INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 8 Batch, Run-to-Run, and Hybrid Control

CHAPTER 3

Reflux Policy on a Batch Column With No Steady State to Tune

Pages 55–82

Edited by John D. Hedengren

How to cite this chapter. "Reflux Policy on a Batch Column With No Steady State to Tune," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 55–82.

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

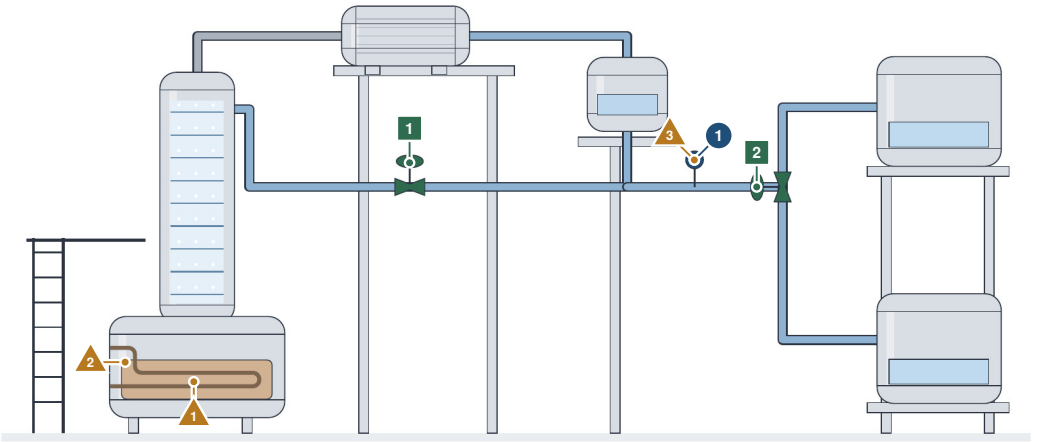
The app for this chapter
https://apmonitor.com/apps/batch-distillation/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v08-batch-hybrid/#batch-distillation

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

CHAPTER 3

Reflux Policy on a Batch Column With No Steady State to Tune



The controller moves 1 reflux split, the share of condensate returned 2 receiver switch, product or slop

The controller reads 1 distillate composition analyzer

What disturbs it 1 pot running out of the light component 2 loss of heater duty 3 analyzer lag and bias

Nothing is fed during the batch, so the reflux returned at the top decides how much of the falling pot reaches the product receiver on specification. Illustration, not to scale.

IN PLAIN TERMS Getting the Most Good Product From One Pot

The main idea. A batch still boils one fixed load of a liquid mixture. The lighter liquid boils off first and is collected at the top. Part of what condenses is poured back down the column, which makes the product purer but slower to come.

The challenge. The pot holds less of the light liquid as it boils, so the same setting makes weaker product as time goes on. The whole tank of product is judged once, at the end, on its average. Very pure early product can pay for weaker product later.

Why it matters. How much is poured back, minute by minute, decides how much good product one batch gives.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history at the start of Section 3.2, the fact on energy in Section 3.7, the two try-its and the fact on analyzer lag in Section 3.6, and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 3.6, through the off-design charge), Technology in Practice (Section 3.8), Limitations (Section 3.9) and the open challenge

The control question A batch column has no operating point, and the batch is judged once, on the amount and the cumulative purity in the receiver at the end. Which reflux policy collects the most product, and why can a composition loop that holds its setpoint closely collect less than an open-loop ramp?

Learning objectives

- Derive the total-reflux separation and explain why the top composition and the process gain fall.
- Explain how the cumulative cut rule creates a purity credit.
- Compare constant reflux, constant composition and a searched ramp on product, energy and reflux demand.
- Predict what analyzer lag and bias do to the loop and to the cut.
- Assess what a trim or re-optimization recovers on an off-design charge.

Prerequisites Sections 1.4.1, 1.4.2, 1.4.3, 1.4.4 and 1.6.2. Section 1.6.3 for the predictive option only.

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

See also Volume 9, Chapter 12 for a time budget met by arcs at bounds and Volume 9, Chapter 6 for predictive control on an economic objective. Volume 13, Chapter 8 for the energy of a continuous separation. Volume 5, Chapter 3 and Volume 12, Chapter 4 for continuous columns through a slow analyzer and an inferred composition.

3.1 Introduction

Regulatory control assumes an operating point. A setpoint is held against disturbances, a step test is read after the response settles, and performance is measured as deviation from a fixed condition [1]. A batch distillation column provides no operating point. The pot is charged once, light component leaves through the top continuously, and no stream replenishes it. The pot composition therefore falls monotonically from the start of collection to the end of the batch, and the separation becomes more difficult with no input change. In the simulator of this chapter the distillate leaves the condenser at 0.9996 mole fraction methanol and, ninety minutes later at an unchanged reflux ratio, at 0.9599. No equipment failed in between. The pot was depleted.

This behavior changes the role of the controller. There is no disturbance in the ordinary sense, because the depletion of the pot is continuous, monotonic and inherent to the operation. There is no operating point, so a tuning identified in the first ten minutes describes a column that no longer exists in the last ten. The quantity the operation is judged on is not a composition at all. It is the moles of on-specification product in the receiver at the terminal time, measured once, on the cumulative average of the whole drum. An objective that exists only at the terminal time is the defining feature of batch operation, and it is why batch plants are run on recipes and schedules rather than on loop tuning [2].

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

The class is broad. Any plant whose feed is its own inventory shares this structure, including batch crystallization to a yield (Chapter 4), a batch reactor driven to a conversion target, evaporative concentration and semi-batch polymerization. The control problem in every case is the choice of an input trajectory over a fixed horizon against a cost evaluated at the end, with path constraints in the middle, and the structure of the solution is known to alternate between constraint-seeking and sensitivity-seeking arcs rather than to hold anything constant [3]. Conclusions drawn on the batch column therefore transfer across the class.

Simulators of distillation columns are usually built around the continuous case, where a feed arrives at a constant rate and the design questions are pairing, interaction and measurement lag [1], as on the continuous column of Volume 5, Chapter 3. However, a student who has only seen continuous columns carries a specific misconception into a batch plant, which is that a composition loop holding its setpoint is performing well. Constant-composition control on the column of this chapter holds the distillate within 0.0012 of 0.990 from 40 to 120 min and finishes second of the three policies, because the nearly flat controlled variable is paid for with a reflux ratio that climbs from 2.53 to its

limit of 40 and a distillate rate that falls by a factor of about twelve. A plot of the controlled variable shows almost none of this.

The column shows the principles that separate batch from continuous control (Section 1.4.1), including an economic objective evaluated only at the terminal time, a process gain that decays over the run, an actuator whose useful range collapses as the plant is depleted, and measurement lag and bias entering a decision that cannot be reversed [1, 4]. The model rests on component material balances, vapor-liquid equilibrium at constant relative volatility as the separation driving force, and a time-scale separation of three orders of magnitude between the tray holdups and the pot. Batch distillation carries most solvent recovery in fine chemicals and pharmaceuticals, so the choice of reflux policy is an industrial decision with a measurable energy and material cost [5].

This chapter compares three reflux policies used in batch operation on one browser column, with the same charge, the same duty and the same time budget behind all three. The column model, the policies and the grid search are all standard. Section 3.2 traces the policies from the simple still to trajectory optimization. Section 3.3 specifies the nineteen-state column down to its cumulative cut rule and its analyzer, with the balances, the constants and the numerics a second implementation would need. Section 3.4 shows why the top composition and the process gain decline over the run, and that the cumulative form of the product specification is what creates the optimization problem at all, since an instantaneous specification would make constant-composition control optimal by construction. Section 3.5 writes the three policies through one discrete controller and describes the grid search behind the reflux ramp. Section 3.6 measures the policies. The optimized trajectory outperforms both classical policies open loop and then loses most of the batch when the charge differs from the one it was computed for. The same section measures analyzer lag and bias, duty, holdup and volatility, the graded challenge and a shrinking-horizon predictive option. Sections 3.7 and 3.9 take up the energy and purity costs and what the model leaves out, and Section 3.8 describes the analyzer, the reflux splitter and the control system of a real batch column.

3.2 Background

The quantitative treatment of batch separation is older than the tray column, and its first balance, for a simple still in 1902, is told in the history segment below. That balance contains the central difficulty of batch separation. The charge is the state, the vapor is always richer than the charge, and integrating the balance forward gives a composition that falls monotonically. Adding trays and a reflux stream raises the enrichment achieved at each instant and does not change the direction of the trajectory.

> ⧖ HISTORY **1902: The Residue Moves One Way**
>
> In 1902 Lord Rayleigh published experiments on a question for which information seemed almost entirely lacking: how the strength of a boiling binary mixture relates to the strength of the vapor in equilibrium with it. If the curve relating the two never crosses the diagonal, the distillate is always richer than the liquid, and the residue of a limited charge moves always in one direction toward purity. If it crosses, as aqueous alcohol does at 96 %, distillation stops having any effect there. His calculation of the residue treated the vapor's composition as the change of one ingredient over the change of the whole. For very weak alcohol in water it found that halving the liquid in the retort cut its alcohol strength in the ratio 2^{11} to 1, so the residue approached pure water rapidly [6].

Two operating policies grew out of that observation and have been the practical choices ever since. In constant reflux the ratio is set once and the distillate composition is allowed to slide until the accumulated product no longer meets the specification, which is what a column with a hand-set reflux divider can do. In constant composition an operator or a controller raises the reflux ratio continuously to hold the overhead composition, which trades distillate rate for purity and requires a measurement. The comparison between the two was posed as an optimal control problem in 1963, when the reflux policy that maximizes distillate over a fixed time was derived and shown to be neither of them [7].

The complementary time-optimal problem, the policy that reaches a specified quantity of a specified purity in the least time, followed shortly afterward and returned a reflux ratio that rises throughout the run [8]. The classical result has been the same ever since. The optimal policy holds nothing constant.

Computational cost kept the optimal policy out of practice for two decades. Rigorous stage-by-stage integration of a batch column was demonstrated in 1968 with the stiffness, the holdup dynamics and the numerical difficulties laid out explicitly [9], and it was expensive enough that the optimization literature turned instead to shortcut models. The shortcut treatment models the column at each instant as a continuous column at steady state whose feed is the current pot, which reduces the dynamics to one differential equation in the pot composition and made optimal reflux policies computable for multicomponent mixtures [10]. The cost of the simplification is holdup. Shortcut models assume none, and holdup is exactly what has to fill with on-specification material before any product can be taken, so the simplification is not neutral for the quantity being optimized [11]. The argument between shortcut and rigorous models organized the field for a generation and is surveyed with its computational history in the standard treatments [12, 13]. Solving the trajectory problem directly on a rigorous model became routine only when orthogonal collocation and nonlinear programming were brought to bear on it [14], and that is the technique the reference behind this application uses.

The measurement side of the problem has its own history. Composition control on a batch column is harder than on a continuous one for reasons that have nothing to do with tuning, because the gain the loop works against changes over the run and the analyzer that reports the overhead composition is slow, so the practical arrangements infer the composition from tray temperatures rather than waiting for the analyzer measurement [15], as the propylene splitter of Volume 12, Chapter 4 does on a continuous column. This application is built directly from a published experimental study of a laboratory column that combines these elements, where the model parameters were estimated from data with nonlinear confidence intervals and a sensitivity ranking before any trajectory was optimized [16]. That apparatus is a 38 tray, vacuum-jacketed Oldershaw column with a 600 W reboiler and a swinging bucket reflux divider that can be moved once every 5 min, charged with a 50 % by weight methanol and ethanol mixture and run to a 99 mol% methanol specification. Its base case is 30 minutes at total reflux and then a constant reflux ratio of 4, and its optimized reflux schedule starts low and rises in a nominally linear pattern, raising production by 14 %. The charge, the duty, the heating efficiency, the specification, the startup and the cut rule used here are all taken from that study.

3.3 Process Model

Column T-401 is a laboratory batch column with a reboiler that holds the charge, ten theoretical trays and a total condenser. The pot is charged with 28 mol of methanol and ethanol at 0.59 mole fraction methanol, a 600 W heater boils it, and all of the vapor leaving the top is condensed. The reflux split, the fraction of that condensate returned down the column against the fraction withdrawn as product, is the only operating decision. Every stream, instrument tag and symbol the equations below use is drawn in Figure 3.1. Time is in minutes throughout, inventories in moles, flows in mol/min and compositions in mole fraction of methanol, which is the light component.

3.3.1 Controlled, Manipulated and Disturbance Variables

Table 3.1 classifies the signals. A single composition is controlled, a single reflux ratio is manipulated, and the three disturbances are all optional, because the pot depletes with no disturbance applied.

Nineteen states are integrated: the pot inventory and composition, ten tray compositions, the condenser composition, four receiver states carrying total and light-component moles for the product and slop drums separately, the analyzer, and a cumulative reboiler energy meter. Integrating the two receivers as moles rather than as running average compositions avoids a zero divided by zero at the first drop and lets both mole balances close to integrator accuracy.

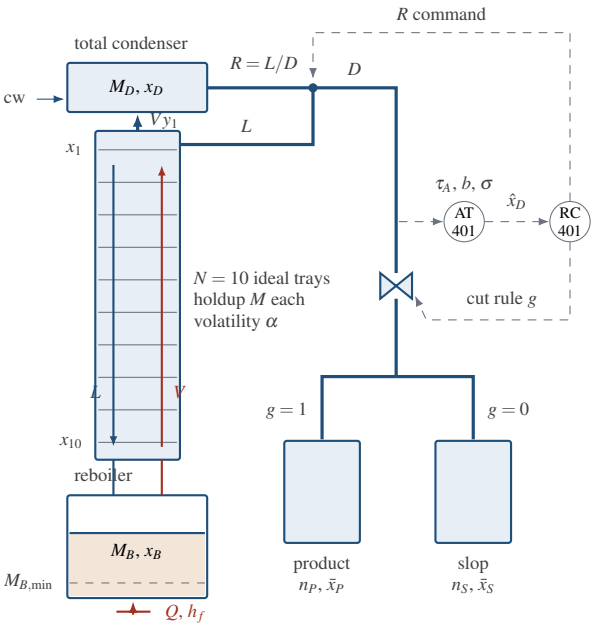


Figure 3.1 Batch column T-401 with its reboiler, ten trays, total condenser and two receivers. The reboiler holds the whole charge M_B at composition x_B and is boiled at rate V set by duty Q and heating efficiency h_f. Vapor rises through trays of holdup M against liquid L falling from above, and the total condenser holds M_D at x_D. The condensate is split at reflux ratio $R = L/D$, which is the only manipulated variable. Analyzer AT-401 reports $\hat{x}_D$ through a lag τ_A with bias b and noise σ, and RC-401 sets R and operates the routing flag g that sends the distillate to the product drum or to slop.

Table 3.1 Controlled, Manipulated, Disturbance and State Variables. Nominal Values Are Those at the Start of Product Takeoff on the Shipped Charge.

Symbol	Role	Nominal	Range
x_D (AT-401)	CV, measured	0.9996	0.90–0.999 SP
R	MV	4	0–$R_{\max}$
duty derate	DV, measured	0 %	0–60 %
b analyzer bias	DV, unmeasured	0	±0.02
σ analyzer noise	DV, unmeasured	0	0–0.01
M_B	state	27.15 mol	4–27.15 mol
x_B	state	0.5787	0–1
$x_1 \ldots x_{10}$	state	0.9993–0.7331	0–1
x_D	state	0.999 64	0–1
n_P, n_{PA}	state	0 mol	≥ 0
n_S, n_{SA}	state	0 mol	≥ 0
x_M	state	0.999 64	0–1
E	state	1080 kJ	≥ 0

3.3.2 Equilibrium and Stage Balances

Stages are ideal, so the vapor leaving a tray is in equilibrium with the liquid standing on it, and Equation (3.1) is that equilibrium written at a fixed relative volatility.

$$y_i^* = \frac{\alpha x_i}{1 + (\alpha - 1) x_i}, \qquad \frac{y^*}{1 - y^*} = \alpha \frac{x}{1 - x} \tag{3.1}$$

where α is the relative volatility, x_i is the liquid mole fraction of methanol on stage i and y_i^* is the vapor mole fraction in equilibrium with it. The second form governs the behavior of the application, because relative volatility multiplies the ratio of light to heavy and not the mole fraction itself. The enrichment available from a lean pot is therefore far smaller than the enrichment from a rich one.

Each tray is a well-mixed holdup receiving liquid from the stage above and vapor from the stage below, with constant molar overflow making L and V the same on every tray, as shown in Equation (3.2).

$$M\frac{\mathrm{d}x_i}{\mathrm{d}t} = L\,(x_{i-1} - x_i) + V\,(y^*_{i+1} - y^*_i), \qquad i = 1 \ldots N \tag{3.2}$$

where M is the liquid holdup per tray in moles, the liquid entering tray 1 arrives from the condenser so that x_{i-1} is x_D there, and the vapor entering tray N leaves the reboiler so that y^*_{N+1} is y^*_B. Both sides carry mol/min. A tray time constant is $M/V = 0.10\,\mathrm{min}$ against a batch that lasts 180 min, so the trays are very nearly quasi-steady and the slow variable is the pot.

The reboiler balances separate the batch column from the continuous case, as shown in Equations (3.3) and (3.4).

$$\frac{\mathrm{d}M_B}{\mathrm{d}t} = L - V = -D \tag{3.3}$$

$$M_B\frac{\mathrm{d}x_B}{\mathrm{d}t} = L\,(x_N - x_B) - V\,(y^*_B - x_B) \tag{3.4}$$

where M_B is the pot inventory, x_B the pot composition and y^*_B the vapor in equilibrium with it. Nothing on the right-hand side replenishes the pot, and everything leaving the top is richer in methanol than the pot is, so x_B falls monotonically for the whole batch. At a constant reflux ratio of 4 it travels from 0.5787 to 0.2337 in ninety minutes.

All the vapor condenses and both outgoing streams leave at the same composition, so the condenser balance does not contain the reflux ratio at all, as shown in Equation (3.5).

$$M_D\frac{\mathrm{d}x_D}{\mathrm{d}t} = V\,(y^*_1 - x_D), \qquad D = \frac{V}{R+1}, \qquad L = V - D \tag{3.5}$$

where M_D is the condenser and reflux drum holdup and R is the reflux ratio. Writing the distillate first and the reflux as the remainder makes $L + D$ equal V exactly, which keeps the total mole balance closed to integrator accuracy, and it accepts $R \to \infty$ for total reflux without a special case. The boil-up follows the reference energy balance in Equation (3.6).

$$V = \frac{Q\,h_f}{\Delta H_{vap}} \tag{3.6}$$

where Q is the reboiler duty expressed in kJ/min, h_f is the heating efficiency and ΔH_{vap} is the heat of vaporization. A 600 W heater, which is 36 kJ/min, at 71.9 % efficiency against 36.6 kJ/mol gives $V = 0.7072\,\mathrm{mol/min}$. Duty sets the throughput and the reflux ratio sets the split, and the two are independent, which is why losing duty costs product without costing any separation.

At total reflux $L = V$ and $D = 0$, the tray balances of (3.2) are satisfied exactly by $y^*_i = x_{i-1}$, and the ratio in (3.1) multiplies by α across each of the ten trays and again across the reboiler. The result is the Fenske relation of Equation (3.7), exact in this model rather than a shortcut.

$$\frac{x_D}{1 - x_D} = \alpha^{N+1}\,\frac{x_B}{1 - x_B} \qquad \text{at total reflux} \tag{3.7}$$

Eleven equilibrium stages at $\alpha = 2$ give a separation factor of $2^{11} = 2048$, and total reflux is the only stationary condition in the model.

3.3.3 Cumulative Cut Rule

The receiver moves through three phases in a fixed order. During the *fore cut* nothing is on specification and the distillate goes to slop. The *main cut* opens when the analyzer first reads the specification. The *tail cut* closes it when the cumulative average of everything already in the product drum falls to the specification, after which everything goes to slop. The receivers integrate as shown in Equation (3.8).

$$\frac{\mathrm{d}n_P}{\mathrm{d}t} = D\,g, \qquad \frac{\mathrm{d}n_{PA}}{\mathrm{d}t} = D\,x_D\,g, \qquad \bar{x}_P = \frac{n_{PA}}{n_P} \tag{3.8}$$
$$\frac{\mathrm{d}n_S}{\mathrm{d}t} = D\,(1-g), \qquad \frac{\mathrm{d}n_{SA}}{\mathrm{d}t} = D\,x_D\,(1-g)$$

where g is 1 during the main cut and 0 otherwise, n_P and n_{PA} are the total and light-component moles in the product drum, and $\bar{x}_P$ is the cumulative average on which the product is graded. The switching rule is Equation (3.9), and it is evaluated on the analyzer rather than on the true composition because the analyzer is the only composition information the operation has.

$$g = \begin{cases} 0 \to 1, & \hat{x}_D \geq x_{\mathrm{spec}} \quad \text{(fore to main)} \\ 1 \to 0, & \dfrac{\int D\,\hat{x}_D\,g\,\mathrm{d}t}{\int D\,g\,\mathrm{d}t} < x_{\mathrm{spec}} \quad \text{(main to tail)} \end{cases} \tag{3.9}$$

The cumulative form is the reference's own rule and it is what creates the optimization problem. Material collected above specification builds a credit, material collected below it spends the credit, and the receiver remains on specification while the credit is positive. Section 3.6 shows the searched policy finishing with an instantaneous distillate composition of 0.8285 and a receiver still at 0.9897.

3.3.4 Analyzer and Energy Meter

One online analyzer measures the distillate. Its lag is a state rather than a display filter, because the lag is large enough to change the outcome of a batch, as shown in Equation (3.10).

$$\tau_A \frac{\mathrm{d}x_M}{\mathrm{d}t} = x_D - x_M, \qquad \hat{x}_D = \mathrm{clamp}(x_M + b + n(t),\, 0,\, 1) \tag{3.10}$$

where x_M is the lagged composition, b is a standing bias and n is band-limited noise, formed by filtering a Gaussian sequence with a 0.5 min time constant at every controller execution and scaling the innovation so that the stationary standard deviation is σ. The bias and the noise are applied in the controller and not in (3.10), because the fourth-order integrator would otherwise sample the noise four times per step at four different values. The energy meter is Equation (3.11).

$$\frac{\mathrm{d}E}{\mathrm{d}t} = 0.06\,Q_{\mathrm{eff}}, \qquad Q_{\mathrm{eff}} = Q\left(1 - \frac{\mathrm{derate}}{100}\right) \tag{3.11}$$

where E is cumulative reboiler energy in kJ, Q_{eff} is the effective duty in watts and the coefficient 0.06 converts watts to kJ/min. The meter accrues at full duty whether or not anything leaves the top, which is why a high reflux ratio held for a long time raises the energy per mole of product.

3.3.5 Parameters, Initialization and Numerical Integration

Table 3.2 lists every constant, the plant parameters above the rule and the operator settings a scenario overrides below it. The identification applications in this series withhold the quantities a step test is meant to recover, while this application discloses every parameter, because the task is the choice of a trajectory against a plant whose parameters are given rather than the recovery of parameters from a response.

A fresh charge opens with the pot, every tray, the condenser and the analyzer all at x_0, and the pot holding $M_B(0) = M_0 - NM - M_D = 27.15$ mol. Scenario 2 instead opens on a column that has already

Table 3.2 Model Parameters and Controller Defaults With Their Adjustable Ranges

Symbol	Meaning	Value	Unit	Range
α	relative volatility	2.0	—	1.2–3.0
N	theoretical trays	10	—	fixed
M_0	charge	28	mol	12–45
x_0	charge composition	0.59	—	0.25–0.90
M	holdup per tray	0.07	mol	0.02–0.30
M_D	condenser holdup	0.15	mol	0.05–0.80
Q	reboiler duty	600	W	150–1000
h_f	heating efficiency	0.719	—	0.40–1.00
ΔH_{vap}	heat of vaporization	36.6	kJ/mol	30–45
x_{spec}	product specification	0.990	—	0.90–0.999
$D_{\min}$	production-stopped threshold	0.02	mol/min	0.002–0.1
τ_A	analyzer time constant	0.5	min	0.05–10
$R_{\max}$	maximum reflux ratio	40	—	5–100
t_{batch}	batch time budget	180	min	30–300
$M_{B,\min}$	minimum pot inventory	4	mol	0.5–15
t_{start}	total-reflux startup	30	min	0–90
$x_{D,\text{SP}}$	composition setpoint	0.990	—	0.90–0.999
R_{set}	constant reflux ratio	4.0	—	0–$R_{\max}$
R_0	hand-set ramp start	2.0	—	0–20
R_1	hand-set ramp end	10.0	—	0–$R_{\max}$
K_c	proportional gain	600	—	0–1500
τ_I	integral time	4.0	min	0.2–60
τ_D	derivative time	0	min	0–5
τ_f	derivative filter	0.2	min	0–2

been brought up at total reflux. That state is constructed from (3.7) by bisecting the light-component balance of Equation (3.12) for the pot composition, which is monotone in x_B.

$$M_B x_B + M \sum_{i=1}^{N} x_i(x_B) + M_D x_D(x_B) = M_0 x_0 \tag{3.12}$$

The result is $x_B = 0.5787$, leaner than the 0.59 charged, because the 0.85 mol of holdup in the trays and the condenser has filled with much richer material and that methanol came out of the pot. The corresponding condenser composition is 0.999 64.

The engine of Section 1.6.1 integrates the column with classical fourth-order Runge-Kutta on a fixed 0.01 min step. RC-401 executes every 0.05 min, which is 3 s of plant time, and the historian keeps ten samples per simulated minute. Noise comes from a generator seeded at 20260829, so a noisy batch is exactly reproducible. The speed setting counts simulated minutes per second of wall time, and at the default of ten a full 180 min batch completes in about eighteen seconds. Every discontinuous event, the end of the startup, the switch between slop and product, and the end of the batch, is implemented in the controller and reaches (3.2) to (3.8) only as R and g, both constant across a step, so the four Runge-Kutta stages never straddle an event.

3.3.6 Assumptions

Six simplifying assumptions define the model, and the condition under which each fails is stated with it.

Constant relative volatility, and not the reference's own value. The reference computes equilibrium from a quartic fit to a physical property database. Across the whole composition range that quartic is a constant relative volatility of 1.725 to an RMS deviation of 4.2×10^{-4} mole fraction, so nothing is lost by using (3.1) instead. What is changed deliberately is the number itself, and Section 3.5.4

gives the reasoning and the cost. Any mixture with a strongly curved equilibrium line is outside a constant-volatility model, and every azeotropic mixture is outside it completely, because the relative volatility passes through 1 at the azeotrope and no number of stages separates anything there.

Constant molar overflow. The internal flows are the same on every tray and there is no energy balance inside the column. Methanol and ethanol differ by about 9 % in molar heat of vaporization, which makes the assumption reasonable here and would not make it reasonable for a water and acetic acid split. The reference validates the same simplification against a detailed model carrying full tray energy balances and finds the difference immaterial for optimization [16].

Ideal equilibrium stages, ten of them. Stage efficiency falls with vapor rate, with fouling and with foaming, and none of that is modeled, so a column that has been in service will not reach these compositions.

Constant liquid holdup and no vapor holdup. Holdup is 0.85 mol against a 28 mol charge, which is 3.0 %, and it is fixed. Holdup is the refinement that separates a rigorous batch model from a shortcut one and it is not a small effect [11]. With no total-reflux startup, raising it from 3.0 % to 10.7 % of the charge lengthens the fore cut from 6.2 min to 21.2 min and costs 2.75 mol of product.

Boil-up follows the duty algebraically. Equation (3.6) has no thermal capacitance, the heat of vaporization is fixed, and the column pressure does not move with the vapor rate. A real reboiler has a thermal time constant of several minutes, so a duty change here appears instantly and the profile answers it within a tray time constant.

The reflux ratio is a continuous variable settable to anything from 0 to 40. The reference apparatus sets it with a swinging bucket that can only be moved every 5 min, which is why its optimizer returns a staircase rather than a smooth ramp. A column with a real bucket cannot execute a constant-composition policy at all, and that is much of why constant reflux is the classical policy.

3.4 Dynamic Analysis

3.4.1 Decline of the Top Composition and the Process Gain

The behavior that organizes the application follows from combining (3.7) with (3.4). The column multiplies the pot ratio by a factor that never changes, and the pot ratio only ever falls, so the top composition falls with it. Over ninety minutes at a constant reflux ratio of 4 the pot goes from 0.5787 to 0.2337, a ratio falling from 1.374 to 0.305 by a factor of 4.5, and the distillate follows from 0.9996 to 0.9599. No disturbance entered and no input changed.

Two consequences follow for the controller. The process gain $\partial x_D/\partial R$ decays toward zero over the run, because raising the reflux ratio buys composition by giving up distillate rate, and as the pot gets leaner the same increment buys less and costs more. And the useful actuator range collapses at the same time, since $D = V/(R+1)$ in (3.5) means the distillate rate falls hyperbolically with R. At $R = 3$ the column makes 0.177 mol/min and at $R = 39$ it makes 0.018 mol/min, so the last three quarters of the actuator's travel buy almost no product at all.

3.4.2 Reflux Demand Under Constant-Composition Control

A constant-composition controller must therefore supply an ever larger reflux ratio to hold a fixed distillate composition, and the demand it faces is set by the plant rather than by the tuning. Figure 3.2 plots the reflux ratio each of the three policies runs against time on the same charge.

The two curves that meet the same specification are separated by a factor of almost four in the reflux ratio at the end of the batch. The constant-composition demand rises without bound because the plant is running out of light component and the loop is being asked for a composition the pot can no longer support at any useful rate. That is not a tuning fault and no tuning removes it.

3.4.3 Effect of the Cumulative Cut Rule on Optimality

The optimization exists because of (3.9) and would not exist without it. Suppose the receiver were switched to slop the moment the instantaneous composition fell below x_{spec}. Then every mole

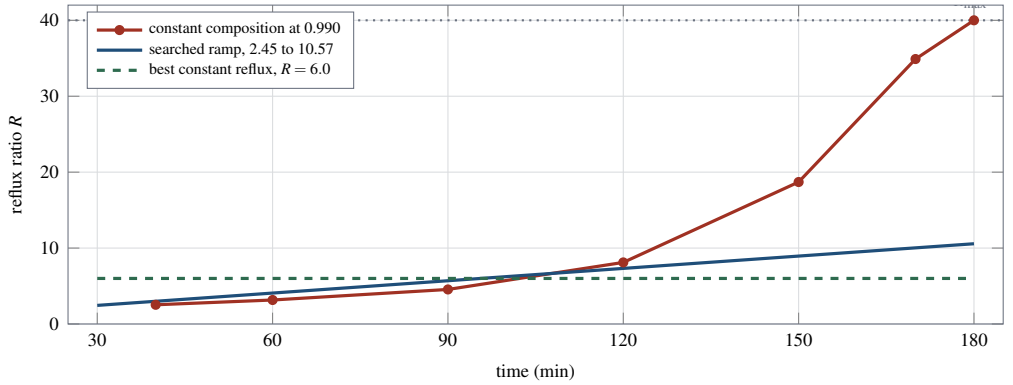


Figure 3.2 Reflux ratio against time for the three policies on the same 28 mol charge and the same 180 min budget: constant composition (solid with dots), the searched ramp (solid) and the best constant reflux ratio (dashed). The constant-composition points are measured on the running application at the shipped tuning. Holding 0.990 costs 2.53 at 40 min and the full 40 by 180 min, while the searched ramp that collects more product never asks for more than 10.57.

collected would be at or above specification, no reflux would be spent making anything purer than the customer pays for, and the best possible policy would be to hold x_D exactly at x_{spec} for as long as possible. That policy is constant-composition control at $x_{D,\text{SP}} = x_{\text{spec}}$, optimal by construction, and there would be nothing left to search for.

The cumulative rule creates a currency instead. Define the purity credit as the integral in Equation (3.13), taken over the main cut.

$$C(t) = \int_{t_{\text{main}}}^{t} D\,\left(x_D - x_{\text{spec}}\right)\,\mathrm{d}\tau \tag{3.13}$$

The receiver is on specification exactly while $C(t) \geq 0$, since $\bar{x}_P - x_{\text{spec}} = C/n_P$, and the tail cut of (3.9) fires when the credit reaches zero. In the terms of Section 1.4.4, the receiver purity is an endpoint constraint and the instantaneous composition carries no path constraint. A policy is therefore free to run the instantaneous composition below specification for as long as earlier material pays for it, and the best policy spends the credit down to exactly zero at the end of the batch. That is what the searched ramp does, and it is why it beats a loop that holds x_D near 0.990 for most of the collection.

🕮 **Pause.** Try Exercise 3.3 before reading on.

3.5 Controller Design

3.5.1 Discrete PID Algorithm

RC-401 runs the velocity-form algorithm of Section 1.6.2, computing an increment to the reflux ratio each scan rather than the ratio itself. Its derivative term acts on the analyzer reading and not on the error, so a composition setpoint change produces no derivative kick at the reflux divider, as shown in Equations (3.14) and (3.15).

$$e_k = x_{D,\text{SP}} - \hat{x}_{D,k} \tag{3.14}$$

$$\begin{aligned}
\Delta R_k^P &= K_c\,(e_k - e_{k-1}) \\
\Delta R_k^I &= \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \\
\Delta R_k^D &= -K_c\,\tau_D\,\frac{\hat{x}_k^f - 2\hat{x}_{k-1}^f + \hat{x}_{k-2}^f}{\Delta t_c}
\end{aligned} \tag{3.15}$$

Table 3.3 The Three Reflux Policies, Expressed Through the Same Controller

Policy	R_{ff}	K_c	τ_I
Constant reflux	—	output discarded, PID tracks	
Constant composition	0	600	4 min
Optimized ramp	Eq. (3.17)	0	∞

The controller is reverse acting because the process gain is positive. More reflux gives a richer distillate, so a measurement below setpoint, which is a positive error in (3.14), calls for a larger reflux ratio. The derivative increment is a second difference on a filtered measurement, with $\hat{x}_k^f = \alpha_f \hat{x}_{k-1}^f + (1-\alpha_f)\hat{x}_{D,k}$ and $\alpha_f = \exp(-\Delta t_c/\tau_f)$, where Δt_c is the controller period and the filter constant τ_f defaults to 0.2 min. The three increments are combined and the demand is clipped to the divider span [17], as shown in Equation (3.16).

$$\begin{aligned} s_k &= s_{k-1} + \Delta R_k^P + \Delta R_k^D \\ c_k &= R_{\mathrm{ff}}(t_k) + s_k + I_{k-1} \\ R_k &= \min(R_{\max},\ \max(0,\ c_k)) \\ I_k &= I_{k-1} + \Delta R_k^I + \frac{\Delta t_c}{\tau_I}\,(R_k - c_k) \end{aligned} \tag{3.16}$$

These are Equations (1.14) to (1.17) written for this loop. The running sum s_k is S_k, the demanded ratio c_k is v_k, the trajectory R_{ff} is the output bias u_b, and the back-calculation that corrects the integral alone has the tracking time $T_t = \tau_I$. Under the trajectory policy τ_I is infinite, so the integral increment and the back-calculation are both absent, and I carries only an offset left by a transfer, which decays with a 2 min time constant. The output is clamped to $0 \le R_k \le R_{\max}$ and carries no rate limit, because a reflux divider changes its split in one step rather than stroking through the intermediate positions. Execution is every 0.05 min, independent of the 0.01 min physics step. Transfers between MANUAL and AUTO are bumpless in the way Section 1.6.2 describes. The tuning is bounded at $0 \le K_c \le 1500$, $0.2\,\mathrm{min} \le \tau_I \le 60\,\mathrm{min}$ and $0 \le \tau_D \le 5\,\mathrm{min}$, with K_c in units of reflux ratio per unit of mole fraction.

Before any of that runs, the startup latch overrides it. For $t < t_{\mathrm{start}}$ the plant is given $R = \infty$ rather than $R_{\max}$, so that (3.5) returns $D = 0$ exactly and the pot does not drain during the startup at all. Behind the latch the loop is open, so in AUTO the controller is not updated there but tracks, as in MANUAL, the ratio its policy will open on, which is R_{set} under constant reflux, R_0 under the trajectory and the operator's reflux station under constant composition. Collection therefore opens bumplessly from that ratio, and the error standing at the opening enters through the integral, or under the trajectory through the transfer offset that decays with the 2 min time constant, rather than as a step. Section 3.6 measures the opening on the shipped analyzer.

3.5.2 The Three Reflux Policies

The three policies differ only in what R_{ff} and τ_I are, as listed in Table 3.3. The plant does not change between them and neither does the actuator.

Under constant reflux the controller output is discarded and the plant is given the operator's constant, while the PID keeps tracking so that selecting another policy later starts from the reflux ratio the column already has. Under constant composition R_{ff} is zero and the loop is an ordinary PI controller on the analyzer. Under the optimized policy R_{ff} is the searched trajectory of Equation (3.17).

$$R_{\mathrm{ff}}(t) = \mathrm{clamp}\left(R_0 + (R_1 - R_0)\,\frac{t - t_{\mathrm{start}}}{t_{\mathrm{batch}} - t_{\mathrm{start}}},\ 0,\ R_{\max}\right) \tag{3.17}$$

where the fraction is itself clamped to $[0, 1]$ so the ramp is held flat outside the collection window. Two details of that arrangement are structural rather than incidental. The trajectory enters as the controller

bias rather than being added on afterward, as the kettle's feedforward does in Section 1.4.3, which puts it inside the clamp and inside the manual tracker, and a gain of zero then gives the pure open-loop trajectory exactly. And reset is switched off, because an integral term would drive the composition onto its setpoint within a few minutes and erase the trajectory, reproducing constant-composition control. Any gain above zero is a proportional trim around the trajectory, and Section 3.6 measures what that trim is worth.

3.5.3 Grid Search for the Reflux Ramp

Selecting the searched ramp runs a forward-simulation grid search at every reset, before the first frame is drawn. Each candidate is integrated with explicit Euler at 0.02 min, a fifth of the tray time constant and a twenty-fifth of the analyzer's, on the same derivative function the plant uses, and reproduces the engine's Runge-Kutta run to within 0.01 mol of product. The search routes the distillate with the same lagged and biased analyzer the operator has, without its noise, because the lag delays the tail cut and changes the collected product.

The structure of the search follows from the objective not being smooth. The tail cut is a discrete event, so a ramp that just avoids taking one collects every mole it made while a ramp a tenth of a unit lower throws the last half hour away. The search therefore runs a diagonal sweep of every constant reflux ratio from 0.5 to 16 in steps of 0.5, which guarantees the result is never worse than the best constant policy on that grid, then a coarse two-dimensional sweep with R_0 from 0.5 to 8 and R_1 from 1.5 to 26 that skips every falling ramp, because the optimal reflux ratio does not decrease over a run in which the separation becomes monotonically harder, then a refinement around the best three candidates and a finer one around the best two, rather than around the single best point. Refining around one candidate instead costs nothing on the shipped case and 0.54 mol at a 0.995 specification, where the best flat ramp and the best rising ramp sit far apart and only one of them is near the coarse grid's winner. The whole search is deterministic, so a reset reproduces the run exactly.

On the shipped charge the search returns $R_0 = 2.45$ and $R_1 = 10.57$. The reference's own optimizer is a nonlinear program that re-picks the reflux ratio every five minutes because its apparatus can only be changed that often, and the trajectory it finds starts low and increases in a nominally linear pattern [16]. A two-parameter straight ramp is a continuous form of that trajectory.

3.5.4 Relative Volatility and Stage Count

One departure from the reference is deliberate and is documented here. The application runs at $\alpha = 2.0$ where the methanol and ethanol chemistry is 1.725. The reason is not the thermodynamics, which the constant-volatility form reproduces to an RMS of 4.2×10^{-4} mole fraction, but the apparatus. The reference column has 38 Oldershaw trays at a fitted Murphree efficiency between 0.37 and 0.45, which is roughly 14 to 17 theoretical stages against the eleven modeled here. Eleven stages at 1.725 give a total-reflux separation factor of $1.725^{11} = 402$ where eleven stages at 2.0 give $2^{11} = 2048$, a factor of 5.1 between them, and the difference at finite reflux is far larger than that. At 1.725 the ten-tray column cannot run the reference's own policy at all, since a constant reflux ratio of 4 collects 1.38 mol before it has to take a tail cut. At 2.0 it reproduces the reference base case, collecting 12.73 mol at a cumulative 0.9909 over the reference's 30 plus 90 minute schedule against the 9.5 mol at 0.992 the reference model reports. Volatility and stage count trade against each other, and here the trade was made in favor of a state vector of 19. The value 1.725 is available as a parameter setting and scenario 13 measures what it costs, so the effect of the choice is measurable inside the application.

One arithmetic correction belongs with this. The quartic equilibrium fit printed in the reference carries a coefficient of -2.016 on the fourth-order term. With that coefficient the expression returns -0.250 at $x = 0.9$ and -0.815 at $x = 1$, which are not compositions. With -0.2016 the same quartic is the constant relative volatility form of (3.1) at $\alpha = 1.725$ to an RMS deviation of 4.2×10^{-4} and a maximum deviation of 8.7×10^{-4} over $0 \le x \le 1$, and it returns 0.9999 at $x = 1$. The printed coefficient is a typographical error of one decimal place, and the model file distributed with the

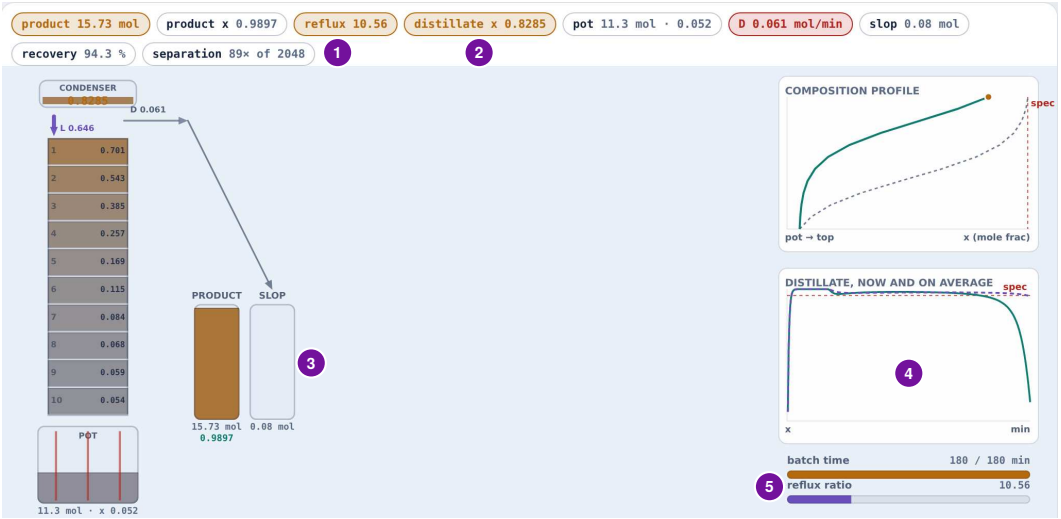


Figure 3.3 The searched reflux ramp at the end of its batch, scenario 6 at $t = 180$ min, cropped to the plant view. (1) The badges: the batch collects 15.73 mol at a cumulative composition of 0.9897, with 0.08 mol of slop and 94.3 % recovery, at 412 kJ/mol (Table 3.5), and the pot is drawn down to 11.3 mol at 0.052 mole fraction methanol. (2) The instantaneous distillate composition ends at 0.8285. (3) The product and slop receivers. (4) The distillate now and on average: the instantaneous trace falls away at the end while the average holds near the specification. (5) The batch-time and reflux-ratio bars.

reference does not use the quartic at all, computing equilibrium from component vapor pressures instead.

3.6 Results

The fifteen scenarios are listed in Table 3.4, ordered by concept rather than by difficulty. Each carries a starting mode, a reflux policy, a tuning, the parameters it overrides, the faults it fires and their timing, its own prediction, experiment, hints and explanation, and a completion rule, as Section 1.6.4 describes for every app. The Exercises at the end of the chapter use them.

All computation, the reflux search included, runs in the browser on the engine of Section 1.6.1, so a slow machine draws fewer frames and integrates the same batch. Every run exports at ten samples per simulated minute, carrying the true composition beside the analyzer reading, the cut phase, both mole balances and every parameter in force, which is enough to reconstruct the purity-credit integral offline or to audit a batch. Twenty-nine verification tests, twelve of them specific to this application, run against the application and every one passes, with both mole balances closing to better than 2×10^{-11} mol at the end of each batch of Table 3.5 against a 1×10^{-9} mol test tolerance. The receiver is graded on specification when its cumulative average is within 0.001 of x_{spec}, because the tail cut closes on the analyzer's running average, which lags a falling distillate, so a correctly run batch closes a few ten-thousandths below the specification. Figure 3.3 is the last frame of the optimized batch of scenario 6.

3.6.1 Comparison of the Three Policies

The central comparison is Table 3.5. All three policies run the same 28 mol charge at 0.59, the same 600 W duty, the same 30 min total-reflux startup and the same 180 min budget against the same 0.990 specification, and every figure is read from the application at the end of the batch. Figure 3.4 sets the product, the energy per mole and the final reflux ratio of each policy side by side.

Table 3.4 The Fifteen Scenarios and the Behavior Each One Shows

#	Scenario	What it shows
1	Total reflux, MANUAL	holdup filling at startup, Fenske
2	The settled column	the total-reflux equilibrium
3	Constant $R = 4$, 120 min	the reference base case
4	Constant $R = 3$ then a sweep	location of the constant-reflux optimum
5	Constant composition at 0.990	actuator saturation with the CV near setpoint
6	Searched reflux ramp	open-loop trajectory, purity credit
7	Hand-set ramp, 2 to 16	purity giveaway
8	Analyzer $\tau_A = 5$ min	lag in a loop and in a cut decision
9	Analyzer $+0.012$ at $t = 34$	bias that improves reported results
10	-35 % duty at $t = 60$	throughput against separation
11	Charge at $x_0 = 0.42$	open-loop policy on an off-design charge
12	Holdup 3.0 mol, no startup	purpose of the total-reflux startup
13	$\alpha = 1.725$	volatility against stage count
14	Replan at $\alpha = 1.85$	nominal predictive feedback
15	Graded challenge	control on a slow, noisy analyzer

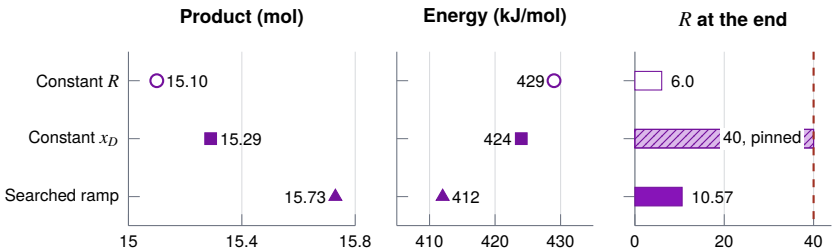


Figure 3.4 The three policies of Table 3.5 on the same charge, duty and 180 min budget: product collected and energy per mole of product (axes not starting at zero), and the reflux ratio at the end against its limit of 40 (dashed). Notice that the policy that collects most and spends least per mole ends at a reflux ratio of 10.57, while the composition loop finishes second with its actuator pinned.

Table 3.5 Terminal Results of a 180 min Batch Against a 0.990 Specification. The Constant Reflux Ratio Is the Best Flat Value Found by Sweeping in Steps of 0.05.

	Constant R $R = 6.0$	Constant x_D SP = 0.990	Searched ramp 2.45 to 10.57
Product (mol)	15.10	15.29	**15.73**
Receiver $\bar{x}_P$	0.9895	0.9899	0.9897
Recovery (%)	90.4	91.6	**94.3**
Energy (kJ/mol)	429	424	**412**
R at the end	6.0	**40, pinned**	10.57
Feedback	none	PI on $\hat{x}_D$	none

The optimized trajectory collects 15.73 mol against 15.29 mol for constant composition and 15.10 mol for the best constant reflux ratio, an improvement of 2.9 % over the loop and 4.2 % over the flat policy. It does so *open loop*, with a proportional gain of zero and a reflux ratio that never exceeds 10.57, while the loop that holds its setpoint within 0.0012 for eighty minutes finishes second with its actuator hard against a limit of 40. The reference reports a 14 % improvement from optimization on its own apparatus, and the difference between 14 % and 4.2 % is the baseline rather than the method, since the reference compares against a constant reflux ratio of 4 on a 90 minute collection and the

comparison here is against the best constant reflux ratio on a budget long enough for the trade to develop. Against a constant reflux ratio of 4 on the same 180 min budget the searched ramp collects 15.73 mol against 13.16 mol, an improvement of 19.6 %.

3.6.2 Sensitivity of the Constant-Reflux Optimum

Sweeping the constant reflux ratio gives 2.17 mol at $R = 2$, 9.16 mol at 3, 13.16 mol at 4, 14.48 mol at 5, 15.10 mol at 6, 14.14 mol at 6.5 and 11.79 mol at 8. Two different failures bound the curve. Below $R = 6$ the tail cut fires before the batch ends and the remaining production is routed to slop, with 13.99 mol of slop at $R = 3$ and a tail cut at 81.8 min. Above $R = 6$ the batch ends with the receiver above specification, and the surplus purity has no value against a 0.990 specification. The best constant reflux ratio is therefore the one whose receiver reaches specification at the moment the batch ends, so the answer moves when the budget moves and is 3.9 on the reference's own 120 min schedule rather than 6, collecting 12.95 mol against 12.73 mol at 4.

Scenario 7 quantifies the same effect on a trajectory. A plausible hand-set ramp of 2 rising to 16 collects 13.15 mol at 0.9957, giving away 16.4 % of the production for 0.0057 of purity above the specification. A hand-set 2.5 rising to 10.5, within a tenth of the searched answer, collects 15.74 mol, 0.006 mol more than the search, because that point lies between the nodes of the search grid. The value of the search is that the optimum moves whenever the charge, the duty, the specification or the batch time moves.

3.6.3 Actuator Saturation Under Constant-Composition Control

The constant-composition policy holds the composition close to setpoint for most of the collection and fails in the actuator. Collection opens on the operator's station at $R = 2$, as described in Section 3.5.1, the loop trims the ratio to 0.88 and back up to 4.66 within three minutes, and the distillate dips only to 0.9896 at 32.5 min. From 40 to 120 min the distillate stays within 0.0012 of 0.990. The reflux ratio the loop demands runs 2.53 at 40 min, 3.16 at 60, 4.55 at 90, 8.11 at 120, 18.7 at 150 and 34.9 at 170, and is hard against 40 from 174.4 min, as plotted in Figure 3.2. The distillate rate falls with it from 0.200 mol/min at 40 min to 0.017 mol/min, a factor of about twelve, and falls below the $D_{\min}$ threshold at 169.5 min.

The demand is a ramp that steepens, and a PI loop follows a ramp with a lag error of about $(\mathrm{d}R/\mathrm{d}t)\,\tau_I/K_c$ in mole fraction. At 150 min the demand rises by 0.58 per minute, which predicts 0.0038 against the 0.0034 measured, and the distillate falls behind the setpoint to 0.9866 at 150 min and 0.9834 at 170 min, and the running average closes the receiver with a tail cut at 162.7 min, leaving 0.34 mol in slop. The policy still performs well. It beats the best constant reflux ratio by 1.3 % because it never over-purifies, and it needs no trajectory and no model. It requires an analyzer and a tuning fast enough to follow a reflux demand that rises throughout the batch. At $K_c = 600$ the loop keeps up until the last hour. At $K_c = 300$ the lag error doubles, the loop takes a tail cut at 119.7 min and collects 13.33 mol, a 12.8 % loss from a tuning change that is barely visible on the controlled variable.

3.6.4 Purity Credit Trajectory of the Optimized Ramp

The searched policy starts at 2.45, near the minimum reflux the column needs, where the distillate rate is high and the pot is still rich enough that the material leaves at 0.9923 at 40 min. That interval builds the purity credit of (3.13), and the remainder of the run spends it. The instantaneous composition falls to 0.9870 at 150 min, 0.9490 at 170 and 0.8285 at the end, while the cumulative average in the receiver is still 0.9897. A policy that never let the instantaneous composition go below specification could not do this and would collect less. The gap between the instantaneous and cumulative traces is the mechanism, and the two compositions are displayed together for that reason.

Table 3.6 Constant-Composition Control at $K_c = 700$ and $\tau_I = 2$ min Against Analyzer Lag, With Noise at Zero. Every Extra Minute of Lag Costs Product, and at 8 min the Dip After the Opening Forces a Tail Cut at 39.9 min.

τ_A (min)	0.5	1	2	3	5	8
Product (mol)	15.51	15.46	15.16	14.61	12.70	1.66
Receiver $\bar{x}_P$	0.9901	0.9902	0.9896	0.9880	0.9834	0.9856

> **TRY IT IN THE APP** **6 · The Optimized Reflux Ramp**
>
> Load the scenario and press Run. Watch the inset of the distillate, now and on average. The instantaneous composition falls to 0.9490 at 170 min and 0.8285 at the end, while the receiver's cumulative average stays at 0.9897. The early, purer product pays for the late, weaker product.

3.6.5 Analyzer Lag and Bias

The analyzer serves two functions and only one of them is feedback. Its effect on the loop is visible. At $\tau_A = 5$ min with a tight tuning of $K_c = 700$ and $\tau_I = 2$ min and the scenario's 0.002 of analyzer noise, the loop drives the ratio from the station's 4 to zero within two minutes of the opening, the distillate falls to 0.884 before the analyzer has seen it, and from then on the ratio hunts from zero to between 25 and 40 in every half hour. Detuning to $K_c = 300$ with $\tau_I = 4$ min stops the hunt, the ratio climbing smoothly between 1.7 and 7.3 from 60 to 110 min. Without the noise the tight tuning still swings by 28.8 over the last seventy minutes. Its effect on the cut is not visible at all, because the tail cut of (3.9) is judged on the same instrument, and while the true composition is falling the analyzer reads high. Table 3.6 measures both failures on noise-free runs at the tight tuning.

At 5 min of lag in the shipped scenario the tail cut never fires at all and the receiver finishes with 13.09 mol at 0.9820, off specification, while the running average computed from the analyzer reads 0.9937. Detuning at the same lag removes the hunt and not the bias. It returns 15.36 mol at 0.9888, still off specification after a tail cut at 159.1 min judged on the same lagging reading. The lag itself is untouched and no tuning touches it.

> **DID YOU KNOW?** **Eight Minutes of Lag Lose the Batch**
>
> From the simulator: under constant-composition control at the tight tuning, an analyzer lag of 0.5 min collects 15.51 mol, 3 min collects 14.61 mol, and 8 min only 1.66 mol (Table 3.6). The cut that ends the collection is judged on the same late instrument.

A bias is more damaging than a lag because it improves the reported result. Drifting the analyzer 0.012 high four minutes into the collection leaves the receiver at 0.9783, which is 0.0117 low and almost exactly the bias. No trend drawn from the analyzer reveals it, because each carries the bias. The controller holds the biased reading within 0.007 of setpoint until 120 min, and the running average computed from the analyzer still reads 0.9907 at 150 min against a true 0.9790. The batch also reports better results than the unbiased one, at 16.33 mol against 15.29 mol and 96.7 % recovery against 91.6 %. Both improvements are fictitious. The extra 1.04 mol is material the cut rule should have sent to slop and the whole drum is now rework. Two measurements reveal the fault, since the true distillate composition falls to 0.8935 by the end and production stops at 151.2 min, eighteen minutes earlier than the unbiased run, with the reflux ratio at its limit of 40 from 154.9 min. Feedback cannot correct a sensor bias, because the controller works to make the biased reading true, so every defense is external to the loop. Both mole balances are recorded at every sample for that reason, and they stay below 1×10^{-9} mol of error even while the analyzer reading is biased.

Table 3.7 A Proportional Trim Around the Wrong Trajectory, on a 0.42 Charge. Re-Optimizing the Trajectory Recovers 10.69 mol on the Same Plant.

K_c	0	100	300	600	1000
Product (mol)	1.26	1.37	1.98	3.10	4.74
Receiver $\bar{x}_P$	0.9882	0.9888	0.9893	0.9896	0.9897
Slop (mol)	14.55	10.53	9.45	8.07	6.24
Recovery (%)	10.6	11.5	16.6	26.1	39.9

TRY IT IN THE APP 9 · A Biased Analyzer Quietly Makes Off-Specification Product

Load the scenario and press Run. Four minutes into the collection the analyzer drifts 0.012 high. The batch then reports 16.33 mol and 96.7 % recovery, better than the unbiased run, while the true receiver ends at 0.9783, off specification. Watch the reflux ratio reach its limit of 40 earlier than in scenario 5.

3.6.6 The Graded Challenge

Scenario 15 grades constant-composition control on an analyzer with a 1 min lag and added noise. It asks for the true distillate within 0.008 of 0.990 from 86 to 130 min and a receiver on specification within 0.001, with the model parameters and disturbances as given. With the shipped tuning, $K_c = 40$ and $\tau_I = 30$ min, the worst error from 86 to 130 min is 0.1198, about fifteen times the band. The receiver closes at 0.9893 with 3.98 mol and 13.31 mol in slop, so receiver quality passes and tracking fails. With the suggested tuning, $K_c = 500$, $\tau_I = 3$ min, $\tau_D = 0.4$ min and $\tau_f = 0.5$ min, the worst error is 0.0047 and the receiver closes at 0.9897 with 15.61 mol and no slop, about 0.0007 above its true-composition minimum of 0.989. PI at the same gain and reset also passes with a worst error of 0.0039, so the derivative is optional. Removing its filter raises the reflux-ratio travel from 997 to 4810, while a gain of 700 with a two-minute reset fails tracking at 0.0328 despite an on-specification receiver. Reset with the suggested settings reproduces the batch.

3.6.7 Sensitivity to an Off-Design Charge

An open-loop trajectory carries no information about the plant it is running on, which is the standing objection to trajectory optimization and the reason the reference spends most of its length on parameter estimation before it optimizes anything. Scenario 11 measures the objection in moles. The ramp searched for a 0.59 charge is applied to a 0.42 charge. Eleven stages multiply a ratio, and a 0.42 charge has a ratio of 0.72 where 0.59 has 1.44, so the top of the column starts a factor of two lower and 2.45 is no longer enough reflux to make specification. The batch collects 1.26 mol at 0.9882, sends 14.55 mol to slop, takes its tail cut at 36.5 min and consumes 5143 kJ/mol against 412 kJ/mol on the design charge. Re-searched against the charge it actually has, the same policy structure returns 3.40 rising to 17.83 and collects 10.69 mol at 0.9898 with 0.004 mol of slop.

Because the trajectory enters as the bias in (3.16), a proportional trim can be added around it without changing anything else, and Table 3.7 measures how much of the loss that recovers.

At $K_c = 1000$ the proportional term drives the reflux ratio to its limit of 40, holds the true distillate at 0.9562 at the end instead of 0.0440, and recovers 4.74 mol. Re-optimizing recovers 10.69 mol. More gain recovers more, because the tail cut moves from 36.5 min open loop to 37.6 min, 42.1 min, 50.8 min and 65.0 min at $K_c = 100$, 300, 600 and 1000, while part of the fall in slop comes only from the reflux ratio later climbing to 40 and the column ceasing to produce. Feedback around a wrong trajectory recovers part of the loss, and a correct trajectory recovers nearly all of it, which is why a batch operation measures its charge before executing a recipe. On the correct charge the same trim is close to free, since $K_c = 100$ collects 15.78 mol at 0.9898 against 15.73 mol open loop while pushing the reflux ratio to 20.6 by the end and putting 0.19 mol in slop against 0.08 mol, and $K_c = 300$ collects 15.73 mol. The trim protects against an off-design charge at a measurable cost in reflux and slop on the design charge.

3.6.8 Effect of Duty, Holdup and Volatility

Three parameters change the outcome in ways that resemble control faults. Losing 35 % of the reboiler duty at 60 min with the reflux ratio held at 6 drops the vapor rate from 0.707 mol/min to 0.460 mol/min and drops L and D with it, so L/D is still 6 and the column separates exactly as well as before on less throughput. Production falls from 15.10 mol to 10.91 mol and recovery from 90.4 % to 65.9 %. The loss is throughput, because the same batch time now passes 35 % less vapor, and the diagnostic is that the product comes out *purer* at 0.9970 against 0.9895. A batch column that ships material well above specification with no setpoint change is running a reflux ratio too high for its remaining throughput.

🕮 **Pause.** Try Exercise 3.2 before reading on.

The total-reflux startup exists to fill the holdup. With no startup at all at $R = 4$, the fore cut runs 6.2 min at the shipped 0.85 mol of holdup and 21.2 min at 3.00 mol, and the product falls from 12.15 mol to 9.40 mol, with $M_B(0)$ following the holdup whether it is set before Run or carried through Reset. Switching 30 min of total reflux back on erases the penalty completely, at 13.17 mol with 3.00 mol of holdup against 13.16 mol with 0.85 mol, because nothing is being thrown away while the holdup fills. That also sets how long a startup should be. At $R = 6$ on the shipped 0.85 mol column, no startup collects 14.59 mol with 3.60 mol of slop, ten minutes collects 15.10 mol with 2.08 mol, thirty minutes collects 15.10 mol with 0.06 mol, and beyond that the startup only consumes batch time, dropping to 13.64 mol at forty-five minutes and 12.12 mol at sixty. The best startup is the shortest one that leaves nothing in the slop drum.

Relative volatility decides whether the separation is possible at all. At $\alpha = 1.725$ on the same ten trays, a constant reflux ratio of 4 collects 1.38 mol and sends 19.84 mol to slop against 13.16 mol collected at 2.0, so a 14 % reduction in relative volatility costs 90 % of the production. The column remains capable but under-refluxed, because $R = 8$ collects 10.89 mol at 0.9899 and the best constant reflux ratio in steps of 0.1, 8.4, collects 11.29 mol at 0.9900. Eleven stages still make 99 mol% methanol at the reference chemistry, but only by sacrificing most of the distillate rate, and even at its best the column produces 25 % less than at $\alpha = 2.0$.

3.6.9 Replanning With an Imperfect Nominal Separator

The MPC option of Section 1.6.3 recalculates reflux every half minute, rather than preserving a ramp searched at reset. Its collection horizon ends at the current batch clock. It uses at most 150 uniformly spaced prediction intervals and twelve graduated move blocks, so the prediction spacing may exceed the decision period. Before takeoff the window spans startup end to batch end, and the first optimization at $t = 0$ describes hypothetical collection from the current estimate. Physical total reflux continues unchanged. This is a local economic MPC [18], not the reference's full nonlinear batch optimization.

The reduced model eliminates the ten tray and condenser transients. For a trial steady overhead x_D^s, its operating line marches downward:

$$y_1 = x_D^s, \quad x_i = \frac{y_i}{2 - y_i}, \quad y_{i+1} = \frac{R x_i + x_D^s}{R + 1}, \tag{3.18}$$

with the shooting condition $y_{11} = 2x_B/(1 + x_B)$. Thus relative volatility stays at two even when the simulated chemistry changes. The fixed boil-up is $V_0 = 0.707213$ mol/min. With $F = x_D^s + b$, confined to physical mole fractions, the four-state nominal balance is

$$\dot{M} = -D_0, \qquad \dot{A} = -D_0 F, \qquad D_0 = \frac{V_0}{1 + R}, \tag{3.19}$$

$$\dot{a} = (F - a)/0.5, \qquad \dot{z} = g D_0 (a - x_{\text{spec}}). \tag{3.20}$$

Here M and A are estimated pot moles and light moles, a is the analyzer, and z is the receiver's measured purity surplus. The current main-cut flag g is held across the forecast. The model is

Table 3.8 Scenario 14: Model Information and Reflux Movement

Policy	Product (mol)	True purity	Travel
Composition PI, 600/4	13.8056	.989937	22.46
Fixed nominal ramp	1.6814	.989013	8.11
Fixed ramp plus P600	11.1138	.989893	23.37
Run-model searched ramp	14.3393	.989882	5.81
MPC	14.2914	.989982	173.47

linearized once at each decision, then its exact affine hold is extrapolated over the shrinking collection window.

The inventory estimator starts from 27.15 nominal pot moles. Its initial composition uses the first analyzer sample, with a nominal Fenske inversion for an already rich settled reading. Nominal distillate flow and analyzer history update its material accounting. The receiver state is initialized from the same operator ledger that the original cut logic uses. That ledger already weights analyzer readings by the app's calculated flow. It is not a new pot or tray instrument and does not update the predictor's fixed boil-up. A filtered analyzer slope compensates a nominal half-minute lag. A gain-.1 residual filter, limited to $\pm.05$ mole fraction, adjusts b. Disturbance correction follows the motivation of offset estimation [19], without identifying volatility, actual charge size, heater performance or calibration. No hidden states, run coefficients or later events enter this predictor.

The QP rewards linearized distillate flow through the positive reflux price $hV_0/(1+\bar{R})^2$. Its other costs are $100h(a-x_{\mathrm{sp}})^2$, terminal weight 100, reflux weight $10^{-6}h$ and block-move weight $.002w$, with default $w=1$. Reflux between zero and the present drive limit, with unit changes between blocks, is hard bounded. Predicted receiver surplus and pot inventory are soft lower constraints. The analyzer floor is $x_{\mathrm{spec}}-.02$ during the main cut and $x_{\mathrm{spec}}+.0005$ otherwise, allowing measured receiver credit to fund some instantaneous dilution. Quadratic slack prices are 10^5, 10^7 and 10^4 for analyzer, surplus and pot, with linear prices 10, 100 and 100. These prices optimize a flow proxy under a local approximation. They do not certify maximum terminal product, true purity or recursive feasibility.

The existing composition PID tracks the applied finite reflux request. Any failed, nonfinite, malformed or invalid plan holds that request on the first scan, then the composition PID resumes until a valid plan returns. Historian marks and notes identify failure and recovery. Total reflux, measured cumulative cuts and the clock or pot stop retain their original authority. A tail cut remains irreversible. Ten appended CSV fields report acceptance, iterations, analyzer slack, fallback, full planning time, horizon, predicted end analyzer and receiver surplus, bias and estimated pot inventory. The badge presents the end analyzer value, not final receiver purity.

Independent operating-line reconstruction gives a maximum 6.53×10^{-11} residual against the full nominal steady balances at twelve pot/reflux points. Implicit differentiation agrees with the implemented Jacobian within 5.71×10^{-8}. A closed scalar optimization using the actual four-state linearization and its economic reflux price has zero first-move error. The reduction nevertheless supplies immediate overhead feedthrough: at $x_B=.4$, changing R from four to five changes its static overhead by .003778, while the full condenser has zero initial direct response to reflux. Tray storage and analyzer dynamics matter when receiver credit is small. A two-minute prototype allowed an early tail cut between decisions, motivating the half-minute period before the final comparison settings were frozen.

Scenario 14 puts actual volatility at 1.85 while MPC retains two. Table 3.8 compares complete batches from the same seed. The fixed nominal ramp uses 2.45 to 10.5667. Adding the original proportional trim at gain 600 recovers part of its loss, but the feedback MPC collects more. The existing Searched option uses the run's actual parameters, so its open-loop row has more model information and slightly exceeds MPC. It must not be confused with a frozen nominal ramp.

The 55 measured rows include six PI tunings, four constant ratios, two move prices, halved physics

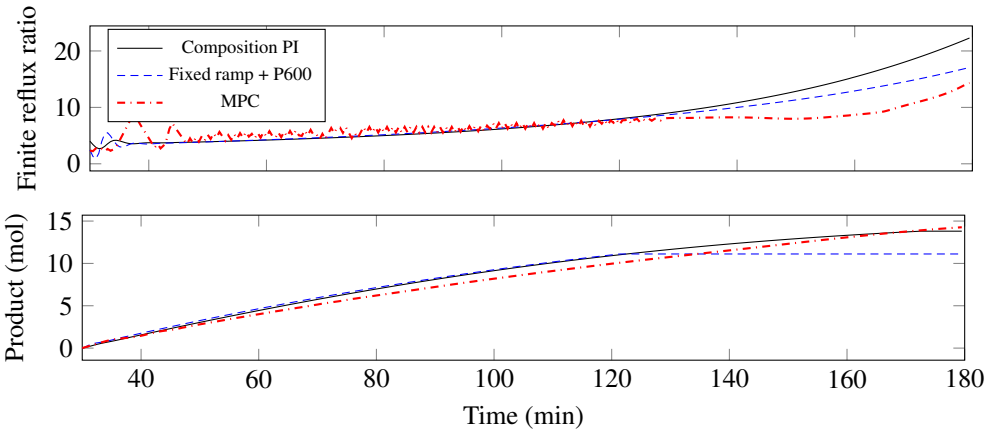


Figure 3.5 Matched volatility-1.85 batches after total reflux: composition PI (solid), fixed ramp with proportional trim (dashed) and MPC (dash-dotted). MPC improves recovery from the fixed nominal policy but requires more reflux movement. Travel excludes startup's infinite reflux and the undefined transition to a finite ratio. All five table rows spend 6480 kJ and meet the app's original receiver tolerance of .001.

steps and the original disturbance cases. The best PI in this small grid collects 13.9663 mol at .990060 with 26.69 of travel. Half-step MPC gives 14.2957 mol at .989990, retaining the comparison order. The nominal searched open-loop case still beats MPC, 15.7349 against 15.4535 mol. At volatility 1.725 the frozen MPC closes its product cut after only .9029 mol at .987859, below the receiver tolerance, while the composition PI collects 9.8613 mol. Larger holdup gives MPC 2.8451 mol against 14.3349 for PI. A positive .012 analyzer offset in scenario 14 produces .978214 true purity under MPC. Current analyzer correction cannot repair a biased reference.

On the unchanged suggested challenge, MPC collects 15.4947 mol at .989950 but fails its tracking target: maximum true error from 86 to 130 min is .01913 against .00472 for the suggested PID. Finite collection travel is 146.71 against 1364.45. The unmodified PID passes, and MPC's smaller movement does not substitute for the challenge objective. Chrome measurements across six viewport/theme captures and 300 plans per capture give .572–.627 ms per full optimization, .114–.125 percent of a core at unit playback and 6.87–7.53 percent at the highest 60-fold speed. This excludes observation, integration and rendering. The reduced horizon is computationally affordable, while its loss cases remain part of the laboratory.

3.7 Discussion

Separation is where a chemical plant spends its energy, and distillation dominates that spending, which is why separations are repeatedly identified as one of the largest available targets for reducing industrial energy use [20]. The dividing-wall column of Volume 13, Chapter 8 meets the same question on a continuous column. A batch column makes the point sharply because the reboiler spends its full duty whether or not anything is leaving the top. The energy per mole in Table 3.5 moves from 429 kJ/mol at the best constant reflux ratio to 412 kJ/mol on the searched ramp, a reduction of 4.0 % bought with no equipment and no extra instrument, and the same number reaches 5143 kJ/mol on a batch whose trajectory was written for the wrong charge.

DID YOU KNOW? A Fifth of the Energy Used in Plants

A 2005 report for the United States Department of Energy put industrial separations at about 4500 trillion Btu a year, about 22 % of all in-plant energy use in the country. Distillation alone took 49 % of the energy spent on separations [21].

The second cost is purity above specification. Scenario 7 ships 0.9957 material against a 0.990 specification and loses 16.4 % of the production doing it, and neither the application nor a plant annunciates over-purity, because no loop is asking for it. The same arithmetic drives solvent recovery, where most batch distillation capacity in fine chemicals and pharmaceuticals sits, and where solvent is the great majority of the mass a batch process handles, so recovering it rather than discarding it is the largest single decision in the material footprint of the route [5]. Measuring the distance between what was shipped and what was specified, and reporting it beside production, quantifies a loss that no control loop annunciates.

3.8 Technology in Practice

What Measures It

Every decision of this chapter's recipe rests on one composition reading. On a production column the reading may come from a process chromatograph, since the monitoring and control of distillation columns is a very common industrial use of chromatographs [22], as on the continuous column of Volume 5, Chapter 3. The laboratory column behind this chapter read refractive index on samples instead. AT-401 stands for either as a first-order lag of 0.5 min with a settable bias and noise, and the same reading closes the loop and switches the receiver.

TECHNOLOGY Sensor: Refractive Index of the Distillate

The column behind this chapter measured its distillate composition through the refractive index of the solution [16]. That index is the speed of light in a vacuum divided by its speed in the liquid, and it sets how far a beam bends on entering it [22]. A laboratory refractometer holds the sample at a controlled temperature and needs about twenty minutes to warm up [23]. A reading becomes a composition only through a calibration curve built from standards of known composition [24]. Instruments on industrial batch columns may report values with drift, noise or outliers [16]. AT-401 of this chapter reports continuously through its lag, and a bias set in a scenario shows on no trend drawn from it.

What Moves It

The only manipulated variable is the split of the condensate. The continuous column of Volume 5, Chapter 3 sets its reflux through a control valve, while a small batch column can split the condensate by time, so that the ratio of two intervals is the reflux ratio. RC-401 applies R as an exact split of V into L and D, with no travel time and no switching cycle, so the staircase of the reference's swinging bucket appears in this chapter only as a smooth ramp.

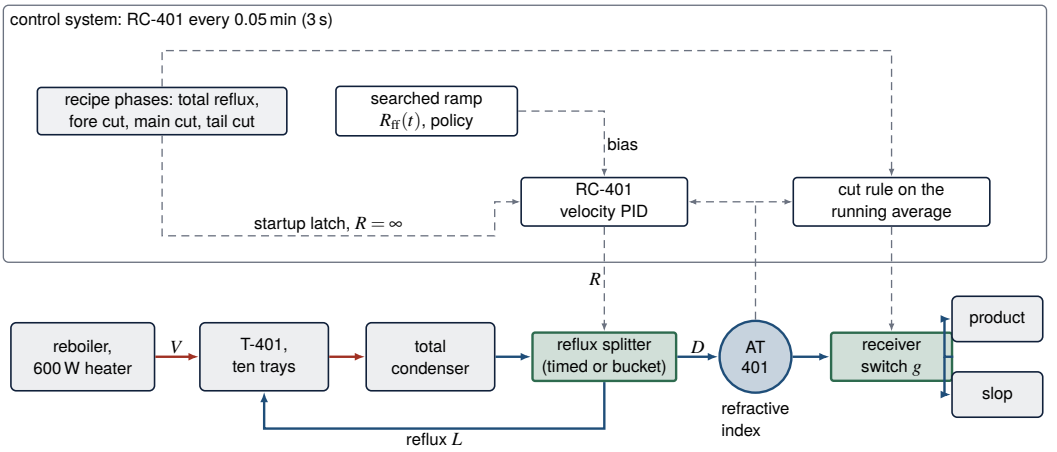


Figure 3.6 The chapter's loop on plant hardware. The analyzer AT-401, a refractive-index measurement on the reference column, reads the distillate, and the same reading drives RC-401, which sets the reflux splitter, and the cut rule, which moves the receiver switch between product and slop. The recipe phases latch total reflux during the startup and order the cuts, and the searched ramp enters RC-401 as its bias. Notice that one late or biased reading reaches both the loop and the cut.

> **TECHNOLOGY** **Final Element: A Timed Reflux Splitter**
>
> On some laboratory columns the reflux ratio is set by time rather than by a valve opening. A reflux valve switches the condensate between the column and the product receiver within a fixed cycle. On one teaching column with a 10 s cycle, a ratio of 4 to 1 sends condensate to the column for 8 s and to the product receiver for 2 s [24]. Frequent switching lets the column more closely approximate a steady state [23]. The column behind this chapter set its ratio with a swinging bucket that could be changed as often as every 5 min [16]. RC-401 instead applies any ratio from 0 to 40 at once, every 0.05 min.

Where the Controller Runs

Batch operations follow recipes developed in the laboratory: a sequence of operations in a prespecified order, repeated on a predetermined schedule to reach the production volume [25]. They are controlled from stand-alone controllers, programmable logic controllers, distributed control systems and personal computers [25]. The batch control committee of ISA describes a functional model that relates recipe management, scheduling, sequential control, regulatory control and safety interlock systems [26]. Read in those terms, the startup at total reflux, the fore cut, the main cut and the tail cut are the sequential control of one recipe, the policy and the searched ramp are recipe parameters, and RC-401 is the regulatory control. Supervisory and lower-level loops run with cycle times from milliseconds to minutes [27]. RC-401 executes every 0.05 min, which is 3 s, as Figure 3.6 draws it on plant hardware.

> **IN PRACTICE** **Running to the Cumulative Specification**
>
> On the column behind this chapter, the base case ran at total reflux for 30 min, then set a constant reflux ratio, usually between 3 and 5, and ran until the cumulative overhead composition reached 99 mol% methanol. Collection usually lasted 60 to 90 min. The instantaneous composition could be measured every 5 min and the total collected was read on a graduated cylinder, so the cumulative composition followed from both through a mass balance [16]. That is the stopping rule of Equation (3.9), kept by hand on readings five minutes apart.

What Is Modernizing It

Two developments use the repetition of batches. The first is learning from run to run. In a simulated binary batch column of ten stages, an iterative learning controller adjusted the reflux trajectory between batches to track a reference for the purity of the accumulated distillate. Its schemes converged after 25 to 30 runs, and a residual error had to be tolerated because the startup phase left the purity at the start of collection uncertain [25]. Most often the path constraints of a batch are met online and its terminal objectives run to run [25], and the receiver purity of this chapter is such a terminal objective. The second is estimation from batch data. The study behind this chapter fitted its model to experiments with an error measure insensitive to outliers, ranked which parameters the data could determine, and only then optimized the reflux schedule [16]. The shrinking-horizon MPC of Section 3.6.9 acts within one batch, and the rework of slop into the next charge, which Limitations sets aside, is where run-to-run learning would act.

3.9 Limitations

No claim about a real column's yield, its cycle time or the economics of a campaign can be founded on this model. The relative volatility was chosen to compensate for a stage count, so the column is a plausible batch column rather than the reference apparatus. There is no pressure, no temperature and no energy balance inside the column, which matters for control rather than only for accuracy, because real columns are frequently run on an inferential tray temperature with an analyzer trimming it and this model cannot show that structure at all [15]. The mixture is binary, where a real batch separates several components and a cut becomes a sequence of cuts with intermediate fractions accumulating between them. One batch is simulated in isolation, so slop is counted and then forgotten, whereas in a campaign it is reworked into the next charge, which changes that charge's composition and therefore the optimal policy for it, and the sequence of batches is the thing actually being optimized. The supervisory search integrates the same equations with the same parameters that the plant does, so it optimizes against a perfect model of a plant whose only unknown is the one a scenario changes, and real trajectory optimization is limited by parameter uncertainty rather than by search [16]. Perfect repeatability is the last simplification and the least obvious one, since a fixed step and a fixed controller period give one trajectory for a given scenario forever, whereas a real batch record from the same recipe never repeats twice. Each of the six modeling assumptions is given the condition under which it gives way in Section 3.3.6.

Two scenario disturbances are stylized. An analyzer that drifts 0.012 high in one step at a nominated minute is not how a calibration fault develops, but a step is the only form of bias whose consequence can be read off a single trend against a known reference run. A charge that is 0.42 instead of 0.59 with no measurement taken would not survive a real batch record, although it isolates what an open-loop trajectory assumes about its plant, which is the whole content of scenario 11.

The added local MPC leaves room for a full nonlinear formulation. Nonlinear Model Predictive Control could re-optimize the reflux trajectory in a receding horizon against the terminal objective, enforcing the cumulative purity constraint and the reflux limit directly rather than through the present frozen local balance or a reset-time ramp [28]. Moving horizon estimation could reconstruct the charge composition from the analyzer and the early trajectory, addressing scenario 11 directly, because a policy re-optimized on the estimated charge recovers 10.69 mol where the fixed ramp collects

1.26 mol. The same estimator structure supports soft sensing of the distillate composition between analyzer updates, which is the practical arrangement on industrial columns [15]. The reference study estimated its parameters and optimized its trajectory in APMonitor [16, 28], and the Gekko Optimization Suite provides the same nonlinear MPC and estimation framework in Python [29].

> ⚠ **OPEN CHALLENGE** **A Model That Predicts the Batch but Not Its Optimum**
>
> A 2017 handbook chapter on batch control calls model adequacy a major challenge in process optimization. A model is trained to predict the plant's outputs, not the gradients that decide optimality, so numerical optimization computes inputs that are optimal for the model and often fails to reach plant optimality. In practice, model-based optimal inputs lead to suboptimal, and often infeasible, operation [25]. Scenario 11 is a small case of it: the ramp searched for a 0.59 charge collects 1.26 mol on a 0.42 charge. Progress would be a policy that finds the plant's own optimum from its measurements as the charge changes.

3.10 Conclusion

This chapter used a ten-tray batch distillation column as an interactive laboratory for control against a plant that is consumed while the controller runs. The simulator integrates nineteen states covering the pot, ten trays, the condenser, two receivers, an analyzer and an energy meter, with equilibrium at a constant relative volatility, constant molar overflow fixing the internal flows, and a boil-up of 0.7072 mol/min set by a 600 W duty independently of the reflux ratio. Eleven equilibrium stages multiply the ratio $x/(1-x)$ by 2048 at total reflux and that factor never changes, while the pot it acts on falls from 0.5787 to 0.2337 in ninety minutes, which takes the distillate from 0.9996 to 0.9599 with no input change. On a 180 min budget against a 0.990 specification the best constant reflux ratio collects 15.10 mol, constant-composition control collects 15.29 mol with its reflux ratio pinned at 40 and its distillate rate down by a factor of about twelve, and a searched ramp from 2.45 to 10.57 collects 15.73 mol open loop. That ordering holds only because the product specification is applied to the cumulative average of the drum, since an instantaneous specification would make constant-composition control optimal by construction. The same ramp collects 1.26 mol on a charge of 0.42, where a proportional trim recovers 4.74 mol and re-optimizing recovers 10.69 mol. Further work would add multicomponent cuts, rework of slop into the following charge, a reboiler with thermal capacitance, and a search that carries parameter uncertainty rather than assuming the model. The new nominal MPC recovers product relative to a mismatched fixed ramp, but extra motion, biased instruments and strong mismatch limit that recovery. The batch column is an accessible browser laboratory for trajectory control against a terminal objective.

KEY POINTS

- A batch column has no operating point: the pot is charged once, its composition falls for the whole batch, and the batch is judged once, on the product drum at the end.
- On a 180 min budget the best constant reflux ratio collects 15.10 mol, constant-composition control 15.29 mol with its reflux ratio pinned at 40, and a searched ramp from 2.45 to 10.57 collects 15.73 mol open loop.
- That ordering holds only because the specification applies to the cumulative average of the drum.
- The same ramp collects 1.26 mol on a 0.42 charge, where a proportional trim recovers 4.74 mol and re-optimizing recovers 10.69 mol.

Where else this applies. Batch crystallization to a yield, in Chapter 4, has the same structure as this column: an input trajectory chosen over a fixed horizon against a result judged at the end, with constraints along the way.

References

[1] 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.

[2] D. Bonvin, "Optimal operation of batch reactors: a personal view," *Journal of Process Control*, vol. 8, no. 5-6, pp. 355–368, 1998.

[3] B. Srinivasan, S. Palanki, and D. Bonvin, "Dynamic optimization of batch processes: I. characterization of the nominal solution," *Computers & Chemical Engineering*, vol. 27, no. 1, pp. 1–26, 2003.

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

[5] D. J. C. Constable, C. Jimenez-Gonzalez, and R. K. Henderson, "Perspective on solvent use in the pharmaceutical industry," *Organic Process Research & Development*, vol. 11, no. 1, pp. 133–137, 2007.

[6] L. Rayleigh, "On the distillation of binary mixtures," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 4, no. 23, pp. 521–537, 1902.

[7] A. O. Converse and G. D. Gross, "Optimal distillate-rate policy in batch distillation," *Industrial & Engineering Chemistry Fundamentals*, vol. 2, no. 3, pp. 217–221, 1963.

[8] I. Coward, "The time-optimal problem in binary batch distillation," *Chemical Engineering Science*, vol. 22, no. 4, pp. 503–516, 1967.

[9] G. P. Distefano, "Mathematical modeling and numerical integration of multicomponent batch distillation equations," *AIChE Journal*, vol. 14, no. 1, pp. 190–199, 1968.

[10] U. M. Diwekar, R. K. Malik, and K. P. Madhavan, "Optimal reflux rate policy determination for multicomponent batch distillation columns," *Computers & Chemical Engineering*, vol. 11, no. 6, pp. 629–637, 1987.

[11] I. M. Mujtaba and S. Macchietto, "Holdup issues in batch distillation-binary mixtures," *Chemical Engineering Science*, vol. 53, no. 14, pp. 2519–2530, 1998.

[12] U. Diwekar, *Batch Distillation: Simulation, Optimal Design, and Control*, 2nd ed. Boca Raton, FL: CRC Press, 2011.

[13] K.-J. Kim and U. M. Diwekar, "New era in batch distillation: Computer aided analysis, optimal design and control," *Reviews in Chemical Engineering*, vol. 17, no. 2, 2001.

[14] J. S. Logsdon and L. T. Biegler, "Accurate determination of optimal reflux policies for the maximum distillate problem in batch distillation," *Industrial & Engineering Chemistry Research*, vol. 32, no. 4, pp. 692–700, 1993.

[15] M. Barolo and F. Berto, "Composition control in batch distillation: Binary and multicomponent mixtures," *Industrial & Engineering Chemistry Research*, vol. 37, no. 12, pp. 4689–4698, 1998.

[16] S. M. Safdarnejad, J. R. Gallacher, and J. D. Hedengren, "Dynamic parameter estimation and optimization for batch distillation," *Computers & Chemical Engineering*, vol. 86, pp. 18–32, 2016.

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

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

[20] D. S. Sholl and R. P. Lively, "Seven chemical separations to change the world," *Nature*, vol. 532, no. 7600, pp. 435–437, 2016.

[21] 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.

[22] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[23] Brigham Young University, Unit Operations Laboratory, "Batch distillation, version 16.2: startup procedure and refractometer instructions," uolab.groups.et.byu.net/files/batchdist/batchdist_manual.pdf, laboratory manual, accessed 22 November 2026.

[24] University of Utah, Department of Chemical Engineering, "UOL distillation column," www.che.utah.edu/sop/uol-distillation-column, standard operating procedure, accessed 22 November 2026.

[25] D. Bonvin and G. François, "Control and optimization of batch chemical processes," in *Coulson and Richardson's Chemical Engineering*. Elsevier, 2017, pp. 441–503.

[26] International Society of Automation, "ISA88: Batch control, committee scope and purpose," www.isa.org, standards committees, 2026, accessed 22 November 2026.

[27] 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.

[28] 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.

[29] 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.

3.11 Exercises

3.1 Concept check · 2 min

A batch of the column of this chapter is run to the end of its 180 min budget. What is the batch graded on?

(a) The largest deviation of the distillate from its 0.990 setpoint during the collection

(b) The moles in the product receiver at the end, provided their cumulative average meets 0.990

(c) The pot composition at the end, because a lean pot shows that the light component was recovered

(d) The reboiler energy spent, because the duty runs whether or not anything leaves the top

3.2 Concept check · 2 min

In scenario 10 the column runs in Manual at a reflux ratio of 6, and at 60 min the reboiler loses 35 % of its duty. What happens to the batch?

(a) The reflux ratio rises, because the reflux falls less than the distillate when the vapor rate drops

(b) The distillate grows leaner, because a lower vapor rate gives each tray less separating power

(c) Vapor, reflux and distillate fall together, so less product is made, and it comes out purer

(d) Nothing changes in the product, because the analyzer sees the fault and the loop restores the duty

3.3 Concept check · 2 min

Suppose the receiver were switched to slop the moment the instantaneous distillate composition fell below the 0.990 specification, instead of when the cumulative average of the drum falls to it. Which policy would then collect the most product?

(a) The searched ramp, because it starts near the minimum reflux and takes product fastest
(b) The best constant reflux ratio, because it needs no analyzer and cannot saturate
(c) Total reflux followed by one high reflux ratio, because every mole would then be pure
(d) Constant-composition control with its setpoint at 0.990, held as long as possible

3.4 Theory and calculation · 6 min

Use the total-reflux relation of Equation (3.7) with ten trays.

(a) At $\alpha = 2$, find the separation factor and the distillate composition when the pot holds $x_B = 0.5787$. *Checkpoint:* $2^{11} = 2048$, $x_D = 0.999\,64$.

(b) After ninety minutes at $R = 4$ the pot has fallen to $x_B = 0.2337$. What would total reflux give at the top now, and how does that compare with the 0.9599 the column actually delivers at $R = 4$? *Checkpoint:* $x_D = 0.9984$ at total reflux.

(c) Repeat part (a) at the reference chemistry, $\alpha = 1.725$. *Checkpoint:* Factor 402, $x_D = 0.9982$.

3.5 Theory and calculation · 5 min

The reboiler supplies 600 W at a heating efficiency of 0.719, and the heat of vaporization is 36.6 kJ/mol. A total-reflux startup of 30 min precedes a collection that runs to 180 min.

(a) Compute the boil-up V. *Checkpoint:* $V = 0.7072\,\text{mol/min}$.

(b) At a constant reflux ratio of 6, find the distillate rate and the total distillate taken over the collection, product and slop together. *Checkpoint:* $D = 0.1010\,\text{mol/min}$, 15.15 mol.

(c) Find the energy per mole of product for the best constant reflux ratio (15.10 mol) and for the searched ramp (15.73 mol). *Checkpoint:* 429 and 412 kJ/mol.

3.6 Theory and calculation · 4 min

A product drum holds 10 mol at a cumulative 0.995, and the specification is 0.990.

(a) Compute the purity credit $C = n_P(\bar{x}_P - x_{\text{spec}})$. *Checkpoint:* $C = 0.050\,\text{mol}$.

(b) Distillate now arrives at a steady 0.980. How many moles can be added before the cumulative rule of Equation (3.9) closes the drum? *Checkpoint:* 5.0 mol.

(c) How many of those moles would an instantaneous specification accept? *Checkpoint:* None.

3.7 App challenge · 8 min

Load scenario 4, Constant reflux $R = 3$: watch the composition slide. Keep Manual, the default seed and the default speed of 10, which runs a 180 min batch in about 18 s.

(a) Change nothing and run until the batch stops. Read the Product collected, Sent to slop and Batch time used tiles, and the time of the tail cut, which is the first `tail` entry in the exported CSV column `cut_phase`. *Checkpoint:* 9.16 mol, 13.99 mol of slop, stops at 161 min, tail cut at 81.8 min.

(b) Press Reset, set the Reflux ratio to 6 before pressing Run, and run to the end of the batch. Read the same tiles and the Product composition tile. *Checkpoint:* 15.10 mol at 0.9895, 0.06 mol of slop, tail cut at 179.5 min.

A ratio of 8 collects 11.79 mol at 0.9978 with no tail cut. Explain why 6 beats both 3 and 8. Give moles to within 0.02 mol and times to within 0.2 min.

3.8 App challenge · 8 min

Load scenario 5, Constant composition, and the reflux runaway. Keep Auto, the default seed, the default speed of 10 and the shipped tuning ($K_c = 600$, $\tau_I = 4\,\text{min}$).

(a) Run to the end of the batch. Read the Product collected and Sent to slop tiles, the reflux ratio at 120 min from the CSV column `mv_reflux_ratio`, and the tail-cut time from the first `tail` entry in `cut_phase`. *Checkpoint:* 15.29 mol, 0.34 mol of slop, $R = 8.11$, tail cut at 162.7 min.

(b) Press Reset, set the Proportional gain to 300 before pressing Run, and run to the end. Read the same quantities. *Checkpoint:* 13.33 mol, tail cut at 119.7 min.

(c) At 150 min the reflux demand rises by about 0.58 per minute. Estimate the ramp lag $(\mathrm{d}R/\mathrm{d}t)\,\tau_I/K_c$ at both gains and explain why the tail cut moves. *Checkpoint:* 0.0039 and 0.0077.

Give moles to within 0.02 mol, times to within 0.2 min and the reflux ratio to within 0.05.

3.9 App challenge · 9 min

Load scenario 11, A leaner charge than the ramp was optimized for. It runs in Auto with Optimal ramp By hand from 2.45 to 10.57, a Proportional gain of 0, the default seed and the default speed of 10.

(a) Run to the end. Read Product collected, Sent to slop and the tail-cut time from the first `tail` entry in the CSV column `cut_phase`. *Checkpoint:* 1.26 mol, 14.55 mol of slop, tail cut at 36.5 min.

(b) Press Reset, set the Proportional gain to 600 before pressing Run, and run to the end. Read the same quantities. *Checkpoint:* 3.10 mol, 8.07 mol of slop, tail cut at 50.8 min.

(c) Press the reload button, which puts the gain back to 0, switch to Advanced, set Optimal ramp to Searched before pressing Run, and run to the end. Read Product collected and the reflux ratio at the end. *Checkpoint:* 10.69 mol, $R = 17.83$.

What does each remedy recover, and what would a plant need to do before executing a stored recipe? Give moles to within 0.02 mol and times to within 0.2 min.

3.10 Application area · 4 min

To shorten the cycle, a production engineer proposes cutting the total-reflux startup of the shipped column from 30 min to 10 min, keeping the 180 min budget and a reflux ratio of 6. What does the chapter predict, and what would you recommend?

3.11 Application area · 4 min

An older plant column has a swinging-bucket reflux divider that can be moved once every 5 min and an online analyzer with about 5 min of lag. The plant asks for constant-composition control at 0.990. Name two problems the chapter predicts and one alternative.

3.12 Review the agent · 5 min

An AI assistant reviews an exported scenario 9 run (the analyzer drifts 0.012 high at 34 min) and writes: “This batch beat the reference run of scenario 5: 16.33 mol against 15.29 mol and 96.7 % recovery against 91.6 %. The receiver average in column `prod_x_meas` ends at 0.9900, so the drum is on specification. Keep the new analyzer calibration.” Find the flaw and say how you would check it.