INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 5 Delays, Slow Measurements, and State Estimation

CHAPTER 3

Teaching Composition Control With a Slow Analyzer on a Thirty-Two Stage Column

Pages 50–81

Edited by John D. Hedengren

How to cite this chapter. "Teaching Composition Control With a Slow Analyzer on a Thirty-Two Stage Column," in *Delays, Slow Measurements, and State Estimation*, John D. Hedengren, Ed., vol. 5 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 50–81.

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

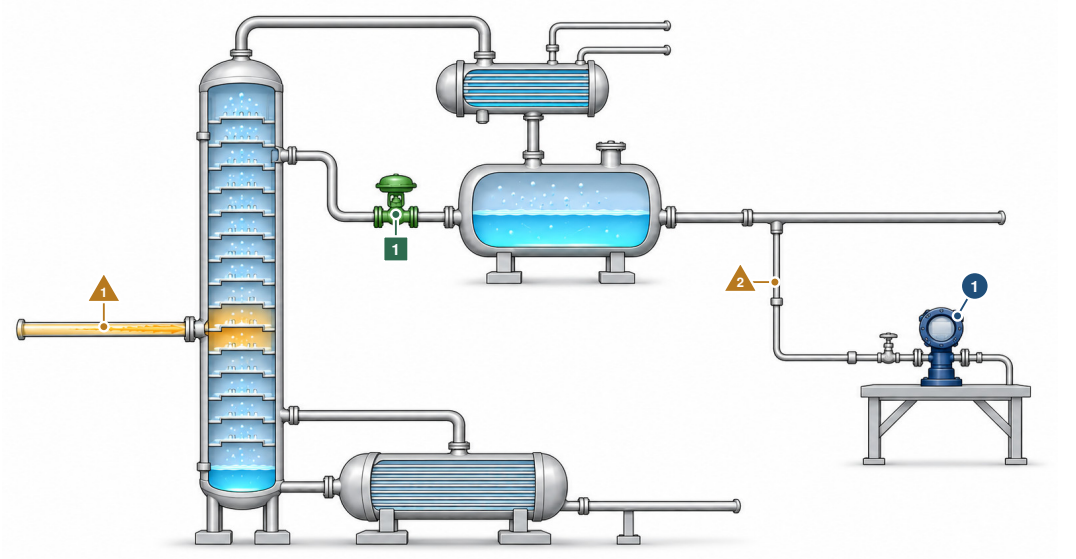
The app for this chapter
https://apmonitor.com/apps/distillation/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v05-delays-estimation/#distillation

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

CHAPTER 3

Teaching Composition Control With a Slow Analyzer on a Thirty-Two Stage Column



The controller moves 1 reflux valve that sets the reflux ratio

The controller reads 1 analyzer of the distillate's purity

What disturbs it 1 feed richer or leaner in the light component 2 a slower analyzer

The reflux reaches the top trays almost at once, while the analyzer reads the distillate only after the sample has traveled its line and been analyzed. Illustration, not to scale.

IN PLAIN TERMS **Steering a Column by a Late Report**

The main idea. A distillation column splits a liquid mixture by boiling it. Some of the liquid made at the top is poured back down the column, and the more is poured back, the purer the top product becomes. A controller sets that flow from a test of how pure the top product is.

The challenge. The test takes minutes, so the controller always acts on old news. A simple controller then overshoots, and a smarter one that forgets the delay can keep swinging back and forth.

Why it matters. Knowing what the late test costs tells engineers whether to slow the controller, predict ahead, or estimate the purity from faster readings.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on distillation (Section 3.2), the history segment on the relay (Section 3.5.6), the try-it on the reflux limit (Section 3.6.8), the try-it and the fact on the late analyzer (Section 3.7) and the key points
Study	80 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Sections 3.6 and 3.7), Technology in Practice (Section 3.9), Limitations (Section 3.10) and the two open challenges

The control question A thirty-two stage column holds its distillate purity with the reflux ratio, and the only measurement is an analyzer that reports minutes late. How much performance does that lag remove from a fixed tuning, and what can detuning or a predictive controller recover when the controller's model leaves the analyzer out?

Learning objectives

- Derive the ceiling that the overall balance places on the distillate purity, and tell an infeasible setpoint from a poor tuning.
- Fit a first-order plus dead-time model to a reflux step, explain why the column shows a dead time it does not contain, and derive the IMC tuning.
- Predict from the phase margin how an analyzer lag erodes damping, and measure its cost in overshoot, settling and margin to the specification.
- Compare the PID with an MPC whose model omits the analyzer lag, and find where that model error ends the MPC's advantage.
- Judge the PI constants a relay test writes when the relay band, not the column, sets the cycle.

Prerequisites Sections 1.4.2 (fitted dead time), 1.4.3 (IMC tuning), 1.4.5 (age of an analyzer result), 1.6.2 (velocity-form PID) and 1.4.7 (disturbance state). Section 1.6.3 is needed for the predictive controller.

The app https://apmonitor.com/apps/distillation/

See also Volume 6, Chapter 3 for a scanning gauge that adds a whole scan of dead time, and Volume 6, Chapter 2 for the pairing question a single-ended column cannot pose. Volume 13, Chapter 8 for interacting purity loops on one shell, and Volume 13, Chapter 9 for a tray temperature loop on a column with more than one steady state.

3.1 Introduction

A distillation column is among the most heavily instrumented vessels in a chemical plant, and its state is a composition profile rather than a single number. Every tray carries a liquid holdup in contact with a rising vapor, every holdup obeys its own component balance, and the plant presented to a composition controller is therefore of very high order. The controller never sees the profile. What arrives is one composition, measured on one product stream, by an instrument that reports minutes after the sample is drawn. Measurement dynamics of this kind are part of the loop and set a limit on achievable feedback performance that no tuning can recover [1].

A single measurement is nonetheless sufficient for control, because the column response is dominated by one slow mode. A column with thirty-two well-mixed stages has thirty-two eigenvalues, and for the simulator of this chapter they span time constants from 0.025 min to 15.24 min. A reflux move nonetheless produces a top-composition response that a first-order plus dead-time model reproduces almost exactly, because one mode carries essentially all of the steady-state gain and the rest decay before the response is half finished. Order reduction of this kind is a standard result for columns [2, 3], and it is the reason a two-parameter tuning rule works on a thirty-two state plant. The application draws all thirty-two liquid compositions as a live profile beside the trend of the distillate composition, so the order reduction is observed rather than asserted.

The reduction does not extend to the measurement path. A gas chromatograph on a product stream reports on a cycle rather than continuously, and the sample line and the analysis time add to the delay between the tray composition and the value the controller receives. Dead time and lag in the measurement path are not attenuated by feedback and cannot be removed by tuning, and the performance they remove is the subject of a literature of its own [4]. The industrial remedy is to infer the composition from tray temperatures, which respond immediately, and to correct the inference against the slow analyzer [5, 6]. The experiments of this chapter measure the loop degradation that justifies that expense.

This chapter quantifies what an adjustable analyzer lag costs a fixed tuning on a thirty-two stage binary column whose distillate composition is held with the reflux ratio. The topics are process order and model order reduction, apparent dead time, measurement dynamics, nonlinear steady-state gain,

integral windup, and the feasibility limit that an overall balance places on a setpoint, which are core topics of a first course in process dynamics and control [7, 8]. The physical content of the model is one component mass balance per stage, with the equilibrium driving force carried by a constant relative volatility and the flows fixed by constant molar overflow. Section 3.2 places the column in the literature on column dynamics and composition control. Section 3.3 sets out the balances, the instrument, the constants and the numerics in enough detail for a second implementation. Section 3.4 decomposes the column response into a symmetric inventory mode and an antisymmetric profile mode, and shows that fixing the distillate flow at half the feed makes every reflux move a pure internal-flow change, which suppresses the slow mode and sets a hard ceiling of $2x_F$ on the distillate composition. It then fits a first-order plus dead-time model to a step test and derives the PI tuning from it. Section 3.5 gives the discrete control law, a linear model predictive controller offered beside the PID, a moving horizon alternative to that controller's Kalman filter and a relay test that tunes the PID. Section 3.6 carries one tuning across five analyzer lags and two feed upsets, and Section 3.7 compares the predictive controller with the PID on five of the app's cases and across the same lags. Section 3.8 measures the estimator option and the relay test on the same column. Section 3.9 sets the loop on plant hardware, and Sections 3.10 to 3.12 state the limitations, the economic argument and the conclusions.

3.2 Background

Stage-by-stage calculation is the oldest quantitative method in separations. The graphical construction that steps between an equilibrium curve and an operating line was published in 1925 and remains the way a first course explains what a tray does [9]. Its central simplifying device, a constant relative volatility that converts the equilibrium calculation into one algebraic expression in the liquid mole fraction, is also what makes a dynamic stage model small enough to integrate. Every dynamic column model used as a control benchmark since then has been built the same way, one component balance and one equilibrium relation per stage, with constant molar overflow removing the energy balance from the problem [3, 10].

Column dynamics entered the control literature as a question about time scales. Measured columns settle far more slowly than any single tray, and the ratio is large enough that early studies looked for an explanation outside the balances. The explanation follows from the balances themselves. A change in the external flows must change the total holdup of light component in the column, and the products carry that inventory change away at a rate that produces a dominant time constant much longer than any residence time [2]. A change in the internal flows at constant external flows rearranges the profile without changing the inventory and is correspondingly faster. The same distinction determines which upsets are slow and which are fast, and it explains why one lag is usually enough to describe the response a controller acts on.

Composition control itself became a benchmark problem with the pilot column whose two-by-two transfer function matrix was identified from step tests in 1973 [11]. That model has been used ever since to pose the central structural question of column control, which is not how to tune a loop but which flows to pair with which composition. The reflux and boil-up pair, the distillate and boil-up pair and their variants behave very differently under disturbances and at high purity, and the choice interacts with the level loops underneath [12, 13]. Volume 13, Chapter 8 meets the same interaction among three purity loops on one shell. Holding both product compositions at once makes the interaction explicit, and high-purity columns are the standard illustration of an ill-conditioned plant, where the gain in one direction of input space is orders of magnitude larger than in another [14].

The measurement is the weakest element of the loop. An online analyzer is expensive to purchase, expensive to maintain and slow to respond, and the composition it reports is minutes old by the time a controller acts on it. Tray temperatures cost little, respond immediately and are only an indirect indication of composition, so the practical arrangement is to control a temperature and let the analyzer correct the temperature setpoint. Selecting which tray to control is a design problem with

Table 3.1 Variable Classification for the Distillate Composition Loop

Symbol	Role	Range	Unit
x_0	CV, distillate composition	0 to 1	mol frac
x_m	measured CV at the analyzer	0 to 1	mol frac
rr	MV, reflux ratio L/D	1.0 to 10.0	—
x_F	DV, feed composition	0.20 to 0.80	mol frac
F	DV, feed rate	0.50 to 1.50	mol/min
τ_a	DV, analyzer lag	0 to 2	min
$x_1 \ldots x_{30}$	tray liquid compositions	0 to 1	mol frac
x_{31}	bottoms composition, not controlled	0 to 1	mol frac
L, V, B, D	internal and product flows	—	mol/min

several competing criteria [6], and regression of composition on several temperatures at once is the standard way to sharpen the inference [5]. A temperature loop is not free of trouble either. Volume 13, Chapter 9 shows a stripping-tray temperature loop that holds one steady state and drives another the wrong way. None of that structure is needed when the analyzer is fast. All of it is justified by how much the loop deteriorates when the analyzer is slow, which is the quantity this chapter makes measurable.

The industrial motivation is energy. Distillation accounts for a large fraction of the separation energy used by the chemical and refining industries, the reboiler duty is the dominant operating cost of most columns, and reducing the purity giveaway of a poorly controlled column is one of the least expensive available savings [15]. A composition loop that is detuned because its analyzer is slow runs further from the specification limit, and the added distance is paid for in reboiler duty.

> ✦ **DID YOU KNOW? Forty Thousand Columns**
>
> A 2005 report for the United States Department of Energy counted about 40,000 distillation columns in operation in the United States. Distillation made up 90 to 95 % of all separations in chemicals and refining and is usually the method of choice when the relative volatility is 1.2 or more [16]. The column of this chapter has 1.6.

3.3 Process Model

Thirty trays, a total condenser with its reflux drum and a partial reboiler make up the thirty-two well-mixed stages of the simulated column, numbered from the top as shown in Figure 3.1. Stage 0 is the condenser and drum together, stages 1 to 15 are the rectifying trays, stage 16 is the feed tray, stages 17 to 30 are the stripping trays, and stage 31 is the reboiler. The single state of each stage is the mole fraction of the light component in its liquid holdup. Time is carried in minutes throughout the model, and the export writes its time column in the unit that suits the run length, minutes for most runs and hours for a run of an hour or more.

3.3.1 Variable Classification

The loop has one controlled variable, one manipulated variable and three disturbances, and each of them is an operator-set input while the simulation runs. Table 3.1 classifies them together with the states and the two outputs that are reported but not controlled. The bottoms composition is drawn, displayed and logged, and there is no bottoms controller, so the column is single ended.

3.3.2 Flows and Vapor-Liquid Equilibrium

Constant molar overflow is assumed, so the vapor rate is the same on every stage and the liquid rate is constant within each section. The distillate flow is fixed at half the feed, which makes the reflux

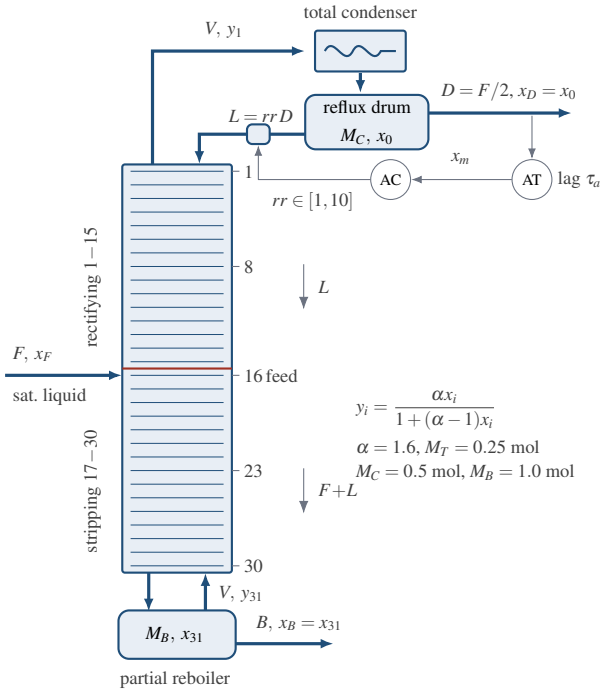


Figure 3.1 Thirty-two stage binary column with the composition loop closed on the distillate. The reflux ratio $rr = L/D$ is the only manipulated variable, the distillate flow is fixed at $D = F/2$, and the analyzer AT reports x_0 through a first-order lag τ_a before the controller AC sees it.

ratio the only degree of freedom, as shown in Equation (3.1).

$$\begin{aligned} D = \tfrac{1}{2}F, \qquad B = F - D, \qquad L = rr\,D \\ V = L + D, \qquad L_S = F + L \end{aligned} \tag{3.1}$$

where D is the distillate flow, B the bottoms flow, L the reflux and the liquid rate through the rectifying section, V the vapor rate through the whole column, and L_S the liquid rate through the stripping section. The feed is a saturated liquid, so it adds to the liquid traffic below the feed tray and leaves the vapor traffic unchanged.

The vapor leaving a stage is in equilibrium with the liquid on it, and the equilibrium is written with a constant relative volatility as shown in Equation (3.2).

$$y_i = \frac{\alpha x_i}{1 + (\alpha - 1)x_i}, \qquad i = 0, 1, \ldots, 31 \tag{3.2}$$

where y_i is the vapor mole fraction of the light component leaving stage i, x_i is the liquid mole fraction on that stage and α is the relative volatility. The default value $\alpha = 1.6$ is a difficult separation. It takes all thirty trays to reach a distillate of 0.9354 at a reflux ratio of 3, and α is adjustable between 1.1 and 3.0 so that the separation difficulty of the same column can be varied over a wide range.

3.3.3 Component Balances

Each stage is a well-mixed liquid holdup in equilibrium with the vapor leaving it, so one component balance per stage closes the model. The condenser receives the overhead vapor and returns liquid of its own composition to the column and to the product, as shown in Equation (3.3).

$$M_C \frac{\mathrm{d}x_0}{\mathrm{d}t} = V\,(y_1 - x_0) \tag{3.3}$$

where M_C is the combined condenser and reflux drum holdup. The condenser is total, so the reflux and the distillate both leave at x_0 and the L and D terms collapse into $V = L + D$.

A rectifying tray receives liquid from the tray above and vapor from the tray below, as shown in Equation (3.4).

$$M_T \frac{\mathrm{d}x_i}{\mathrm{d}t} = L\left(x_{i-1} - x_i\right) - V\left(y_i - y_{i+1}\right), \quad i = 1, \ldots, 15 \tag{3.4}$$

where M_T is the liquid holdup of one tray. The feed tray adds the feed and passes the larger stripping liquid rate downward, as shown in Equation (3.5).

$$M_T \frac{\mathrm{d}x_{16}}{\mathrm{d}t} = F x_F + L x_{15} - L_S x_{16} - V\left(y_{16} - y_{17}\right) \tag{3.5}$$

A stripping tray is the rectifying tray with the larger liquid rate, as shown in Equation (3.6).

$$M_T \frac{\mathrm{d}x_i}{\mathrm{d}t} = L_S\left(x_{i-1} - x_i\right) - V\left(y_i - y_{i+1}\right), \quad i = 17, \ldots, 30 \tag{3.6}$$

The reboiler receives the liquid leaving tray 30, returns vapor of composition y_{31} to the column and withdraws the bottoms product, as shown in Equation (3.7).

$$M_B \frac{\mathrm{d}x_{31}}{\mathrm{d}t} = L_S x_{30} - B x_{31} - V y_{31} \tag{3.7}$$

where M_B is the reboiler holdup. The reboiler is partial, so the vapor it returns is in equilibrium with the sump liquid and the bottoms product leaves at x_{31}. Every term in Equations (3.3) to (3.7) is a molar flow of light component in mol/min, and dividing by a holdup in mol leaves a rate of change of mole fraction in 1/min.

Adding Equations (3.3) to (3.7) cancels every internal term and leaves the overall balance of Equation (3.8).

$$\sum_{i=0}^{31} M_i \frac{\mathrm{d}x_i}{\mathrm{d}t} = F x_F - D x_0 - B x_{31} \tag{3.8}$$

This is the constraint that governs the whole chapter, and it is derived again in Section 3.4.

3.3.4 Instrument

The controller does not act on x_0. It acts on an analyzer reading that carries a first-order lag and additive Gaussian noise, as shown in Equation (3.9).

$$\tau_a \frac{\mathrm{d}x_a}{\mathrm{d}t} = x_0 - x_a, \qquad x_m = \mathrm{clip}_{[0,1]}(x_a + \sigma\varepsilon) \tag{3.9}$$

where x_a is the analyzer state, τ_a is the analyzer lag in min, x_m is the reading passed to the controller, σ is the noise standard deviation in mole fraction and ε is drawn from the unit normal distribution independently at every controller execution, through a Box-Muller transform of the seeded generator of Section 1.6.1, so a run with noise is as reproducible as one without. Setting $\tau_a = 0$ collapses the first expression and the analyzer reports x_0 directly. The lag is implemented as a disturbance rather than as a fixed property of the instrument, so it can be changed while the column is running and the column itself is untouched by the change.

3.3.5 Parameters, Limits and Initial Conditions

Table 3.2 lists every constant of the column, the analyzer and the PID, with the three settings of the predictive controller, whose model and filter are given in Section 3.5.4. The three holdups, the relative volatility and the analyzer noise are adjustable, which allows the tray time constant and the separation difficulty to be varied independently. The reflux ratio limits are hard clamps applied both to the manual input and to the controller output.

Table 3.2 Model Parameters, Limits and Defaults

Sym.	Meaning	Value	Unit
Adjustable			
α	relative volatility (1.1-3.0)	1.6	—
M_T	tray holdup (0.10-1.00)	0.25	mol
M_C	condenser holdup (0.20-2.00)	0.50	mol
M_B	reboiler holdup (0.40-3.00)	1.00	mol
σ	analyzer noise SD (0-0.010)	0.0015	m.f.
Fixed by the model			
N	modeled stages	32	—
n_F	feed stage index	16	—
$rr_{\min}$	lower reflux ratio limit	1.0	—
$rr_{\max}$	upper reflux ratio limit	10.0	—
$\sum M_i$	total liquid holdup	9.00	mol
Operator settings and defaults			
F	feed rate (0.50-1.50)	1.00	mol/min
x_F	feed composition (0.20-0.80)	0.50	m.f.
τ_a	analyzer lag (0-2, step 0.05)	0	min
x_0^{SP}	setpoint (0.700-1.000)	0.970	m.f.
—	band about the setpoint	$\pm$0.010	m.f.
rr	manual reflux ratio, case 1 (step 0.1)	3.0	—
Controller			
K_c	prop. gain (0-300, step 0.6)	60	1/m.f.
τ_I	integral time (0.002-100, 100 is off)	4	min
τ_D	derivative time (0-2)	0	min
u_{bias}	bias, reflux ratio at the start	rr_{man}	—
τ_t	anti-windup tracking time	τ_I	min
Δt	controller sample period	0.02	min
h	integration step	0.002	min
Δt_s	history sample interval	0.2	min
T_s	MPC decision period	0.2	min
T_p	MPC horizon (4-60)	30	min
λ	MPC move suppression (0-3)	0.3	—

The abbreviation m.f. stands for mole fraction.

Every case starts at steady operation. The application solves the thirty-two balances of Equations (3.3) to (3.7) for their steady profile by Newton's method, with the parameters and disturbances in force. A case in manual starts at the steady profile of its manual reflux ratio, and a case in automatic starts at the steady profile of its setpoint, with the manual reflux ratio set to the value that holds it. The reference operating point case starts in manual at $rr = 3$, $F = 1$ and $x_F = 0.5$, where the profile runs from $x_0 = 0.93542$ to $x_{31} = 0.06458$ and stays there exactly, and the other nine cases start at the 0.970 setpoint with $rr = 3.7319$. A setting changed before the run starts, whether the mode, the setpoint, the manual reflux ratio, a parameter or a disturbance, re-initializes the case exactly as a reset would, with the Reset semantics of Section 1.6.1. It therefore changes where the run starts rather than producing a step, and a step or an upset is made after the run has started.

3.3.6 Numerical Integration and Timing

Integration is explicit Euler on a step of $h = 0.002\,\mathrm{min}$, in place of the Runge–Kutta step of Section 1.6.1. The simulation clock advances the column one step of h at a time, every state is clipped to the physical interval $[0, 1]$ after each step, and the analyzer state of Equation (3.9) is advanced on the same clock. The integration step is not a tuning choice. Linearizing the thirty-two balances about

the $rr = 3$ profile gives a fastest eigenvalue of −39.50/min, so the explicit Euler stability limit is $2/39.50 = 0.0506$ min and the integration step sits twenty-five times inside it. The slowest eigenvalue is −0.0656/min, a time constant of 15.24 min, so the stiffness ratio across the state vector is 602.

The controller executes every $\Delta t = 0.02$ min, which is ten integration steps, and one sample is written to the history every $\Delta t_s = 0.2$ min, which is every tenth execution. A sample stamped t carries the analyzer reading the controller acts on at t and the reflux ratio held over the execution period that ends at t, so a reflux move made at an execution first appears in the next sample stamped after it. A change the operator makes between executions reaches the column at the next execution. The speed setting runs from 0.5 min/s to 10 min/s and, as Section 1.6.1 explains, changes the pace of a run and never its result.

3.3.7 Assumptions

The model is deliberately the smallest one that reproduces the behavior under study, and each simplification has a limited domain of validity.

Constant molar overflow. The vapor rate is taken as constant up the whole column and the liquid rate as constant within each section, which removes the energy balance from the problem. This requires equal molar latent heats and negligible heat loss and sensible heat effects. It breaks down for pairs with very different latent heats, and it hides the fact that a real reflux change alters the internal traffic only after the condenser and reboiler duties have responded.

Constant holdup. Every M_i is a constant, so there is no hydraulic dynamics, no downcomer lag, no weir relationship and no level loop on the drum or the sump. Real columns show a liquid hydraulic lag of a few seconds per tray that delays a reflux change on its way down, and thirty of those in series are not negligible against a 7 min composition response. The model therefore reaches the lower trays faster than a real column would.

Fixed distillate split. The distillate flow is $D = F/2$ at all times, as though a perfect ratio controller held it there. This is what makes the reflux ratio the only degree of freedom, and it is also what pins the sum of the two product compositions. A real column has an independent draw, and the choice of which flows to pair with which composition is the configuration question the model does not pose.

Saturated liquid feed. The feed adds its whole rate to the liquid traffic below the feed tray and nothing to the vapor traffic, which is the $q = 1$ case. A subcooled or partially vaporized feed would change the split between L_S and V and move the feed tray optimum.

Constant relative volatility and ideal stages. Equation (3.2) assumes α does not vary with composition, temperature or pressure, and each stage is assumed to reach equilibrium exactly. Real trays have Murphree efficiencies well below unity, and a column running under a pressure profile has a volatility that changes from top to bottom.

Binary mixture and no energy balance. There are two components and no temperatures anywhere in the model. Inferential temperature control, which is what most columns actually run on, cannot be demonstrated on this simulator for that reason.

Instrument. The analyzer is a first-order lag with white measurement noise. A real chromatograph is closer to a sampler with a dead time equal to its cycle (Section 1.4.5), and a pure dead time is harder for a loop to tolerate than a lag of the same magnitude. The results in Section 3.6 therefore understate the damage a real analyzer does.

3.4 Dynamic Analysis

3.4.1 Steady-State Coupling of the Product Compositions

Setting the derivatives in Equation (3.8) to zero and substituting $D = B = F/2$ gives the steady-state relation of Equation (3.10).

$$x_0 + x_{31} = 2x_F \tag{3.10}$$

The relation holds exactly and independently of the reflux ratio, the number of trays, the holdups and the relative volatility, because it is nothing more than the overall component balance with a fixed split.

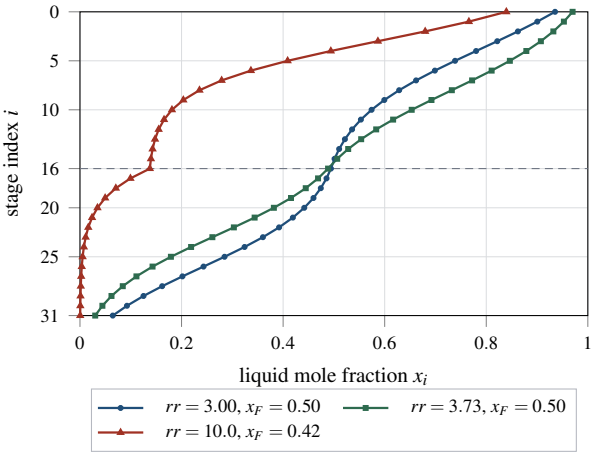


Figure 3.2 Converged liquid composition profiles from Equations (3.3) to (3.7), with circles at $rr = 3.00$, squares at $rr = 3.73$ and triangles at $rr = 10.0$ with $x_F = 0.42$. The dashed line marks the feed stage. Raising the reflux ratio from 3.00 to 3.73 lifts the top and drops the bottom by the same 0.0346, while a feed of $x_F = 0.42$ leaves the distillate pinned at 0.8395 even with the reflux ratio at its 10.0 limit.

The reflux ratio therefore does not change the amount of light component available. It determines only how that amount is divided between the two products.

Two consequences follow. The first is a ceiling. Since $x_{31} \geq 0$, the distillate composition can never exceed $2x_F$ at steady state, so the default setpoint of 0.970 stops being reachable as soon as the feed composition falls below 0.485. The second is a check on every steady state the application produces. At $rr = 3$ the profile settles at $x_0 = 0.9354$ and $x_{31} = 0.0646$, and at $rr = 3.5$ it settles at 0.9622 and 0.0378, both pairs summing to 1.0000 against a feed of 0.5000.

3.4.2 Modal Decomposition of the Linearized Column

The steady-state profiles at $rr = 3$ and $rr = 3.5$ differ by $\Delta x_0 = +0.0267$ and $\Delta x_{31} = -0.0267$, and the holdup-weighted sum of the changes across all thirty-two stages is -5.1×10^{-4} mol, which is 0.012 % of the 4.244 mol of light component the column holds at $rr = 3$. Raising the reflux ratio therefore rearranges the profile without changing the inventory, as shown in Figure 3.2. That is a direct consequence of fixing $D = F/2$. The external flows never move, so the light component that leaves the top faster must leave the bottom slower by the same amount.

Linearizing the balances about the $rr = 3$ profile makes the same point in the frequency domain. All thirty-two eigenvalues are real, so the open-loop plant cannot oscillate, and the two slowest are 15.24 min and 8.17 min. Their eigenvectors are different in kind. The 15.24 min eigenvector has the same sign on every stage, so it is the inventory mode in which the whole column becomes richer or leaner together. The 8.17 min eigenvector changes sign once, near the feed stage, so it is the profile mode in which the top gains what the bottom loses. Decomposing the steady-state gain of x_0 with respect to rr across the thirty-two modes assigns 109 % of it to the 8.17 min profile mode and −2.8 % to the 15.24 min inventory mode, with the remaining thirty modes contributing the small negative balance between them. The slowest dynamics in the column are essentially invisible from the reflux ratio, and a feed composition upset, which does change the inventory, is what brings them out.

3.4.3 Origin of the Apparent First-Order Response

The liquid residence time of a single stage is its holdup divided by the liquid flow through it. At $rr = 3$ and $F = 1$ those are 0.167 min on a rectifying tray, 0.100 min on a stripping tray, 0.250 min in the

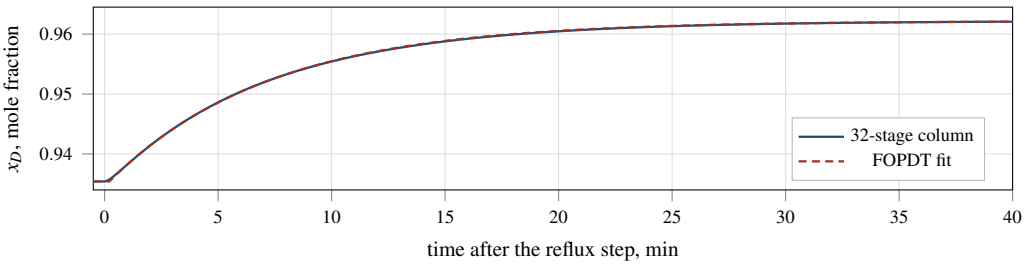


Figure 3.3 Open-loop response of the distillate composition to a reflux ratio step from 3.0 to 3.5 with the analyzer noise set to zero. The solid curve is the column, and the dashed curve is the fitted first-order plus dead-time model, with $K_p = 0.0535$, $\tau_p = 7.07\,\text{min}$, $\theta_p = 0.21\,\text{min}$.

condenser and 0.400 min in the reboiler. No stage in the column holds liquid for longer than 0.40 min, and the distillate responds to a reflux step with an apparent time constant of 7.07 min. The factor of eighteen is not a delay accumulating down a chain. It is the profile mode, whose time constant is set by the holdup that must be redistributed against the separating action of each tray, and it is the reason a column responds far more slowly than any individual stage.

3.4.4 Step-Test Identification

The step test follows the standard identification procedure of a first course [7]. In the reference operating point case, with the analyzer noise set to zero and the loop in manual, the reflux ratio is held at 3.0 on the already converged profile for 5 min and then stepped to 3.5, which gives the response of Figure 3.3. The distillate moves from 0.9354 to 0.9622 in the 60 min after the step for a reflux ratio change of 0.5, so the fitted process gain is $K_p = 0.0535$ mole fraction per unit of reflux ratio. The response reaches 28.3 % of its final value at 2.56 min and 63.2 % at 7.28 min, so the two-point fit gives $\tau_p = 7.07\,\text{min}$ and $\theta_p = 0.21\,\text{min}$. The fitted curve stays within 3.3×10^{-4} mole fraction of the column, 1.2 % of the step.

The apparent dead time of 0.21 min is not a transport delay. Nothing in Equations (3.3) to (3.7) delays anything, and the composition on stage 0 starts moving on the first integration step after the reflux changes. What the two-point fit reads as dead time is the thirty fast eigenvalues rounding the foot of the response, and it is of the same order as the sum of the fastest stage residence times. This is the fitted dead time that is not transport of Section 1.4.2, here on a plant that contains no delay at all. Chapter 2 fits a dead time to a plant that does carry a transport delay and finds the fit far longer than the delay, for the same reason.

Pause. Try Exercise 3.2 before reading on.

3.4.5 Steady-State Gain Nonlinearity

Holding a reflux ratio to convergence at the base feed maps the manipulated variable onto the steady distillate composition, and the map is strongly curved. The model source tabulates it to four decimal places in steps of 0.5 from $rr = 1$ to $rr = 10$, and the table runs 0.7306, 0.8431, 0.9354, 0.9765, 0.9883 and 0.9971 at reflux ratios of 1.0, 2.0, 3.0, 4.0, 5.0 and 10.0. The exact steady states show half a unit of reflux raising the distillate by 0.0571 mole fraction at $rr = 1$ and by 0.00022 at $rr = 9.5$, a factor of 256, which the rounded table would put at 190. The secant gain of 0.0535 measured over the 3.0 to 3.5 step is already 22 % below the local gain of 0.0687 at $rr = 3$, so even a half-unit step test averages over a visible amount of curvature. The same half-unit step test made at $rr = 6$, with the manual reflux ratio set to 6.0 before the run so that the case starts at that steady state, returns a gain of 0.0024, more than an order of magnitude below the 0.0535 at $rr = 3$, on a plant whose equations contain no explicit nonlinearity other than Equation (3.2). A tuning computed at 0.970 is far too slow

at 0.90 and far too aggressive near 0.99, which is the standard argument for controlling a purity loop on the logarithm of the impurity rather than on the composition itself.

3.4.6 Controller Tuning From the Fitted Model

Internal model control for a first-order plus dead-time plant (Section 1.4.3) sets the proportional gain and reset time as shown in Equation (3.11).

$$K_c = \frac{\tau_p}{K_p\,(\tau_c + \theta_p)}, \qquad \tau_I = \min(\tau_p,\, 4\,(\tau_c + \theta_p)) \tag{3.11}$$

with τ_c the specified closed-loop time constant [17]. Specifying $\tau_c = 2\,\text{min}$ with the fitted model gives $K_c = 59.8$ and $\tau_I = 7.07\,\text{min}$. The default tuning is $K_c = 60$ and $\tau_I = 4\,\text{min}$, which is the same proportional gain with the integral time shortened to 4 min, and the exercises of Section 3.13 compare the two.

3.5 Controller Design

3.5.1 Velocity-Form Algorithm

The loop is the velocity-form PID of Section 1.6.2, with the derivative on the measurement. The complete algorithm as coded is given by Equation (3.12), in the loop's own symbols.

$$e_k = x_0^{\text{SP}} - x_{m,k} \tag{3.12a}$$

$$\Delta s_k = K_c\,(e_k - e_{k-1}) - K_c\,\tau_D\,\frac{x_{m,k} - 2x_{m,k-1} + x_{m,k-2}}{\Delta t} \tag{3.12b}$$

$$s_k = s_{k-1} + \Delta s_k \tag{3.12c}$$

$$c_k = u_{\text{bias}} + s_k + I_{k-1} \tag{3.12d}$$

$$rr_k = \text{clip}_{[1,10]}\,(c_k) \tag{3.12e}$$

$$I_k = I_{k-1} + \frac{K_c}{\tau_I}\,e_k\,\Delta t + \frac{\Delta t}{\tau_t}\,(rr_k - c_k) \tag{3.12f}$$

where e_k is the error in mole fraction, s_k is the accumulated proportional and derivative contribution, c_k is the demanded output before the limits are applied, rr_k is the controller output in reflux ratio units at execution k, I_k is the integral state, $\Delta t = 0.02\,\text{min}$ is the sample period, u_{bias} is the manual reflux ratio in force when the case starts or is reset and $\tau_t = \tau_I$ is the tracking time of the back-calculation. The integral used at execution k was accumulated through execution $k-1$, and the derivative acts on the unfiltered reading. Setting τ_I to the top of its range, 100 min, switches integral action off. The integral gain is then zero, the back-calculation term is dropped with it, and the loop is proportional only. In the symbols of Equations (1.15) to (1.18), $s_k = S_k$, $c_k = v_k$, $rr_k = u_k$, $u_{\text{bias}} = u_b$, $\tau_t = T_t$ and $a = +1$. The code forms each increment as the difference of the positional terms for the gains in force, so a gain changed during a run takes effect at the next execution, as in Section 1.6.2.

The controller is reverse acting in the instrument-engineering sense. Raising the reflux ratio raises the distillate composition, so the process gain is positive, and a rise in the measured composition lowers the error and with it the reflux ratio.

3.5.2 Initialization and Bumpless Transfer

A velocity-form recursion has no bias of its own, and its starting output is set by the bias and the integral state. When a case starts or is reset, the bias is set to the manual reflux ratio in force and the integral to zero, so the first execution in automatic returns the positional expression of Equation (3.13).

$$rr_0 = \text{clip}_{[1,10]}\,(rr_{\text{man}} + K_c\,e_0) \tag{3.13}$$

where rr_{man} is the manual reflux ratio. A case that starts in automatic starts at the steady state of its setpoint with rr_{man} at the reflux ratio that holds it, so e_0 is zero apart from the analyzer noise

and the first execution leaves the reflux ratio where it is. Transfers between manual and automatic are bumpless in both directions, as in Section 1.6.2. A switch to automatic adds only $K_c(e_k - e_{k-1})$, which is zero for an unchanged setpoint and a steady reading and is K_c times the setpoint change when the setpoint is raised at the same instant.

A setpoint raised in automatic therefore moves the reflux ratio at the next execution by K_c times the change, the proportional kick that opens the setpoint runs of Section 3.6, and integral action then carries the output to the reflux ratio that holds the new setpoint, which is $rr = 3.7319$ for a distillate of 0.970 at the base feed.

3.5.3 Anti-Windup by Back-Calculation

Anti-windup is the back-calculation term of Equation (3.12). While the output is clipped, $(\Delta t/\tau_t)(rr_k - c_k)$ pulls the integral toward the value that reproduces the applied reflux ratio with a tracking time of $\tau_I = 4\,\text{min}$, so each execution removes 1/200 of the clipped excess and the accumulator s_k is never truncated. With $\tau_D = 0$ and a steady error, the integral settles where the demand exceeds the limit by exactly $K_c e_k$, because $(K_c/\tau_I)\,\tau_t = K_c$. The demand therefore winds no further past the limit than the proportional term, and once the integral has settled the output leaves the limit as the error changes sign. Run XF of Section 3.6 holds the reflux ratio on its upper limit for more than 1200 min, and the demand there settles at 17.83, the limit of 10.0 plus $K_c e = 7.83$. Away from the limits the recursion is algebraically identical to the positional law with the same gains.

The upper limit has a second effect that has nothing to do with windup. At $rr = 10$ the steady distillate is 0.9971 and the local gain is close to zero, so a controller operating near saturation is close to open loop even before the clip engages.

3.5.4 Model Predictive Control

A linear model predictive controller (MPC) is available as a second automatic controller for the distillate. The PID stays the default, and only case 9 selects the MPC. Every $T_s = 0.2\,\text{min}$, or ten controller executions, the MPC optimizes the reflux ratio over a prediction horizon, sends only the first move to the column, keeps that move in force for the period and then optimizes again, the arrangement of Section 1.6.3 [18, 19]. Its model was identified on the simulator at the operating point of the automatic cases, $\bar{x}_0 = 0.970$ with $\overline{rr} = 3.7319$ and the base feed, by averaging the responses to reflux ratio steps of $+0.2$ and -0.2 with the analyzer noise and lag at zero and fitting two first-order lags in series, as shown in Equation (3.14).

$$\begin{aligned} \tau_1 \frac{\mathrm{d}z_1}{\mathrm{d}t} &= -z_1 + K_m\,(rr - \overline{rr} + \delta) \\ \tau_2 \frac{\mathrm{d}z_2}{\mathrm{d}t} &= z_1 - z_2, \qquad \tilde{x}_0 = \bar{x}_0 + z_2 \end{aligned} \tag{3.14}$$

where z_1 and z_2 are model states in mole fraction, $\tilde{x}_0$ is the predicted distillate composition, δ is an unmeasured disturbance in units of reflux ratio, $K_m = 0.02928$ per unit of reflux ratio, $\tau_1 = 7.115\,\text{min}$ and $\tau_2 = 0.566\,\text{min}$. The fit needs no dead time, and its root-mean-square residual of 1.4×10^{-5} is 0.23 % of the response.

The fit restates Section 3.4 at a higher purity. K_m lies 1.4 % above the local gain of 0.02889 that the thirty-two balances give at 0.970, because a ± 0.2 step averages over the curvature of Section 3.4.5, and 45 % below the 0.0535 measured at 0.9354 in Section 3.4.4. The second lag carries the fast modes that the fit of Section 3.4.4 read as 0.21 min of dead time, and a first-order plus dead-time fit of the same response reads 0.48 min instead, with a residual 1.6 times larger. The model contains neither the analyzer lag nor the feed composition, so an analyzer lag is a model error and a feed upset reaches the controller only through the analyzer.

The estimate $\hat{\delta}$ removes offset. A steady-state Kalman filter on Equation (3.14) treats δ as a random walk, the disturbance state of Section 1.4.7 that makes a predictive controller offset-free [20], and it

updates the estimate at every execution whichever controller is acting, as shown in Equation (3.15).

$$\begin{aligned} \bar{z}_k &= \Phi\,\hat{z}_{k-1} + \gamma\left(rr_{k-1} - \overline{rr} + \hat{\delta}_{k-1}\right) \\ \eta_k &= x_{m,k} - \bar{x}_0 - \bar{z}_{2,k} \\ \hat{z}_k &= \bar{z}_k + \ell_z\,\eta_k, \qquad \hat{\delta}_k = \hat{\delta}_{k-1} + \ell_\delta\,\eta_k \end{aligned} \tag{3.15}$$

which is the prediction and correction of Equation (1.11) with steady gains, where Φ and γ are Equation (3.14) sampled with a zero-order hold over $\Delta t = 0.02\,\text{min}$, rr_{k-1} is the reflux ratio applied over the period just ended, $x_{m,k}$ is the analyzer reading that the PID also receives, η_k is the innovation, and $\ell_z = [0.0547 \;\; 0.0360]^\mathsf{T}$ and $\ell_\delta = 10.60$ per unit of mole fraction are the converged Riccati gains for a measurement variance $\sigma^2 = 2.25 \times 10^{-6}$, a state noise variance of $10^{-6}\sigma^2$ and a disturbance noise variance $q_\delta = 0.1\,\sigma^2/K_m^2 = 2.62 \times 10^{-4}$. A richer feed enters $\hat{\delta}$ as the added reflux ratio that would raise the distillate as much.

At each decision a quadratic program on the model of Equation (3.14), sampled at T_s over a horizon $T_p = 30\,\text{min}$ of $N_p = 150$ steps, chooses twelve reflux ratios v_j, held over blocks of $(1,1,1,1,2,2,3,4,6,9,15,105)$ steps, to minimize Equation (3.16) within the hard limits $1.0 \le v_j \le 10.0$.

$$J = \frac{T_s}{\tau_n}\sum_{k=1}^{N_p}\left(\frac{\tilde{x}_{0,k} - x_0^{\text{SP}}}{y_s}\right)^2 + \lambda\sum_{j=1}^{12}\left(\frac{v_j - v_{j-1}}{u_s}\right)^2 \tag{3.16}$$

where $\tilde{x}_{0,k}$ is predicted k steps ahead with $\hat{\delta}$ held, v_0 is the reflux ratio applied now, $\tau_n = 7.7\,\text{min}$, near $\tau_1 + \tau_2$, keeps the first sum independent of the horizon, $y_s = 0.01$ is the half-width of the specification band, $u_s = 9$ is the reflux span and $\lambda = 0.3$ is the move suppression, for weights of 259.7 per squared mole fraction of predicted error and 3.7×10^{-3} per squared move. The symbol λ keeps the app's name for the move suppression here and is not the bias weight of Chapter 1. The program has no specification band, rate limit, terminal weight or feed preview. Its decisions are departures from $\overline{rr}$, a ridge of 10^{-9} relative to the weights keeps it strictly convex, and a dual active-set method [21] solves it. One tuning rule fixed λ and the factor 0.1 in q_δ for each of the seven single-loop studio apps before any comparison was run. Among the settings that move the reflux no more than the PID does in cases 3 and 4, it keeps the largest λ whose true IAE is within 10 % of the lowest.

A plan is rejected on any of three outcomes (a solver status other than optimal, a first move that is not a finite number, or a first move outside the reflux limits). The PID then takes the reflux ratio, and because it follows the applied value at every execution, in the same way that it follows the operator's setting in manual (Section 3.5.2), the transfer is bumpless, and the PID keeps the reflux until a later plan succeeds. When a rejection is imposed 6 min into case 9 with the noise removed, the handover changes the reflux ratio by 0.0006, less than the 0.0008 of each following PID move, and the plan that resumes control changes it by 0.067. The first plan is solved as a case loads, so the controller acts from $t = 0$. It receives only the analyzer reading, its own reflux ratio, the setpoint and the reflux limits, and its constants are fixed, so the thirty-two balances, the fit behind the PID tuning and Equation (3.14) remain three separate models, and a parameter the operator changes reaches only the first.

3.5.5 Moving Horizon Estimation

Each plan may instead start from a moving horizon estimator (MHE), chosen in place of the filter of Equation (3.15) and selected by no case. It solves the window problem of Equation (1.13), which can keep the fitted states inside hard bounds [22]. Writing $\zeta = [z_1 \;\; z_2 \;\; \delta]^\mathsf{T}$ for the augmented state of the

filter, execution k fits the readings $j = s, \ldots, k$, with $s = k - N_e$, by solving Equation (3.17).

$$\begin{aligned} \min_{\zeta_s,\, \omega} \; & \| \zeta_s - \bar{\zeta}_s \|^2_{\bar{P}^{-1}} + \sum_{j=s}^{k-1} \omega_j^{\mathsf{T}} Q^{-1} \omega_j \\ & + \frac{1}{\sigma^2} \sum_{j=s}^{k} \left(x_{m,j} - \bar{x}_0 - z_{2,j} \right)^2 \\ \text{s.t. } & \zeta_{j+1} = \Phi_e \, \zeta_j + \Gamma_e \, (rr_j - \overline{rr}) + \omega_j \\ & 0 \le \bar{x}_0 + z_{i,j} \le 1, \qquad i = 1, 2 \end{aligned} \tag{3.17}$$

In Equation (3.17) $\Phi_e = \left[\begin{smallmatrix} \Phi & \gamma \\ 0 & 1 \end{smallmatrix}\right]$ and $\Gamma_e = [\gamma^{\mathsf{T}} \;\; 0]^{\mathsf{T}}$ repeat the prediction of Equation (3.15), rr_j is the reflux ratio applied after reading j, ω_j perturbs each of the three states, and $Q = \mathrm{diag}(10^{-6}\sigma^2, 10^{-6}\sigma^2, q_\delta)$ and σ^2 are the variances the filter was designed with. The first term ties the window's first state to $\bar{\zeta}_s$, the MHE's own estimate from execution $s-1$ advanced one period, through the steady prior covariance $\bar{P}$ of the filter, held constant. With $N_e = 10$ the window spans 0.2 min, one decision period, and eleven readings, so the dual active-set method of Section 3.5.4 receives 33 decisions and 44 bound rows.

Only declared values set the bounds. A mole fraction lies between 0 and 1, and the same interval is imposed on z_1, the composition that the second lag follows, so in deviation form both states are confined to -0.97 to 0.03. Nothing declared limits δ, which stays free. With no bound active the window problem carries the filter's model, weights and prior and returns the filter's estimate. Should a window stop short of an optimum, return a non-finite value or break a bound, the estimate falls back to the MHE's one-period prediction corrected as the filter would correct it, clipped to the bounds, and a flag in the exported history shows each such case. Whenever the MHE is chosen, at the start or in mid-run, its window opens from the filter's own prediction. The filter keeps running under either choice, and the MHE reads the analyzer, the applied reflux ratio and declared constants only.

3.5.6 Relay Autotuning

A button starts the relay feedback experiment of [23] while the PID holds the distillate in automatic, and none of the ten cases runs it. The relay takes the place of the PID and holds the reflux ratio at $u_0 + d_r$ until the reading rises more than h_r above the setpoint, then at $u_0 - d_r$ until it falls more than h_r below, and so on. The center u_0 is the mean reflux ratio applied over the last 100 executions, 2 min or a tenth of the trend window, so that no single noisy output fixes it. The amplitude $d_r = 0.45$ is five percent of the reflux span, or 0.9 of the distance from u_0 to the nearer limit if that is smaller, and the band $h_r = 3\sigma_0 = 0.0045$ is three times the declared analyzer noise σ_0 whatever noise the run carries. A cycle with a half shorter than four executions is not counted. Once the last two counted cycles match to 5 %, both in period and in the swing $\pm a_r$ of the reading, the first harmonic of the two applied levels, with a correction for the band, yields the ultimate gain K_{cu}, and the Tyreus–Luyben rule converts it and the period P_u into the PI constants of Equation (3.18).

$$K_{cu} = \frac{4 d_r \sin(\pi f_h)}{\pi \sqrt{a_r^2 - h_r^2}}, \qquad K_c = \frac{K_{cu}}{3.2}, \qquad \tau_I = 2.2\, P_u \tag{3.18}$$

Here f_h is the fraction of a cycle spent at the upper level, one half for a symmetric cycle, and K_{cu} has the unit of K_c, the reciprocal of a mole fraction. The written constants carry three significant figures, with $\tau_D = 0$, each inside its input range and τ_I below the 100 min that switches integral action off, and they stay editable. An unconverged test ends after 30 cycles or 40 min, and a second press or a switch to manual or to the MPC cancels it. The PID follows the applied reflux ratio throughout, as in manual (Section 3.5.2), so it resumes without a bump. Because the band delays every switch, the cycle settles where the loop's phase falls $\psi = \arcsin(h_r / a_r)$ short of -180°, and the result note flags any cycle with $\psi \ge 10^\circ$ as one that the band, not the column, has fixed.

Table 3.3 Runs Built on the Reference Operating Point Case

Run	Mode	Change applied	Horizon
STEP	manual	rr 3.0 to 3.5 at 5 min	65 min
L0	auto	SP 0.970 at 2 min, $\tau_a = 0$	62 min
L0.5	auto	L0 with $\tau_a = 0.5$ min	62 min
L1	auto	L0 with $\tau_a = 1$ min	62 min
L1.5	auto	L0 with $\tau_a = 1.5$ min	62 min
L2	auto	L0 with $\tau_a = 2$ min	62 min
D2	auto	L2 with $K_c = 36$, $\tau_I = 5$ min	62 min
N0	auto	L0 with $\sigma = 0.0015$	62 min
XF	auto	x_F 0.50 to 0.42 at 2 min	1296 min
FR	auto	F 1.00 to 1.50 at 2 min	62 min
SF	manual	STEP repeated at $F = 1.50$	65 min

⧖ HISTORY **1984: A Relay Finds the Critical Point**

Most regulators in industry are of the PID type, a large plant may have hundreds of them, and a paper of 1984 observed that many are poorly tuned in practice, in part because simple and robust tuning methods had not been available. The classic Ziegler–Nichols procedure raises the gain of a proportional controller until the loop oscillates and reads the critical gain and period, an experiment that is hard to automate with the swing kept under control. The 1984 answer put a relay in the loop instead. A process with a phase lag of at least 180° at high frequencies can then oscillate at its critical period, and the first harmonic of the relay output gives the critical gain from the relay amplitude and the swing of the output. A relay with hysteresis is less sensitive to measurement noise [23]. The relay of this chapter has such a band.

3.6 Results

3.6.1 Protocol

Every run below starts from the reference operating point case, which loads the column at its steady profile in manual at a reflux ratio of 3.0. The noise, the analyzer lag, the tuning and the mode listed in Table 3.3 are chosen before the run starts, so the case starts at steady operation with those settings, and every step or upset is made after the run has started, as an operator would at the controls. Runs L0 to N0 switch to automatic before the run and so hold the case's own setpoint of 0.9354 with the reflux ratio at 3.0 until the setpoint is raised to 0.970 at 2 min. Runs XF and FR set the 0.970 setpoint before the run, start at the steady state that holds it and step the feed at 2 min. A change made at 2 min reaches the column at the controller execution at 2.00 min, and times below are measured from that execution. The runs were driven through the application's own simulation engine, the one the browser runs, with every change applied through the same calls the controls make. The application's performance panel, the scorecard of Section 1.6.4, reports the integral absolute error, the overshoot, the settling time, the reflux travel and the time at a limit, computed on the analyzer reading from the 0.2 min history against a band of ± 0.01 mole fraction. The measures below are defined on the true distillate composition instead and computed at the 0.002 min integration step over the 60 min after the step. Overshoot is quoted as a percentage of the setpoint step. Settling is the last time the true distillate composition is further than 0.002 mole fraction from setpoint. Travel is $\sum_k |rr_k - rr_{k-1}|$ over every controller execution from the step on. The exported history gives integral absolute errors 2 to 4 % lower and settling times up to 0.2 min earlier, because it is sampled ten times less often. Except where stated, the analyzer noise was set to zero. The settings chosen before the run and the same settings carried through a reset give the same run, as Section 1.6.1 states for every app.

Table 3.4 One Tuning ($K_c = 60$, $\tau_I = 4$ min) Across Five Analyzer Lags

Run	τ_a min	peak m.f.	over %	settle min	IAE m.f. min	travel —
L0	0	0.9726	7.6	13.5	0.0945	3.54
L0.5	0.5	0.9738	10.9	13.7	0.0928	3.87
L1	1	0.9758	16.8	13.9	0.1050	4.24
L1.5	1.5	0.9783	24.0	14.4	0.1312	4.73
L2	2	0.9806	30.7	23.6	0.1687	5.32
D2	2	0.9752	15.0	20.1	0.1487	2.43

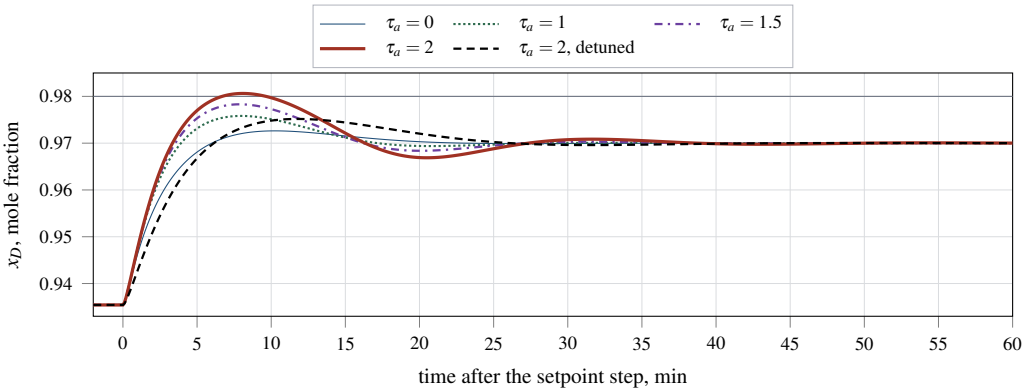


Figure 3.4 Capture of the 0.970 setpoint by the default tuning at four analyzer lags in minutes, after the setpoint is raised from 0.9354 at 2 min: thin solid at no lag, dotted at 1 min, dash-dotted at 1.5 min and thick solid at 2 min. The 0.980 specification limit is the thin gray line. The detuned run, dashed, lowers K_c to 36 and raises τ_I to 5 min at 2 min of lag, which brings the peak back inside the band.

3.6.2 Setpoint Capture With the Default Tuning

Run L0 holds the starting profile in automatic with the default tuning and raises the setpoint from 0.9354 to 0.970 at 2 min, so the loop has to move the distillate by 0.0346 against no disturbance at all. The execution at the step moves the reflux ratio from 3.0 to 5.07, it peaks at 5.10 at the execution 0.10 min later, the composition peaks at 0.9726 10.25 min after the step, and the loop settles inside 0.002 mole fraction of setpoint 13.53 min after the step. The reflux ratio then approaches the 3.7319 that holds 0.970 at steady state and reads 3.7314 60 min after the step. The overshoot is 7.6 % of the step and the composition spends 2.76 min outside the 0.960 to 0.980 specification band, all of it below the band on the way up.

3.6.3 Effect of the Analyzer Lag

Runs L0.5 to L2 repeat L0 with the analyzer lag raised in steps of 0.5 min to the 2 min maximum the application allows, and nothing else changed. The column is identical in all five runs. The only difference is the age of the measurement the controller acts on, and Table 3.4 quantifies the cost. Overshoot rises from 7.6 % to 30.7 %, the settling time rises from 13.5 min to 23.6 min, the integral absolute error rises by a factor of 1.79 and the reflux ratio travels 1.50 times as far. Half a minute of lag lowers the integral absolute error slightly (0.0928 against 0.0945), because the stale reading lets the proportional term push the reflux ratio higher early in the rise. The peak composition crosses the upper specification limit between 1.5 min and 2 min of lag, and at 2 min the loop that was designed to hold 0.970 peaks at 0.9806 and spends 2.95 min above the band, as shown in Figure 3.4. Figure 3.5 draws the same table as bars.

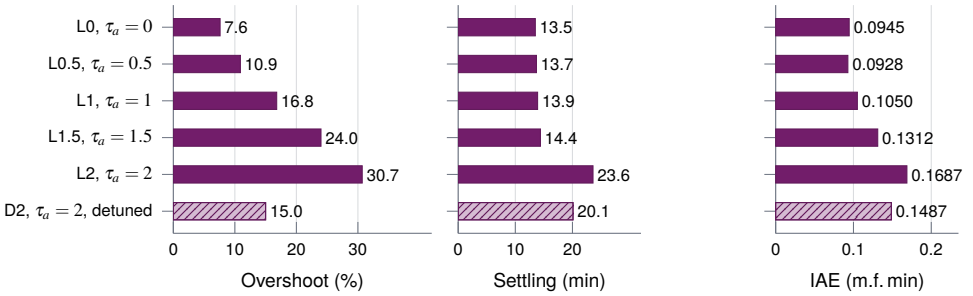


Figure 3.5 One tuning across five analyzer lags, from Table 3.4: the overshoot of the setpoint step (left), the settling time within 0.002 of setpoint (middle) and the integral absolute error (right), with the lag τ_a in minutes. Dark bars are the default tuning, the hatched bar the detuned run D2. Notice that the overshoot grows with every half minute of lag while the settling time barely moves until 2 min, and that detuning about halves the overshoot at 2 min without recovering the settling time of the fast loop.

The identified model explains the whole table. Adding a first-order lag of τ_a to the plant of Figure 3.3 and closing the loop with $K_c = 60$ and $\tau_I = 4$ min gives a phase margin of 73.2° at $\tau_a = 0$, falling to 59.7°, 48.7°, 40.5° and 34.3° at lags of 0.5, 1, 1.5 and 2 min. The gain margin falls from 16.5 to 8.32 over the same range, so the loop remains nominally stable throughout and the lag removes damping rather than stability. That is exactly the shape of Figure 3.4, where the 2 min case swings 0.0106 above setpoint 8.1 min after the step, 0.0031 below at 20.5 min and 0.0009 above at 31.6 min, decaying without diverging.

3.6.4 Detuned Controller at Two Minutes of Lag

Run D2 keeps the 2 min analyzer lag and detunes the loop to $K_c = 36$ and $\tau_I = 5$ min, chosen so that the predicted phase margin at 2 min of lag is 49.9°, close to the 48.7° the default tuning had at 1 min of lag. Every quantity in Table 3.4 improves. The peak falls from 0.9806 to 0.9752 and back inside the specification band, the overshoot falls from 30.7 % to 15.0 %, the settling time falls from 23.6 min to 20.1 min, the integral absolute error falls from 0.1687 to 0.1487 and the reflux travel falls from 5.32 to 2.43. Run D2 is outside the band for 3.42 min, all of it below on the slower rise, against 5.01 min for run L2, of which 2.95 min was above. Detuning further does not help, because $K_c = 24$ with the same integral time settles 0.22 min later than run L2 and raises the integral absolute error to 0.1743. The detuned loop is still 6.5 min slower to settle than the default tuning was with an instantaneous analyzer, and that gap is the part of the performance the lag removes permanently rather than trades against damping.

3.6.5 Effect of the Analyzer Noise on the Manipulated Variable

Run N0 repeats L0 with the default analyzer noise of $\sigma = 0.0015$ mole fraction kept. The trajectory is barely affected. The overshoot moves from 7.6 % to 7.8 % and the settling time from 13.53 min to 13.95 min, while five other noise seeds give 7.3 to 7.8 % and 13.5 to 13.8 min. The reflux ratio is affected far more strongly. Its travel over the 60 min after the step rises from 3.54 to 304.9, a factor of 86, and after the loop has settled the reflux ratio still varies with a standard deviation of 0.091 over a range of 0.62. The plant view at the end of a closed-loop run at 2 min of lag with this noise level is shown in Figure 3.6.

The velocity form makes the arithmetic transparent. The proportional term contributes $K_c(e_k - e_{k-1})$ to every increment, and for white measurement noise the difference of two independent samples has a standard deviation of $\sigma\sqrt{2}$, so the increment has a standard deviation of $K_c\sigma\sqrt{2} = 0.127$ reflux ratio units. The mean absolute value of a zero-mean normal variable is $\sqrt{2/\pi}$ times its standard

deviation, which gives 0.102 per execution and 305 over the 2999 executions that follow the step, against the 303.0 measured over the same executions. The exported history shows a travel of only 31.5 over the 60 min after the step and the performance panel 32.2 over the whole run, because between two samples 0.2 min apart the proportional increments telescope into a single difference of two readings. Raising K_c trades valve movement at this rate for speed, and the derivative time is left at zero on this loop because the second difference of the measurement would amplify the same noise by another factor of $1/\Delta t$.

3.6.6 Feed Composition Disturbance and the Purity Ceiling

Run XF applies the feed composition upset. The case starts in automatic at the steady state that holds 0.970 under the default tuning with no analyzer lag, the reflux ratio at 3.7319, and at 2 min the feed composition is stepped from 0.50 to 0.42. Equation (3.10) sets the new ceiling at $2x_F = 0.840$, well below the lower specification limit, so the setpoint is infeasible at steady state rather than merely difficult. The controller responds as feedback requires. The distillate leaves the band 15.8 min after the step, the reflux ratio reaches its 10.0 limit 28.7 min after the step and never leaves it, and the distillate reads 0.9377 at 30 min after the step, 0.8851 at 40 min and 0.8435 at 60 min. Left to run for 1296 min the column settles at $x_D = 0.8395$ and $x_B = 0.00045$, a sum of 0.8400 against the $2x_F = 0.840$ that Equation (3.10) requires.

Run XF demonstrates that no tuning changes the outcome. The reflux ratio is not saturated because the controller is badly tuned. It is saturated because the setpoint demands more light component in the distillate than the feed contains, and the only remedies are to raise the vapor rate independently of the reflux, to cut the distillate draw below $F/2$, or to lower the target.

3.6.7 Feed Rate Disturbance

Run FR steps the feed rate from 1.00 to 1.50 mol/min, the top of its range, at 2 min with everything else unchanged, and it has no closed-loop effect at all. Over the 60 min after the step the distillate stays at 0.9700 and the reflux ratio at 3.7319, and the run is identical to the same run without the step to the last digit the simulation carries. The reason is in Equation (3.1). Multiplying F by a factor multiplies D, L, V and L_S by the same factor at a fixed reflux ratio, so every right-hand side in Equations (3.3) to (3.7) is multiplied by that factor while the holdups stay where they are. The steady-state profile is therefore independent of the feed rate altogether, and what a feed rate change does is rescale time. Repeating the step test of Figure 3.3 in the application at $F = 1.5\,\text{mol/min}$ returns the same gain of 0.0535 measured at $F = 1$, with the time constant reduced from 7.07 min to 4.71 min and the apparent dead time from 0.208 min to 0.139 min, both within 0.2 % of the factor of 1.5 the scaling requires. The feed rate therefore changes the plant dynamics without displacing the steady state, a behavior easily mistaken for an ordinary well-rejected load change. The tray holdup separates the same two quantities from the other side. Raising it from 0.25 to 1.00 mol divides every tray derivative by four and leaves every steady state exactly where it was, because the holdups enter Equations (3.3) to (3.7) only as divisors. The same step test then returns the same gain with a much longer time constant, 27.1 min against 7.07 min.

🕮 **Pause.** Try Exercise 3.1 before reading on.

3.6.8 The Shipped Cases Under the Default Tuning

The app carries ten cases in a fixed order, each with its own completion rule (Section 1.6.4). Every case starts at steady operation, so in cases 2 to 10 the scripted event at 2 min is the first thing that moves the column, apart from the loop's response to analyzer noise in automatic. The upset of cases 4, 5, 9 and 10 is a richer feed, x_F from 0.50 to 0.58, which raises the $2x_F$ ceiling of Equation (3.10) instead of lowering it as run XF does. In automatic the reflux ratio falls toward the 2.90 that holds 0.970 on the richer feed. In manual with the reflux ratio left at 3.73 the distillate climbs to 0.9888 by 30 min. A cut to 2.90 made at once in case 5 acts on the top before the richer feed does, so the distillate dips to 0.9588 at 9.5 min. Case 5 credits a reflux ratio that holds the distillate in the band

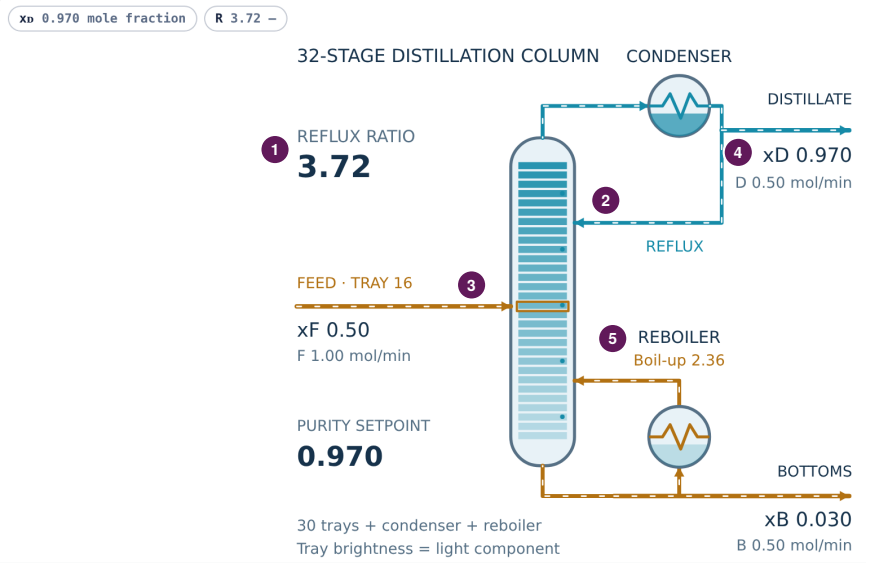


Figure 3.6 The Distillation Column Control Studio holding the 0.970 setpoint with a 2 min analyzer lag and analyzer noise of 0.0015 mole fraction, at 60 min of the reference operating point case, switched to automatic before the run and with the setpoint raised from 0.9354 to 0.970 at 2 min, cropped to the plant view (the header, control panel, trends and performance panel are cut). The default tuning of $K_c = 60$ and $\tau_I = 4$ min has brought the distillate back to 0.970 with the reflux ratio at 3.72. The callouts mark (1) the reflux ratio the controller sets, (2) the reflux returned to the top tray, (3) the feed on tray 16, (4) the distillate whose composition the analyzer reports and (5) the reboiler, whose boil-up follows the reflux because $V = L + D$.

once settled, kept there for 10 min, longer than the 7.07 min time constant, and the cut to 2.90 is credited at 23.2 min.

The ± 0.01 band equals the 0.010 setpoint step of case 3, so settling there is judged within ± 0.002 of the new setpoint, held for 1 min, and the case is credited at 8.6 min. The setpoint of 1.000 in case 8 lies above the 0.9971 the column reaches at the reflux limit, so the reflux ratio works up past 6, 8 and 10 at 4.5, 31.2 and 69.2 min while each unit buys less purity, and the distillate stops at 0.9971. The challenge starts at 0.970 with $K_c = 12$ and $\tau_I = 4$ min and asks for the analyzer reading to stay within ± 0.010 of 0.970 from 19.8 to 30 min through the richer feed at 2 min. With the shipped $K_c = 12$ the worst logged error over that window is 0.0143, so the attempt fails by 0.0043, and the true distillate is still 0.0112 above 0.970 at 20.9 min. The suggested tuning changes only the proportional gain, to $K_c = 60$, and gives 0.0052, 52 % of the tolerance, of which the true deviation is 0.0021 and the rest analyzer noise. Revealing the suggested tuning before the run and carrying it through a Reset give the same run.

> **TRY IT IN THE APP** **8 · Actuator Limit**
>
> Load the case and press Run at 10 min/s. Watch the reflux ratio on the lower strip work up past 6, 8 and 10 at 4.5, 31.2 and 69.2 min, while each unit buys less purity. Once the reflux sits on its limit, the purity stops at 0.9971, short of the 1.000 setpoint, and no tuning closes the gap.

3.7 Predictive Control Through a Lagging Analyzer

Table 3.5 runs cases 3, 4, 8 and 9 and the challenge once under each controller. The controller is chosen before the run starts, so the two runs of a case see identical events and identical noise draws.

Table 3.5 Cases 3, 4, 8, 9 and 10 Run Once Under Each Controller With the Same Analyzer Noise

Case	Controller	IAE (m.f. min)	IAE_x	e_{max} (m.f.)	$t_{0.002}$ (min)	Travel —	Limit (%)
3, setpoint	PID	0.048	0.038	0.0006	5.8	11.7	0
	MPC	0.031	0.016	0.0007	2.4	11.6	0
4, feed	PID	0.061	0.050	0.0045	—	11.4	0
	MPC	0.026	0.008	0.0008	0	6.5	0
8, limit	PID	0.519	0.534	—	—	37.6	36.8
	MPC	0.372	0.386	—	—	12.9	94.2
9, $\tau_a = 1$ min	PID	0.064	0.055	0.0052	17.4	11.4	0
	MPC	0.028	0.019	0.0022	5.8	7.2	0
10, challenge	PID, $K_c = 12$	0.225	0.216	0.0112	—	3.1	0
	PID, $K_c = 60^a$	0.076	0.057	0.0045	18.2	15.3	0
	MPC	0.040	0.010	0.0008	0	9.8	0

IAE uses the analyzer reading and IAE_x the true distillate. e_{max} is the largest true overshoot in case 3 and the largest true departure from setpoint after the feed step otherwise, $t_{0.002}$ the time from the step until the true distillate stays within ±0.002 (0 where it never leaves, a dash where it has not settled), Travel the summed reflux change on the log and Limit the time on a reflux limit. [a]Suggested, $\tau_I = 4$ min.

The PID keeps the shipped $K_c = 60$, $\tau_I = 4$ min and $\tau_D = 0$, the challenge loading $K_c = 12$ and suggesting 60, and the predictive controller keeps the settings of Section 3.5.4 throughout. Unlike Section 3.6, each case starts at 0.970 with the default noise and is read over its whole length on the 0.2 min log, as the performance panel reads it. The noise alone adds about $20\sigma\sqrt{2/\pi} = 0.024$ to a 20 min IAE on the reading, so the table also gives the IAE of the true distillate.

In case 3 the plan made at the setpoint step to 0.980 asks for reflux ratios of 5.74, 6.61 and 6.64 in its first three blocks, against a proportional kick to about 4.3 under the PID, and the true distillate stays within ±0.002 from 2.4 min after the step against 5.8 min, with overshoots of 0.0007 and 0.0006 and a true IAE 59 % lower. The equal logged travel, 11.6 against 11.7, hides that without noise the predictive controller moves the reflux ratio by 6.9 against 0.84, while with noise the PID moves it by 103 over its executions (Section 3.6.5).

In case 4 the estimate $\hat{\delta}$ passes half of its 20 min value, 0.63 or 0.64, at 5.0 min after the feed step under either controller, so the difference lies in the plan. The predictive controller keeps the true distillate within 0.0008 of 0.970, while the PID peaks at 0.9745 at 12.8 min and is still 0.0020 above at 20 min, with a true IAE 6.2 times higher. With the noise removed the true IAEs are 0.0053 and 0.0494 for reflux travels of 0.80 and 0.86, so the gain comes from acting earlier, not from moving more.

Case 8 asks for 1.000, above the 0.9971 that the column holds at the reflux limit (Section 3.4.5). The model places 1.000 at a reflux ratio of 4.76, where the column gives 0.9865, because the secant gain from 3.73 to 10 is 0.0043 and the local gain at 10 is 0.00041, 6.8 and 72 times below K_m. The first plan moves the reflux ratio to 9.17 and the second, 0.2 min later, to 10.0. Each of these plans predicts an overshoot, to 1.0013 in the first, and from 2.8 to 7.6 min the plans hold the reflux below the limit, down to 7.22, while $\hat{\delta}$ books the missing purity as a loss of reflux ratio that reaches −5.3 by 10 min, after which the reflux stays on the limit. The PID passes 6, 8 and 10 at 4.5, 31.2 and 69.2 min. Both end at 0.9971, the predictive controller with an IAE 28 % lower, 0.120 against 0.239 before 30 min and 0.250 against 0.278 after it, because the error in its gain, not any knowledge of the ceiling, drives it to the limit early.

Case 9 is case 4 seen through an analyzer lag of 1 min that Equation (3.14) omits, as shown in Figure 3.7. The predictive controller still scores less than half the PID's IAE, 0.028 against 0.064, with a true peak of 0.9722 at 7.2 min and the ±0.002 band held from 5.8 min after the step, against 0.9752 at 12.0 min and 17.4 min. Against case 4, however, the lag raises the true IAE of the PID by 9 %, from 0.050 to 0.055, and that of the predictive controller by 140 %, from 0.008 to 0.019, whose distillate then cycles about the setpoint with a period near 7 min, down to 0.9681. With the noise removed the cycle shrinks by a factor of only 0.66 to 0.83 per period, and with the noise present the true error between 90 and 120 min still reaches 0.0019 against 0.0002 for the PID.

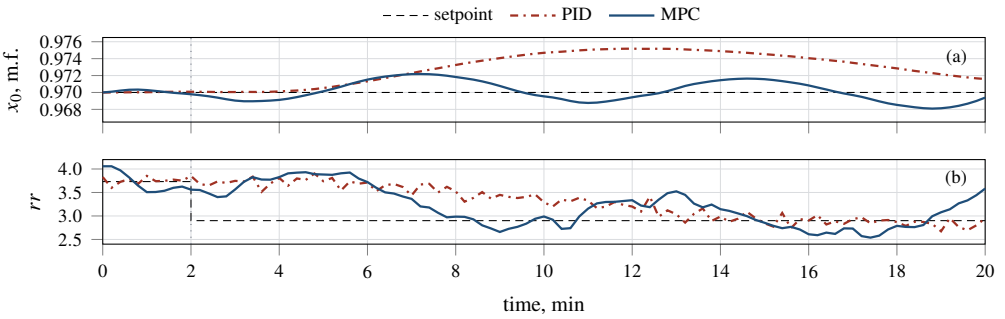


Figure 3.7 Distillate and reflux ratio of case 9 under each controller with the same analyzer noise, sampled on the 0.2 min log. The feed composition steps from 0.50 to 0.58 at 2 min (dotted) and the analyzer lags by 1 min. The PID is dash-dotted and the MPC solid. (a) True distillate composition and setpoint (dashed). (b) Reflux ratio and the steady value that holds 0.970 (dashed), 3.73 on the base feed and 2.90 on the richer one.

> **TRY IT IN THE APP** **9 · MPC With a Late Analyzer**
>
> Press Run and watch the distillate purity after the feed turns richer at 2 min. The MPC answers within minutes, then the purity keeps swinging about the setpoint with a period near seven minutes, because its model leaves out the one-minute analyzer lag. Switch the Controller to PID, press Reset and run again. The PID peaks later and higher, then settles.

Repeating case 9 at the lags of Table 3.4 measures how much of this model error the predictive loop tolerates. Over 20 min with the noise present its true IAE is 0.008, 0.010, 0.019, 0.034 and 0.042 at lags of 0, 0.5, 1, 1.5 and 2 min, against 0.050, 0.052, 0.055, 0.058 and 0.061 for the PID, lower at every lag but growing 5.2 times where that of the PID grows by 21 %. Over 180 min without noise the cycle shrinks per period by a factor of 0.82 at 1 min, 0.88 at 1.05 and 0.94 at 1.10 min, changes by less than 1 % per period at 1.15 min and grows at longer lags, from 1.5 min until the reflux ratio reaches its lower limit of 1.0. At 2 min the distillate cycles between 0.9525 and 0.9851 with a period near 10 min, while the swings of the PID decay by a factor of 0.10 to 0.14 per period, as its 34.3° phase margin implies. The predictive controller answers the upset faster and so has less phase to give up to a lag it does not model, and its largest move suppression, $\lambda = 3$, moves the stability limit only to between 1.25 and 1.5 min. Written into the model as a third first-order lag, outside the application, the analyzer lag restores the margin. With 2 min in its model the controller settles at 2 min of lag with a true IAE of 0.017 over 20 min, and with 1 min it lowers that of case 9 from 0.019 to 0.013, although 2 min in the model against a 1 min analyzer leaves a sustained swing of 0.0059.

Pause. Try Exercise 3.3 before reading on.

> **DID YOU KNOW?** **The Better Controller Loses More to the Lag**
>
> From the simulator: a one-minute analyzer lag added to the richer feed of case 4, which makes case 9, raises the true integral absolute error of the PID from 0.050 to 0.055 and that of the MPC from 0.008 to 0.019 (Table 3.5). The MPC still scores lower, but it gives up far more of its lead.

The predictive controller does not read K_c, so over the first 20 min the challenge reproduces case 4 exactly under the predictive controller and under the suggested tuning. On the judged worst reading error from 19.8 to 30 min it scores 0.0056 against 0.0052 for the suggested $K_c = 60$, both inside the 0.010 tolerance that $K_c = 12$ misses at 0.0143, although its largest true deviation there is 0.0008 against 0.0021, because the judged error is mostly analyzer noise. All 1050 plans were accepted,

and the reflux limit was binding in 566 of the 600 plans of case 8 and in no other, with at most 12 active-set iterations. Averaged over a run, a plan took 0.08 to 0.26 ms under Node.js in three repetitions and 229 and 244 µs in a pair of Chrome runs of case 9, the longest 2.1 ms, which on one desktop computer is 0.12 % of one processor core at one simulated minute per second.

3.8 Estimator Option and Relay Test

3.8.1 The Moving Horizon Estimate

Every case was run twice with the controller set to the MPC, once with each estimator chosen before the run, with the shipped noise seed, for 60 min or, for the challenge, its judged 30 min, and cases 1, 2 and 5 stayed in manual as given. Across the ten cases the two reflux records differ by at most 2.6×10^{-8} and the two readings by at most 7.7×10^{-11} mole fraction, no window was rejected and no plan failed. No bound was ever active, because the filter's purity estimate never left the interval from 0 to 1. Case 8 brought it closest, at 0.9977 under the setpoint of 1.000. Over the 2999 executions of case 9 the two purity estimates differ by at most 1.0×10^{-8} of the noise standard deviation and the two values of $\hat{\delta}$ by 4.8×10^{-9} units of reflux ratio. That residue does not come from rounding. The Riccati iteration that supplies the filter's gain stops after 526 steps with its covariance 2.6×10^{-8} (relative) short of its fixed point. The window performs those Riccati steps afresh at every execution, which separates the two estimates slightly, and with the converged covariance the difference drops to 7.3×10^{-14} of the noise.

Judged against the true distillate, which only the measuring script reads, the two estimators are equally accurate, and no case shows a bound that the filter violates and the MHE holds. Over the 20 min of case 9 the purity estimate has a root-mean-square error of 8.6×10^{-4} and a largest error of 1.7×10^{-3}, while against the analyzer state x_a of Equation (3.9) its error is 3.9×10^{-4}, so the estimate tracks the lagged composition that Equation (3.14) mistakes for the distillate. In case 4, without the lag, the two errors coincide at 3.0×10^{-4}, with a largest error of 1.6×10^{-3}. On a loaded computer under Node.js an estimate of case 9 took a median of 76 µs and a 99th percentile of 244 µs, against 85 and 515 µs for a plan, and each 0.2 min decision period holds ten estimates and one plan. In three Chrome runs on the application's frame loop the estimates and plans of case 9 occupied 1.2 to 1.5 % of a processor core at the 1 min/s speed setting and 5.3 to 8.5 % at 10 min/s, inside the tenth of a core allowed, though within a factor of two of it at the top speed.

3.8.2 The Relay Test and the Constants It Writes

The relay test was pressed 2 min into case 3, its setpoint step removed, with the purity steady at 0.970 in automatic and the default noise. Its center was 3.745, the mean of the last 100 executions, against 3.850 applied at that instant and the 3.732 that holds 0.970, so d_r took its full 0.45 and h_r was 0.0045. The test converged after three cycles and 12.6 min with $a_r = 0.00503$, $f_h = 0.492$, $P_u = 5.36$ min and $K_{cu} = 256$, and it wrote $K_c = 79.9$ and $\tau_I = 11.8$ min, 1.33 times the shipped gain and 2.95 times the shipped integral time. The swing exceeded the band by a factor of only 1.12, so $\psi = 63.5^\circ$ and the note marked the cycle as set by the band. Of the seeds 1001 to 1010, seven converged, with K_{cu} from 116 to 288, P_u from 3.26 to 6.46 min and ψ from 42.5° to 66.2°, and three reached the 40 min limit after 11 to 14 cycles without two agreeing.

The cycle reveals little about the column's ultimate point. The two lags of Equation (3.14) approach -180° without reaching it, and only the zero-order hold over Δt creates an ultimate point for the sampled model, with a period of 0.456 min and a gain of 26,400. The measured period is twelve times longer and the measured gain a hundredth of it. The same relay returns 514 at 11.9 min on the noise-free model, and 767 at 12.3 min on the simulator with the analyzer noise set to zero before the run, where $f_h = 0.518$ and $\psi = 80.6^\circ$. On this loop the swing barely clears the band, and Equation (3.18) then amplifies any change in it, since $\mathrm{d}K_{cu}/K_{cu} = -\left[a_r^2/(a_r^2 - h_r^2)\right] \mathrm{d}a_r/a_r$. The factor is 5.0 at the measured $a_r/h_r = 1.12$ and 37 at the noise-free 1.014. Noise of $\sigma/a_r = 0.30$ widens the swing from

Table 3.6 Shipped, Relay-Written and Two Counterfactual PI Pairs on Cases 3 and 4 Over 60 min With the Shipped Noise Seed

Case	PI pair	IAE (m.f. min)	IAE_x	$e_{\max}$ (m.f.)	$t_{0.002}$ (min)	Travel —
3, setpoint	shipped	0.103	0.042	0.0006	5.8	30.0
	relay	0.109	0.066	0.0001	7.8	39.9
	PI[a]	0.120	0.087	0.0001	11.5	30.0
	PI[b]	0.098	0.036	0.0006	4.6	39.9
4, feed	shipped	0.117	0.059	0.0045	18.2	29.7
	relay	0.154	0.120	0.0051	29.8	39.7
	PI[a]	0.187	0.160	0.0063	35.6	29.8
	PI[b]	0.106	0.044	0.0037	15.6	39.6

Shipped: $K_c = 60$, $\tau_I = 4$ min. Relay: $K_c = 79.9$, $\tau_I = 11.8$ min. [a] $K_c = 60$, $\tau_I = 11.8$ min. [b] $K_c = 79.9$, $\tau_I = 4$ min. Neither counterfactual is shipped or written. Columns as in Table 3.5, over 60 min instead of the case's 20 min.

0.00456 to 0.00503 and so cuts K_{cu} from 767 to 256. The uncorrected $4d_r/(\pi a_r)$ reads 114 and 125, so the band term, not the column, supplies most of either result.

Table 3.6 runs cases 3 and 4 for 60 min under four PI pairs, the shipped and the written ones and two counterfactual pairs that each exchange one constant. The written pair loses both cases. On the setpoint step its scorecard IAE is 0.109 against 0.103 and its true IAE 0.066 against 0.042. On the richer feed they are 0.154 against 0.117 and 0.120 against 0.059, and the true purity returns within ± 0.002 of 0.970 at 29.8 min after the feed step instead of 18.2 min. In both cases the reflux travels 33 % further, the ratio of the two proportional gains, because the analyzer noise sets the travel (Section 3.6.5). The counterfactual pairs locate the loss in the integral time. The shipped gain with the written τ_I does worse still, a true IAE of 0.160 on the richer feed, while the written gain with the shipped τ_I wins both cases, at 0.036 and 0.044. Taken as $2.2P_u$ from a cycle twelve times slower than the ultimate one, the written τ_I lowers the integral gain K_c/τ_I from 15.0 to 6.77 per mole fraction per minute. Across the 14 of 21 tests that converged (the shipped seed and the seeds 1001 to 1020) the written pairs scored 0.101 to 0.148 on the setpoint step, below the shipped 0.103 in 3 of them, and 0.124 to 0.272 on the richer feed, above the shipped 0.117 in all 14. With the noise removed from both case runs the written pair still loses the feed upset, with a true IAE of 0.120 against 0.056.

3.9 Technology in Practice

What Measures It

Chromatographs very commonly monitor and control the composition of distillation products, since one instrument quantifies several species, and an analyzer with its sample system needs much more maintenance than other instruments, with harder fault diagnosis [24]. In a survey reported in a 2013 review, the problems named most often with hardware sensors were maintenance and calibration, and one answer in ten named long dead time [25]. The chapter lumps the sample line and the analysis into the lag τ_a of Equation (3.9).

TECHNOLOGY Sensor: The Process Gas Chromatograph

A process chromatograph injects a precise small volume of sample into a long, narrow column, and a steady flow of carrier gas washes it through. Each species travels at its own rate, so the time it leaves tells it apart, and a detector that only tells carrier gas from anything else is enough. The area under each peak gives the quantity, and the timing is calibrated with a sample of known composition. The separation takes time, so the result arrives late, and ramping the column temperature or switching columns shortens the wait. A leak in the sample line lets outside gas dilute the sample, which a calibration gas sent through the whole line reveals [24].

What Moves It

The manipulated variable is the reflux ratio L/D, with the distillate draw held at half the feed as though by a perfect ratio controller. In a plant control system loops are often nested or cascaded, the setpoint of one taken from the output of another [26], and Figure 3.8, which draws the loop of Section 3.5 on field devices, sets the reflux valve under a flow controller whose setpoint the composition loop writes. The model applies each reflux ratio at once, with no flow loop and no valve lag.

TECHNOLOGY Final Element: The Reflux Control Valve

The reflux runs from the drum back to the top tray through a control valve. Tested at a constant pressure drop, a valve with linear trim passes a flow in proportion to its stem position. In a plant the pressure drop across the valve usually falls as the flow rises, so the installed flow droops at large openings, and equal-percentage trim brings the installed curve closer to a straight line. No trim makes up for pressure lost in restrictive piping, and a large valve can add a lag of tens of seconds to a loop [24]. The application applies each reflux ratio exactly, so all the curvature of Section 3.4.5 comes from the column.

Where the Controller Runs

A column's loops run in a distributed control system, which places a supervisory level over field controllers that hold their process conditions around set points, with cycle times from milliseconds to minutes [26]. A survey of predictive control sketches the usual hierarchy: basic PID loops in that system execute about every second, and a predictive controller above them about every minute [18]. The application sits close to that picture. Its PID executes every 0.02 min, or 1.2 s, its MPC every 0.2 min, and both are short beside an analyzer lag of up to 2 min and the 7.07 min time constant of the column. The velocity-form PID of Equation (3.12) is the function block of the standard loop, and the MPC with its Kalman filter or its moving horizon estimator replaces it when selected.

What Is Modernizing It

Three developments bear on this loop. The first is the analyzer's link to the control system. Older chromatographs give one analog output per species, while most modern ones can send several concentrations over a digital bus and compute derived values from them [24]. The second is the soft sensor, the estimator the Discussion proposes. In one company's survey it served to steady product quality, to save energy by operating close to the specification and to check the online analyzers, and distillation is its largest target in refining and chemicals [25]. The third is tuning that the controller does itself. Adaptive and on-demand tuning of PID controllers has been very successful in the marketplace, while adaptive predictive controllers remain rare, and predictive control has been moving into the distributed control systems themselves [18]. The relay test of Section 3.5.6 is such an on-demand tuner.

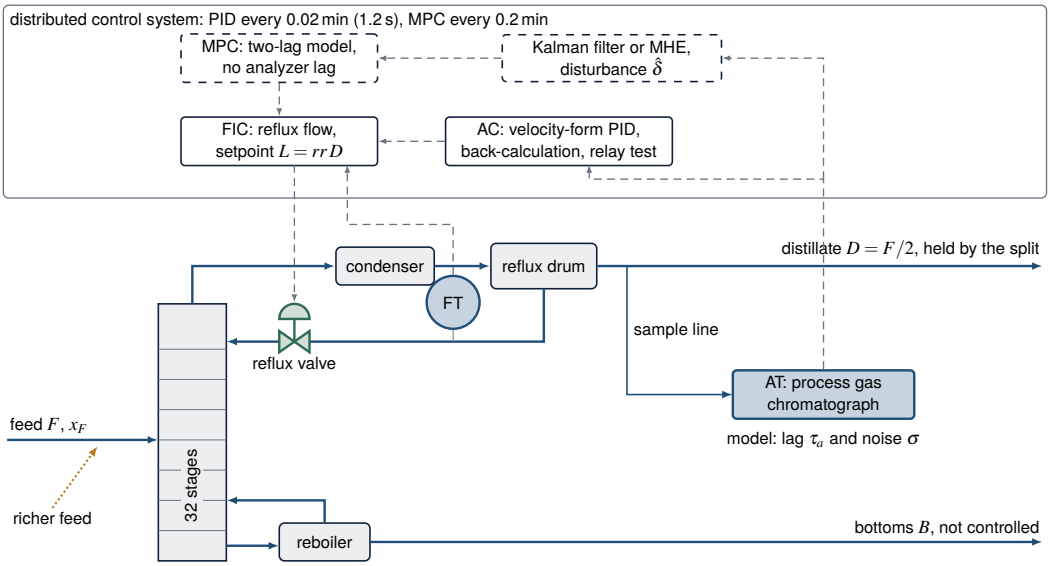


Figure 3.8 The composition loop on plant hardware. A process gas chromatograph (AT) reads a sample drawn from the distillate line, and its result feeds the composition controller (AC) and the estimator of the dashed predictive option. Either controller writes the reflux setpoint $L = rr\,D$ to a reflux flow loop, the plant's way of applying a reflux ratio that the model applies exactly. Notice that in the model the reflux acts on the top tray at once, while the measurement reaches the controller only after the sample line and the analysis, the lag that the MPC's model leaves out.

> **IN PRACTICE When the Analyzer Has Nothing New**
>
> Industrial predictive controllers meet an analyzer whose result is not available at every control execution. A typical solution skips the bias update of that controlled variable for a number of control intervals, and a counter switches control of the variable off if too many executions pass without feedback. A variable whose measurement status is good, and which the operator has enabled, stays under control [18]. The filter of Equation (3.15) corrects its estimate at every execution from a continuous reading, so a chromatograph that reports once per cycle would need such logic first.

3.10 Limitations

The analyzer is a lag and not a sampler. A real chromatograph delivers a new number once per cycle and the number describes the sample taken at the start of that cycle, which is closer to a dead time than to a first-order lag, with the mean age of Equation (1.8). A dead time of 2 min would leave the default tuning a predicted phase margin of 17.1°, half the 34.3° the 2 min lag of run L2 leaves, so Table 3.4 is an optimistic account of what a slow analyzer does.

The column is single ended by construction. Fixing $D = F/2$ removes the second degree of freedom, so dual composition control, the pairing question and the interaction that the relative gain array measures cannot be demonstrated. Volume 6, Chapter 3 measures that interaction on a two-by-two plant read by a scanning gauge. The bottoms composition is reported and logged, and Equation (3.10) means it carries no information the distillate composition does not already carry.

There is no energy balance and there are no temperatures, so inferential control, which is what most industrial columns actually run on, has nothing to be built from. There is no hydraulic lag, no flooding limit, no weeping limit and no pressure, so the reflux ratio can be moved between 1.0 and 10.0 without any of the physical consequences that would stop an operator doing the same thing.

The plant is stiff and the integrator is explicit, so the integration step is fixed at 0.002 min and

cannot be raised. Reducing the tray holdup to its lower limit of 0.10 mol shortens the fastest time constant from 0.0253 min to 0.0101 min and the stability limit from 0.0506 min to 0.0203 min, which is the reason the holdup ranges do not extend further down.

The predictive controller's model is two linear lags fitted at 0.970 on the base feed. It omits the analyzer lag, which brings the loop to the edge of stability at 1.15 min, the flattening of the gain toward high purity, absorbed as -5.3 units of reflux ratio in case 8, and the feed composition, seen only through the analyzer. There is no specification band or reflux rate limit, and the filter takes the reading for the distillate itself, so it is not the soft sensor of the Discussion. The MHE of Section 3.5.5 bounds the states of that same two-lag model, not the column's, so it shares the model's blindness to the analyzer, and with no bound reached on any case it adds nothing to the filter here. The relay test meets a loop with no dead time, where the swing exceeds the band by only 1.01 to 1.48 times, so its K_{cu} depends on the band and on the noise more than on the column, and the Tyreus–Luyben rule is one conservative choice among many.

> **OPEN CHALLENGE** **How Accurate Must the Model Be?**
>
> A survey of industrial predictive control found model uncertainty not addressed adequately. Every product could be detuned for robustness, but the trade-off between performance and robustness was not clear, so it could not be decided when a model was accurate enough for an application. Nearly all products judged model mismatch by brute-force simulation, and robust stability guarantees would greatly cut the time spent tuning and testing [18]. This chapter meets the question in small. The MPC's model leaves out the analyzer lag, its loop stops settling at 1.15 min of lag, and the largest move suppression moves that limit only a little. A test that predicted the limit before the run would be progress.

Finally, the noise is white and additive on the reading. Analyzer drift, a calibration offset and a failed sample line are all more common in practice than Gaussian noise, and none of them is represented.

> **OPEN CHALLENGE** **Trusting an Estimate Between Analyses**
>
> An estimate of composition is only as good as the measurements it is built from. It fails when an input sensor fails, and a statistical estimator should interpolate, never extrapolate. A blocked sample line makes the analyzer itself wrong. The simplest check compares the estimate with the analyzer when a result arrives, so its validity can be judged only then, and a 2013 review calls for another way to validate online. The hard part is to tell normal changes in the process from faults [25]. The analyzer drift, calibration offset and failed sample line that the Limitations name are such faults.

3.11 Discussion

The economics of a column are dominated by the reboiler. Separation work has to be supplied as heat at the bottom and rejected at the top, distillation is responsible for a large share of the process energy used by the chemical and refining industries, and the efficiency of the separation is set as much by how the column is operated as by how it was designed [15]. Composition control decides the operating part. A column held close to its specification limit meets the specification at the lowest energy use, while a column held well inside the limit gives away purity that the customer does not pay for.

How close a loop can be run to a limit is set by its variability, and Table 3.4 quantifies the analyzer directly in those terms. At $\tau_a = 0$ the peak excursion of the default tuning is 0.0026 mole fraction past setpoint. At $\tau_a = 2$ min it is 0.0106, so a setpoint that was safely inside the band at the fast analyzer has to be moved 0.0080 mole fraction away from the limit to give the same protection at the slow one. The detuned run D2 recovers most of that with an excursion of 0.0052, at the cost of settling 6.5 min later than the fast loop, which is time spent off target after every upset. This is the calculation that justifies a temperature inference trimmed by a slow analyzer rather than direct control on the

analyzer itself [5, 6], and it is the same calculation whether the loop is a single tray composition or a full multivariable scheme on both products [13].

The model of this chapter also supports control and estimation methods beyond the implemented PID loop and the linear predictive controller of Section 3.5.4. The thirty-two balances of Equations (3.3) to (3.7) form a complete nonlinear state-space model, so Model Predictive Control can enforce the reflux ratio limits and the specification band as explicit constraints rather than through detuning, and an added independent vapor rate would let one controller coordinate both product compositions. The lagged, noisy analyzer of Equation (3.9) is the standard setting for state estimation, where a moving horizon estimator or a Kalman filter reconstructs the current distillate composition from the delayed reading, a soft sensor (Section 1.4.8) that recovers part of the phase the analyzer removes. APMonitor implements nonlinear MPC and moving horizon estimation for models of exactly this differential algebraic form [27], and the Gekko Optimization Suite provides the same capabilities in Python [28]. The equations and parameters of Section 3.3 are sufficient to reproduce the column in either tool.

3.12 Conclusion

This chapter studied a thirty-two stage binary distillation column as an interactive laboratory for composition control through a slow measurement. The simulator integrates one component balance on each of a total condenser, fifteen rectifying trays, a saturated liquid feed tray, fourteen stripping trays and a partial reboiler, with a constant relative volatility of 1.6, tray holdups of 0.25 mol, explicit Euler on a 0.002 min step and a controller executing every 0.02 min. Fixing the distillate flow at $D = F/2$ makes the reflux ratio the only degree of freedom and forces $x_D + x_B = 2x_F$ at steady state, so a reflux move rearranges the profile without changing the 4.244 mol of light component the column holds, and the 8.17 min profile mode carries 109 % of the steady-state gain while the 15.24 min inventory mode carries −2.8 %. That is why a plant with thirty-two states and no stage slower than 0.40 min responds to a reflux step as a single 7.07 min lag with 0.21 min of apparent dead time and a gain of 0.0535 mole fraction per unit of reflux ratio.

Carrying one tuning across analyzer lags from 0 to 2 min moves the overshoot from 7.6 % to 30.7 % and the settling time from 13.5 min to 23.6 min without touching the column, and the phase margin of the identified model falls from 73.2° to 34.3° over the same range. Detuning to $K_c = 36$ and $\tau_I = 5$ min brings the peak back inside the 0.960 to 0.980 band at 2 min of lag and still settles 6.5 min later than the fast loop. A feed of $x_F = 0.42$ puts the 0.970 target above the $2x_F$ ceiling, and the reflux ratio sits on its 10.0 limit for the rest of the run because no tuning can supply light component the feed does not contain.

The linear predictive controller added beside the PID rejects the richer feed with an integral absolute error of 0.026 against 0.061 and reaches the reflux limit at once when 1.000 is requested, but its two-lag model omits the analyzer, and at 1.15 min of lag its loop stops settling. Selecting the moving horizon estimator for that controller instead of its Kalman filter reproduces the filter's run on every case to 2.6×10^{-8} in the reflux ratio, since no purity bound is reached, and a relay test on the PID writes an integral time from a cycle twelve times slower than the ultimate one, which raises the hour-long integral absolute error of the feed upset from 0.117 to 0.154.

A fuller model would replace the analyzer lag with a sampled dead time, add an independent vapor rate so that both product compositions can be controlled, and add tray temperatures so that an inferential scheme can be built and compared. As it stands, the column is a high-order plant whose limiting element is the measurement rather than the plant itself.

KEY POINTS

- A column of thirty-two stages answers a reflux move like one slow lag, which is why a two-parameter tuning rule works on a thirty-two state plant.
- The fitted dead time of 0.21 min is not a transport delay. It is the thirty fast eigenvalues rounding the foot of the response.
- Carrying one tuning across analyzer lags from 0 to 2 min moves the overshoot from 7.6 to 30.7 % and the settling time from 13.5 to 23.6 min without touching the column.
- The predictive controller rejects the richer feed better than the PID, but its model omits the analyzer, and at 1.15 min of lag its loop stops settling.

Where else this applies. Volume 12, Chapter 4 holds a far taller column through a tray temperature under a slow chromatograph, and Chapter 4 cooks hours ahead of its kappa analyzer.

References

[1] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[2] S. Skogestad and M. Morari, "Understanding the dynamic behavior of distillation columns," *Industrial & Engineering Chemistry Research*, vol. 27, no. 10, pp. 1848–1862, 1988.

[3] E. F. Wahl and P. Harriott, "Understanding and prediction of the dynamic behavior of distillation columns," *Industrial & Engineering Chemistry Process Design and Development*, vol. 9, no. 3, pp. 396–407, 1970.

[4] A. Ingimundarson and T. Hägglund, "Performance comparison between PID and dead-time compensating controllers," *Journal of Process Control*, vol. 12, no. 8, pp. 887–895, 2002.

[5] T. Mejdell and S. Skogestad, "Estimation of distillation compositions from multiple temperature measurements using partial-least-squares regression," *Industrial & Engineering Chemistry Research*, vol. 30, no. 12, pp. 2543–2555, 1991.

[6] W. L. Luyben, "Evaluation of criteria for selecting temperature control trays in distillation columns," *Journal of Process Control*, vol. 16, no. 2, pp. 115–134, 2006.

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

[8] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

[9] W. L. McCabe and E. W. Thiele, "Graphical design of fractionating columns," *Industrial & Engineering Chemistry*, vol. 17, no. 6, pp. 605–611, 1925.

[10] S. Skogestad, "Dynamics and control of distillation columns," *Chemical Engineering Research and Design*, vol. 75, no. 6, pp. 539–562, 1997.

[11] R. K. Wood and M. W. Berry, "Terminal composition control of a binary distillation column," *Chemical Engineering Science*, vol. 28, no. 9, pp. 1707–1717, 1973.

[12] S. Skogestad and M. Morari, "LV-control of a high-purity distillation column," *Chemical Engineering Science*, vol. 43, no. 1, pp. 33–48, 1988.

[13] S. Skogestad, "The dos and don'ts of distillation column control," *Chemical Engineering Research and Design*, vol. 85, no. 1, pp. 13–23, 2007.

[14] S. Skogestad, M. Morari, and J. C. Doyle, "Robust control of ill-conditioned plants: high-purity distillation," *IEEE Transactions on Automatic Control*, vol. 33, no. 12, pp. 1092–1105, 1988.

[15] A. A. Kiss and R. Smith, "Rethinking energy use in distillation processes for a more sustainable chemical industry," *Energy*, vol. 203, p. 117788, 2020.

[16] BCS, Incorporated and Oak Ridge National Laboratory, "Materials for separation technologies: Energy and emission reduction opportunities," U.S. Department of Energy, Industrial Technologies Program, Tech. Rep., 2005.

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

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

[19] 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.

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

[21] 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.

[22] 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.

[23] K. J. Åström and T. Hägglund, "Automatic tuning of simple regulators with specifications on phase and amplitude margins," *Automatica*, vol. 20, no. 5, pp. 645–651, 1984.

[24] 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.

[25] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

[26] 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.

[27] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[28] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, "GEKKO optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.

3.13 Exercises

3.1 Concept check · 2 min

Run FR steps the feed rate from 1.00 to 1.50 mol/min at 2 min with the loop in automatic. Over the next hour the distillate stays at 0.9700 and the reflux ratio at 3.7319. Why does the step move nothing?

(a) At a fixed reflux ratio every flow scales with the feed, so the steady profile does not depend on the feed rate and only the time scale changes

(b) The PID rejects the step within one execution, faster than the analyzer can register the change in the distillate

(c) The feed rate enters only the feed-tray balance, and the reflux ratio overrides that tray before the change reaches the top

(d) The distillate draw takes up all of the extra feed, so the liquid and vapor rates inside the column do not change

3.2 Concept check · 1 min

The reflux step from 3.0 to 3.5 is fitted with a first-order plus dead-time model, which gives $\theta_p = 0.21$ min. The analyzer lag and noise are zero during the test. Where does this dead time come from?

(a) The analyzer, which reports the distillate through a first-order lag

(b) The thirty fast modes of the stage balances, which round the foot of the response

(c) The time the reflux liquid takes to run down the trays to the reboiler

(d) The controller period and the integration step, which hold each reflux move

3.3 Concept check · 2 min

With the analyzer lag at 1.15 min the MPC loop of case 9 stops settling, while the PID with the shipped tuning still settles at a lag of 2 min. Which explanation fits the chapter's measurements?

(a) The MPC has no integral action, so the lag leaves an offset that it cannot remove
(b) The analyzer noise excites the plans, and without the noise the MPC settles at every lag
(c) The MPC's model omits the analyzer lag, and its faster answer leaves it less phase to lose
(d) The reflux ratio reaches its lower limit early, and the clipped plans wind up the estimate

3.4 Theory and calculation · 6 min

The column runs at $rr = 3$ with $F = 1\,\text{mol/min}$ and the holdups of Table 3.2.

(a) Compute D, B, L, V and L_S from Equation (3.1). *Checkpoint:* $D = B = 0.5$, $L = 1.5$, $V = 2.0$ and $L_S = 2.5\,\text{mol/min}$.

(b) Compute the liquid residence time of a rectifying tray, a stripping tray, the condenser and the reboiler, and the ratio of the apparent time constant of 7.07 min to the longest of them. *Checkpoint:* 0.167, 0.100, 0.250 and 0.400 min. The ratio is 17.7, about eighteen.

(c) The feed rate is raised to 1.5 mol/min. Predict the gain, the time constant and the apparent dead time of the same reflux step, from 0.0535, 7.07 min and 0.208 min at $F = 1$. *Checkpoint:* Gain 0.0535, $\tau_p = 4.71\,\text{min}$, $\theta_p = 0.139\,\text{min}$.

3.5 Theory and calculation · 6 min

The step test gives $K_p = 0.0535$ per unit of reflux ratio, $\tau_p = 7.07\,\text{min}$ and $\theta_p = 0.21\,\text{min}$. Use the IMC rule of Equation (3.11).

(a) Compute K_c and τ_I for $\tau_c = 2\,\text{min}$. *Checkpoint:* $K_c = 59.8$ per mole fraction, $\tau_I = 7.07\,\text{min}$.

(b) Compute K_c and τ_I for $\tau_c = 1\,\text{min}$. *Checkpoint:* $K_c = 109$, $\tau_I = 4.84\,\text{min}$.

(c) The shipped tuning is $K_c = 60$ and $\tau_I = 4\,\text{min}$. Compare its integral gain K_c/τ_I with that of part (a), and say which of the two removes a sustained load error faster. *Checkpoint:* 15.0 against 8.46 per mole fraction per minute. The shipped tuning, by a factor of 1.77.

3.6 Theory and calculation · 6 min

Run N0 repeats the setpoint step with the analyzer noise at $\sigma = 0.0015$ mole fraction, $K_c = 60$ and $\tau_D = 0$, and the reflux ratio travels far more than without noise. Treat the reading noise as white.

(a) Find the standard deviation of the proportional increment $K_c(e_k - e_{k-1})$ that the noise alone produces. *Checkpoint:* $K_c\sigma\sqrt{2} = 0.127$ units of reflux ratio.

(b) Find the expected travel $\sum |rr_k - rr_{k-1}|$ over the 2999 executions that follow the step, ignoring the small integral increments. *Checkpoint:* $0.127\sqrt{2/\pi} = 0.102$ per execution, 305 in total.

(c) The exported history samples the reflux ratio every 0.2 min, every tenth execution. Estimate the travel it shows over the same hour, and explain why it is so much smaller. *Checkpoint:* About 30.5 over 300 samples. The ten increments between two samples telescope into one difference of two readings.

3.7 App challenge · 9 min

Load scenario 9, MPC with a late analyzer, with the default seed. The MPC is in charge and the analyzer lags by 1 min. Run to 20 min at a speed of 1 and read the IAE and MV travel tiles. Set Controller to PID, press Reset and repeat.

(a) Report the two tiles for each controller. *Checkpoint:* MPC: IAE 0.028, MV travel 7.22. PID: IAE 0.064, MV travel 11.4. Within 0.002 and 0.1.

(b) Reload scenario 9, set Analyzer lag to 1.5 min before Run, and run to 60 min at a speed of 5. Read the IAE and Saturated tiles. Set Controller to PID, press Reset and repeat. *Checkpoint:* MPC: IAE 0.218, Saturated 12.7 %. PID: IAE 0.122, Saturated 0 %. Within 0.005 and 0.5 %.

(c) Explain why the order of the two controllers reverses between parts (a) and (b). *Checkpoint:* The MPC's model omits the lag. At 1.5 min its swings grow until the reflux reaches its lower limit, while the PID's decay.

3.8 App challenge · 7 min

Load scenario 8, Actuator limit, with the default seed. The setpoint steps from 0.970 to 1.000 at 2 min. Run to 120 min at a speed of 10.

(a) Export the CSV file and find the first row in which `saturated` is 1. Read the Saturated and IAE tiles at 120 min. Why does the LIMIT badge flicker for several minutes before it stays lit? *Checkpoint:* 73.4 min, within one row. Saturated 36.8 % and IAE 0.519, within 0.5 % and 0.005.

(b) The feed composition is 0.50. What ceiling does Equation (3.10) put on the distillate, and is it the limit that stops this run? *Checkpoint:* $2x_F = 1.000$. No. The reflux limit stops the distillate at 0.9971.

(c) Run XF also ends with the reflux on its limit. Name the remedy that each run calls for. *Checkpoint:* Case 8: lower the setpoint, since even a higher reflux limit buys almost nothing. Run XF: cut the draw below $F/2$, raise the vapor independently, or lower the target.

3.9 App challenge · 8 min

Load the challenge, Challenge · return to target, the tenth case in the list, with the default seed. It starts at 0.970 with $K_c = 12$ and $\tau_I = 4$ min, and the feed composition steps to 0.58 at 2 min. Run to 30 min at a speed of 5 each time. The challenge result is the worst $|\texttt{pv} - \texttt{sp}|$ logged from 19.8 min on, which the case's progress list reports and the exported CSV also gives.

(a) Run as loaded and read the challenge result. *Checkpoint:* Worst error 0.0143 against ± 0.01: not met.

(b) Press Reveal suggested tuning, press Reset and run again. Then set Controller to MPC, press Reset and run a third time. Read the challenge result of each. *Checkpoint:* Suggested PID ($K_c = 60$): 0.00518, met. MPC: 0.00555, met. Within 0.0003.

(c) Export the CSV of the last two runs and find the largest $|\texttt{true_pv} - \texttt{sp}|$ from 19.8 min on. Which controller holds the true distillate closer, and why does the judged error not show it? *Checkpoint:* MPC 0.0008, PID 0.0021, within 0.0002. The judged error is read on the noisy analyzer, and the noise dominates it.

3.10 Application area · 4 min

The analyzer lag of Table 3.4 decides how far the setpoint must sit from the specification limit. Under constant molar overflow the reboiler duty is proportional to the vapor rate $V = (rr + 1)D$.

(a) From the peaks of runs L0, L2 and D2, find how far the setpoint must move away from the limit when the lag rises from 0 to 2 min, and what share of that the detuned run D2 recovers. *Checkpoint:* 0.0106 against 0.0026: 0.0080. D2 peaks 0.0052 past setpoint, recovering 0.0054, about 68 %.

(b) The steady map gives 0.9765 at $rr = 4.0$, against $rr = 3.7319$ for 0.970. By what percentage must the vapor rate rise to hold 0.9765 instead of 0.970 at the base feed? *Checkpoint:* V from 2.37 to 2.50 mol/min, 5.7 % more.

3.11 Application area · 4 min

A plant chromatograph on the distillate runs a 2 min cycle. Each result describes the sample drawn at the start of a cycle, is published at its end and is held through the next cycle.

(a) Use Equation (1.8) to find the mean age of the result the controller sees. *Checkpoint:* $t_d = T_s = 2$ min, so $t_d + T_s/2 = 3$ min.

(b) The chapter's Limitations section finds that a pure dead time of 2 min leaves the default tuning a phase margin of 17.1°, against 34.3° for a lag of 2 min. What do you expect for this chromatograph, and what would you do before closing the loop on it? *Checkpoint:* Worse than 17.1°: on the fitted model the loop is unstable, gain margin about 0.85. Detune, or control a tray temperature and trim it with the analyzer.

3.12 Review the agent · 5 min

An AI assistant is asked to improve the composition loop. It replies: "Run the relay autotune in case 3 and keep the constants it writes, $K_c = 79.9$ and $\tau_I = 11.8$ min. A relay test measures the column's ultimate point directly, so Tyreus–Luyben settings from it are better founded than the hand-picked 60 and 4 min, and the longer integral time can only make the loop more robust." Find the flaw, and say how you would check the advice in the app.