Furnace Severity, Coking and the Run Length

A cracking coil absorbs heat through its wall, and coke laid down inside the wall adds a resistance between the metal and the gas. With the outlet held by the firing loop, the heat flux q through each square meter is fixed by the cracking duty, and the metal runs above the gas by the flux times the resistance. Coke formation and run length are studied by experiment and model [16], and a simple form lets the coking rate grow by a factor e for every θ_c kelvin of metal temperature. Together these give Equation (1.1).

$$\mathrm{TMT} = \mathrm{COT} + q\,(R_w + R_f), \qquad \frac{\mathrm{d}R_f}{\mathrm{d}t} = \dot{R}_0\,\mathrm{e}^{(\mathrm{TMT}-\mathrm{TMT}_0)/\theta_c} \tag{1.1}$$

where R_w is the clean wall, $\dot{R}_0$ the coking rate at the start of the run and TMT_0 the metal temperature then. Each unit of coke raises the metal by q, and the hotter metal cokes faster. Integrating from a

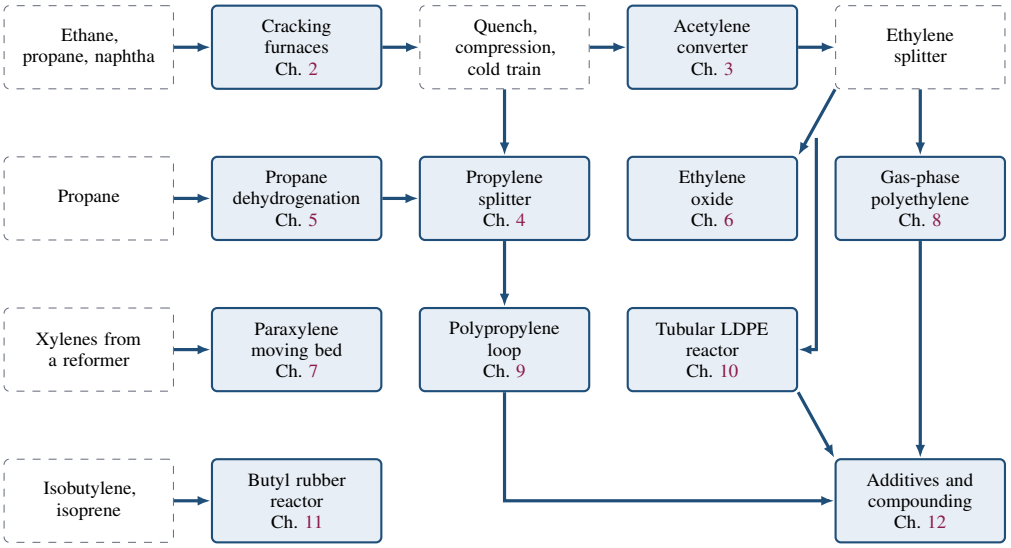


Figure 1.2 The Olefin, Aromatics and Polymer Chain. Filled Blocks Are the Plants of Chapters 2 to 12, Dashed Blocks Lie Outside the Volume. The Two-Carbon Cut Passes the Acetylene Converter, the Three-Carbon Cut Goes to the Propylene Splitter, and Ethylene From the Splitter Feeds the Oxidation Reactor and Both Polyethylene Reactors. Light Ends, Fuel Gas and Most Products Are Omitted.

clean tube gives the metal rise and the run length of Equation (1.2).

$$\Delta\mathrm{TMT}(t) = -\theta_c \ln\!\left(1 - \frac{t}{t^*}\right), \qquad t^* = \frac{\theta_c}{q\dot{R}_0}, \qquad t_{run} = t^*\left(1 - \mathrm{e}^{-\Delta_M/\theta_c}\right) \tag{1.2}$$

where Δ_M is the margin from the clean metal to its limit. The rise is a blow-up in finite time at t^*, not an exponential, and the limit is reached before it.

P-1's furnace cracks 30 t/h of ethane at a COT of 840 °C to 65 % conversion at 80 % selectivity, 15.6 t/h of ethylene. The clean metal reads 1000 °C against a limit of 1090 °C, the flux is 100 kW/m^2, $\dot{R}_0 = 2.5 \times 10^{-7}\,\mathrm{m^2\,K/(W\,h)}$ and $\theta_c = 30\,\mathrm{K}$. The metal starts by rising 0.025 K/h, so $t^* = 1200\,\mathrm{h}$ and the run lasts 1140 h, or 47.5 days, as shown in Figure 1.3a. A straight line through the first day predicts 3600 h, more than three times too long. Half the run spends only 21 % of the margin, and the last 30 K take 103 h. With a constant fouling rate, the limit $\theta_c \to \infty$ of Equation (1.2), the run is simply $\Delta_M/(q\dot{R}_0)$.

A firing loop that holds the COT hides all of this. The outlet stays on setpoint while the firing walks up, and the metal limit is a consequence of the firing, not a variable any loop sets. Chapter 2 derives Equation (1.2) for its own furnace and measures a run of 1073 h. It also shows that four passes balanced to the same outlet temperature leave 16.1 K of metal spread, because a coked pass reads cold at the outlet and hot at the skin. On a furnace whose fourth pass cokes faster, balancing on the outlets gives a run of 841 h against 888 h with no balancing at all.

Severity trades rate against run. A hotter outlet converts more and selects less, and it raises the metal and the coking rate together. The furnace must also stop to decoke, so the quantity to maximize is the production per elapsed hour of Equation (1.3).

$$Y = \frac{P\,t_{run}}{t_{run} + t_{down}} \tag{1.3}$$

where P is the production rate while running and t_{down} the time out of service. On P-1, with selectivity falling 0.8 for each unit of conversion, conversion rising 0.004 per kelvin of COT, the metal following

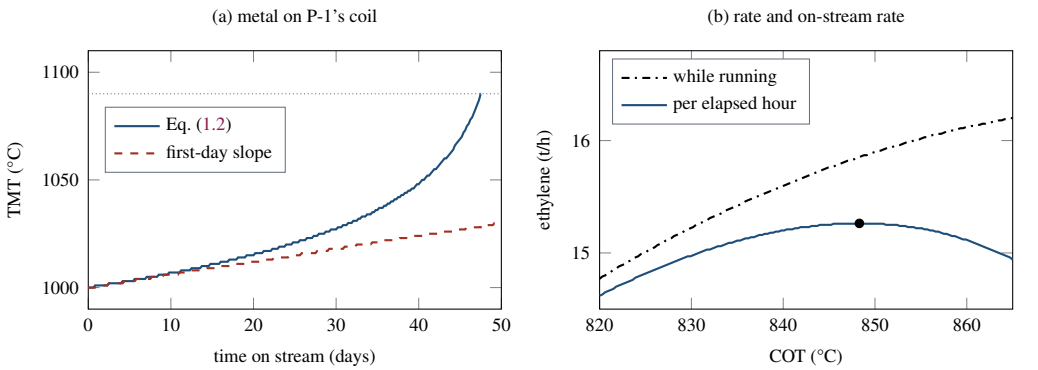


Figure 1.3 P-1's Furnace. (a) Tube Metal Over the Run (Solid), Reaching the 1090 °C Limit (Dotted) at 1140 h, Against the Straight Line Through the First Day (Dashed). (b) Ethylene While Running (Dash-Dotted) and Per Elapsed Hour With a 30 h Decoke (Solid), Whose Flat Maximum (Dot) Lies at 848.3 °C.

the outlet one for one, the coking rate growing as the 1.8 power of conversion and a decoke of 30 h, the design gives 15.20 t/h per elapsed hour. The optimum lies at 848.3 °C with 15.26 t/h and a run of 777 h, as shown in Figure 1.3b. The top is flat, 0.41 % above the design value, while the instantaneous rate keeps rising past it. Neither the conversion nor the rate can find this optimum, because both are still improving where it has passed. Only the run's total, read after the decoke, can find it.

The same arithmetic recurs. Chapter 2 measures an on-stream optimum whose top is flat between 845 and 855 °C. Chapter 11 fouls its jacket at a constant rate, reads its run of 48.9 h from the linear form, and weighs the run against an eight-hour wash with Equation (1.3). Its optimum reactor temperature is warmer than intuition expects, because a colder reactor starts closer to the refrigeration floor, which shortens the run more than the slower fouling lengthens it. Cleaning schedules against fouling are an optimization subject of their own [17]. Chapter 3 loses catalyst activity to green oil over months and Chapter 8 lets its cooler foul. In each case the controlled variable stays on setpoint and the actuator, or a margin computed from it, is the clock.

> **✦ DID YOU KNOW? The Largest-Volume Organic Chemical**
>
> Ethylene is the organic material that industry makes in the largest volume, and most of it becomes polyethylene. In 2013 the olefin plants of Europe had a capacity of about 25 million tonnes of ethylene, about a quarter of the world's [3].

1.4.3 Selectivity and Runaway in Fixed Beds

Many petrochemical reactions run beside an unwanted one on the same catalyst. Ethylene is oxidized to ethylene oxide and burned to carbon dioxide, and acetylene is hydrogenated to ethylene while ethylene is hydrogenated to ethane. When the two rates share their dependence on composition, the ratio of the unwanted to the wanted rate depends on temperature alone, as Equation (1.4) states.

$$\frac{r_2}{r_1} = \frac{r_{2,0}}{r_{1,0}} \exp\!\left[-\frac{E_2 - E_1}{R}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right], \qquad S = \frac{1}{1 + r_2/r_1} \tag{1.4}$$

If the unwanted reaction has the larger activation energy, selectivity falls as the bed warms. If it also releases more heat, its share of the heat is larger than its share of the conversion. The heat generated then rises with temperature as one reaction with the effective activation of Equation (1.5).

$$\frac{E_{eff}}{R} = \varphi_1 \frac{E_1}{R} + \varphi_2 \frac{E_2}{R}, \qquad \varphi_2 = \frac{(1-S)\,\Delta H_2}{S\,\Delta H_1 + (1-S)\,\Delta H_2} \tag{1.5}$$

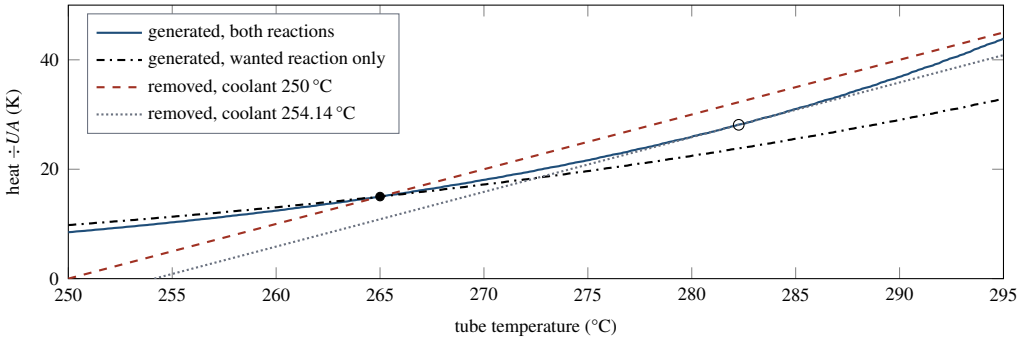


Figure 1.4 P-1's Oxidation Tube. Heat Generated by Both Reactions (Solid) and by the Wanted Reaction Alone With the Same Heat at Design (Dash-Dotted), Against Heat Removed at the Design Coolant (Dashed) and at the Coolant Where the Line Becomes Tangent (Dotted). The Dot Is the Design State and the Circle the Last Cool State.

where $\varphi_2 = 1 - \varphi_1$ is the unwanted reaction's share of the heat.

P-1's oxidation tube has $E_1/R = 8000\,\mathrm{K}$ and 100 kJ/mol for the wanted reaction and $E_2/R = 12\,000\,\mathrm{K}$ and 1300 kJ/mol for the combustion, at a selectivity of 0.85. The 15 % that burns makes 69.6 % of the heat, so $E_{eff}/R = 10\,786\,\mathrm{K}$. Each kelvin raises the ratio of the rates by 1.4 %, and 10 K hotter the selectivity is 0.832.

A cooled tube balances heat generated against heat removed to its coolant. The generation is exponential in temperature and the removal is linear, and a steady state exists only where the two curves cross. The lowest crossing is lost when the removal line becomes tangent to the generation curve, the classical construction of heat generated against heat removed [18]. With the generation rising as one reaction of activation E_{eff}, tangency requires Equation (1.6).

$$T - T_c = \frac{T^2}{E_{eff}/R} \tag{1.6}$$

A tube that runs closer to its coolant than this has a stable cool state, and the ratio of the two sides measures how much of the room is used. Parametric sensitivity studies of tubular reactors and fixed beds with axial profiles give criteria of the same kind [19–21].

P-1's tube runs at 265 °C, 15 K above a coolant at 250 °C, while Equation (1.6) allows about 27 K. A slightly warmer coolant raises the tube more each time, as Figure 1.4 shows. The cool state ends at a coolant of 254.14 °C, only 4.14 K above design, where the tube has reached 282.3 °C and its selectivity has fallen to 0.818. With the same heat at design made by the wanted reaction alone, the margin would be 13.05 K. The side reaction therefore costs two thirds of the coolant margin as well as 15 % of the feed.

Chapter 6 measures the same structure on a twelve-cell tube, where combustion makes 65 % of the heat at 87.2 % selectivity and the cool state ends 4.2 K above the design coolant. Its recycle loop stabilizes the tube, because a hotter tube burns its own oxygen, but the loop turns over in 7.7 min and the peak overshoots on the way. Its chloride moderator has an optimum, where the sign of the process gain changes and a PID started on the wrong side walks the reactor to its trip. Chapter 3 adds a reagent window open on one side [22]. Too little hydrogen puts the product off grade, too much hydrogenates ethylene, and where a bed runs out of acetylene the released ethylene reaction lights the bed off. Its bed rise first falls and reaches the trip only 3.95 min after the hydrogen step, so a temperature limit cannot see a light-off coming. The peak of Chapter 10 is a small runaway lit on purpose and limited by its initiator. Endothermic beds do not run away. The cracking coil of Chapter 2 and the dehydrogenation beds of Chapter 5 cool themselves as they react, and their conversion is limited by the heat supplied or stored. Runaway as a whole is reviewed in [23].

Table 1.1 How Each Polymer Reactor Removes Its Heat, the Limit, and What Spends the Margin

Plant	Heat leaves through	Limit	What spends the margin
P-1 reactor	recycle gas, 15 MW sensible and 5 MW latent	cooler capacity	more production, less condensing agent
Chapter 8	recycle gas, 19.85 and 5.90 MW	cooling at 88.3 % use, sticking 8.0 K away	condensing agent, low-density grades
Chapter 9	jacket on a liquid-full loop	bubble point 3.43 bar below the pressure	hydrogen at 2.44 bar per unit of ratio, inerts, 0.63 bar/K
Chapter 10	four peaks and the jackets	decomposition at 335 °C, peaks at 290	excursions, the pulse valve, a strong initiator lot
Chapter 11	jacket on a slurry	refrigeration floor −106 °C	polymer fouling the wall, reached at hour 49

1.4.4 Heat Removal in Polymer Reactors

A polymerization releases its heat in proportion to the polymer made, so a reactor's capacity is the heat it can remove, as Equation (1.7) states.

$$P_{max} = \frac{Q_{max}}{\Delta H_p}, \qquad \Delta T_{ad} = \frac{\Delta H_p}{c_p} \tag{1.7}$$

where Q_{max} is the heat the reactor can remove, ΔH_p the heat of polymerization per unit mass of polymer and ΔT_{ad} the rise the monomer would see if it converted completely without cooling. P-1 takes $\Delta H_p = 3.6\,\mathrm{MJ/kg}$, so each tonne an hour of polymer is a megawatt. For ethylene with $c_p = 2.8\,\mathrm{kJ/(kg\,K)}$ the adiabatic rise is 1286 K, or 12.9 K for each percent converted. One front that converted 25 % would heat the gas 321 K.

P-1's gas-phase reactor makes 20 t/h and so must remove 20 MW. A fluidized bed loses little heat through its wall, so its heat leaves in the recycle gas. P-1 circulates 150 kg/s at 2.0 kJ/(kg K), entering at 35 °C and leaving the bed at 85 °C, which carries 15 MW. Dry, the reactor could make only 15 t/h. The other 5 MW come from liquid condensed in the recycle cooler that boils again in the bed. At 0.35 MJ/kg that is 14.3 kg/s, 9.5 % of the recycle, and each further percent condensed adds 0.525 t/h of capacity. The gas-phase processes, their cooling and their property correlations are reviewed in [24].

Each polymer reactor of the volume removes its heat a different way, and each way has its own limit and its own consumer of the margin, as Table 1.1 lists. In Chapter 8 the condensing agent that adds cooling also lowers the temperature at which the powder sticks. The loop reactor of Chapter 9 is liquid full, and the hydrogen that sets its melt flow rate raises the bubble point of the liquid toward the pressure the circulation pump needs. The tube of Chapter 10 cannot pass its heat through the wall fast enough, so it converts in four separate peaks, each a few tens of kelvin below the temperature at which ethylene decomposes [25]. The slurry of Chapter 11 fouls its jacket, and the run ends when the jacket reaches its refrigeration floor.

Two lessons follow from the table. The variable that sets quality often spends the cooling margin, as hydrogen does in the loop reactor and the condensing agent and comonomer do in the fluidized bed. The quality loop and the limit then want the same valve in opposite directions, and Chapter 9 lets them meet at a low selector. The second lesson is that the margin is read in a computed quantity. The cooling utilization of Chapter 8, the bubble-point margin of Chapter 9 and the cooling margin of Chapter 11 are each computed from several measurements, and each chapter shows that the computed margin warns earlier than the controlled temperature.

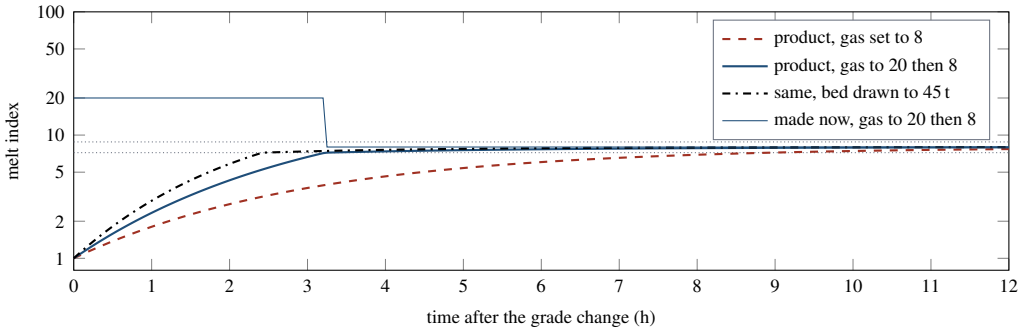


Figure 1.5 A Grade Change on P-1's Reactor From Melt Index 1 to 8. The Product Leaves the Off-Grade Region Below the Band (Dotted, 7.2 to 8.8) at 8.95 h When the Gas Is Set to the Target (Dashed), at 3.23 h When It Is First Driven to 20 (Thick Solid, the Thin Solid Line Being the Melt Index Made at Each Moment), and at 2.42 h With the Bed Also Drawn Down (Dash-Dotted).

> ✦ **DID YOU KNOW? A Reactor More Than a Kilometer Long**
>
> Commercial tubular reactors for low-density polyethylene are typically 1000 to 2500 meters long, built of high-pressure tubes 10 to 15 meters long laid in a serpentine inside a concrete bay. They convert up to 36 percent of their ethylene to polymer, against about 20 percent in an autoclave [26].

1.4.5 Grade Transitions

A polymer property is set by the conditions at the moment a chain grows, but what leaves the reactor is the average of everything the inventory holds [10]. A well-mixed inventory M_b that makes and discharges P obeys the mixing balance of Equation (1.8) for the logarithm of melt index, which averages additively in the usual correlations.

$$\tau_s \frac{\mathrm{d}\bar{q}}{\mathrm{d}t} = q_{inst} - \bar{q}, \qquad \tau_s = \frac{M_b}{P} \tag{1.8}$$

where q_{inst} is the logarithm of the melt index being made now and $\bar{q}$ that of the product discharged. If the gas composition holds q_{inst} constant from the change, the product crosses the edge q_{spec} of the new band after the time and tonnage of Equation (1.9).

$$t_{off} = \tau_s \ln \frac{q_{inst} - q_0}{q_{inst} - q_{spec}}, \qquad W_{off} = P t_{off} = M_b \ln \frac{q_{inst} - q_0}{q_{inst} - q_{spec}} \tag{1.9}$$

where q_0 is the starting value. The production rate cancels. A slower reactor makes less per hour and takes longer by the same factor, so the off-grade tonnage depends only on the inventory and on how far past the target the instantaneous property is driven.

P-1's reactor holds 60 t of powder, so $\tau_s = 3$ h. A change from melt index 1 to 8, with a band of 10 %, must bring the product above 7.2. Setting the gas to make melt index 8 at once leaves 8.95 h and 179 t off grade. Driving the gas to make melt index 20 until the product reaches 7.2, and then to 8, cuts this to 3.23 h and 64.5 t, as shown in Figure 1.5. At half the production rate the same overshoot still costs 64.5 t, over 6.46 h. Drawing the bed down to 45 t before the change cuts it to 48.4 t. Optimal grade transitions have been computed for gas-phase and multistage polyolefin reactors [27, 28].

Chapter 8 derives Equation (1.9), tests it against ten runs of one transition and finds 0.77 to 0.84 t of off-grade resin per tonne of bed against the 0.72 the logarithm gives, because its hydrogen inventory takes 142 min to turn over and the instantaneous property cannot step. Retuning its quality cascade with a drawdown cuts the off-grade tonnage from 179.1 to 77.5 t. The loop reactor of Chapter 9 averages over 52.4 min, and its transition is limited by the bubble-point margin the overshoot in

hydrogen spends. Chapter 10 changes grade through the chain-transfer agent and its peak setpoints, and Chapter 12 times each recipe switch 240 s ahead of the change, the residence of its blender, extruder and transport.

1.4.6 Inferred Product Quality

Few qualities of the volume are measured continuously. A chromatograph reports every few minutes and a laboratory every few hours, so each unit infers its quality between reports from what it does measure, a tray temperature, a gas ratio, a bed temperature or a weighed flow [29, 30]. Three practices recur.

A tray temperature is a bubble point, so it reads pressure as well as composition. The compensation of Equation (1.10) removes the pressure term arithmetically.

$$T_{comp} = T - \frac{\mathrm{d}T}{\mathrm{d}P}\left(P - P_{ref}\right) \tag{1.10}$$

In the propylene splitter of Chapter 4 the composition profile spans 5.80 K over 151 stages, 0.038 K a stage, while pressure moves every tray 3.10 K per bar. A tenth of a bar therefore reads like eight stages of the profile. The same chapter shows that compensating the measurement does not compensate the process, because pressure also changes the relative volatility.

> ✦ DID YOU KNOW? **Purifying Olefins Takes a Share of the World's Energy**
>
> A 2020 review reports that purifying propylene and ethylene accounts for about 0.3 percent of the world's energy use, and it counts the task among seven chemical separations that could change the world [31]. The column of Chapter 4 separates propylene from propane over 151 stages.

A model carries the quality between laboratory results, and each result corrects the model's bias as in Equation (1.11).

$$b_k = b_{k-1} + \lambda\left(y_{lab,k} - \hat{y}(t_{s,k}) - b_{k-1}\right) \tag{1.11}$$

where $\hat{y}(t_{s,k})$ is the model's value at the instant the sample was drawn, not when the result returned, and $0 < \lambda \leq 1$. For a polymer the model is the instantaneous correlation passed through the inventory lag of Equation (1.8). P-1's laboratory draws a sample every 2 h and reports 1 h later, so a result is 1 to 3 h old when it is used. During the direct grade change of Section 1.4.5, a result drawn at 1 h and compared with the model at 2 h would put 0.42 into the bias of ln MI, a factor of 1.53 in melt index, although model and plant agree. Chapter 12 compares each additive assay with its prediction averaged over the sample's own collection interval, and the density estimate of Chapter 10 moves the laboratory value by the change in the mean peak since the sample was drawn.

An inference is only as good as the measurement that corrects it. A frozen chromatograph republishes its last result, and a cascade that trusts it integrates an error nobody is measuring. In Chapter 4 that cascade ends worse than the tray loop alone until a validity hold is armed. When the inlet analyzer of Chapter 3 fails, a hydrogen balance on the outlet analyzer infers the feed acetylene. The weighed flow of Chapter 12 cannot see powder held below the scale, so its loop holds 1.00 kg/h while the product receives 0.600 kg/h, and only the laboratory finds it. Quality control of polymer plants has long been organized around such delayed and inferred measurements [32].

1.4.7 Cyclic Units

Two kinds of cycle appear in the volume. The slow ones are campaigns ended by a decoke, a wash or a catalyst regeneration, and Equation (1.3) prices them. The fast ones are part of the process. The dehydrogenation beds of Chapter 5 store heat during regeneration and spend it on the endothermic reaction, a bed of 36.3 MJ/K cooling 65 K in nine minutes on line. The simulated moving bed of Chapter 7 steps its ports every 90 s, and its purities exist only as twelve-minute averages. In both, a quality is reported once a cycle, and a move of the handle shows only in later reports [33].

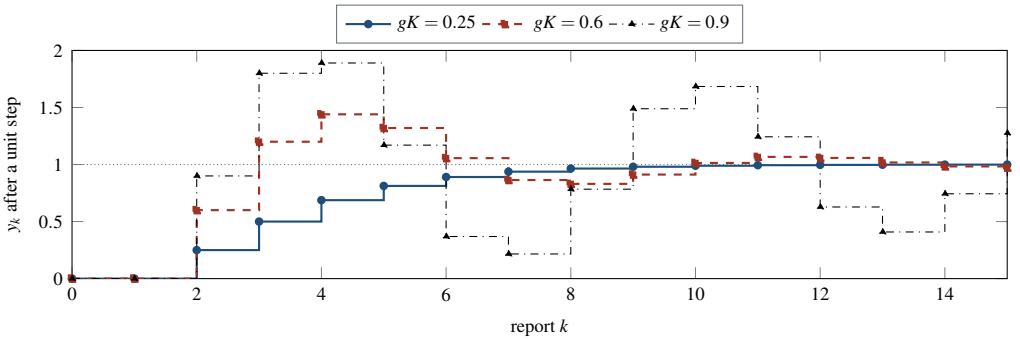


Figure 1.6 A Cycle-to-Cycle Integral Loop on a Plant That Answers Two Reports Late, Equation (1.12) With $d = 2$, After a Unit Setpoint Step: the Double Root (Solid, Dots), a Faster Setting (Dashed, Squares) and One Near the Edge of Stability (Dash-Dotted, Triangles).

A loop on such a unit acts once a report. If the reported quality y_k answers the handle u with gain g after d reports, an integral law that moves the handle by K times each report's error gives the closed loop of Equation (1.12).

$$y_k = g\,u_{k-d}, \qquad u_k = u_{k-1} + K\,(y_{SP} - y_k), \qquad z^d - z^{d-1} + gK = 0 \tag{1.12}$$

For $d = 1$ the loop settles in one report at $gK = 1$. For the delay of two reports that both chapters meet, the characteristic equation is $z^2 - z + gK = 0$. Its roots are a double root at 0.5 for $gK = 0.25$, the fastest setting without overshoot, and lie on the unit circle at $gK = 1$.

P-1's dehydrogenation bank reports its conversion once a cycle and answers the regeneration air by 0.3 points per kelvin two reports later. The integral gain is therefore 0.83 K per point for the double root and 3.33 K per point at the edge of stability. As Figure 1.6 shows, $gK = 0.25$ reaches 5 % of the setpoint at the eighth report, $gK = 0.6$ overshoots 44 % and settles by the thirteenth, and $gK = 0.9$ is still swinging.

The real units add a tail to the pure delay. A regeneration move in Chapter 5 reaches the first bed two slots late, at 0.293 points per kelvin, and a smaller second rise arrives a cycle later as the beds start their next regeneration warmer. Its cycle-to-cycle PI is therefore tuned slower than the pure delay allows, and a predictive controller that knows the delay and the measured feed holds the target sooner. Chapter 7 runs two such loops, on the logarithm of each impurity, with an integral time of two cycles, and draws the region of flow ratios inside which both outlets can be pure, the triangle of equilibrium theory [34]. Its real unit needs a ratio well inside the ideal edge, 0.568 against 0.422, because the beds disperse what the ideal theory sharpens.

1.4.8 Rarely Measured Quality Against Safety Limits

Every chapter of the volume holds a quality that it sees seldom, and every chapter runs near a limit that it cannot afford to cross. The two differ in how they are known. The quality arrives late, every few minutes to every few hours. The limit is measured continuously, but often in the wrong place: behind a coke layer, at a thermocouple that a moving peak has passed, or at a transmitter upstream of an impeller eye. The setpoint that the quality loop may use is what the limit leaves after the largest excursion and the error of its measurement, as in Equation (1.13).

$$\mathrm{SP}_{max} = y_{lim} - e_{max} - e_{meas} - \Delta_{res} \tag{1.13}$$

where y_{lim} is the limit, e_{max} the largest excursion of the controlled variable from its setpoint, e_{meas} how far the true value can run ahead of its measurement and Δ_{res} a declared reserve. Chapter 10 walks its peak setpoints toward exactly this value. Its gas runs up to 10.9 K ahead of every thermocouple while

Table 1.2 For Each Chapter, the Quality Held and How Often It Is Known, the Limit, and the Structure That Guards It

Ch.	Quality, and how it is known	Limit	Guard
2	Conversion on an analyzer every 21 min	Tube metal 1100 °C, behind the coke	One-way override on the COT, decoke trigger
3	Outlet acetylene under 1 ppm on a 5 min chromatograph	Bed rise trip at 70 K	Override at 55 K through a low select on the hydrogen
4	Propylene purity on a 20 min chromatograph, a tray temperature	Reflux and pressure limits	Pressure compensation, analyzer validity hold
5	Conversion on a chromatograph once a 3 min slot	Bed temperature limit	Clamp on the regeneration, a constraint in the MPC
6	Selectivity on a 30 min analyzer	Hot spot, flammable oxygen	One-way override on the production rate
7	Both purities as twelve-minute composites	Edges of the operating triangle	Ratio clamps, purity floors in the MPC
8	Melt index and density from a laboratory, 45 to 105 min old	Sticking temperature, cooling capacity	Six overrides that no switch removes
9	Melt flow rate on a 6 min analyzer behind a 52 min average	Pump cavitation	Bubble-point controller at a low selector, interlocks
10	Melt index every 6 min, density from the laboratory	Decomposition at 335 °C	High-high trip at 325 °C, a supervisor that backs off
11	Mooney viscosity from the laboratory every 2 h	Refrigeration floor	Wash on a schedule or on the cooling margin
12	Additive content from a composite every 3 h, 45 min late	Screw speed, a plugged sock	Level switch, torque diagnostics, a bounded trim

the pressure recovers from a pulse, and its supervisor keeps a reserve of 20 K to the decomposition line.

On P-1's tube the coolant may rise only 4.14 K before the cool state is lost. A coolant loop that swings 1 K leaves 3.14 K, and 1 K of reserve leaves a coolant ceiling of 252.14 °C. A selectivity loop on a half-hourly analyzer that asks for warmer coolant must be stopped well before that, and it learns of its mistake only at the next report. That is why the guards of the volume are fast and act in one direction only. The overrides of Chapter 2 and Chapter 6 can only lower the outlet setpoint and the production rate. Chapter 3 and Chapter 9 take the reagent away from the quality loop through a low selector, and the loop that loses is told what was applied so that it does not wind up.

Table 1.2 lists, for each chapter, the quality held, how and how often it is known, the limit that binds and the structure that guards it. Chapter 13 adds the measured results and asks how each unit holds a quality that is rarely measured against its safety limits.

1.5 Instruments and Final Elements in This Volume

Every loop in this volume begins at a sensor and ends at a final element, the device that changes a flow of material or of energy. Figure 1.7 lists what the controllers of each chapter read and what they move. A temperature is read in ten of the eleven chapters and an online analyzer in nine, and seven chapters act on a quality that is inferred or estimated rather than measured. A heating or cooling duty is moved in nine chapters, a valve or flow loop in seven and a feed of catalyst, initiator or additive in six.

Temperature carries most of the limits of the volume. A thermocouple is a junction of two different metals that behaves like a temperature-sensitive battery [35], as drawn in Figure 1.8. It sits in a thermowell, whose mass adds a first-order lag, and a lag severe enough can compromise the stability of a feedback loop [35]. The gas of Chapter 10 runs up to 10.9 K ahead of every thermocouple while the pressure recovers from a pulse, and Chapter 2 is the only chapter that reads a tube metal.

The analyzers set the pace of quality control. A chromatograph separates the components of a sample in time, as shown in Figure 1.9, so cycles of minutes are not uncommon and some dead time

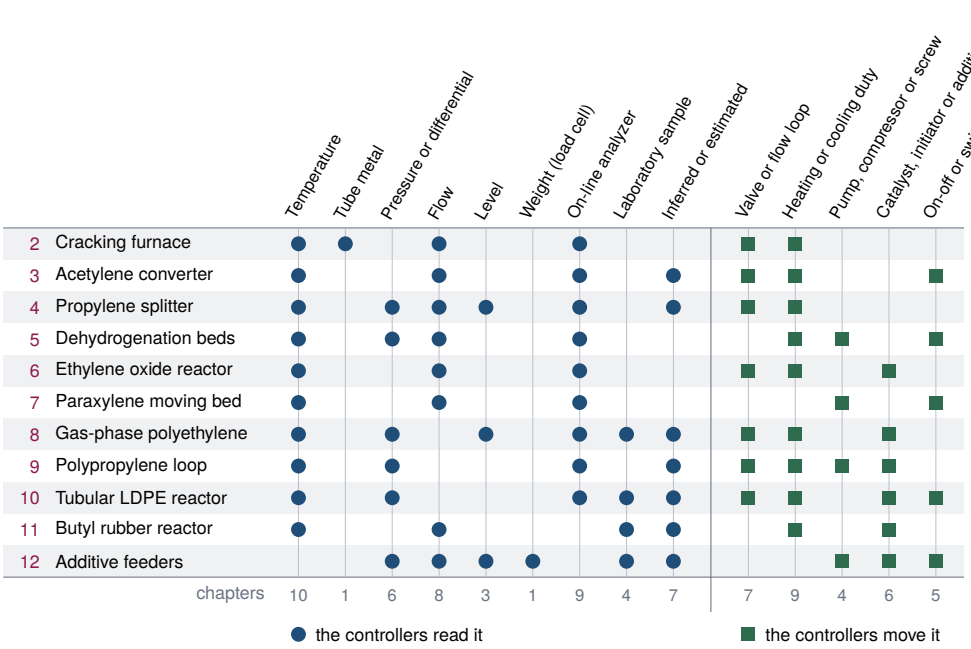


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers or operator read the device, a square mark that they move it. A mark is set only where the chapter's variables table, its instrument text or its section on technology in practice names the device or its reading, so a blank says that the model leaves the device out and not that a real plant lacks it. Notice that only one chapter reads a tube metal and only one weighs its product.

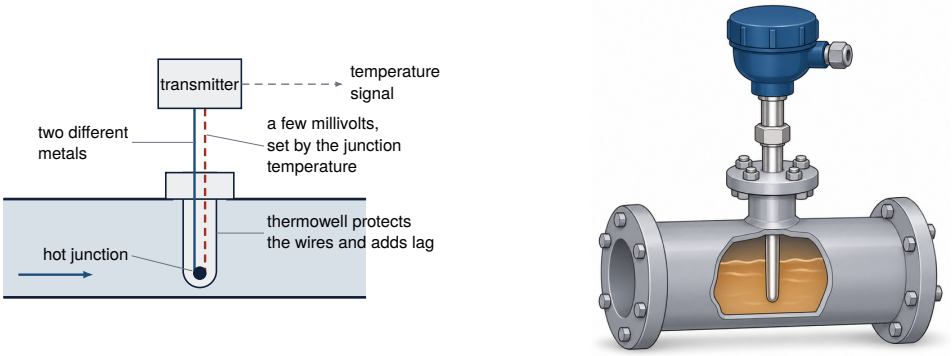


Figure 1.8 Principle of a thermocouple in a thermowell. Notice the metal between the process and the junction. It lets the element be replaced without opening the process, and it delays the reading [35], which is why a moving peak or a hot spot can run ahead of it.

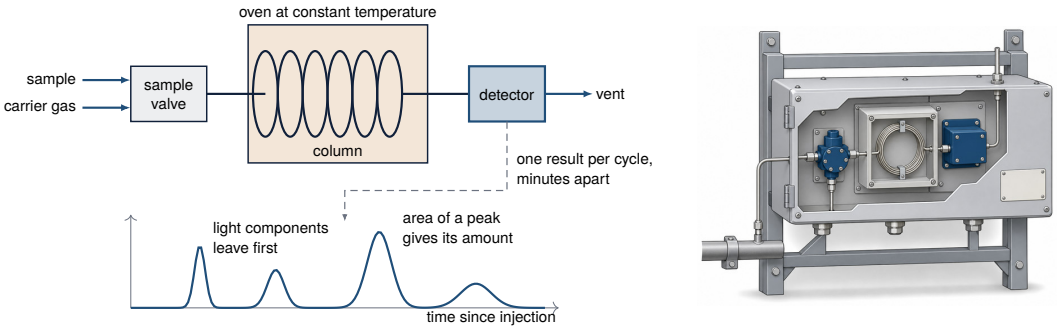


Figure 1.9 Principle of a process gas chromatograph, with a view of the device as installed. Notice that one sample gives one result at the end of the cycle, so a loop on this instrument acts on a value that is always old.

Layer	What it does in this volume	Where it runs	How often
Optimization	picks severity, grade sequence and the time to decoke or wash	**Supervisory computers** send setpoints to the controllers below	every hour for a unit, every day for the plant
Predictive control	moves several handles together against several limits		every minute; 2 s to 0.1 h in the apps
Quality and cycle to cycle control	trims setpoints from analyzers, the laboratory and inferentials	**Distributed control system** field controllers with PID, lead-lag and select blocks; trips in a separate safety instrumented system	at each analyzer, laboratory or cycle report
Overrides, selectors and trips	one-way overrides, low selectors, interlocks and trips		at each scan of the loops below
Regulatory control	holds flows, pressures, levels and temperatures		every second; 1 s to 72 s in the apps
Field devices	thermocouples, analyzers, load cells, valves, pumps and feeders	point-to-point wiring or a fieldbus	continuously

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The periods of the regulatory, predictive and optimization layers are those of [9, 37], the platforms those of [9, 36], and the ranges in the apps those of the chapters. Notice that each layer sends setpoints to the layer below and acts less often than it.

is inevitable [35]. The chromatographs of Chapter 3 report 3 min after a sample drawn every 5 min, and that of Chapter 4 reports every 20 min. Four chapters wait for a laboratory. The laboratory of Chapter 8 posts each result 45 min after its hourly sample, and the Mooney viscosity of Chapter 11 arrives every 2 h.

The final elements are mostly valves on coolant, fuel and feed, and the polymer reactors add the feeders of catalyst and initiator. Where a limit must never be crossed, a trip acts beside the loops. A safety instrumented system is built of sensors, logic solvers and final elements that bring the plant to a safe state, and it is often independent of the basic process control system [36]. Figure 1.10 places the layers of the volume on their platforms. Regulatory loops act about every second, predictive control every minute and optimization every hour to every day [9, 37]. The apps run their loops every 1 s in Chapter 10 and Chapter 12 and every 72 s in Chapter 2 and Chapter 11, whose predictive controllers plan every 0.1 h.

Table 1.3 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

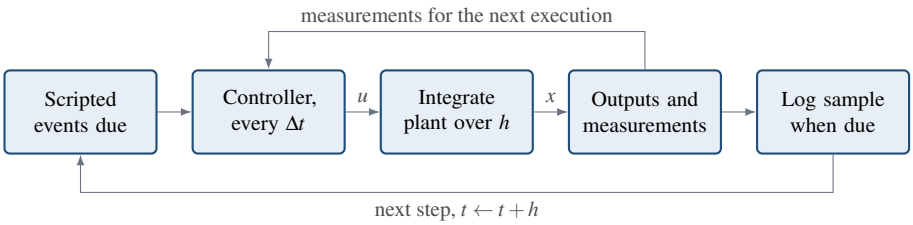


Figure 1.11 Work of the Engine in One Integration Step

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{dx}{dt} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.11. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings

are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.14) to (1.17) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.14}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.15}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.16}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.17}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.17) drives I toward the value that reproduces the output applied, which is back-calculation [38]. With no integral action, Equation (1.17) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.18).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.18}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.19) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.19}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [39]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [40]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [41].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according

to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed (the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 13.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Ethane cracking furnace	Balancing the measured variable does not balance the constraint, and the run length is the objective.	intermediate	2.5 h
3	Acetylene converter	A reagent window open on one side, and a light-off a thermocouple cannot see.	advanced	2 h
4	Propylene splitter	A tray temperature is a pressure gauge first, and the analyzer is fast enough.	intermediate	2.5 h
5	Propane dehydrogenation beds	The handle acts on the next cycle's stored heat.	intermediate	2 h
6	Ethylene oxide reactor	The minority reaction owns the heat, and an optimum has no sign.	advanced	2.5 h
7	Paraxylene moving bed	Purities exist only as cycle averages.	advanced	2 h
8	Gas-phase polyethylene	Two time scales two decades apart, and grade changes priced in tonnes.	intermediate	3 h
9	Polypropylene loop	The quality handle spends the pump's margin, so a selector arbitrates.	intermediate	3 h
10	Tubular LDPE reactor	Excursions, not setpoints, spend the margin to decomposition.	advanced	2.5 h
11	Butyl rubber reactor	A flat controlled variable and an actuator counting down.	introductory	2.5 h
12	Additive feeders	A held weighed flow is not a held concentration.	introductory	2.5 h

1.8 Chapter Roadmap

Table 1.4 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 11 and 12 form a first course: one loop whose actuator is the clock, and a ratio loop whose product is seen only by a laboratory. They rest on Sections 1.4.2 and 1.4.6. Chapters 2, 4, 5, 8 and 9 add a hidden constraint, an inferential, a loop that acts once a cycle, a grade change and a selector. Chapters 3, 6, 7 and 10 combine a selectivity or purity that is measured late with a limit that a fast excursion can reach, and are best left until last. A reader who follows the molecules can read the volume in order, from the cracker to the pellet.

KEY POINTS

- Each unit holds a quality it sees only every few minutes or hours, against a limit it must never cross.
- A firing loop that holds the outlet hides the coking, and only the run's total, read after the decoke, finds the severity optimum.
- A side reaction that releases more heat costs P-1's tube two thirds of its coolant margin.
- The off-grade tonnage of a grade change depends on the inventory and on how far past the target the property being made is driven.
- The guards are fast and act in one direction, because a quality loop learns of its mistake only at the next report.

Where else this applies. Inferred quality returns in Volume 5, overrides and selectors in Volume 4, cycle-to-cycle control in Volume 8 and predictive control in Volume 9.

References

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

[2] American Chemical Society, "Discovery of polypropylene and the development of a new high-density polyethylene," National Historic Chemical Landmark booklet, Office of Communications, Washington,

DC, 1999, designated at Phillips Petroleum Company, Bartlesville, Oklahoma, 12 November 1999; acs.org/education/whatischemistry/landmarks/polypropylene.html, read November 2026.

[3] 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.

[4] R. M. Thomas, I. E. Lightbown, W. J. Sparks, P. K. Frolich, and E. V. Murphree, "Butyl rubber. a new hydrocarbon product," *Rubber Chemistry and Technology*, vol. 14, no. 1, pp. 177–195, 1941.

[5] K. Ziegler, "Consequences and development of an invention," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/ziegler-lecture.pdf, read November 2026.

[6] D. B. Broughton and C. G. Gerhold, "Continuous sorption process employing fixed bed of sorbent and moving inlets and outlets," U.S. Patent 2,985,589, assigned to Universal Oil Products Co., 1961, filed May 22, 1957, issued May 23, 1961.

[7] Y. Liu, "Natural intelligence principle of supervisory-level control for condensed-mode polyethylene fluidized bed reactors: A proportional-integral-derivative perspective," *ACS Omega*, vol. 11, no. 39, pp. 60 064–60 072, 2026.

[8] J. M. Jenkins, III, R. L. Jones, and T. M. Jones, "Fluidized bed reaction systems," U.S. Patent 4,543,399, assigned to Union Carbide Corp., 1985, filed April 3, 1984, issued September 24, 1985.

[9] 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.

[10] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[11] 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.

[12] R. M. Thomas and W. J. Sparks, "Mixed olefinic polymerization process and product," U.S. Patent 2,356,128, filed 20 October 1939 (continuation in part of an application filed 29 December 1937), granted 22 August 1944, 1944.

[13] G. Natta, "From the stereospecific polymerization to the asymmetric autocatalytic synthesis of macromolecules," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/natta-lecture.pdf, read November 2026.

[14] H. Zimmermann and R. Walzl, "Ethylene," in *Ullmann's Encyclopedia of Industrial Chemistry*. Wiley-VCH, 2009.

[15] J. J. H. B. Sattler, J. Ruiz-Martinez, E. Santillan-Jimenez, and B. M. Weckhuysen, "Catalytic dehydrogenation of light alkanes on metals and metal oxides," *Chemical Reviews*, vol. 114, no. 20, pp. 10 613–10 653, 2014.

[16] 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.

[17] 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.

[18] C. van Heerden, "Autothermic processes," *Industrial & Engineering Chemistry*, vol. 45, no. 6, pp. 1242–1247, 1953.

[19] O. Bilous and N. R. Amundson, "Chemical reactor stability and sensitivity: II. Effect of parameters on sensitivity of empty tubular reactors," *AIChE Journal*, vol. 2, no. 1, pp. 117–126, 1956.

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

[21] 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.

[22] A. N. R. Bos and K. R. Westerterp, "Mechanism and kinetics of the selective hydrogenation of ethyne and ethene," *Chemical Engineering and Processing: Process Intensification*, vol. 32, no. 1, pp. 1–7, 1993.

[23] A. Kummer and T. Varga, "What do we know already about reactor runaway? – A review," *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

[24] T. Xie, K. B. McAuley, J. C. C. Hsu, and D. W. Bacon, "Gas phase ethylene polymerization: Production processes, polymer properties, and reactor modeling," *Industrial & Engineering Chemistry Research*, vol. 33, no. 3, pp. 449–479, 1994.

[25] C. Kiparissides, G. Verros, and J. F. MacGregor, "Mathematical modeling, optimization, and quality control of high-pressure ethylene polymerization reactors," *Journal of Macromolecular Science, Part C: Polymer Reviews*, vol. 33, no. 4, pp. 437–527, 1993.

[26] European IPPC Bureau, "Reference document on best available techniques in the production of polymers," European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[27] K. B. McAuley and J. F. MacGregor, "Optimal grade transitions in a gas phase polyethylene reactor," *AIChE Journal*, vol. 38, no. 10, pp. 1564–1576, 1992.

[28] M. Takeda and W. H. Ray, "Optimal-grade transition strategies for multistage polyolefin reactors," *AIChE Journal*, vol. 45, no. 8, pp. 1776–1793, 1999.

[29] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. Steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[30] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[31] Y. Ren, X. Liang, H. Dou, C. Ye, Z. Guo, J. Wang, Y. Pan, H. Wu, M. D. Guiver, and Z. Jiang, "Membrane-based olefin/paraffin separations," *Advanced Science*, vol. 7, no. 19, p. 2001398, 2020.

[32] M. Ohshima and M. Tanigaki, "Quality control of polymer production processes," *Journal of Process Control*, vol. 10, no. 2–3, pp. 135–148, 2000.

[33] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009.

[34] M. Mazzotti, G. Storti, and M. Morbidelli, "Optimal operation of simulated moving bed units for nonlinear chromatographic separations," *Journal of Chromatography A*, vol. 769, no. 1, pp. 3–24, 1997.

[35] 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.

[36] 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.

[37] S. Skogestad, "Plantwide control: The search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[38] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

[39] 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.

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

[41] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

1.9 Exercises

1.1 Theory and calculation · 5 min

A cracking coil starts a run with its tube metal 75 K below the limit, and in the first days the metal rises 0.03 K/h. Its coking rate grows by a factor e for every 30 K of metal temperature, and the outlet and the flux are held, as in Section 1.4.2.

(a) Find t^* and the run length. *Checkpoint:* 1000 h and 917.9 h.

(b) Find the run a straight line through the first days predicts, and the metal rise after 500 h. *Checkpoint:* 2500 h and 20.8 K.

1.2 Concept check · 2 min

A four-pass furnace has one pass more coked than the others. The passes are balanced by trimming their flows until the four coil outlet temperatures agree. What happens to the coked pass?

(a) its metal falls to match the others, because equal outlets mean equal heat flux

(b) its flow is cut and its metal runs hottest, because coke makes it read cold at the outlet and hot at the skin

(c) its flow is raised and its metal cools, because the trim sends more feed to the cold outlet

(d) nothing changes, because balancing moves only the total feed and not its split

1.3 Theory and calculation · 3 min

P-1's furnace can run at 15.6 t/h of ethylene for a run of 1140.3 h, or hotter at 16.0 t/h for a run of 700 h. Each decoke takes 30 h. Which gives more ethylene per elapsed hour, and by how much?

1.4 Theory and calculation · 5 min

P-1's oxidation tube, with the activation temperatures and heats of Section 1.4.3, runs at 265 °C on a less selective catalyst, $S = 0.80$.

(a) Find the combustion's share of the heat and E_{eff}/R. *Checkpoint:* 0.765 and 11 059 K.

(b) Find the room $T^2/(E_{eff}/R)$ at the design temperature and compare it with the 0.85 catalyst. *Checkpoint:* 26.2 K against 26.9 K.

1.5 Concept check · 2 min

An acetylene converter adds hydrogen in ratio to the acetylene it must remove. Too little hydrogen leaves the product off grade, and every mole beyond the need hydrogenates ethylene. Where should the ratio be held?

(a) in the middle of the range between off grade and the trip, to keep equal margins

(b) well above the need, because more hydrogen is always the safe side

(c) at the need exactly, because any error either way costs the same

(d) just on the clean side of the breakthrough edge, because the two failures cost very different amounts

1.6 Theory and calculation · 4 min

A cooled tube carries one exothermic reaction with $E/R = 10\,000$ K and runs at 500 K, 20 K above its coolant.

(a) Find the room $T^2/(E/R)$ and the share of it the tube uses. *Checkpoint:* 25 K and 80 %.

(b) A side reaction raises E_{eff}/R to 12 500 K. Find the room again. *Checkpoint:* 20 K, all of it used.

1.7 Theory and calculation · 4 min

P-1's gas-phase reactor carries 15 MW in its recycle gas as sensible heat, and its condensed liquid takes 0.35 MJ/kg to boil, with $\Delta H_p = 3.6$ MJ/kg and 150 kg/s of recycle. Find the condensed flow and its share of the recycle at 24 t/h.

1.8 Theory and calculation · 6 min

A reactor holds 80 t of polymer and changes grade from melt index 2 to 6 with a band of 10 %, so the product must rise above 5.4.

(a) Find the off-grade tonnage when the gas is set to make melt index 6 at once. *Checkpoint:* 187.6 t.

(b) Find it when the gas first makes melt index 12 until the product reaches 5.4. *Checkpoint:* 64.7 t.

(c) At 16 t/h, how long is the product off grade in (b)? *Checkpoint:* 4.04 h.

1.9 Theory and calculation · 4 min

A cycle-to-cycle loop moves a regeneration temperature once a report, and the reported conversion answers by 0.4 points per kelvin two reports later. With the integral law of Equation (1.12), find the gain K for the double root and the gain at the edge of stability.

1.10 Concept check · 2 min

A propylene splitter runs a tray temperature loop under an analyzer cascade. The chromatograph freezes and keeps republishing its last result, which happens to be slightly high, and the feed then changes. What does the cascade do without a validity check?

(a) it keeps integrating the stale error and walks the tray setpoint away, ending worse than the tray loop alone

(b) it holds the tray setpoint, because a constant reading produces no error

(c) it trips the column, because the analyzer reading stops changing

(d) it tracks the feed change, because the tray loop below it sees the change