INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 2 Practical PID Control

CHAPTER 3

Stirred Blending Tank Simulator for Composition Control Against an Actuator Limit

Pages 55–86

Edited by John D. Hedengren

How to cite this chapter. "Stirred Blending Tank Simulator for Composition Control Against an Actuator Limit," in *Practical PID Control*, John D. Hedengren, Ed., vol. 2 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 55–86.

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

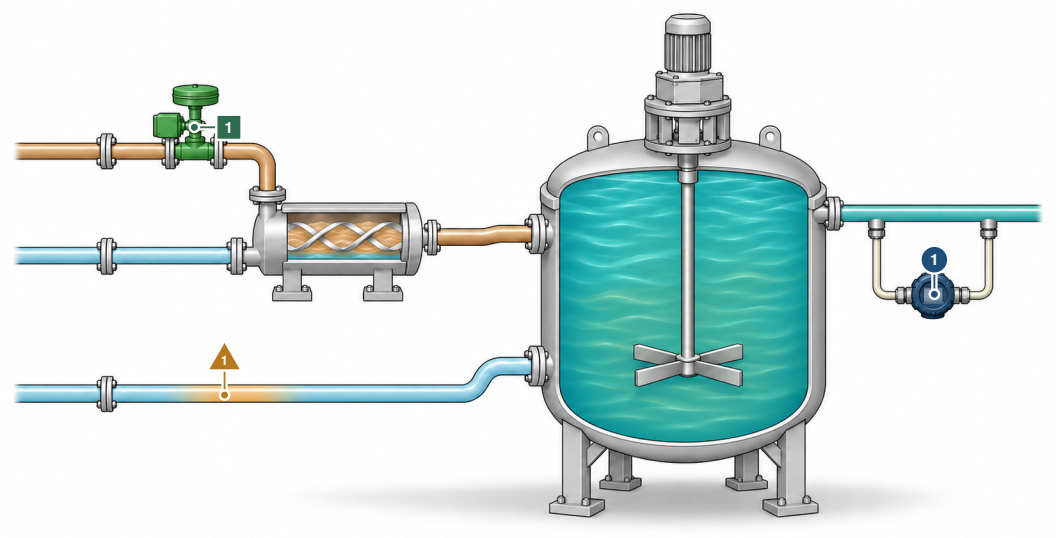
The app for this chapter
https://apmonitor.com/apps/blend/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v02-practical-pid/#blend

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

CHAPTER 3

Stirred Blending Tank Simulator for Composition Control Against an Actuator Limit



The controller moves 1 the valve that sets how strong the feed is

The controller reads 1 an analyzer that reads the strength of the product

What disturbs it 1 plain diluent that waters the product down

Notice that the diluent enters on its own line, so no setting of the valve can make the feed stronger than the concentrate it draws on. Illustration, not to scale.

IN PLAIN TERMS Holding a Mixture's Strength With a Limited Supply

The main idea. A blending tank mixes a strong feed with a plain stream that carries none of the ingredient. A controller reads the product and makes the feed stronger or weaker to hold a chosen strength as more of the plain stream comes in.

The challenge. The feed can never be stronger than the supply it comes from. When too much of the plain stream arrives, even the strongest feed falls short, and no setting of the controller can close the gap.

Why it matters. Before blaming the tuning for a lasting error, check whether the supply could reach the target at all.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history of reset (Section 3.1), the history of model-based tuning (Section 3.4), the try-it segment of case 8 (Section 3.6), the try-it segment of case 4 and the two facts (Section 3.7) and the key points
Study	70 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 3.7), Technology in Practice (Section 3.8), Limitations (Section 3.9) and the open challenge

The control question A blending loop manipulates a feed concentration that is bounded by what the upstream supply contains. Which standing errors can a controller remove, and at what dilution flow does a composition setpoint pass out of reach of any tuning?

Learning objectives

- Derive the gain, the residence time and the steady-state inverse of the tank from the species balance.
- Explain why a dilution change moves the integral gain of an internal model control tuning and leaves the proportional gain alone.
- Decide from the steady-state balance whether a setpoint is reachable, and compute the offset when it is not.
- Distinguish the offset of a proportional loop from the offset that an actuator bound imposes.
- Describe how back-calculation holds the stored output at the bound, and what a model predictive controller adds at the limit.

Prerequisites Sections 1.4.2 and 1.6.2 (velocity form and offset), 1.4.4 (step test), 1.4.5 (tuning rule) and 1.4.9 (windup). Sections 1.4.7 and 1.6.3 are needed only for the relay test and the predictive controller. Chapter 2 uses the same algorithm.

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

See also Volume 1, Chapter 2 and Volume 1, Chapter 3 for step tests on tanks whose model moves with the operating point. Volume 4, Chapter 9 for cases that separate a tuning problem from missing capacity. Volume 13, Chapter 5 for feeds held in ratio with analyzer trims and Volume 13, Chapter 10 for a second handle that keeps the first off its limit.

3.1 Introduction

Composition is the product property that a specification or a permit constrains, and blending is the operation that delivers it. Streams are combined so that the mixture meets a number written into a product specification or a permit, and the combination is enforced continuously on gasoline in a refinery header [1–3], on coagulant in a water treatment works [4], and on active ingredient in a continuous tablet line [5, 6]. The stirred vessel that performs the mixing is among the earliest systems studied in chemical engineering dynamics, and its residence time behavior was formalized as a distribution rather than a single number in 1953 [7]. A composition loop built on such a vessel is the standard first example of process dynamics in the textbooks, because a species balance on a well-mixed inventory reduces conservation of mass to a single linear ordinary differential equation [8, 9]. There is one state, no dead time when the mixing is ideal, a gain that is close to unity, and a first-order response that any proportional-integral controller can be made to follow.

What separates a blending loop from the level, temperature and flow loops taught beside it is the nature of its manipulated variable. A temperature loop moves a valve, and the travel of that valve is a sizing decision that can be revisited. A blending loop moves a composition or a ratio, and its ceiling is set by what the upstream tank contains. Raising the concentrate strength is a procurement decision rather than a tuning decision. The consequence is that a large enough load disturbance does not degrade the loop. It removes the setpoint from the reachable set, and the resulting steady-state error is a property of the mass balance rather than of the controller.

The usual treatment presents offset as a deficiency with a known remedy. A proportional controller leaves an error, integral action removes it, and the integrator that removes it is then protected against saturation by an anti-windup mechanism so that the loop recovers promptly when the limit clears [8, 10]. Sections 1.4.2 and 1.4.9 present the two steps in that order. Each of those statements is correct, and taken together they suggest that any standing error is a tuning failure. However, an actuator bound makes some standing errors a property of the steady-state balance, and the literature that treats saturation as a design problem is concerned with recovery after the limit clears rather than with whether the setpoint was reachable at all [11–14]. A reader who has only encountered the removable offset may respond to the irremovable one by raising the gain.

> ⧖ HISTORY **Early Process Controllers and Reset, 1930s**
>
> By the mid 1930s many companies supplied sensors, actuators and controllers to the process industry. A proportional controller holds its setpoint without error only if one more number is right: the steady output that the process needs, called the reset value. In early controllers it was adjusted by hand. Automatic reset, one of the early inventions of integral control, produced it by passing the controller output through a low-pass filter and feeding it back with positive gain, which was easier to build than an integrator [10]. Integral action is in this sense a self-resetting bias, which is why it is also called reset. An older name is floating [15]. The stored output of this chapter's PID is that reset value.

This chapter works with a stirred blending tank simulation in which the manipulated feed concentration is bounded at 0 to 2 mol/L. It quantifies the dilution flow at which a composition setpoint passes from a tuning limitation to a capacity limitation. The experiments built on the simulation cover process gain and time constant identification from a step test, proportional offset, integral action, actuator saturation, anti-reset windup and steady-state reachability, all of which are standard topics in a first course on process dynamics and control [8–10].

The chapter is organized as follows. Section 3.2 places the composition loop in blending practice. Section 3.3 specifies the simulator, covering the balance, the bound, the instrument, the constants and the integration scheme in enough detail for a second implementation to reproduce the same trajectories. Section 3.4 shows that the dilution stream moves the gain and the time constant together. As a result, internal model control leaves the proportional gain of the loop untouched by the dilution while moving its integral gain, and both results are checked against step tests run in the application. Section 3.5 describes the velocity-form PID with its back-calculation, a linear model predictive controller offered beside it, a moving horizon estimator that may stand in for the filter of that controller, and a relay test that sets the PID constants from the running loop. Section 3.6 derives the reachability inequality for this plant in closed form. Section 3.7 reports the measured runs. They are built around two standing errors of similar size, 0.149 mol/L left by a proportional-only loop against a 33 m^3/min dilution and 0.167 mol/L imposed by the bound at 50 m^3/min, only the first of which is removable by feedback. The same section lists the ten shipped cases with their completion requirements, compares the predictive controller with the PID to separate what a model of the load contributes from what an explicit treatment of the bound contributes, and assesses the estimator and the relay on the same cases. Section 3.8 describes the devices and the platform of a real blend loop, Sections 3.9 and 3.10 state the limitations of the model and its economic context, Section 3.11 summarizes the results, and the exercises follow in Section 3.12.

3.2 Background

Blending entered the engineering literature as an optimization problem before it entered as a control problem. The aviation gasoline blend written as a linear program in 1952 is one of the first industrial applications of linear programming of any kind, and it was posed exactly as a blending question, which combination of available stocks meets every octane, volatility and vapor pressure specification at the lowest cost [1]. That economic framing remains current. The economic incentive in a modern blend is giveaway, the quality delivered above the specification that the customer pays for, and moving the average blend closer to its limit has value in direct proportion to how tightly the composition can be held [2]. Refinery blend planning is still posed as a simultaneous optimization over recipes and schedules, with the nonlinearity of the quality models and the discrete structure of the tank assignments as the two difficulties [3]. Volume 11, Chapter 15 treats an in-line gasoline blender with a recipe optimizer, whose product is certified from the average in its tank.

Feedback control of blending became practical when the in-line analyzer arrived. An offline laboratory assay supports recipe adjustment between batches, while an in-line analyzer supports feedback within a blend, and the second changes what the optimization is allowed to assume. Model-based real-time optimization of a gasoline blend header reduces giveaway by driving the blend toward

Table 3.1 Classification of Variables. Nominal Values Are Those of the 1.5 mol/L Operating Point With the Diluent Stream Closed.

Symbol	Role	Nominal	Range
C_A	state, true concentration	1.50 mol/L	≥ 0
$\hat{C}_A$ (AT)	CV, analyzer reading	1.50 mol/L	≥ 0
C_A^{SP}	setpoint	1.50 mol/L	0 to 2, step 0.01
C_{Af}	MV, feed concentration	1.50 mol/L	0 to 2, step 0.01
q_d	DV, diluent flow	0 m^3/min	0 to 100, step 1
q	parameter, feed flow	100 m^3/min	20 to 200, step 5
T_f	DV, feed temperature	300 K	270 to 370, step 1
$q+q_d$	output, outlet flow	100 m^3/min	≥ 0
τ	output, residence time	1.00 min	$V/(q+q_d)$
limit flag	output, feed at a bound	false	—

the specification and relying on the regulatory layer to hold it there, so the achievable margin is set by the variance of the composition loop [2]. Water treatment carries the same structure with a different specification. Coagulant dose is manipulated against a raw water quality that changes with weather and season, and the reason online monitoring of the coagulation process has attracted sustained attention is that dosing too little fails the treatment objective while dosing too much wastes chemical and loads the sludge [4]. Continuous pharmaceutical manufacturing has adopted the same architecture most recently, with blend uniformity held by feedback around a continuous mixer inside a plantwide control structure that mixes model-based and proportional-integral-derivative loops [5].

The modeling side of the subject is older and is dominated by one question, how far a real vessel departs from perfect mixing. The residence time distribution was introduced to answer it, with the perfectly mixed tank and the plug flow tube as the two limiting distributions and every real vessel somewhere between [7]. The same tool is now the standard characterization of continuous solids equipment, where the width of the distribution sets how much an upstream feeder fluctuation is smoothed before it reaches the tablet press [6]. A perfectly mixed vessel is the limiting case in which the outlet composition equals the inventory composition and the transfer function is a single first-order lag with no dead time, which is the case the application implements.

3.3 Process Model

The simulated vessel is a closed, liquid-full, agitated tank fed by two streams. The first is a concentrated stream of fixed flow whose concentration is set by the controller, and the second is a pure diluent stream of variable flow whose concentration is zero. Because the vessel is liquid full and the liquid is incompressible, the outlet flow is the sum of the two inlet flows at all times and the volume is not a state. The controlled variable is the outlet concentration reported by analyzer AT, the manipulated variable is the feed concentration, and the diluent flow is left to the operator whichever mode the loop is in. Figure 3.1 places every symbol used below on the equipment.

Table 3.1 sorts the signals into their roles. The feed flow q is a model parameter that the operator can change during a run, and it appears in every model coefficient. The feed temperature is measured, recorded in the exported file, and deliberately absent from every equation. Concentrations are carried in mol/L, flows in m^3/min, volume in m^3 and time in min, and no conversion between m^3 and L appears in the source because the factor divides out of every term of the balance.

3.3.1 Species Balance

A solute balance over the liquid-full vessel, with the inlet solute carried entirely by the concentrated stream and the outlet composition equal to the inventory composition, is shown in Equation (3.1).

$$V\frac{\mathrm{d}C_A}{\mathrm{d}t} = qC_{Af} - (q+q_d)C_A, \qquad 0 \leq C_{Af} \leq C_{Af}^{\max} \tag{3.1}$$

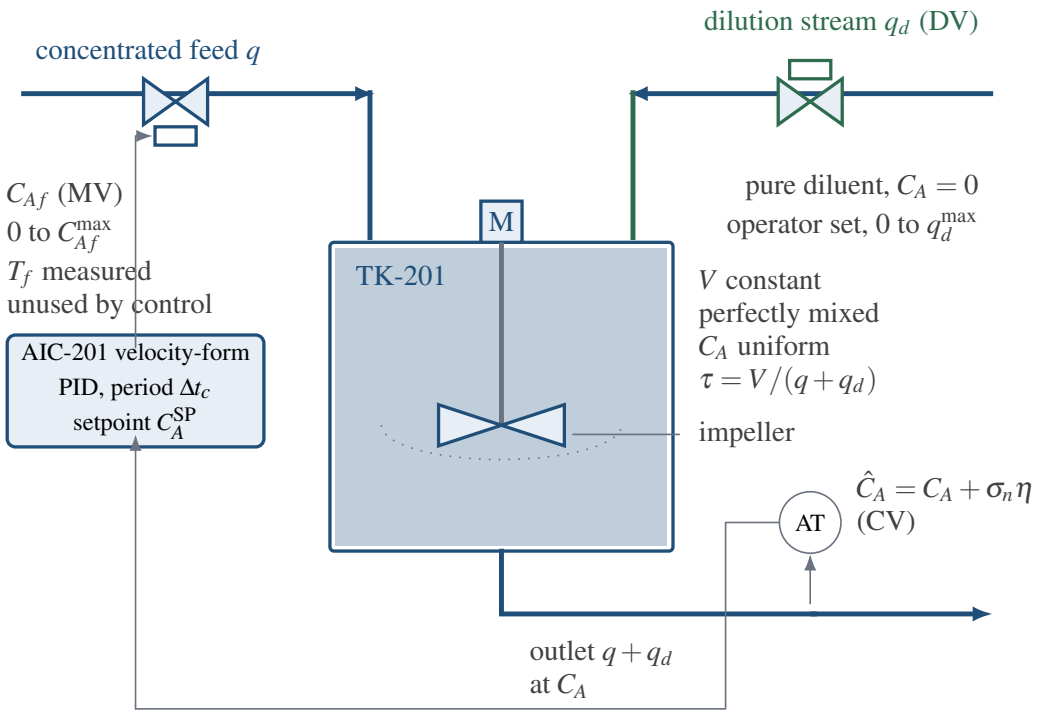


Figure 3.1 Blending tank TK-201 with its concentrated feed, its diluent stream and the outlet composition analyzer. The feed concentration C_{Af} is the manipulated variable and is bounded at $C_{Af}^{\max}$, the diluent flow q_d is the load disturbance and carries no solute, and the vessel is liquid full so the outlet flow is $q+q_d$. Controller AIC-201 reads the noisy analyzer value $\hat{C}_A$ and sets the feed concentration.

where V is the liquid volume in m^3, C_A is the outlet and inventory concentration in mol/L, C_{Af} is the concentration of the feed stream in mol/L, q is the feed flow and q_d is the diluent flow, both in m^3/min. Every term carries units of m^3/min times mol/L. The diluent contributes no term on the right because its concentration is zero, and it acts only through the outlet term, which is the mechanism by which it disturbs the loop.

Two structural facts follow from the shape of Equation (3.1). The volume V is a constant and not a state, because the vessel is liquid full and the outlet flow is defined as the sum of the inlets rather than by a valve or a pump. There is no total mass balance to integrate and no level loop. The equation is also linear in C_A at fixed flows, so the nonlinearity of the plant is entirely in the coefficients. Both the coefficient of C_{Af} and the coefficient of C_A move when a stream moves, which is what makes a diluent flow change a retuning of the plant rather than an additive load.

Holding the feed concentration and the two flows constant gives the first-order response of Equation (3.2).

$$C_A(t) = \bar{C}_A + \left[C_A(0) - \bar{C}_A\right] \mathrm{e}^{-t/\tau}, \qquad \bar{C}_A = \frac{q}{q+q_d} C_{Af} \tag{3.2}$$

where the residence time is $\tau = V/(q+q_d)$ in min. The steady-state gain from feed concentration to outlet concentration is $K_p = q/(q+q_d)$, which is dimensionless and is exactly unity when the diluent valve is shut. There is no dead time anywhere in the plant, because a perfectly mixed vessel begins to respond in the first instant and the outlet is drawn from the same well-mixed inventory that the feed enters.

Table 3.2 Model and Controller Parameters

Symbol	Meaning	Default	Range	Unit
Configuration parameters				
V	liquid volume	100	20–300	m^3
q	feed flow	100	20–200	m^3/min
σ_n	analyzer noise	0.01	0–0.05	mol/L
Operating parameters				
C_{Af}	manual feed concentration	1.50	0–2	mol/L
C_A^{SP}	concentration setpoint	1.50	0–2	mol/L
q_d	diluent flow	0	0–100	m^3/min
T_f	feed temperature	300	270–370	K
K_c	proportional gain	2.0	0–10	—
τ_I	integral time (20 = off)	1.0	0.01–20	min
τ_D	derivative time	0.0	0–2	min
—	controller	PID	PID, MPC	—
T_p	MPC prediction horizon	4	0.4–8	min
λ	MPC move suppression	30	0–100	—
—	MPC estimator	Kalman	or MHE	—
—	speed multiplier	1	0.5–10	—
Set when a run starts				
$C_A(0)$	initial concentration, steady	1.50	—	mol/L
u_b	controller bias, starting feed	1.50	—	mol/L
Fixed in the source				
C_{Af}^{max}	feed concentration limit	2.0	—	mol/L
Δt	integration step (Euler)	0.01	—	min
Δt_c	controller sample period	0.02	—	min
Δt_h	logging interval	0.1	—	min
T_s	MPC decision period	0.04	—	min
N_e	MHE window	10	—	executions
ε_r	relay band	0.03	—	mol/L

3.3.2 Analyzer

What the control law sees is not the state of Equation (3.1) but the analyzer output of Equation (3.3).

$$\hat{C}_A = \max(0,\, C_A + \sigma_n\, \eta) \tag{3.3}$$

where σ_n is the analyzer noise standard deviation in mol/L and η is a fresh standard normal deviate at each controller execution, obtained through a Box-Muller transformation from the seeded generator of the engine (Section 1.6.1), which is reseeded with the same seed whenever a run starts. Nothing else stands between the state and the reading. There is no analyzer lag, no sample transport delay, no calibration bias and no drift, and the noise is white at the control rate. The true concentration and the measured value are both displayed and both exported, which puts the instrument error beside the state it corrupts.

Table 3.2 holds every number a second implementation would need. Two numbers of the default configuration are deliberately unphysical. A feed of 100 m^3/min carrying 2 mol/L of solute is a throughput no blend drum of 100 m^3 would ever see, but it makes the residence time exactly 1 min and fits a complete step response into a window that can be watched at ordinary speed. A diluent stream reaching half the feed flow would not occur in a physical system, but it is scaled that way so that the boundary of Equation (3.19) can be crossed with one input change.

Configuration values are guarded before use, with V raised to at least 1 m^3 and both q and σ_n raised to at least zero, and the feed concentration is clipped to $[0, C_{Af}^{max}]$ as it enters the model whichever mode produced it. The proportional gain is dimensionless because the manipulated variable and the controlled variable are both concentrations, which is uncommon among process loops. The integral

and derivative gains of the parallel form (Section 1.4.2) follow from the two times as $K_I = K_c/\tau_I$ in 1/min and $K_D = K_c\tau_D$ in min, so the default tuning carries $K_I = 2.0/\text{min}$, and the top of the integral-time range, 20 min, switches integral action off. Every case starts at steady operation. In automatic the tank is placed at the steady state of the setpoint in force, and in manual at the steady state of the manual feed, with the flows and parameters in force, so the defaults of the second group in Table 3.2 are those of the reference case.

3.3.3 Numerics and Clocks

A fixed 0.01 min step and explicit Euler advance Equation (3.1), with the state floored at zero afterwards. The app supplies this step in place of the Runge–Kutta default of the framework (Section 1.6.1). The most demanding configuration the input ranges allow is $q = 200\,\text{m}^3/\text{min}$ with $q_d = 100\,\text{m}^3/\text{min}$ and $V = 20\,\text{m}^3$, giving $\tau = 0.067\,\text{min}$ and a ratio $\Delta t/\tau$ of 0.15, far inside the stability bound of the method across the whole parameter range. Control acts on every second step, so the sample period is 0.02 min of process time, and one historian row is written every 0.1 min, at every fifth execution. Neither clock is driven by wall time. The engine keeps them apart from the drawing of the screen, as Section 1.6.1 describes, which leaves the trajectory and the exported time base identical at every speed. One second of wall time corresponds to 1 min of process time at $1\times$.

Within one step the order is that of Figure 1.10, and it is visible in the exported file. Scheduled events are applied first, then, on every second step, the feed concentration is chosen by the control law or copied from the manual input and applied, and the balance is advanced across the step. The analyzer reading for an execution is drawn when the step ending at that instant completes, from the state at that instant, and a row is written at the end of the step. Rows fall on every fifth execution, so each row holds the noise-free concentration and the analyzer reading the controller acts on at its time stamp, beside the feed applied over the 0.02 min period that ends there. A time constant read graphically from the reading column is therefore not displaced, while the feed column trails its time stamp by one sample period.

3.3.4 Assumptions

A mechanism appears in the model when it changes a commissioning decision and not otherwise. Every item listed here is given with the conditions that make it acceptable and the conditions that break it.

1. *Perfect mixing.* The inventory is uniform, so the outlet concentration equals the tank concentration at every instant and the residence time distribution collapses to the single exponential of Equation (3.2). Real agitated vessels deviate from that limit, and the deviation is normally reported as the spread of a measured residence time distribution [6, 7]. The consequence here is that the simulated plant has no dead time at all, while a real blend header carries transport delay from the mixing point to the analyzer and a finite analyzer cycle. A loop tuned on this simulator is therefore tuned for the most forgiving composition plant that exists.

2. *Constant volume.* The vessel is liquid full and incompressible, so the outlet flow equals the sum of the inlets instantaneously and the volume never appears as a state. A real blend drum with a level controller would pass an inlet flow change to the outlet only after the level loop had acted, which adds a second dynamic element between the diluent valve and the outlet composition. None of that is represented.

3. *Ideal streams and constant density.* The two streams mix without volume change, the solute is dilute enough that the concentrated feed carries the same volumetric flow whatever its concentration, and no density is used anywhere in the source. A concentrate strong enough to matter would change the volumetric balance as well as the solute balance.

4. *Feed concentration as a direct actuator.* The manipulated variable is the concentration of the feed itself, applied with no valve, no blend dynamics upstream and no lag. In practice that

concentration is produced by a ratio of two streams or by a dosing pump, and the loop would be a cascade whose inner element has its own time constant, rangeability and stiction. Volume 13, Chapter 5 treats such a ratio station, with feeds held in ratio to the cell current and trimmed by analyzers. The simulator collapses that inner element to an ideal instantaneous setting, so the only actuator imperfection retained is the bound itself.

5. *Ideal analyzer apart from noise.* AT has no lag, no dead time, no drift, no calibration error and no resolution limit, and the noise it adds is Gaussian and white. Real composition measurements are dominated by sampling delay and by low-frequency calibration error, neither of which is white, and the cycle time of a real analyzer is frequently long enough to be the dominant dynamic element in the loop. Volume 5, Chapter 3 measures what an analyzer lag removes from a composition loop.

6. *Isothermal operation.* No energy balance is integrated. The feed temperature is measured and recorded because the model signature that the application inherited carries it, and it enters no equation. Any real blend in which viscosity, solubility or reaction rate depends on temperature would need that balance.

7. *Omitted entirely.* Upstream flow disturbances on the concentrated stream, a second solute, reaction in the vessel, agitator failure, level control, heat effects and any second vessel are absent. None of them alters the one behavior this application is built to show, which is what a bounded feed composition does to a setpoint, so all of them are left out.

3.4 Dynamic Analysis

3.4.1 Dependence of Gain and Time Constant on Throughput

Both constants of the first-order model in Equation (3.2) carry the same denominator, as collected in Equation (3.4).

$$K_p = \frac{q}{q+q_d}, \qquad \tau = \frac{V}{q+q_d}, \qquad \theta_p = 0 \tag{3.4}$$

where K_p is dimensionless, τ is in min and θ_p is the dead time. Opening the diluent valve shrinks both. At the nominal $q = 100\,\mathrm{m^3/min}$ and $V = 100\,\mathrm{m^3}$ the plant is a unity-gain lag of 1.00 min, and at $q_d = 50\,\mathrm{m^3/min}$ the same vessel is a plant of gain 0.667 and time constant 0.667 min. The response becomes faster and weaker at the same time, while the more common pattern in process loops pairs a slower response with a stronger one. The gravity-drained tank of Volume 1, Chapter 2 and the vehicle of Chapter 2 show the same pairing, a gain and a time constant that move together at a fixed ratio.

The feed flow acts differently, because it appears in the numerator of the gain as well as in the shared denominator. With the diluent valve shut, K_p is unity whatever q is, so a feed flow change moves only the time constant. Doubling q from 100 to $200\,\mathrm{m^3/min}$ halves τ from 1.00 to 0.50 min and leaves the resting concentration unchanged. The residence time is set by the vessel volume and the throughput rather than by the controller, so a faster composition response requires a higher throughput or a smaller inventory.

3.4.2 Internal Model Control Tuning

For a first-order plant carrying no dead time, internal model control tuning gives $K_c = \tau_p/(K_p\tau_c)$ and $\tau_I = \tau_p$, with τ_c the specified closed-loop time constant [16, 17]. This is the rule of Equation (1.11) with $\theta_p = 0$ and τ_c large enough that the integral time stays at τ_p (Section 1.4.5). Substituting Equation (3.4) cancels the throughput term from the proportional gain but not from the integral gain, as shown in Equation (3.5).

$$K_c = \frac{V}{q\,\tau_c}, \qquad K_I = \frac{K_c}{\tau_I} = \frac{q+q_d}{q\,\tau_c} \tag{3.5}$$

Table 3.3 Step Tests From $C_{Af} = 1.0$ to 1.5 mol/L With $\sigma_n = 0$ and $V = 100\,\mathrm{m}^3$. Predictions Are Equation (3.4).

q (m^3/min)	q_d	$\hat{K}_p$	K_p	$\hat{\tau}$	τ (min)	t_{63}	R^2
100	0	1.0000	1.0000	0.9950	1.0000	0.9952	1.00000
100	25	0.8000	0.8000	0.7950	0.8000	0.7953	1.00000
100	50	0.6667	0.6667	0.6617	0.6667	0.6634	1.00000
200	0	1.0000	1.0000	0.4950	0.5000	0.4954	1.00000

where K_c is dimensionless and K_I is in 1/min. The required proportional gain depends on the vessel volume and the feed flow and not at all on the diluent stream, because gain and time constant shrink in step and their quotient is a constant of the vessel. The integral gain depends only on the dilution ratio $(q+q_d)/q$ and on τ_c. With the default $V = 100\,\mathrm{m}^3$, $q = 100\,\mathrm{m}^3$/min and a moderate $\tau_c = 0.5$ min, Equation (3.5) gives $K_c = 2.0$ and $K_I = 2.0$/min, which are the default $K_c = 2.0$ and $\tau_I = 1.0$ min, and opening the diluent valve to 50 m^3/min raises the correct K_I to 3.0/min ($\tau_I = 0.667$ min) while leaving K_c at 2.0. A dilution change therefore requires retuning of the integral gain only.

Pause. Try Exercise 3.1 before reading on.

> HISTORY **Tuning From a Model, 1986 to 2003**
>
> In 1986 a paper on internal model control derived PI and PID settings analytically from a process model. These IMC-PID rules achieved widespread industrial acceptance. They proved robust and gave very good responses to setpoint changes, but a poor response to load disturbances on integrating processes. A paper of 2003 kept the model-based start and changed the rule for the integral term, so that a single rule serves a first-order or second-order model with time delay, for setpoints and for load disturbances. Its integral time is the smaller of the process time constant and four times the sum of the closed-loop time constant and the time delay [17].

3.4.3 Identification From the Running Application

Four manual step tests are collected in Table 3.3, and Section 1.4.4 gives the step test and its fit in general. Each starts from the reference case with the analyzer noise switched off and the flows and a 1.0 mol/L feed set before the run, so the tank starts at the steady state of that feed, and raises the feed concentration by 0.5 mol/L 1 min into the run. Gain is the secant between the starting state and the state 10 min after the step. The time constant comes from a least-squares exponential fitted to the logged noise-free concentration with the initial value, the final value and τ all free, and t_{63} is read from the logged analyzer column by linear interpolation.

The secant gains reproduce $q/(q+q_d)$ to four decimals in all four tests, because with the noise switched off the starting state is the exact steady state of Equation (3.2). The largest shortfall, 0.00004 at $\tau = 1$ min, is the part of the step that a lag of that length has not completed 10 min after it. The fitted time constants sit below the analytic values by 0.50 to 1.00 %, and the explicit Euler integrator accounts for all of it, because a deviation decays by a factor $1 - \Delta t/\tau$ per step against the $\mathrm{e}^{-\Delta t/\tau}$ of the exact solution and the discrete response therefore runs slightly fast. The equivalent discrete time constant, $-\Delta t/\ln(1-\Delta t/\tau)$, is 0.9950, 0.7950, 0.6617 and 0.4950 min over these four flows, equal to the fitted values to four decimals. The graphically read t_{63} exceeds the fitted time constant by 0.0002 to 0.0017 min, which is the error of interpolating linearly between logged rows 0.1 min apart, because each logged reading is taken at its own time stamp. Both discrepancies are smaller than the difference between any two rows of the table, so neither obscures the scaling result.

Three distinct models appear in this chapter, and two of them consist of one or two numbers. Equation (3.1) with its bound is the simulation, exact inside the assumptions of Section 3.3.4. The pair $(\hat{K}_p, \hat{\tau})$ in Table 3.3 is the identified model, and it describes the plant at one pair of flows and nowhere else. Whatever model the PID holds is carried by the sum of its bias and its integral state, a

single number that starts at the feed of the starting steady state and settles at the steady-state inverse of the plant derived in Section 3.6. The PID gains carry no model of the tank. The predictive controller of Section 3.5.4 carries the second model itself, the pair of the first row of Table 3.3, and keeps it at every dilution.

3.5 Controller Design

3.5.1 Velocity Form With Derivative on the Measurement

The PID is the velocity-form algorithm of the framework, which Section 1.6.2 gives in general, and this subsection writes it in the symbols of the composition loop. Each cycle the algorithm returns an adjustment to the feed concentration rather than the concentration itself, and the operating point is held in an integral state rather than in a bias. Error is defined in Equation (3.6) as the setpoint less the analyzer reading, and Equations (3.7) to (3.9) give the three contributions to the increment.

$$e_k = C_A^{\mathrm{SP}} - \hat{C}_{A,k} \tag{3.6}$$

$$\Delta u_k^P = K_c\,(e_k - e_{k-1}) \tag{3.7}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \tag{3.8}$$

$$\Delta u_k^D = -K_c\,\tau_D\,\frac{\hat{C}_{A,k} - 2\hat{C}_{A,k-1} + \hat{C}_{A,k-2}}{\Delta t_c} \tag{3.9}$$

where K_c is dimensionless, τ_I and τ_D are in min and $\Delta t_c = 0.02\,\mathrm{min}$. The proportional and derivative increments are accumulated, the demand is clipped to the bound, and the integral is updated afterwards, as shown in Equation (3.10).

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= u_b + s_k + I_{k-1} \\ C_{Af,k} &= \min\!\left(C_{Af}^{\max},\ \max(0,\ c_k)\right) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{\tau_I}\left(C_{Af,k} - c_k\right) \end{aligned} \tag{3.10}$$

Here s_k accumulates the proportional and derivative increments without limit, c_k is the feed concentration the algorithm demands, the bias u_b is the feed concentration that holds the starting steady state, the integral starts at zero, and the integral used at execution k is the one accumulated through execution $k-1$. In the notation of Equations (1.23) to (1.26), s_k is the running sum S_k and c_k is the output before limiting v_k. The last term is a back-calculation correction with a tracking time equal to τ_I. Away from the bound $C_{Af,k} = c_k$, the correction vanishes, and the recursion is algebraically identical to the positional law $C_{Af,k} = u_b + K_c e_k + (K_c/\tau_I)\sum_{j<k} e_j \Delta t_c$ plus the derivative term. At the bound the two differ, because only the integral is corrected. At the 20 min top of its range the integral time switches integral action off, and then neither increment nor correction is applied in automatic, so the law is proportional and derivative action about $u_b + I$. The derivative acts on the measurement rather than on the error [18], as Section 1.4.8 explains, so a setpoint change reaches the feed only through the proportional term. The controller is reverse acting in the ISA sense (Section 1.4.3), because more solute in the feed means more solute in the outlet. With $e_k = C_A^{\mathrm{SP}} - \hat{C}_{A,k}$ a positive K_c delivers that action, and no action switch is offered.

3.5.2 Initialization and Transfer

The bias and the integral state carry the operating point that a positional law carries in its bias alone. A case that opens in automatic places the tank at the steady state of its setpoint and sets u_b to the feed concentration that holds it, the value of Equation (3.18), with the integral at zero, so its first execution issues Equation (3.11).

$$C_{Af,0} = \min\!\left(C_{Af}^{\max},\ \max(0,\ u_b + K_c\,e_0)\right) \tag{3.11}$$

where u_b is 1.50 mol/L at the default setpoint with the diluent shut, and the error e_0 is the analyzer noise alone. A case that opens in manual places the tank at the steady state of the manual feed u_m and sets $u_b = u_m$. A setting changed before a run starts, whether a setpoint, a flow, a tuning value or the mode, is part of these starting conditions, so the run that follows is the one the same setting gives when carried through Reset. Transfer between manual and automatic is bumpless in both directions by the tracking of Sections 1.4.10 and 1.6.2. In manual the controller sets $I_k = u_m - u_b - s_k$ at each execution so that its demand equals the manual feed, whatever error stands at that moment. The sum $u_b + I$ at the moment of transfer therefore plays the role of the operator-entered bias of a positional law, and of the bias in the offset analysis of Section 3.7. The transfer back to manual holds the last controller output to its full precision. Edits made in automatic act through the increments alone, because the increments are differences of the complete proportional and derivative terms, so a setpoint change moves the feed by $K_c \Delta C_A^{\mathrm{SP}}$ and a gain change on a settled loop moves nothing.

3.5.3 Back-Calculation as the Anti-Windup Mechanism

Anti-windup is the last term of Equation (3.10) rather than a separate stage, and it is one of the remedies that Section 1.4.9 compares on a reachable setpoint. While the feed is clipped, the correction drives the integral toward the value that reproduces the applied feed with a tracking time of τ_I, and the running sum s_k is never truncated [11, 12]. With $\tau_D = 0$ the sum s_k telescopes to $K_c e_k$, so on the upper bound $c_k - C_{Af}^{\max} = u_b + K_c e_k + I_{k-1} - C_{Af}^{\max}$, and the error term of the correction cancels Δu_k^I exactly. What remains is Equation (3.12), the sampled counterpart of Equation (1.18).

$$u_b + I_k - C_{Af}^{\max} = \left(1 - \frac{\Delta t_c}{\tau_I}\right)\left(u_b + I_{k-1} - C_{Af}^{\max}\right) \tag{3.12}$$

The stored output $u_b + I$ therefore relaxes to the bound itself with time constant τ_I whatever the error, and the demand exceeds the bound only by the proportional term $K_c e_k$. Once the stored output has relaxed, the feed leaves the bound when the error changes sign, which is when the outlet crosses the setpoint, and no stored integral has to be worked off. The application offers no free integrator, so it cannot be used to demonstrate the failure mode itself, only its absence, and the size of the failure it avoids has to be computed from the measured error history rather than observed.

That computation quantifies the failure the anti-windup literature addresses [18]. During the constrained upset of Section 3.7 the feed sat on its 2.00 mol/L bound for 9.34 min and accumulated an error integral of 1.559 mol min/L. A positional algorithm with a free integrator and the same $K_I = 2.0/\mathrm{min}$ would have carried 3.12 mol/L of integral demand above the bound at the instant the disturbance was removed, and would have had to accumulate the same 1.559 mol min/L of negative error area before the feed could begin to come down. With the outlet at 2.0 mol/L against a 1.5 mol/L setpoint that unwinding would take 3.1 min of full overshoot. The measured bias and integral sum to 2.0000 mol/L (1.5000 and 0.5000) when the dilution closes, which is the bound itself, and the feed leaves the bound 0.30 min later, when the outlet crosses the setpoint. The comparison is a calculation and not a measurement, because the algorithm it describes is not implemented.

3.5.4 Model Predictive Control

A linear model predictive controller (MPC) can hold the feed concentration in automatic mode in place of the PID, which remains the default of every case except case 9. It follows the general arrangement of Section 1.6.3, and this subsection gives the model, the estimator and the settings it carries for the tank. Every $T_s = 0.04$ min, two executions of the control law, it solves for a sequence of feed concentrations over a horizon, applies the first of them and keeps it until the next decision [19, 20]. The prediction model is Equation (3.13), identified on the simulator at the design point $\bar{C}_A = \bar{C}_{Af} = 1.50\,\mathrm{mol/L}$ with the diluent shut and the analyzer noise off, from the averaged responses to feed steps of $\pm 0.1\,\mathrm{mol/L}$.

$$\tau_m \frac{\mathrm{d}x}{\mathrm{d}t} = -x + K_m\left(C_{Af} - \bar{C}_{Af} + d\right), \qquad \tilde{C}_A = \bar{C}_A + x \tag{3.13}$$

Here x is the model state in mol/L, $\tilde{C}_A$ is the outlet concentration the controller predicts, d is an unmeasured disturbance expressed as an equivalent feed concentration in mol/L, $K_m = 1.000$ and $\tau_m = 0.995$ min. The fit finds no dead time and leaves a root-mean-square residual below 1×10^{-6} mol/L.

The fitted pair repeats the first row of Table 3.3, and the model sampled at Δt_c shrinks a deviation by $0.99^2 = 0.9801$ per execution, exactly as the Euler step of the simulator does with the diluent shut. The model contains no flow. Against 20 m³/min of diluent the tank has $K_p = 0.833$ and $\tau = 0.833$ min, but their ratio $K_p/\tau = q/V$ does not depend on q_d, which is also why the dilution leaves the proportional gain of Equation (3.5) alone. The model is therefore right about the start of a feed move at any dilution and wrong only about where the move settles. A feed step of 1 mol/L raises the predicted outlet by 0.0394 mol/L after one decision period and by 0.982 mol/L at the end of the horizon, against 0.0393 and 0.827 mol/L for the diluted tank.

Offset-free control comes from the estimator. A steady-state Kalman filter on Equation (3.13), sampled with a zero-order hold at Δt_c and extended with d as a random walk, which is the input-disturbance model of offset-free predictive control [21], runs at every execution whichever controller holds the feed. Its update is Equation (3.14).

$$\begin{aligned} \bar{x}_k &= a\,\hat{x}_{k-1} + b\left(C_{Af,k-1} - \bar{C}_{Af} + \hat{d}_{k-1}\right) \\ \varepsilon_k &= \hat{C}_{A,k} - \bar{C}_A - \bar{x}_k \\ \hat{x}_k &= \bar{x}_k + \ell_x\,\varepsilon_k, \qquad \hat{d}_k = \hat{d}_{k-1} + \ell_d\,\varepsilon_k \end{aligned} \tag{3.14}$$

In Equation (3.14) $a = 0.9801$ and $b = K_m(1-a) = 0.0199$, $C_{Af,k-1}$ is the feed applied over the period just ended, $\hat{C}_{A,k}$ is the reading of Equation (3.3) that the PID also receives, and $\ell_x = 0.166$ and $\ell_d = 0.913$ are the converged gains for a measurement variance $r = \sigma_n^2$, a state noise variance of $10^{-6}r$ and a disturbance noise variance $Q_d = r/K_m^2$. A load is booked in $\hat{d}$ as the feed deficit that would cause it, −0.30 mol/L when 1.80 mol/L of feed holds 1.50 mol/L against 20 m³/min of diluent, and the plan holds $\hat{d}$ over its horizon, so a constant load leaves no offset.

Each decision solves a quadratic program on Equation (3.13) sampled at T_s, over a horizon $T_p = 4$ min of $N = 100$ steps, about four model time constants. Twelve feed concentrations v_j, held over blocks of $1,1,1,1,2,2,3,4,6,9,15$ and 55 steps, minimize Equation (3.15) subject to $0 \le v_j \le C_{Af}^{\max}$.

$$J = \frac{T_s}{\tau_n}\sum_{k=1}^{N}\left(\frac{\tilde{C}_{A,k} - C_A^{\mathrm{SP}}}{y_s}\right)^2 + \lambda\sum_{j=1}^{12}\left(\frac{v_j - v_{j-1}}{u_s}\right)^2 \tag{3.15}$$

Here $\tilde{C}_{A,k}$ is the outlet predicted k steps ahead with $\hat{d}$ held, v_0 is the feed applied now, $\tau_n = 1$ min turns the first sum into a time integral, $y_s = 0.05$ mol/L is the settling band of the scorecard, $u_s = 2$ mol/L is the span of the feed and $\lambda = 30$ is a dimensionless move suppression, so the two weights are 16 and 7.5. Only the feed bound constrains the plan, as a hard limit. With no output band the program carries no slack, so an unreachable setpoint shows only as an active bound, and there is no rate limit, terminal weight or preview. Neither controller reads the diluent flow. The decisions are departures from $\bar{C}_{Af}$, a ridge of 10^{-9} relative to the weights keeps the program strictly convex, and a dual active-set quadratic programming code [22] returns its optimum. A rule common to the seven single-loop studio apps fixed λ and Q_d before any comparison, keeping, among the settings whose actuator travel on cases 3 and 4 stays within that of the shipped PID, the largest λ within 10 % of the lowest true-value IAE.

A plan is rejected when the solver reports no optimum or when its first move is not finite or falls outside the bound. The PID is the fallback of Section 1.6.3. It follows the applied feed at every execution as it follows the manual feed, so it takes the feed without a bump and keeps it, with the history marked, until the solver again returns an acceptable plan. Forced at 1.2 min of case 9 with the noise off, a rejection changes the feed by 0.006 mol/L, less than each of the next three PID moves of 0.007 to 0.008 mol/L, and the returning plan moves it by 0.069 mol/L. Because the first plan is

made while a case loads, the MPC is in charge from the first logged sample. It uses the analyzer reading, the feed it applied, the setpoint and the feed limits and nothing else, which keeps the plant of Equation (3.1), the instrument of Equation (3.3) and the model of Equation (3.13) apart, and the exported history gains its estimates and the outlet each plan predicts at the end of its horizon.

3.5.5 Moving Horizon Estimation

An estimator option decides the state from which each plan of the MPC starts. The default is the filter of Equation (3.14). The alternative is a moving horizon estimator (MHE), which at every execution refits the same augmented model to a sliding record of analyzer readings under box bounds on the model state, with the arrival cost of [23]. The filter keeps running and keeps being logged under either choice. At execution k the MHE chooses the augmented state at the start of its window and the noise sequence $w_{k-N_e},\dots,w_{k-1}$ that solve Equation (3.16).

$$\begin{aligned} \min\ & \|z_{k-N_e}-\bar{z}_{k-N_e}\|^2_{\bar{P}^{-1}}+\sum_{j=k-N_e}^{k-1} w_j^{\mathsf{T}} Q^{-1} w_j \\ & +\frac{1}{r}\sum_{j=k-N_e}^{k}\left(\hat{C}_{A,j}-\bar{C}_A-x_j\right)^2 \\ \text{s.t. } & z_{j+1}=\begin{bmatrix} a & b \\ 0 & 1 \end{bmatrix} z_j+\begin{bmatrix} b \\ 0 \end{bmatrix}\left(C_{Af,j}-\bar{C}_{Af}\right)+w_j \\ & 0\le \bar{C}_A+x_j\le C_{Af}^{\max} \end{aligned} \tag{3.16}$$

In Equation (3.16) $z_j=[x_j\ \ d_j]^{\mathsf{T}}$ stacks the state and the disturbance of Equation (3.13), a and b are the coefficients of Equation (3.14), $C_{Af,j}$ is the feed applied over the period that ends at execution $j+1$, w_j is the process noise with covariance $Q=\operatorname{diag}(10^{-6}r,\,Q_d)$, and $r=\sigma_n^2$, all in the units of Equation (3.14). The arrival term ties the start of the window to $\bar{z}_{k-N_e}$, the estimate the MHE made one execution earlier carried one step through the model, through the steady a-priori covariance $\bar{P}$ of the filter, which is held fixed. With $N_e=10$ the window spans 0.2 min and eleven readings, so each solve has 22 variables and 22 bound rows and goes to the dual active-set code that solves the plan.

The upper bound is the feed limit $C_{Af}^{\max}$, a constant of the source, because an outlet fed at most that strength and diluted only by a stream free of solute cannot exceed it, and a nonnegative feed keeps the outlet of Equation (3.1) at or above zero. No declared value limits d, so it is left free. Unless a bound is active, the window reproduces the recursive estimate of the filter, and the option can change a plan only where the filter would put the outlet outside $[0,C_{Af}^{\max}]$. A window that the solver does not bring to an optimum, or whose estimate is not finite or leaves the bounds, is replaced by the Kalman correction of the one-step prediction of the MHE, clipped to the bounds, and the exported history marks the replacement. The estimator reads the analyzer, the applied feed and declared constants and nothing else, so a parameter edited during a run reaches it only through the reading.

3.5.6 Relay Autotuning

The PID constants can also be taken from an experiment on the running loop, the relay feedback test of Åström and Hägglund [24], which Section 1.4.7 introduces. A button starts it from the present state whenever the PID holds the feed in automatic, and no case starts with it engaged. The relay sets the feed to $u_0-\delta_r$ when the reading exceeds the setpoint by more than ε_r and to $u_0+\delta_r$ when it falls short by more than ε_r. Here u_0 is the mean feed applied over the last 0.5 min, a tenth of the trend window, δ_r is 0.1 mol/L, five percent of the feed range, unless 0.9 of the distance from u_0 to the nearer bound is smaller, and $\varepsilon_r=3\sigma_{n,0}=0.03\,\text{mol/L}$ is three times the default analyzer noise $\sigma_{n,0}$, whatever noise the run carries. A cycle counts only if each of its halves lasts four executions or more, and the test converges once two successive counted cycles have periods, and swings $\pm\alpha$, within 5 % of each other. The describing function of the hysteretic relay then supplies the ultimate gain,

which the Tyreus–Luyben rule of Table 1.2 turns, together with the period P_u, into the PI settings of Equation (3.17), in which D is the fraction of each cycle spent at $u_0+\delta_r$, one half for a symmetric cycle. The first expression is Equation (1.16) with the factor $\sin(\pi D)$ added for a cycle that is not symmetric.

$$K_{cu}=\frac{4\delta_r\sin(\pi D)}{\pi\sqrt{\alpha^2-\varepsilon_r^2}},\qquad K_c=\frac{K_{cu}}{3.2},\qquad \tau_I=2.2\,P_u \tag{3.17}$$

Both gains are dimensionless, and P_u and τ_I are in min. The settings are rounded to three significant figures with $\tau_D=0$, each is held inside its input range, τ_I below the 20 min that switches integral action off, and they remain editable. A test that does not converge ends after 30 cycles or 10 min, and a second press, a change to manual or a change to the MPC cancels it. The PID follows the applied feed during the test as it does in manual, so it resumes without a bump. The result note gives $\psi=\arcsin(\varepsilon_r/\alpha)$, the phase lag that the band adds to the relay, so that the cycle settles where the loop phase is $-180^\circ+\psi$, and from $\psi\geq 10^\circ$ on it warns that the band rather than the tank has set the cycle.

3.6 Setpoint Reachability

3.6.1 The Steady-State Inverse and Its Range

Setting the derivative in Equation (3.1) to zero and solving for the feed concentration gives the exact steady-state inverse of the plant, as shown in Equation (3.18).

$$C_{Af}^{\star}=\frac{q+q_d}{q}\,C_A^{\mathrm{SP}} \tag{3.18}$$

Equation (3.18) is the value at which the stored output u_b+I of Equation (3.10) comes to rest, and it is the only model of the tank the PID ever holds. Holding a composition against dilution requires feed strength in direct proportion to the dilution ratio, and the expression contains no dynamic term. The application does not display the required value, but once a reachable case has settled the applied feed concentration equals it. The feed concentration is flagged whenever it sits on a bound, which turns the test of Equation (3.19) into a visible condition rather than an inference from a trend.

The manipulated variable is bounded, so Equation (3.18) defines a reachable set rather than a solution. A setpoint is attainable at steady state only when the required feed concentration falls inside the actuator range, which gives the three equivalent statements of Equation (3.19).

$$C_A^{\mathrm{SP}}\leq\frac{qC_{Af}^{\max}}{q+q_d},\quad q_d\leq q\left(\frac{C_{Af}^{\max}}{C_A^{\mathrm{SP}}}-1\right),\quad \frac{q_d}{q}\leq\frac{C_{Af}^{\max}}{C_A^{\mathrm{SP}}}-1 \tag{3.19}$$

With the default $q=100\,\mathrm{m^3/min}$, $C_{Af}^{\max}=2.0\,\mathrm{mol/L}$ and $C_A^{\mathrm{SP}}=1.5\,\mathrm{mol/L}$, the loop rejects any diluent flow up to 33.3 $\mathrm{m^3/min}$ and no more. Beyond that the highest outlet concentration the plant can produce is $qC_{Af}^{\max}/(q+q_d)$, and the standing error is the difference of Equation (3.20).

$$e_\infty=C_A^{\mathrm{SP}}-\frac{qC_{Af}^{\max}}{q+q_d} \tag{3.20}$$

At $q_d=50\,\mathrm{m^3/min}$ the ceiling is 1.333 mol/L and the standing error is 0.167 mol/L, which is 11.1 % of the setpoint. Neither expression contains a controller gain, a sample period or a tracking time.

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

3.6.2 Open-Loop Behavior During Saturation

A pinned actuator opens the loop. With C_{Af} held at its bound, Equation (3.1) is the same first-order equation it is in manual mode, so the approach to the constrained steady state has the open-loop

residence time $\tau = V/(q+q_d)$ regardless of how the controller is tuned. The trajectory during saturation is a plant property. Fitting the pinned portion of the two constrained runs of Section 3.7 returns $\hat{\tau} = 0.6617\,\text{min}$ at $K_c = 2.0$ and $\hat{\tau} = 0.6617\,\text{min}$ at $K_c = 8.0$, against the 0.6667 min of $V/(q+q_d)$ and against the $\tau_p/(K_pK_c)$ of 0.50 and 0.125 min that the same two tunings would have imposed had the feed been free to move. The two fits agree to five decimals at 0.66165 min, which is the discrete time constant of the Euler step at $\tau = 0.6667\,\text{min}$, so the controller leaves no trace on the pinned trajectory.

> **TRY IT IN THE APP** **8 · Actuator Limit**
>
> Load the case and run it to 5 min. Against a standing dilution of 20 m^3/min the setpoint steps to 1.9 mol/L at 1 min. Watch the feed trend reach its 2.0 mol/L limit and stay there for 80 % of the run, while the outlet stops at 1.6653 mol/L. The predictive controller ends at the same outlet concentration.

3.7 Results

Each experiment starts from the reference case, configured by hand from the operating inputs of Table 3.2 before the run, so the run starts at the steady state those settings define, and every step or upset is applied 1 min into the run, the instant at which the shipped cases apply their first. The exported file (Section 1.6.4) records, every 0.1 min, the time, mode, case, setpoint, analyzer reading, noise-free outlet concentration, applied feed, limit flag and both disturbances. Times, peaks and integrated errors quoted below are evaluated at the 0.01 min step of the same runs, which the exported rows reproduce to within their spacing and 0.007 mol min/L. Table 3.5 lists the sequence.

The application also ships ten cases, listed with their completion requirements in Table 3.4, and each hand-configured experiment of Table 3.5 starts from the first of them with its settings changed before the run. Every case starts at steady operation, and each case with a scripted step or upset applies the first one 1 min into the run. Completion is judged as Section 1.6.4 describes. The challenge is judged on the case as given, so changing a model parameter or a disturbance during the attempt voids it. A setting changed before the run gives the same run as the same setting carried through Reset, while reloading the case restores its shipped settings.

3.7.1 Startup From an Empty Tank

Every case starts at steady operation, so a startup is built from the reference case by setting the manual feed to zero before the run, which starts the tank empty at the steady state of zero feed, with the default $K_c = 2.0$, $\tau_I = 1.0\,\text{min}$, $\tau_D = 0$ and the noise switched off. How the loop is then closed decides the approach. Switching to automatic 1 min into the run with the 1.5 mol/L setpoint already in force is a bumpless transfer onto a standing error, so the feed starts from zero and is raised by integral action alone. It never reaches the bound, the outlet comes within 5, 2 and 1 % of setpoint 3.63, 4.54 and 5.23 min after the switch with no overshoot, and the feed comes to rest at exactly 1.5000 mol/L. Setting that feed by hand instead reaches the same bands in 2.98, 3.89 and 4.58 min, so this transfer is slower than no feedback at all. Holding the setpoint at zero until the switch and then raising it to 1.5 mol/L applies the proportional kick of Equation (3.7) instead. The demand of $2.0 \times 1.50 = 3.00\,\text{mol/L}$ puts the feed on its upper bound at the first execution, where it stays for 0.70 min while the stored output relaxes toward the bound as Equation (3.12) describes and reaches 1.03 mol/L, and the outlet comes within 5, 2 and 1 % in 1.61, 2.05 and 2.38 min, passing the setpoint by less than 0.0001 mol/L. Even that time exceeds the four closed-loop time constants that $\tau_c = 0.5\,\text{min}$ would suggest, because the first part of the approach is made at the bound at the open-loop rate. A nominal startup can touch the constraint, and the way automatic is engaged rather than the tuning decides how fast it goes.

Table 3.4 Scenarios and Completion Requirements

	Scenario	Kind	Required to complete
1	Reference operating point	Experiment	Hold the tank in manual at steady operation, outlet and feed at 1.50 mol/L, for 1.3 min and note which quantities stay stationary.
2	Open-loop actuator step	Experiment	In manual from steady operation, run through the feed step from 1.50 to 1.90 mol/L at 1 min and on to 3.5 min, reading the direction, size and speed of the outlet response.
3	Closed-loop setpoint step	Experiment	In automatic from steady operation, run through the setpoint step from 1.5 to 1.7 mol/L at 1 min until the outlet stays within ±0.05 mol/L of the new setpoint for 0.25 min.
4	Disturbance rejection	Experiment	In automatic, run through the dilution step from 0 to 20 m^3/min at 1 min and on to 3.5 min.
5	Recover by hand	Exploration	After the same dilution step in manual, move the feed by hand to a setting that holds the outlet, once settled, within half of its largest error since the step (or within ±0.05 mol/L), and keep the outlet there for 0.25 min in manual.
6	Conservative tuning	Experiment	With $K_c = 0.8$, run through the setpoint step from 1.5 to 1.66 mol/L at 1 min and on to 3.5 min.
7	Aggressive tuning	Experiment	With $K_c = 4.0$, run through the same setpoint step and on to 3.5 min.
8	Actuator limit	Demonstration	Against a standing dilution of 20 m^3/min, with the feed at 1.80 mol/L, run through the setpoint step from 1.5 to 1.9 mol/L at 1 min until the feed reaches its 2.0 mol/L limit and on to 1.7 min, the outlet then able to reach only 1.67 mol/L.
9	MPC at the edge of the feed	Experiment	Under the MPC from $t = 0$, run 5 min through the dilution step from 0 to 20 m^3/min at 1 min, the setpoint step to 1.70 mol/L at 2 min and its return to 1.60 mol/L at 3.5 min, then run the case again from $t = 0$ under the PID and compare the two runs.
10	Challenge	Challenge	Starting at steady operation from $K_c = 0.4$ and $\tau_I = 1.0$ min, with the dilution stepping from 0 to 20 m^3/min at 1 min, hold the analyzer reading within ±0.12 mol/L of 1.5 mol/L from 2.64 to 4.0 min, with the model parameters and disturbances as given.[a]

[a] With the shipped $K_c = 0.4$ the worst reading error over the scored window is 0.159 mol/L, 1.33 times the band, and the outlet is still 0.114 mol/L low at 4.0 min. With the suggested $K_c = 2.0$ and $\tau_I = 1.0$ min, revealed before the run or carried through Reset, which give one run, it is 0.053 mol/L, inside the band with 0.067 mol/L to spare. Most of it is the end of the recovery rather than analyzer noise, because the true concentration is still up to 0.048 mol/L below setpoint over the same window.

Table 3.5 Experiment Sequence and the Mechanism Each One Exposes

Experiment	Setup	What it exposes
Step test	manual, C_{Af} 1.0 to 1.5 at four flow settings	$K_p = q/(q+q_d)$ and $\tau = V/(q+q_d)$
Startup	empty tank, automatic engaged on the error or by a setpoint step	how engaging the loop sets the approach
Reachable upset	q_d 0 to 25 m^3/min	integral action removes the offset, feed rests at 1.875
Boundary	q_d = 33 then 34 m^3/min	Equation (3.19) closes between two adjacent settings
Unreachable upset	q_d 0 to 50 m^3/min at K_c = 2 and 8	the same offset at two tunings
Recovery	q_d back to 0	off the bound as the outlet crosses setpoint
Proportional only	integral off, q_d = 33 and 50 m^3/min	an offset a controller does remove
Analyzer noise	$\sigma_n = 0.01$ mol/L, K_c = 2 and 8	feed activity proportional to K_c
Edits in automatic	C_A^{SP} 1.50 to 1.45, K_c 2.0 to 2.5	increments act, nothing is reseeded

3.7.2 Reachable and Unreachable Dilution

Table 3.6 reports four dilution upsets applied to a settled loop at 1.5 mol/L, with the analyzer noise switched off so that the offsets reported are the deterministic ones. Each column is the 10 min

Table 3.6 Dilution Upsets Applied to a Settled Loop at $C_A^{SP} = 1.5\,\text{mol/L}$ With $\tau_D = 0$ and $\sigma_n = 0$, Measured Over the 10 min Following the Step

q_d (m³/min)	25	33	34	50
K_c, τ_I (min)	2, 1	2, 1	2, 1	2, 1
$C_{Af}^{\star}$ required (mol/L)	1.875	1.995	2.010	2.250
Reachable by Eq. (3.19)	yes	yes	no	no
Final C_{Af} (mol/L)	1.8750	1.9949	2.0000	2.0000
Final C_A (mol/L)	1.4999	1.4999	1.4925	1.3333
Offset (mol/L)	0.0001	0.0001	0.0075	0.1667
Predicted offset, Eq. (3.20)	—	—	0.0075	0.1667
Maximum $\lvert e \rvert$ (mol/L)	0.0898	0.1165	0.1198	0.1703
IAE (mol min/L)	0.187	0.247	0.288	1.643
Time on the bound (min)	0.00	0.00	6.48	9.34

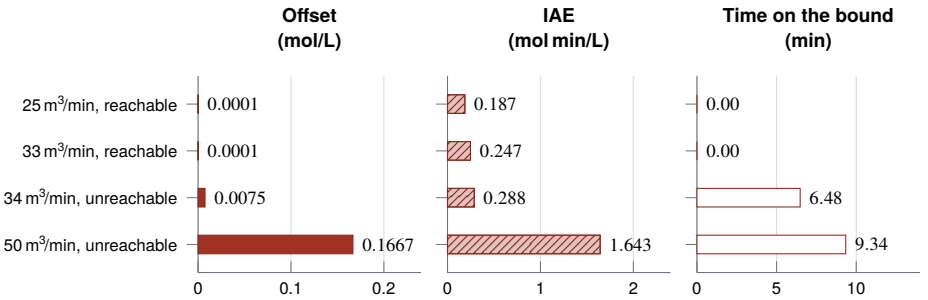


Figure 3.2 The dilution upsets of Table 3.6 under one tuning: the offset left after ten minutes (solid bars), the integrated absolute error (hatched bars) and the time the feed spends on its bound (open bars). Notice that one more m³/min of diluent, from 33 to 34, puts the feed on its bound for most of the window and leaves an offset, and that the offset at 50 is the bound's, not the controller's.

following a step in the diluent flow. The disturbance rejection case ships a reachable upset of 20 m³/min, shown in Figure 3.3.

> **TRY IT IN THE APP** **4 · Disturbance Rejection**
>
> Load the case and press Run. Watch the outlet trend when the diluent opens at 1 min. The outlet falls to 1.428 mol/L at 1.66 min and stays within ±0.05 mol/L of the setpoint after 2.55 min, while the feed trend climbs toward the 1.80 mol/L that the new dilution requires without reaching its bound.

The reachable set closes between two adjacent values of a diluent flow whose resolution is 1 m³/min. At $q_d = 33\,\text{m}^3/\text{min}$ the required feed concentration is 1.995 mol/L, the feed comes to rest at 1.9949 mol/L and the outlet returns to 1.4999 mol/L with an integrated absolute error of 0.247 mol min/L. At 34 m³/min the requirement is 2.010 mol/L, the feed spends 6.48 min of the ten on its bound and the outlet settles 0.0075 mol/L low, matching Equation (3.20) to four decimals. The loop has not changed and its tuning has not changed. The inequality of Equation (3.19) has simply stopped being satisfied. Figure 3.2 draws the table as bars.

Table 3.7 repeats the largest upset at two tunings that differ by a factor of four in both gains, with τ_I held at 1.0 min so that K_I rises with K_c.

Quadrupling both gains changes the final outlet concentration by nothing at all. Both tunings settle at 1.3333 mol/L against a 1.5 mol/L setpoint, both leave an offset of exactly 0.1667 mol/L, and the integrated absolute error over the ten minutes improves only from 1.643 to 1.578 mol min/L, which is

Table 3.7 The Same $q_d = 50\,\mathrm{m}^3/\mathrm{min}$ Upset and Its Removal at Two Tunings, $\tau_I = 1.0\,\mathrm{min}$, $\sigma_n = 0$. The Upset Window Is 10 min and the Recovery Window Is 8 min.

	$K_c = 2.0$	$K_c = 8.0$
Diluent opened to 50 m³/min		
Final C_A (mol/L)	1.3333	1.3333
Offset (mol/L)	0.1667	0.1667
Maximum $\lvert e \rvert$ (mol/L)	0.1703	0.1667
IAE (mol min/L)	1.643	1.578
Time on the bound (min)	9.34	9.86
Fitted $\hat{\tau}$ while pinned (min)	0.6617	0.6617
Diluent closed again		
Delay off the bound (min)	0.30	0.30
Peak C_A (mol/L)	1.6266	1.5476
Overshoot (%)	8.44	3.17
Return within 1 % (min)	3.72	1.83
IAE (mol min/L)	0.272	0.084

3.9 % for four times the gain. The benefit of the additional gain is concentrated in the initial transient and in the recovery. The aggressive tuning reaches the bound sooner and spends 9.86 min of the ten there against 9.34 min, and after the diluent valve is closed it returns to within 1 % of setpoint in 1.83 min against 3.72 min, with 3.17 % overshoot against the 8.44 % of the default tuning. Both tunings hold the feed on the bound for 0.30 min after the dilution closes, until the outlet crosses the setpoint, because Equation (3.12) has left both stored outputs at the bound. The recovery overshoot is therefore not windup. The reachable upset at 33 m³/min, which never touches the bound but leaves the stored output at 1.995 mol/L, gives a recovery peak of 1.6253 mol/L against the 1.6266 mol/L that follows the unreachable one.

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

3.7.3 Two Offsets of Similar Size

A second standing error, close in size to the first, arises by a mechanism a controller does remove. A proportional loop whose stored output $u_b + I$ is frozen at b has the steady state of Equation (3.21) while its demand stays inside the range. Equation (1.3) is the same balance for a setpoint step taken from a bias that held the old setpoint.

$$C_A^{\infty} = \frac{K_p\,(b + K_c\,C_A^{\mathrm{SP}})}{1 + K_p\,K_c}, \qquad e_{\infty} = \frac{C_A^{\mathrm{SP}} - K_p\,b}{1 + K_p\,K_c} \tag{3.21}$$

Setting the integral time to its 20 min top before the run switches integral action off, and a run started in automatic at the 1.5 mol/L setpoint then has $b = u_b = 1.50\,\mathrm{mol/L}$, the value of Equation (3.18) with the diluent shut, so Equation (3.21) predicts no offset at all until a load arrives, which is the purpose of the bias in a positional law. With $K_c = 2.0$ and the noise switched off, opening the diluent to 33 m³/min gives $K_p = 0.752$, and Equation (3.21) predicts 1.3514 mol/L with the feed at 1.797 mol/L. The run settles there to four decimals, an offset of 0.149 mol/L that the same upset with $\tau_I = 1.0\,\mathrm{min}$ removes entirely (Table 3.6). At 50 m³/min the proportional loop stands at 1.2857 mol/L, an offset of 0.214 mol/L, with the feed at 1.929 mol/L and still inside its range. Restoring $\tau_I = 1.0\,\mathrm{min}$ 5 min after the step puts the feed on its bound 0.20 min later, and the outlet rises only to the 1.3333 mol/L of Equation (3.20). Quadrupling the gain instead lowers the offset that Equation (3.21) predicts to 0.079 mol/L, but its 2.13 mol/L demand exceeds the bound, and the outlet again stops at 1.3333 mol/L.

The two experiments therefore report offsets of similar size against the same 2.0 mol/L feed ceiling, and they are not the same phenomenon. The offset of Equation (3.21) is the finite loop gain



Figure 3.3 A reachable dilution upset rejected by feedback, in the disturbance rejection case as shipped, with the default $K_c = 2.0$, $\tau_I = 1.0\,\text{min}$, $\tau_D = 0$ and the default 0.01 mol/L analyzer noise, at $t = 5.00\,\text{min}$. The loop starts in automatic at steady operation, outlet and feed at 1.50 mol/L, and the diluent stream opens to 20 m^3/min at $t = 1.0\,\text{min}$. Equation (3.18) then requires 1.80 mol/L of feed, inside the range, so the feed never reaches its bound (its peak is 1.87 mol/L). The outlet falls to 1.428 mol/L at 1.66 min and stays within ±0.05 mol/L of setpoint after 2.55 min. At the instant shown the outlet is 1.494 mol/L, the analyzer reads 1.504 mol/L, the feed is 1.80 mol/L and the scorecard reads an IAE of 0.150 mol min/L, 1.17 mol/L of feed travel and no saturation. The residence time is 0.83 min because the throughput has risen to 120 m^3/min. The screenshot is cropped to the plant view, with the scorecard tiles moved into a frame at its right, above the trends. Callouts: (1) the outlet readout, (2) the analyzer, which the app labels CT, (3) the scorecard, (4) the low point of the outlet trend after the dilution step at the dotted marker, and (5) the feed trend climbing to carry the new dilution.

of a proportional controller and is removed by integral action or reduced by more gain, each of which the application demonstrates within minutes, until the feed reaches its bound. The offset of Equation (3.20) contains no controller parameter and is removed only by more feed strength, less dilution or a lower setpoint. At 50 m^3/min the two meet, and integral action removes the 0.048 mol/L by which the proportional offset exceeds the reachability limit and nothing more.

Table 3.8 Feed Activity of a Settled Loop With $\sigma_n = 0.01\,\text{mol/L}$, $q_d = 0$ and $\tau_D = 0$

K_c, τ_I (min)	2.0, 1.0	8.0, 1.0
σ of $\hat{C}_A$ (mol/L)	0.01068	0.01103
σ of C_A (mol/L)	0.00121	0.00297
σ of C_{Af} measured (mol/L)	0.02119	0.08739
σ of C_{Af} from $K_c\sigma_n$ (mol/L)	0.02000	0.08000
Feed travel (mol/L per min)	1.22	5.14

DID YOU KNOW? More Gain Buys No Feed Strength

From the simulator: against 50 m^3/min of diluent, a gain of 8.0 leaves the outlet at 1.3333 mol/L, exactly where a gain of 2.0 leaves it, and trims the integrated absolute error only from 1.643 to 1.578 mol min/L (Table 3.7).

3.7.4 Analyzer Noise, With and Without the Bound

Summing Equation (3.7) over a run telescopes the proportional increments, so with $\tau_D = 0$ the instantaneous noise contribution to the feed concentration is $-K_c n_k$ with n_k the noise sample at step k, and the standing feed noise is $\sigma_u = K_c\sigma_n$ rather than a random walk. Table 3.8 reports a loop started in automatic at 1.5 mol/L with the default 0.01 mol/L analyzer noise over the first 6 min of the run with no disturbance applied, sampled at every controller execution, because the exported rows see one execution in five and record only about a quarter of the feed travel.

The telescoped estimate falls 5.6 % below the measured feed noise at $K_c = 2.0$ and 8.5 % below it at $K_c = 8.0$. Part of that deficit belongs to the noise itself, because the 300 draws in the window have a standard deviation of 0.01054 mol/L, and with that value the estimate falls only 0.5 and 3.5 % below. The remainder underpredicts because the loop has started to move the tank in response to its own noise. The evidence is in the second row, where the true concentration standard deviation rises from 0.00121 to 0.00297 mol/L while the noise draws are identical. Quadrupling the gain therefore multiplies the feed travel by 4.22 and the concentration variability the loop is meant to suppress by 2.46 (2.45 to 2.70 over three 6 min windows of the same run), which is the ordinary argument against aggressive tuning on a measurement-limited loop.

Noise does not move the saturated offset. Repeating the $q_d = 50\,m^3/\text{min}$ upset with the analyzer noise left at its default puts the feed exactly on 2.000 mol/L at every controller execution from 0.70 min after the step (93.2 % of the 10 min window, the rest being the approach), and the mean outlet concentration over the last 6 min is 1.33333 mol/L, the noise-free value. The stored output sits at the bound by Equation (3.12), and over those 6 min the proportional term holds the demand at least 0.27 mol/L above it, so no noise sample of the default size brings the demand back inside the range and nothing is clipped on one side only.

3.7.5 Edits in Automatic

Editing the setpoint or a gain while the loop is in automatic acts through the increments of Equations (3.7) to (3.9) alone, as described in Section 3.5, and nothing is reseeded. Lowering the setpoint from 1.50 to 1.45 mol/L on a settled loop with the noise switched off lowers the feed concentration from 1.5000 to 1.4000 mol/L at the next execution, the $K_c\,\Delta C_A^{\text{SP}}$ that Equation (3.7) gives, and the outlet then moves to the new setpoint without passing it, coming within 0.005 mol/L of it in 1.12 min. Raising K_c from 2.0 to 2.5 with the setpoint untouched leaves the feed at 1.5000 mol/L and the outlet where it was, because the error is zero when the gain changes. Two tunings can therefore be compared within one run without an excursion caused by the edit itself.

Table 3.9 Five Cases Under the PID and Under the Predictive Controller, Each Pair Run With Identical Analyzer Noise

Case	Controller	IAE (mol min/L)	IAE_C	e_{max} (mol/L)	Travel (mol/L)	Limit (%)	Bound
3, setpoint	PID	0.111	0.089	0.162	1.44	0	—
	MPC	0.082	0.053	0.163	1.30	4	8
4, dilution	PID	0.150	0.143	0.071	1.17	0	—
	MPC	0.061	0.032	0.048	0.91	0	0
8, limit	PID	1.066	1.062	0.379	0.33	80	—
	MPC	1.068	1.063	0.381	0.25	80	100
9, edge	PID	0.281	0.272	0.240	1.20	30	—
	MPC	0.207	0.197	0.184	0.99	30	38
10, challenge	PID, $K_c = 0.4$	0.378	0.372	0.149	0.24	0	—
	PID, $K_c = 2^a$	0.138	0.133	0.071	0.94	0	—
	MPC	0.053	0.030	0.048	0.79	0	0

IAE uses the analyzer reading, as the scorecard does, and IAE_C the true concentration. e_{max} is the largest true deviation from setpoint after the first event, Travel is the summed absolute change of the feed, Limit is the fraction of the run spent with the feed at a bound, and Bound counts the plans, of 125 (100 in the challenge), with the feed bound active. [a]Suggested, $\tau_I = 1.0$ min.

3.7.6 Predictive Control on the Shipped Cases

Table 3.9 runs cases 3, 4, 8 and 9 for 5 min and the challenge for its 4 min under each controller. The controller is selected before each run starts, so both runs of a case meet the same events and the same analyzer noise draws, and the noise columns of the two runs of case 9 agree to the last digit. The PID keeps the shipped $K_c = 2.0$, $\tau_I = 1.0$ min and $\tau_D = 0$ in cases 3, 4, 8 and 9, the challenge ships $K_c = 0.4$ and suggests $K_c = 2.0$ with $\tau_I = 1.0$ min, while the MPC runs every case with the settings of Section 3.5.4. Unlike the preceding subsections, the numbers here are read on the 0.1 min log, the samples from which the scorecard of the application computes them.

At the setpoint step of case 3 the model reproduces the plant apart from the noise, and the MPC uses the bound. Its first eight plans after the step have the upper bound active, the feed rests on 2.00 mol/L for 0.20 min, and the outlet stays within ±0.01 mol/L of 1.70 mol/L from 0.60 min after the step, where the PID needs 1.60 min. The MPC passes the new setpoint by 0.005 mol/L and the PID by 0.001 mol/L, and the true-value IAE is 0.053 against 0.089 mol min/L with 9 % less feed travel.

The dilution of case 4 activates the bound in none of the 125 plans, so the whole difference there belongs to the model and its estimate. The estimate $\hat{d}$ reaches 90 % of the −0.30 mol/L deficit within 0.5 min of the upset and ends the case at −0.297 mol/L. The true outlet falls to 1.452 mol/L under the MPC and to 1.429 mol/L under the PID, and it returns for good within ±0.02 mol/L of setpoint 0.70 min after the upset against 2.70 min. The scorecard IAE falls from 0.150, the value of Figure 3.3, to 0.061 mol min/L, and the feed travel from 1.17 to 0.91 mol/L. A PID at $K_c = 8.0$, the aggressive tuning of Table 3.7 and not a shipped one, reaches a similar 0.058 only with 4.53 mol/L of travel, five times that of the MPC.

Case 8 is a tie. Both controllers reach the bound at the execution at 1.00 min and hold it for 80 % of the run, both leave the outlet at 1.6653 mol/L at 5 min, and their scorecard IAEs, 1.066 and 1.068 mol min/L, differ by 0.1 %. All 100 plans after the step have the bound active, $\hat{d}$ ends at −0.331 mol/L against the −0.333 mol/L that the feed at its bound implies at the 1.667 mol/L ceiling of Equation (3.19), and the last plan predicts 1.671 mol/L at the end of its horizon. The MPC therefore plans with the setpoint out of reach and gains nothing from doing so. Its lower travel accrues in the quiet minute before the step, 0.08 against 0.21 mol/L.

Case 9 joins the dilution of case 4 to a setpoint of 1.70 mol/L at 2 min, which Equation (3.18) prices at 2.04 mol/L of feed, and to a reachable 1.60 mol/L at 3.5 min, which needs 1.92 mol/L, as Figure 3.4 shows. Before the setpoint step the MPC rejects the dilution as in case 4, with a true-value IAE over the first 2 min of 0.024 against 0.061 mol min/L, and meets the step with the outlet at 1.497 mol/L, where the PID has 1.434 mol/L. Both feeds reach the bound at the execution at 2.00 min

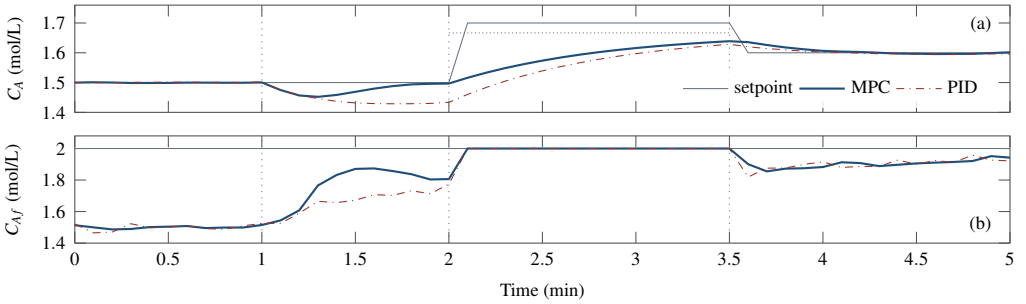


Figure 3.4 The two controllers on case 9, with identical analyzer noise draws, from the 0.1 min log. The MPC is the thick solid line and the PID the dash-dotted line. The diluent opens to 20 m^3/min at 1 min, the setpoint steps to the unreachable 1.70 mol/L at 2 min and returns to 1.60 mol/L at 3.5 min (vertical dotted lines). (a) True outlet concentration and setpoint (thin gray line), with the 1.667 mol/L ceiling of Equation (3.19) dotted between 2 and 3.5 min. (b) Feed concentration and its 2.0 mol/L bound (thin gray line).

and stay on it until the setpoint returns, so for 1.5 min the tank runs open loop under either controller, and the 0.063 mol/L gap present at 2 min shrinks by the Euler factor $1 - \Delta t\,(q + q_d)/V = 0.988$ per step. That decay, summed over the logged samples, is 0.0411 mol min/L, which is all of the difference between the true-value IAEs of the two runs over the interval, 0.161 and 0.202 mol min/L. Every one of the 38 plans in the interval has the bound active, and at 3.0 min the plan already predicts 1.658 mol/L at the end of its horizon, close to the ceiling, yet the MPC gains nothing while the feed is pinned.

When the setpoint returns, the PID leaves the bound at the execution at 3.50 min and the MPC at its next decision, 0.02 min later. The stored output $u_b + I$ of the PID is then 1.916 mol/L, below the bound, because Equation (3.12) had been drawing it up from 1.624 mol/L after the execution at 2.00 min and predicts the measured value to the fifth decimal. No integral demand above the bound remains to be worked off, so the windup that a constraint inside the plan is meant to prevent is already absent from the PID. The PID drops the feed to 1.886 mol/L at that execution, while the move suppression $\lambda = 30$ makes the MPC lower it by 0.069, 0.030 and 0.026 mol/L over its first three decisions. The PID is within ±0.02 mol/L of 1.60 mol/L for good 0.20 min after the return and the MPC after 0.30 min, and over the last 1.5 min the MPC loses, 0.012 against 0.009 mol min/L of true-value IAE. Over the whole run it wins, 0.207 against 0.281 mol min/L on the scorecard, and the whole margin traces to the rejection of the dilution before the bound was reached.

The MPC reads none of the PID settings, so its challenge run repeats the first 4 min of its case 4 run. Its worst reading error over the judged window from 2.64 to 4.0 min is 0.018 mol/L, against 0.159 mol/L for the shipped $K_c = 0.4$, which fails the ±0.12 mol/L band, and 0.053 mol/L for the suggested tuning, which passes.

All 600 plans were optimal, so no fallback occurred, and plans with the bound active needed up to 12 active-set iterations. The mean logged solve time was 0.05 to 0.30 ms per case in four Node.js runs and 237 and 253 µs in two Chrome runs of case 9 (largest 2.1 ms), 0.6 % of one processor core at one process minute per second and at most 6.3 % at the 10× top speed, wall-clock values from one desktop computer that vary between runs. Across the five cases the MPC wins the setpoint step, the dilution and the challenge, ties case 8, and in case 9 loses only the 1.5 min after the setpoint returns, so at the feed bound its advantage over this PID is its model of the load rather than its handling of the bound.

Table 3.10 Cases 3 and 4 Over 15 min Under the Shipped PI, Two Counterfactual PI Settings, the Relay-Tuned PI and the MPC, With One Noise Seed

Case	Controller	IAE (mol min/L)	IAE_C	$e_{\max}$ (mol/L)	Travel (mol/L)	Travel^0
3, setpoint	PI, shipped	0.194	0.108	—	3.86	0.54
	PI^a	0.205	0.123	—	3.33	0.45
	PI^b	0.243	0.183	—	3.86	0.54
	PI, relay	0.267	0.209	—	3.33	0.46
	MPC	0.167	0.088	—	2.60	0.82
4, dilution	PI, shipped	0.233	0.159	0.072	3.59	0.30
	PI^a	0.253	0.181	0.078	3.10	0.30
	PI^b	0.325	0.272	0.079	3.59	0.30
	PI, relay	0.363	0.313	0.086	3.10	0.30
	MPC	0.145	0.058	0.048	2.19	0.50

PI, shipped: $K_c = 2.0$, $\tau_I = 1.0$ min. PI^a: the written $K_c = 1.73$ with the shipped τ_I. PI^b: the shipped K_c with the written $\tau_I = 1.83$ min. PI^a and PI^b are counterfactuals that the application neither ships nor writes. PI, relay: $K_c = 1.73$, $\tau_I = 1.83$ min. Columns as in Table 3.9, with IAE_C on the 0.01 min step and $e_{\max}$ after the dilution step. Travel^0 repeats Travel with $\sigma_n = 0$. Either estimator gives the MPC rows.

3.7.7 Estimator Choice and Relay Autotuning

The runs of this subsection last 15 min, three trend windows, with the challenge at its 4 min, the shipped noise seed and every setting made before the run, so their integrals are larger than those of Table 3.9. In a run of case 9 under the MPC, the outlet estimates of the MHE and of the filter differ by no more than $2.6 \times 10^{-10}\sigma_n$ over 749 executions, and the disturbance estimates by 4.1×10^{-11} mol/L. This residue does not come from rounding. The Riccati iteration that produced ℓ_x and ℓ_d stopped after 111 steps at a relative distance of 2.6×10^{-10} from its fixed point, the window repeats the covariance update that the fixed gain of the filter leaves out, and with the fixed-point covariance the two estimates agree to $1.4 \times 10^{-14}\sigma_n$. Run under the MPC with one estimator and then the other, the ten cases give feed records within 2.1×10^{-11} mol/L of each other, with no bound active, no window replaced and no plan rejected. The true outlet stays between 1.25 and 1.90 mol/L in every case and the filter never places it outside the bounds, so no case separates the two estimators and no trend of them is drawn.

Read against the true concentration by the measuring script, and by nothing the controller sees, the two estimates have the same error to the digits shown, a root-mean-square value of 0.0038 to 0.0047 mol/L across the ten cases, under half of σ_n, and a largest value of 0.013 to 0.019 mol/L. Under Node.js an estimate took a median of 37 to 43 µs and a 99th percentile of 160 to 222 µs in two runs. On the animation loop of the application in Chrome, during case 9, the estimator and the planner together used at most 1.3 % of one core at normal speed and 6.0 % at the tenfold speed, the higher of two runs at a load average of 2.1 to 2.8, against a budget of one tenth of a core. A window of 15 executions, the first choice, needs a median of 70 µs per estimate against 38 µs for 10 in the same Node.js session, a ratio of 1.84, and the same ratio applied to the share of the estimates in Chrome, 3.7 of the 6.0 percent, places it at about 9 % of a core at the tenfold speed. That margin was judged too narrow under load, so the window is 10, and the feed records of the two windows differ by 5.2×10^{-12} mol/L on case 9.

For the relay, case 3 was started without its setpoint step, so the loop held 1.50 mol/L with the feed at 1.50 mol/L, and the button was pressed at 0.5 min. With $\delta_r = 0.1$ mol/L the test converged after two cycles and 1.49 min, with $\alpha = 0.0378$ mol/L, $D = 0.49$, $P_u = 0.830$ min and $K_{cu} = 5.55$, and it wrote $K_c = 1.73$ and $\tau_I = 1.83$ min over the shipped 2.0 and 1.0 min. The ratio ε_r/α is 0.794, so ψ is 53° and the note flags the result. Under ten other noise seeds, 1001 to 1010, every test converged, with K_{cu} between 3.3 and 11.5, P_u between 0.558 and 1.09 min and ψ between 38° and 70°.

The tank has no ultimate point anywhere near these values. Its model, Equation (3.13), is a lag without dead time, and with a zero-order hold at Δt_c its phase first reaches −180° at the Nyquist frequency, where $K_{cu} = (1 + a)/b = 99.5$ and $P_u = 2\Delta t_c = 0.04$ min. The same relay run on that sampled model without noise cycles with $K_{cu} = 37.0$ and $P_u = 1.24$ min at $\varepsilon_r/\alpha = 0.993$, a cycle

placed by the band and the sampling. The analyzer noise widens the swing, and Equation (3.17) amplifies that error, since $\mathrm{d}\ln K_{cu}/\mathrm{d}\ln\alpha = -1/[1-(\varepsilon_r/\alpha)^2]$ is -2.71 at $\varepsilon_r/\alpha = 0.794$. The measured K_{cu} is therefore 0.15 of the noise-free relay value and 0.056 of the sampled ultimate gain.

The relay settings lose to the shipped ones in both cases of Table 3.10. The true-value IAE rises from 0.108 to 0.209 mol min/L on the setpoint step and from 0.159 to 0.313 mol min/L on the dilution step, where the largest deviation grows from 0.072 to 0.086 mol/L. The integral time carries 74 percent of each loss, because the shipped K_c with the written τ_I already reaches 0.183 and 0.272 mol min/L, while the written K_c with the shipped τ_I costs only 0.123 and 0.181 mol min/L. Equation (3.5) explains the direction. For a lag without dead time, internal model control sets τ_I to the residence time, 1.00 min with the diluent shut, whereas $2.2P_u$ gives 1.83 min and lowers K_I from 2.0 to 0.945/min, although the 20 m^3/min of case 4 calls for 2.4/min. The written settings move the feed 14 percent less, 3.33 against 3.86 and 3.10 against 3.59 mol/L, which is the ratio $1.73/2.0$ of the gains, because most of the travel is analyzer noise passed through K_c, as Table 3.8 showed for a settled loop. With $\sigma_n = 0$ the travel falls to 0.54 and 0.46 mol/L on case 3 and to 0.30 mol/L under both settings on case 4. The MPC, with either estimator, reaches 0.088 and 0.058 mol min/L on the same runs.

> **DID YOU KNOW? The Autotuner Loses to the Model Rule**
>
> From the simulator: the relay test writes an integral time of 1.83 min against the 1.0 min residence time of the undiluted tank, and over 15 min of the dilution step of case 4 the true-value integrated error rises from 0.159 to 0.313 mol min/L (Table 3.10). The tank has no ultimate point near the cycle the relay finds.

A relay started away from the reference point shows how strongly the swing sets the result. Once the setpoint step of case 3 has settled at 1.70 mol/L, a relay started at 5 min converges after four cycles and 2.41 min on $K_{cu} = 4.87$ and $P_u = 0.960$ min with $\alpha = 0.0398$ mol/L and $\psi = 49°$, and it writes $K_c = 1.52$ and $\tau_I = 2.11$ min. Run from the shipped settings with $\sigma_n = 0$, the same test needs two cycles, returns $K_{cu} = 37.6$ at $P_u = 1.24$ min with $\alpha = 0.0302$ mol/L and $\psi = 84°$, and writes $\tau_I = 2.73$ min with K_c held at the top of its range, 10, in place of the 11.7 the rule gives. Swings that differ by a factor of 1.32 yield ultimate gains that differ by a factor of 7.7, because both swings lie close to the band ε_r under the square root of Equation (3.17), and both gains stay far below the 99.5 of the sampled model.

3.8 Technology in Practice

The loop of Section 3.3 reads an ideal analyzer and sets a feed concentration directly. Figure 3.5 draws it with the devices a real blend would use in the place of the schematic's symbols.

What Measures It

Analyzer AT of Equation (3.3) adds white noise of 0.01 mol/L and nothing else. A real composition measurement brings what Section 3.3.4 leaves out: a sample line, an analysis cycle and calibration error. Continuous drug lines, for example, have relied heavily on near-infrared spectroscopy to monitor blend composition [25]. The range, accuracy and speed of such instruments depend on the principle and the sample system, and no single figure applies.

> **TECHNOLOGY Sensor: The Process Analyzer and Its Sample System**
>
> A composition analyzer either measures inside the process stream, in situ, or receives a sample carried to it through tubes, filters, pumps and regulators. Problems with process analyzers often stem from sample systems that are poorly built or poorly maintained. If debris plugs the sample line, the analyzer goes on measuring old, stale sample with no way of knowing it, and reports values that no longer follow the process. Some analyzers are slow by principle. An on-line chromatograph separates species in time, so cycle times of minutes are common, and enough of that dead time makes a feedback loop oscillate [15].

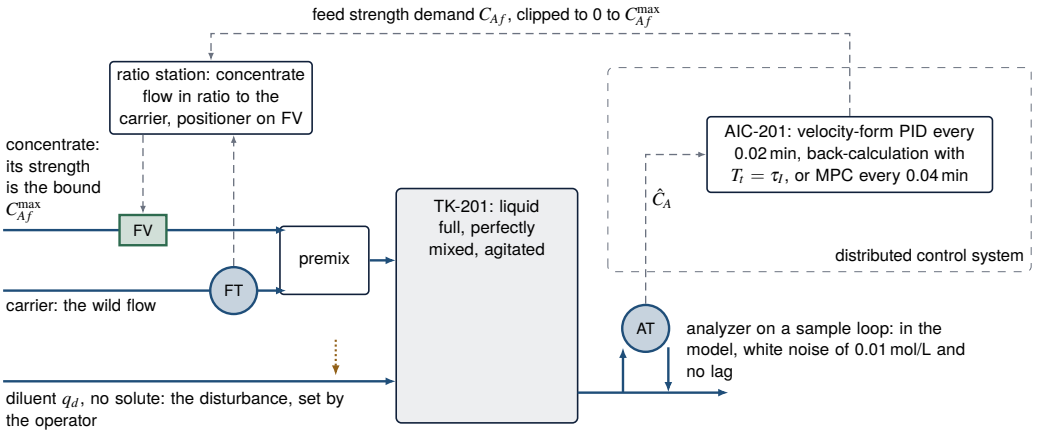


Figure 3.5 The loop of the chapter with real devices in place of the schematic's symbols. An analyzer AT on a sample loop reads the outlet, the controller AIC-201 on a distributed control system computes a feed strength, and a ratio station sets it by holding the concentrate flow through valve FV in ratio to the carrier flow measured by FT. The model collapses the ratio station into an ideal, instantaneous setting. Notice that the diluent is not measured, and that no valve opening can make the feed stronger than the concentrate.

What Moves It

The manipulated feed concentration of Equation (3.1) is applied with no valve and no lag. In a plant it is made by a ratio of two streams or by a dosing pump, so the composition loop is the outer loop of a cascade whose inner element has its own time constant, rangeability and stiction, as the fourth assumption of Section 3.3.4 states. The bound $C_{Af}^{\max}$ stands for the strength of the supply itself, which no valve opening can raise.

> **TECHNOLOGY** **Final Element: The Concentrate Valve of a Ratio Station**
>
> A feed of set strength is usually made by holding two flows in ratio. The transmitter on the uncontrolled, or wild, flow supplies the setpoint of a flow loop on the other line, whose flow is held captive to it [15]. That flow is set by a control valve in a cascade, with an inner loop, the valve positioner, that reduces the effects of friction and nonlinearity [10]. Friction the positioner cannot overcome makes the stem stick and then jump, and integral action then hunts above and below the setpoint without settling [15].

Where the Controller Runs

A distributed control system runs production at one site in refineries, water treatment, chemical manufacturing and pharmaceutical processing, with a supervisory level over the local controllers. Its loops are often nested, with the setpoint of one taken from another, and they run with cycle times from milliseconds to minutes [26]. The distributed control system replaced the cabinets of analog controllers and is now the standard tool of the process industry. Valve positioners close the innermost loops, while loops on quality variables sit at the higher levels [10]. AIC-201 is such a quality loop. Its PID executes every 0.02 min, and the predictive controller decides every 0.04 min. Model predictive control originated in the process industries, where the processes are often slow enough to allow the computation [10].

> **IN PRACTICE** **Windup Where the Setpoint Is Out of Reach**
>
> Reset windup happens when outside conditions make the setpoint impossible to reach. The integral term accumulates for as long as the error lasts, and the measurement must later pass the setpoint long enough to cancel what was stored. Controllers may limit how far the integral term can grow. The surest fix is the operator, who places the loop in manual. Because these conditions usually arise at shutdown, when operators run the process by hand anyway, windup often never shows [15]. A weak supply can bring it on in normal running, and there the back-calculation of Section 3.5 holds the stored output at the bound with the loop left in automatic.

What Is Modernizing It

Three developments touch this loop. The first is the smart positioner. An electronic positioner records actuator pressure against stem position, and this valve signature identifies excessive packing friction, trim interference and poor seating [15], the faults that make an inner flow loop stick. The second is the model of the mixing itself. Residence time models fed with real-time upstream data predict the downstream concentration and allow material that does not conform to be removed from the product [25]. The third is estimation and prediction in the controller, which Sections 3.5.4 and 3.5.5 already carry: a disturbance estimate that books the dilution as a feed deficit and a planner that sees the bound.

3.9 Limitations

The simulator is not intended for design work. A perfectly mixed vessel with no transport delay is the most favorable composition plant that can exist, and the dead time, the analyzer cycle and the sample line that dominate a real blend loop are all absent, so a tuning obtained here is optimistic by the amount of delay a real header carries. The manipulated variable is a feed concentration applied directly, with no inner ratio loop, no dosing pump dynamics, no rangeability limit, no stiction and no stroking time, so the only actuator imperfection represented is the bound itself. The vessel is liquid full with the outlet flow defined as the sum of the inlets, so there is no level loop and no inventory dynamic between the diluent valve and the outlet composition. A single solute is tracked, no reaction occurs, no energy balance is integrated, and the feed temperature is carried through the model and the exported file without entering any equation. Every case starts at steady operation, so a startup from an empty tank has to be configured by hand through the manual feed. The noise generator is reseeded whenever a run starts, so a case repeats sample for sample at any speed multiplier, noise included, while the seed itself is not an input.

Four things were kept deliberately, the bound on the feed, a diluent stream that shifts the gain and the time constant at once, white analyzer noise and a short sample period, because deleting any of them deletes something the loop must be commissioned around. Where each assumption breaks is stated in Section 3.3.4.

The predictive controller plans on one first-order model fixed at the design point. It omits the dilution, whose effect on the gain and time constant (0.833 and 0.833 min at 20 m^3/min) the estimate $\hat{d}$ must absorb, and it omits q and V, so a change in either is a mismatch the estimate also has to absorb. The measured diluent flow is not used as feedforward, the Kalman gain assumes the default analyzer noise, no output band reports an unreachable setpoint through a slack, $\lambda = 30$ comes from a rule shared with six other plants, and the comparison rests on one noise realization.

The moving horizon estimator bounds the state of that same fixed lag, so wherever the lag is wrong its estimate is wrong inside the bounds, and in the shipped cases the outlet stays 0.1 mol/L or more inside both. The relay records the switching band, the analyzer noise and the sampling more than the vessel, because a lag without dead time reaches -180° only at the Nyquist frequency, and the Tyreus–Luyben conversion is one published rule among several. Chapter 2 reaches the same finding on a second lag without dead time.

3.10 Discussion

The reachable set of Equation (3.19) has a direct economic reading. Product giveaway is the quality delivered above the specification the customer pays for, and it exists because a blend has to be aimed away from its limit by enough to cover the variability of the composition loop. Real-time optimization of a gasoline blend header is profitable precisely to the extent that the regulatory layer beneath it can hold a tighter band, so every reduction in the standard deviation of the composition converts into a shift of the target toward the specification [2]. The planning layer above it assigns components to blends on the assumption that the recipes it computes can actually be delivered, and a recipe outside the reachable set of the blend header is an infeasible plan rather than a difficult one [1, 3].

Water treatment inverts the direction of the same argument. Coagulant dosing has an asymmetric cost, because underdosing fails the treatment objective while overdosing wastes chemical and increases sludge production, and raw water quality changes on a timescale no dosing pump can be tuned out of [4]. Continuous pharmaceutical manufacturing adds a specification that is legally rather than commercially binding, since blend uniformity is a release criterion and the plantwide control structure around a continuous mixer exists to keep the product inside it [5]. In each of those settings the useful question at a commissioning meeting is not how the loop is tuned but whether the requested composition lies inside what the streams can produce, and that question is answered by a steady-state balance rather than by a step test. A reader who has watched a well-tuned loop sit 0.167 mol/L short of setpoint with its feed on a bound has encountered that distinction in its simplest form.

The model of this chapter also supports control methods beyond those implemented, the PID loop, its relay autotune, the predictive controller and its two estimators. A model predictive controller built on Equation (3.1) rather than on the fixed lag of Section 3.5.4 would carry the dilution in its model instead of in a disturbance estimate, and a measured diluent flow would enter its prediction as feedforward information rather than as an unannounced load [27]. The two estimators of Section 3.5.5 already recover the true concentration from the noisy analyzer of Equation (3.3) for the planner, and reported as a soft sensor they would turn the offset the bound imposes into a displayed quantity rather than one read from a noisy trend. Economic operation follows directly, because giveaway minimization against the reachable set is an objective the one-equation balance can carry. The APMonitor framework implements these estimation and predictive control methods for differential algebraic models of this class [27], and the Gekko Optimization Suite exposes them in Python against the same exported data [28].

> **OPEN CHALLENGE** **Measuring a Blend Fast Enough to Control It**
>
> Continuous drug manufacture must hold its active ingredient tightly, and tracking that ingredient through the line is a major hurdle. In-line near-infrared spectroscopy watches blend composition, and residence time models can predict the downstream concentration from upstream data, but such models have served mainly for monitoring rather than as part of the control strategy [25]. Offline checks of blend uniformity still stop the process to take a sample, and an in-line probe must look at several locations for an accurate estimate [29]. Progress would be a composition measurement fast and representative enough to close a loop on.

3.11 Conclusion

This chapter studied a stirred blending tank whose manipulated variable is a composition and whose composition has a ceiling. The simulator integrates one species balance on a liquid-full, perfectly mixed vessel with a concentrated feed whose concentration is limited to 0 to 2 mol/L, a pure diluent stream as the load disturbance, and an analyzer carrying 0.01 mol/L of Gaussian noise.

The steady-state gain $q/(q+q_d)$ and the residence time $V/(q+q_d)$ share a denominator, so step tests in the running application return 1.0000 and 0.9950 min at zero dilution, against analytic values of 1.0000 and 1.0000 min, and 0.6667 and 0.6617 min at 50 m^3/min, against analytic values of 0.6667 and 0.6667 min. The internal model control proportional gain $V/(q\tau_c)$ carries no dilution term at all while the integral gain $(q+q_d)/(q\tau_c)$ carries it in full.

The reachability condition $q_d \le q(C_{Af}^{\max}/C_A^{\mathrm{SP}} - 1)$ closes at 33.3 m³/min for the default setpoint, and the measured loop crosses it between 33 m³/min, where the offset is 0.0001 mol/L, and 34 m³/min, where it is the 0.0075 mol/L the balance predicts. At 50 m³/min the offset is 0.1667 mol/L at both $K_c = 2.0$ and $K_c = 8.0$, and the trajectory while the feed is pinned fits a time constant of 0.6617 min at both gains, which is the open-loop value of the discrete plant.

Feedback is applied by an incremental PID whose back-calculation holds the stored output at the bound while the feed is pinned, so the feed leaves the bound 0.30 min after the disturbance clears, as the outlet crosses the setpoint, rather than after working off the 3.12 mol/L of integral demand a free integrator would have stored. A predictive controller planning on the zero-dilution model with a disturbance estimate rejects a 20 m³/min dilution with an integrated error of 0.061 against 0.150 mol min/L and no active bound, ties the PID while the feed is pinned and leaves the bound 0.02 min later, so its advantage lies in its model of the load rather than in its treatment of the constraint. Its moving horizon estimator, bounded by the feed range, never reaches a bound and changes the feed by at most 2.1×10^{-11} mol/L, and a relay autotune writes an integral time of 1.83 min against the residence time of 1.0 min and raises the true-value IAE of the dilution step over 15 min from 0.159 to 0.313 mol min/L. With integral action switched off, a 33 m³/min dilution leaves the outlet 0.149 mol/L below setpoint from a completely different cause, which integral action removes entirely, and at 50 m³/min it removes only the 0.048 mol/L by which that offset exceeds the reachability limit.

A fuller model would add an analyzer dead time, a cascade to an inner ratio loop, a level loop under the vessel and a second solute so that the reachable set becomes a region rather than an interval. The experiment to carry away from the blending tank is the one that separates the offset a controller removes from the offset an actuator bound imposes.

KEY POINTS

- The feed can never be stronger than its supply, so too much dilution leaves an error that no tuning removes.
- The gain and the residence time share a denominator, so internal model control keeps the proportional gain and moves only the integral gain.
- The reachable set closes between 33 and 34 m³/min of diluent at the default setpoint.
- A proportional loop leaves a standing error of similar size that integral action removes entirely.
- Back-calculation holds the stored output at the bound, so the feed leaves the bound as the outlet crosses the setpoint.

Where else this applies. A velocity-form PID with back-calculation against a hard stop also holds the car of Chapter 2.

References

[1] A. Charnes, W. W. Cooper, and B. Mellon, "Blending aviation gasolines—a study in programming interdependent activities in an integrated oil company," *Econometrica*, vol. 20, no. 2, p. 135, 1952.

[2] A. Singh, J. F. Forbes, P. J. Vermeer, and S. S. Woo, "Model-based real-time optimization of automotive gasoline blending operations," *Journal of Process Control*, vol. 10, no. 1, pp. 43–58, 2000.

[3] C. A. Méndez, I. E. Grossmann, I. Harjunkoski, and P. Kaboré, "A simultaneous optimization approach for off-line blending and scheduling of oil-refinery operations," *Computers & Chemical Engineering*, vol. 30, no. 4, pp. 614–634, 2006.

[4] H. Ratnaweera and J. Fettig, "State of the art of online monitoring and control of the coagulation process," *Water*, vol. 7, no. 11, pp. 6574–6597, 2015.

[5] R. Singh, M. Ierapetritou, and R. Ramachandran, "System-wide hybrid MPC–PID control of a continuous pharmaceutical tablet manufacturing process via direct compaction," *European Journal of Pharmaceutics and Biopharmaceutics*, vol. 85, no. 3, pp. 1164–1182, 2013.

[6] Y. Gao, F. J. Muzzio, and M. G. Ierapetritou, "A review of the residence time distribution (RTD) applications in solid unit operations," *Powder Technology*, vol. 228, pp. 416–423, 2012.

[7] P. V. Danckwerts, “Continuous flow systems,” *Chemical Engineering Science*, vol. 2, no. 1, pp. 1–13, 1953.

[8] 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.

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

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

[11] K. J. Åström and L. Rundqwist, “Integrator windup and how to avoid it,” in *1989 American Control Conference*. IEEE, 1989, pp. 1693–1698.

[12] R. Hanus, M. Kinnaert, and J.-L. Henrotte, “Conditioning technique, a general anti-windup and bumpless transfer method,” *Automatica*, vol. 23, no. 6, pp. 729–739, 1987.

[13] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, “A unified framework for the study of anti-windup designs,” *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

[14] S. Tarbouriech and M. Turner, “Anti-windup design: an overview of some recent advances and open problems,” *IET Control Theory & Applications*, vol. 3, no. 1, pp. 1–19, 2009.

[15] 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.

[16] D. E. Rivera, M. Morari, and S. Skogestad, “Internal model control: PID controller design,” *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

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

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

[19] 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.

[20] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

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

[22] D. Goldfarb and A. Idnani, “A numerically stable dual method for solving strictly convex quadratic programs,” *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

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

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

[25] S. Hurley, A. Tantuccio, M. S. Escotet-Espinoza, M. Flamm, and M. Metzger, “Development and use of a residence time distribution (RTD) model control strategy for a continuous manufacturing drug product pharmaceutical process,” *Pharmaceutics*, vol. 14, no. 2, p. 355, 2022.

[26] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, “Guide to operational technology (OT) security,” National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

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

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

[29] E. J. Kim, J. H. Kim, M.-S. Kim, S. H. Jeong, and D. H. Choi, “Process analytical technology tools for monitoring pharmaceutical unit operations: A control strategy for continuous process verification,” *Pharmaceutics*, vol. 13, no. 6, p. 919, 2021.

3.12 Exercises

3.1 Concept check · 2 min

The loop runs with the default tuning, $K_c = 2.0$ and $\tau_I = 1.0\,\text{min}$, which is the internal model control tuning for $\tau_c = 0.5\,\text{min}$ with the diluent shut. The diluent flow then rises to $50\,\text{m}^3/\text{min}$ and stays there, with $q = 100\,\text{m}^3/\text{min}$ and $V = 100\,\text{m}^3$. Which retuning keeps the same τ_c?

(a) Raise the gain to 3.0 and keep the integral time, because the process gain has fallen to 0.667

(b) Raise the gain to 3.0 and shorten the integral time to 0.667 min, because both constants moved

(c) Keep the gain at 2.0 and shorten the integral time to 0.667 min, because both constants fell together

(d) Keep the gain and the integral time, because the gain and the time constant fell by the same factor

3.2 Concept check · 1 min

With the default tuning and a setpoint of 1.5 mol/L, the diluent flow steps to $50\,\text{m}^3/\text{min}$. The outlet settles at 1.333 mol/L with the feed on its 2.00 mol/L limit. Which change brings the outlet back to the setpoint?

(a) A proportional gain of 8.0 in place of 2.0, which quadruples the loop gain

(b) An integral time of 0.667 min, which matches the new residence time

(c) A switch to the predictive controller, which carries the bound in its plan

(d) A diluent flow of $33\,\text{m}^3/\text{min}$ or less, which puts the required feed in range

3.3 Concept check · 2 min

The feed has been on its limit for most of 10 min against a diluent flow of $50\,\text{m}^3/\text{min}$. The diluent valve is then closed. Under the default tuning the outlet peaks at 1.6266 mol/L, an overshoot of 8.44 %. Is this overshoot integrator windup?

(a) No. The stored output sat at the bound itself, and a reachable upset that never saturates gives nearly the same peak

(b) No. The overshoot comes from analyzer noise, which the proportional term passes to the feed while it is on the bound

(c) Yes. The integral grew for 9.34 min on the bound and had to be worked off before the feed could move

(d) Yes. The feed stayed on the bound for 0.30 min after the valve closed, the time the integral took to unwind

3.4 Theory and calculation · 5 min

The tank holds $V = 100\,\text{m}^3$. The feed flow is raised to $q = 200\,\text{m}^3/\text{min}$ and the diluent runs at $q_d = 50\,\text{m}^3/\text{min}$.

(a) Find the process gain and the residence time. *Checkpoint:* 0.80 and 0.40 min.

(b) The feed concentration is held at 1.5 mol/L. Find the steady outlet concentration. *Checkpoint:* 1.20 mol/L.

(c) Find the internal model control settings for $\tau_c = 0.5\,\text{min}$, and say which of them changes when the diluent valve is shut. *Checkpoint:* $K_c = 1.0$ and $K_I = 2.5/\text{min}$. Only the integral gain changes, to 2.0/min.

3.5 Theory and calculation · 6 min

A customer asks for 1.5 mol/L from the default tank, with $q = 100\,\text{m}^3/\text{min}$ and a supply that ends at $C_{Af}^{\max} = 2.0\,\text{mol/L}$.

(a) Find the largest diluent flow at which this setpoint is reachable. *Checkpoint:* 33.3 m^3/min.

(b) The diluent runs at 40 m^3/min. Find the highest outlet concentration the tank can hold and the standing error. *Checkpoint:* 1.429 mol/L and 0.071 mol/L.

(c) Operations proposes to restore the setpoint against 50 m^3/min of diluent by raising the feed flow. Find the feed flow needed and the residence time at that flow. *Checkpoint:* At least 150 m^3/min, with 0.50 min.

3.6 Theory and calculation · 7 min

Integral action is switched off, the gain is $K_c = 2.0$ and the stored output is frozen at $b = 1.50\,\text{mol/L}$. The setpoint is 1.5 mol/L and $q = 100\,\text{m}^3/\text{min}$.

(a) The diluent opens to 20 m^3/min. Find the steady outlet concentration, the offset and the feed concentration. *Checkpoint:* 1.406 mol/L, 0.094 mol/L and 1.69 mol/L.

(b) Find the gain that would lower this offset to 0.05 mol/L, and check that the feed stays inside its range. *Checkpoint:* $K_c = 4.8$, with the feed at 1.74 mol/L.

(c) The diluent opens to 50 m^3/min instead and the gain is 8.0. Show that Equation (3.21) does not apply, and give the offset that results. *Checkpoint:* The demand of 2.13 mol/L is past the bound. The offset is 0.167 mol/L.

3.7 App challenge · 9 min

Load scenario 2, Open-loop actuator step. The loop is in Manual and the feed concentration steps from 1.50 to 1.90 mol/L at 1 min. Use the default seed, run each case to 6 min and export the CSV file. Read the true outlet concentration, column `true_pv`.

(a) Run the scenario as shipped. Read the value at 6 min and compute the gain of the step. *Checkpoint:* 1.897 mol/L, within 0.003 mol/L, and a gain of 0.99.

(b) Find the time after the step at which the true outlet has covered 63.2 % of the change to 1.90 mol/L. *Checkpoint:* 1.0 min, within 0.05 min.

(c) Press the reload button, set the Dilution flow to 50 m^3/min before Run and run again. Read the values at the start and at 6 min, and repeat part (b) for the settled change. *Checkpoint:* 1.000 and 1.267 mol/L, a gain of 0.67, and 0.66 min, within 0.05 min.

3.8 App challenge · 9 min

This exercise sets an offset that a controller removes beside one that it cannot remove. Use the default seed and run each case to 5 min. At 5 min read the IAE tile of the scorecard, and in the exported file the true outlet concentration (`true_pv`) and the feed concentration (`mv`).

(a) Load scenario 4, Disturbance rejection, and run it as shipped. The diluent opens to 20 m^3/min at 1 min. *Checkpoint:* IAE 0.150 mol min/L, outlet 1.494 mol/L, feed 1.80 mol/L.

(b) Press the reload button, set the Integral time to its top, 20 min, which switches integral action off, and run again. *Checkpoint:* Outlet 1.406 mol/L, feed 1.69 mol/L.

(c) Load scenario 8, Actuator limit, and run it as shipped. The diluent stands at 20 m^3/min and the setpoint steps from 1.5 to 1.9 mol/L at 1 min. *Checkpoint:* IAE 1.066 mol min/L, outlet 1.665 mol/L, feed 2.00 mol/L.

(d) Press the reload button, set the Proportional gain to 8 and run again. *Checkpoint:* IAE 1.065 mol min/L, outlet 1.665 mol/L, feed 2.00 mol/L.

Accept an IAE within 0.005 mol min/L, an outlet within 0.005 mol/L and a feed within 0.05 mol/L. Name the equation that predicts each standing error.

3.9 App challenge · 7 min

Load scenario 9, MPC at the edge of the feed. The diluent opens to 20 m^3/min at 1 min, the setpoint steps to 1.70 mol/L at 2 min and returns to 1.60 mol/L at 3.5 min. Use the default seed and run each case to 5 min. For each run read the IAE tile of the scorecard and, in the exported file, the true outlet concentration at 2.0 and at 3.5 min.

(a) Run the scenario as shipped, with the Controller at MPC. *Checkpoint:* IAE 0.207 mol min/L, 1.497 and 1.639 mol/L.

(b) Set the Controller to PID, press Reset and run again. *Checkpoint:* IAE 0.281 mol min/L, 1.434 and 1.629 mol/L.

Accept an IAE within 0.005 mol min/L and an outlet within 0.005 mol/L. State the highest outlet concentration either controller could hold between 2 and 3.5 min, and say where the MPC earns its lower IAE.

3.10 Application area · 4 min

A blend is sold against a minimum concentration, and the setpoint is aimed above that minimum by three standard deviations of the true outlet concentration. Table 3.8 gives this standard deviation as 0.00121 mol/L at $K_c = 2.0$ and 0.00297 mol/L at $K_c = 8.0$, with the same analyzer noise. The outlet flow is 100 m^3/min.

(a) Find the giveaway for each tuning, in mol/L and in solute per day. *Checkpoint:* 0.0036 mol/L and 523 kmol per day, 0.0089 mol/L and 1283 kmol per day.

(b) State from Table 3.7 what the higher gain buys in return, and whether it buys anything while the feed is on its bound. *Checkpoint:* A return within 1 % in 1.83 min against 3.72 min. Nothing on the bound.

3.11 Application area · 3 min

Section 3.9 states that a real blend header carries a transport delay and an analyzer cycle that the model omits. Suppose the analyzer adds a dead time of 0.5 min to the default tank, with $K_p = 1$ and $\tau = 1.0$ min.

(a) Find the PI settings of Equation (1.11) with $\tau_c = \theta_p$. *Checkpoint:* $K_c = 1.0$ and $\tau_I = 1.0$ min.

(b) Find the dead time at which the same rule returns the default gain of 2.0, and state what this means for a tuning carried from the app to a plant. *Checkpoint:* 0.25 min.

3.12 Review the agent · 5 min

An AI assistant is given the result of the first relay test of Section 3.7.7 and replies as follows. “The test measured an ultimate gain of 5.55 and an ultimate period of 0.830 min. The Tyreus–Luyben rule gives $K_c = 1.73$ and $\tau_I = 1.83$ min, which should replace the shipped 2.0 and 1.0 min. The loop becomes unstable at a gain of 5.55, so the proportional gain must stay below about half of that.” Find the flaw in the reply and say how it could be checked.