INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 2 Practical PID Control

CHAPTER 6

pH Neutralization Virtual Laboratory for Gain Scheduling and Measurement Linearization

Pages 147–170

Edited by John D. Hedengren

How to cite this chapter. "pH Neutralization Virtual Laboratory for Gain Scheduling and Measurement Linearization," in *Practical PID Control*, John D. Hedengren, Ed., vol. 2 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 147–170.

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

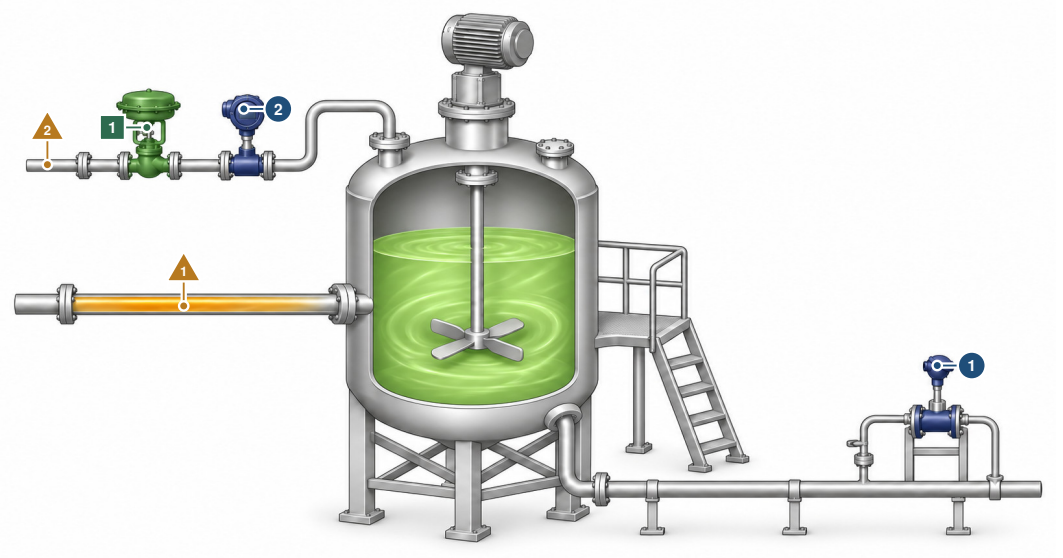
The app for this chapter
https://apmonitor.com/apps/ph-neutralization/

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

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

CHAPTER 6

pH Neutralization Virtual Laboratory for Gain Scheduling and Measurement Linearization



The controller moves 1 the caustic valve

The controller reads 1 the pH electrode on the discharge 2 the caustic flow meter of the inner loop

What disturbs it 1 acid and buffering in the influent 2 caustic supply pressure and strength

Notice that the caustic is dosed into the top of the tank while the pH that decides the permit is read on the discharge line downstream, after mixing and transport. Illustration, not to scale.

IN PLAIN TERMS **Neutralizing Acid Waste Before It Leaves**

The main idea. Acid waste water from a factory flows through a stirred tank. A valve adds caustic soda, a strong base, to bring the water close to neutral. The water may leave only inside a legal band of acidity, read by a probe in the outlet.

The challenge. Near neutral, a tiny extra opening of the valve swings the water from too acidic to too alkaline. Far from neutral, a wide opening barely moves it. A controller set for one place is too jumpy or too sleepy at the other.

Why it matters. Water outside the band breaks the plant's permit, so the controller must change how hard it pushes as the reading moves.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on the permit and the history of gain scheduling (Section 6.1), the try-it segments of cases 1 and 4 (Section 6.6), the fact on caustic and cost, and the key points
Study	50 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 6.6), Technology in Practice (Section 6.7), Limitations (Section 6.8) and the open challenge

The control question A discharge permit of pH 6 to 9 occupies 0.155 % of the reagent valve's travel, and the gain from valve to pH changes by four orders of magnitude across one stroke. When does such a loop need a different structure instead of a better tuning, and which structure repairs which defect?

Learning objectives

- Derive the reagent demand and the process gain at the setpoint from the invariant balances and the charge balance.
- Explain why two correct step tests on this loop return gains that differ by a factor of 691.
- Compare a fixed gain, a gain schedule and control on the inverted titration curve after a load change and after a buffering change.
- State which disturbances an inner reagent flow loop rejects and which it cannot see.
- Judge when valve size, reagent strength or tank volume is the remedy that no control structure supplies.

Prerequisites Sections 1.4.4 and 1.4.5 (step test and tuning), 1.4.13 (scheduling a gain), 1.4.9 (windup) and 1.6.2 (the PID algorithm of the apps). Chapter 4 for tuning from an identified model.

The app https://apmonitor.com/apps/ph-neutralization/

See also Volume 7, Chapter 3, Volume 7, Chapter 2 and Volume 7, Chapter 4 for plants whose gain moves or reverses with the operating point. Volume 3, Chapter 2 for matching a structure to the path of a disturbance and Volume 3, Chapter 6 for a dose computed from an upstream analyzer. Volume 13, Chapter 5 for ratio control with analyzer trim on the plant that makes the caustic.

6.1 Introduction

The discharge permit on the neutralization tank of this chapter is pH 6 to 9, which is three whole units and sounds generous. The reagent valve holds pH 7.00 at 19.899 % of its travel, reaches the bottom of the permit at 19.791 % and passes the top of it at 19.947 %. The entire legal operating range is 0.155 % of one valve, which is finer than the resolution of a typical positioner and far finer than an operator can hold by hand. That single arithmetic fact is what a block diagram of a pH loop cannot deliver, and it is the reason the loop has the industrial reputation it has [1, 2].

DID YOU KNOW? The Same Band in a Federal Rule

The permit of this chapter, pH 6 to 9, is the band that the United States secondary treatment rule sets for publicly owned treatment works. Their effluent pH must stay within 6.0 to 9.0, unless the works adds no inorganic chemicals in treatment and its industrial contributors do not push the pH outside that band [3].

The misconception the simulator corrects is that pH control is hard because the chemistry is hard. It is not. Written in reaction-invariant form, the tank carries two states that blend exactly the way any other conserved quantity blends in a stirred vessel, with one time constant and no coupling [4]. The pH is then recovered from those two states by an algebraic charge balance. The plant is linear and the measurement is not, and a later study by the originators of that form sets out how much of the folklore around pH loops follows from confusing the two [5]. A student who watches the invariant trend travel a clean first-order path while the pH trace falls off a cliff has seen the distinction, and a transfer function from valve position to pH hides it by construction, because the transfer function is only valid at the point where it was linearized.

The general problem is control of a plant whose gain is a strong function of where the plant is sitting rather than a constant. Level in a horizontal drum, temperature through a phase change, composition near a distillation pinch, and flow through a butterfly valve at low lift all have the same shape, and the control literature treats the class rather than the instance. There are two standard answers. The gain can be scheduled on a measured variable so that the product of controller gain and process gain stays fixed, as Section 1.4.13 introduces. The theory and limits of the technique are surveyed for the general case [6]. Alternatively the nonlinearity can be undone by a transformation,

so that the controller works on a variable in which the plant is linear, which is the input-output linearization construction [7] and, applied to this plant, the strong acid equivalent [8]. pH is the textbook case for both because its transformation is the titration curve, which can be computed from equilibrium constants rather than fitted. Chapter 2 is the mild case of the same class, a process gain that falls by a factor of three across the speed range of a car.

> ⧖ HISTORY **Scheduled Gains in Flight Control, 1959**
>
> Aircraft gave gain scheduling its first large field. Wide variations of operating conditions in flight stimulated the development of gain scheduling and of adaptive control, and a high-performance aircraft schedules its controller on Mach number and dynamic pressure, both of them measured. Adaptive control grew from the wish to avoid scheduling in flight control and from the hope of a process controller that adjusts itself, and early work on it was gathered at a symposium on self-adaptive flight control in 1959. Scheduling still dominates flight control, in part because air-data sensors improved so much [9]. In this chapter the measured pH plays the part of Mach number, and the slope of an assumed titration curve sets the gain.

Simulation studies of neutralization are usually built to test an advanced controller, and the model is presented as the setting for that controller [10, 11]. However, a student handed the finished nonlinear controller never performs the step test that returns the wrong number, never sees a correctly tuned loop go to sleep, and never has to decide which of three structural remedies the plant in front of them actually needs. This chapter works with an interactive pH neutralization tank in which the process gain, the reagent demand and the permit band are all measurable from the running plant, so that the case for scheduling, for inversion and for cascade is quantified rather than asserted. The invariant formulation, the titration-curve inverse and the tuning arithmetic are all established methods.

The chapter is organized as follows. Section 6.2 gives the history of the control problem. Section 6.3 states the five-state reaction-invariant model with its instrument chain and its actuator model, in the detail needed to rebuild it. Section 6.4 computes the process gain across one valve stroke and shows why two correct step tests return incompatible answers. Section 6.5 describes the four control architectures, a fixed-gain PI loop, a gain schedule, a loop on the inverted titration curve and a cascade with an inner flow loop. Section 6.6 reports a scenario sequence in which each architecture repairs exactly one defect and the last disturbance defeats all of them. Section 6.7 sets the real instruments and control platform beside the model, and Sections 6.8 and 6.9 state the limitations of the model and weigh the three structural remedies against the equipment decisions. Section 6.10 summarizes the results, and the exercises follow in Section 6.11.

6.2 A History of pH Control

Neutralization became a control problem rather than a batch chemistry problem when effluent limits were written into discharge permits and the treatment had to run continuously alongside production. The first monograph devoted to the subject collected what plants had learned by then, that the reagent valve is almost always oversized for its duty, that a single stirred tank rarely holds a tight permit and two or three in series are the ordinary remedy, that mixing supplies most of the apparent dead time, and that equal-percentage or split-range trim is the cheapest available linearization [1]. Those conclusions were reached from plant experience before there was a dynamic model to support them.

The dynamic model arrived in the early 1970s. A continuous stirred tank receiving a strong acid and a weak acid against a strong base was written down, solved and compared with laboratory data, and it established the shape of the problem, an ordinary mixing lag followed by an equilibrium relation that converts a small change in composition into a large change in pH [12]. That vessel became the benchmark the field re-used, and it was extended to a multivariable version with acid, base and buffer streams that is still the standard test plant [13]. The chemistry in the present application is that benchmark, acetic acid at pK_a 4.76 with a strong mineral acid against sodium hydroxide, scaled from the laboratory rig to a plant-sized vessel.

The decisive modeling step was made in 1983. Rather than integrating concentrations of the ionic species, whose equilibria are microseconds fast against a five-minute tank, two reaction invariants are carried, and neither is changed by any acid-base reaction [4]. The states become linear, the stiffness disappears, and the nonlinearity is isolated in an algebraic output map. The same construction supplies the controller with an exact inverse, and it was later shown that the resulting nonlinear and adaptive schemes are the ones that survive a change in the influent composition [14]. An equivalent transformation, the strong acid equivalent, was derived independently from input-output linearization and demonstrated on an industrial stream [7, 8].

What engineers did before that answer became standard was adapt. Because the titration curve is what the controller needs and nobody measures buffering on line, the curve was estimated on line from the reagent demand the loop turned out to require, and the controller was retuned or re-characterized from the estimate [15]. Adaptive nonlinear schemes that estimate the unmeasured invariants from the pH measurement itself followed and were tested against changes in buffering [10]. Later work identified the curve as the static block of a Wiener model and inverted it inside a feedforward-feedback structure [11]. All of these share one structure, an inverse of an assumed curve placed between the measurement and a linear controller, and they differ in where the assumed curve comes from.

The industrial reputation of the loop did not improve as much as the theory did. Practitioner references written thirty years apart give the same list of causes, an oversized reagent valve, an unmeasured and moving buffering, badly mixed vessels, and a glass electrode that drifts and whose lag grows by an order of magnitude as it fouls [1, 2]. The measurement problem is specific to pH, because the controller is holding a logarithm and a small analog error at the electrode becomes a large error in the quantity the permit is written on. The application of this chapter carries all four of those causes as adjustable features rather than as remarks.

6.3 Process Model

Neutralization tank TK-401 is a 400 L agitated vessel treating 80 L/min of plant effluent that carries 5.0 mmol/L of a strong mineral acid and 0.6 mmol/L of acetic acid at pH 2.301. Reagent is 0.50 mol/L sodium hydroxide through control valve FCV-401, which passes 4.5 L/min at full travel. A small acid trim pump is available for the opposite direction. Electrode AIT-401 measures the discharge, and the permit is pH 6 to 9. Figure 6.1 shows the arrangement, the instrument tags and every symbol that appears in the equations. The residence time at the design point is 4.94 min, and the reagent demand that holds pH 7.00 is 0.8955 L/min, which is 19.899 % of the valve.

Table 6.1 lists what the loop interacts with. The clock is in minutes, concentrations in mol/L, flows in L/min, and valve travel and reagent demand in percent. Five states are integrated and not one of them is the concentration of a reacting species.

6.3.1 Two Reaction Invariants

Writing differential equations for $[H^+]$, $[OH^-]$, [HAc] and $[Ac^-]$ is unnecessary because the equilibria among them complete in microseconds against a 4.94 min tank. The model instead carries the two quantities of Equation (6.1), neither of which any acid-base reaction changes [4].

$$W_a = [H^+] - [OH^-] - [Ac^-], \qquad W_b = [HAc] + [Ac^-] \tag{6.1}$$

where W_a is the excess of strong anion over strong cation in mol/L and W_b is the total weak acid however it is dissociated. Because no reaction term is needed, each invariant blends in the vessel exactly as an inert tracer does, and the balances are Equations (6.2) and (6.3).

$$V \frac{\mathrm{d}W_a}{\mathrm{d}t} = Q_1 c_a - Q_2 C_b + Q_3 C_{\mathrm{HCl}} - Q_t W_a \tag{6.2}$$

$$V \frac{\mathrm{d}W_b}{\mathrm{d}t} = Q_1 c_b - Q_t W_b \tag{6.3}$$

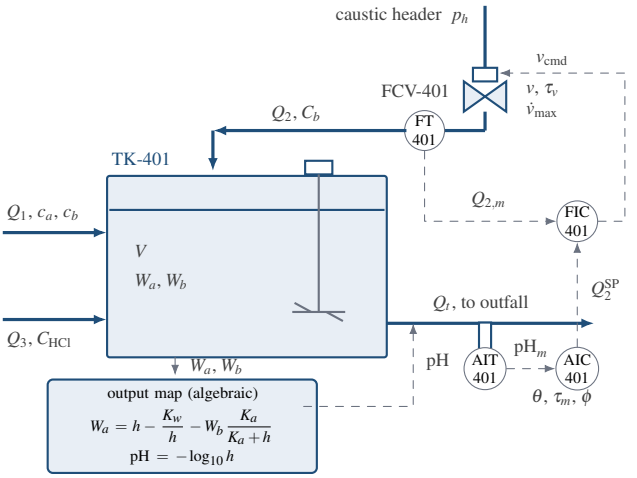


Figure 6.1 Neutralization tank TK-401 with its reagent, effluent and trim streams. Influent Q_1 carries strong acid c_a and weak acid c_b into working volume V, which holds the two reaction invariants W_a and W_b and discharges Q_t. Caustic of concentration C_b enters as Q_2 through valve FCV-401, whose stem v follows the command v_{cmd} with lag τ_v and rate limit $\dot{v}_{\max}$ against header pressure p_h. The pH leaving the tank is not a state and is recovered from W_a and W_b by the charge balance shown. AIT-401 reads that pH through transport delay θ, electrode lag τ_m and fouling fraction ϕ, AIC-401 sets a reagent flow demand, and FIC-401 closes the inner loop on the transmitter FT-401.

Table 6.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Closed-Loop Scenario Starts.

Symbol	Role	Nominal	Range
pH (AIT-401)	CV, measured	7.00	4–11 SP
u demand	MV, outer loop	19.899 %	0–100 %
Q_2^{SP}	MV, cascade only	0.8955 L/min	0–4.5 L/min
v travel	MV, inner loop	19.899 %	0–100 %
c_b weak acid	DV, unmeasured	0.6 mmol/L	0–20 mmol/L
c_a strong acid	DV, unmeasured	5.0 mmol/L	1–15 mmol/L
Q_1 influent	DV, measured	80 L/min	20–170 L/min
p_h header	DV, unmeasured	1.0×	0.4–1.6×
reagent strength	DV, unmeasured	1.0×	0.6–1.4×
Q_3 acid trim	DV, measured	0 L/min	0–0.4 L/min
fouling rate	DV, unmeasured	0 %/h	−30–40 %/h
W_a	state	−0.590 mmol/L	±1000 mmol/L
W_b	state	0.593 mmol/L	0–500 mmol/L
v	state	19.899 %	0–100 %
Q_1^f pump	state	80 L/min	≥ 0.1 L/min
ϕ fouling	state	0	0–1

where V is the working volume in L, Q_1, Q_2 and Q_3 are the influent, reagent and acid trim flows in L/min, $Q_t = Q_1 + Q_2 + Q_3$ is the discharge, c_a and c_b are the influent strong and weak acid concentrations, C_b is the caustic concentration and C_{HCl} is the trim acid concentration. Both sides carry units of mol/min. The caustic appears in (6.2) with a negative sign because a strong base raises the strong cation, and it does not appear in (6.3) at all because it adds no weak acid. Equations (6.2) and (6.3) are two uncoupled first-order lags of time constant V/Q_t, and there is no nonlinearity in either of them.

6.3.2 The Output Map

The pH is recovered from the two invariants by the charge balance of Equation (6.4), solved for the hydrogen ion concentration h.

$$W_a = h - \frac{K_w}{h} - W_b \frac{K_a}{K_a + h}, \qquad \mathrm{pH} = -\log_{10} h \tag{6.4}$$

where $K_a = 10^{-\mathrm{p}K_a}$ and $K_w = 10^{-\mathrm{p}K_w}$. The right side of (6.4) is strictly increasing in h, because its derivative is Equation (6.5) and every term is positive for physical W_b.

$$J = \frac{\partial W_a}{\partial h} = 1 + \frac{K_w}{h^2} + \frac{W_b K_a}{(K_a + h)^2} > 0 \tag{6.5}$$

A bisection on $\log_{10} h$ over $[10^{-15.5}, 10]$ therefore cannot fail. The implementation takes 48 halvings, which places the answer within 6×10^{-14} of a decade, and a check in the application's test suite confirms that inverting (6.4) and evaluating it again reproduces a specified pH to better than 10^{-9} over $1.5 \le \mathrm{pH} \le 12.5$ at weak acid levels from 0 to 20 mmol/L. This one algebraic line is where the entire difficulty of the application lives.

Differentiating (6.4) gives the sensitivity of the measurement to each state, as shown in Equation (6.6).

$$\frac{\partial \mathrm{pH}}{\partial W_a} = \frac{-1}{h \ln 10 \, J}, \qquad \frac{\partial \mathrm{pH}}{\partial W_b} = \frac{-1}{h \ln 10 \, J} \cdot \frac{K_a}{K_a + h} \tag{6.6}$$

The leading h in the denominator moves by ten decades across the pH scale while J moves by two or three, so the sensitivity is dominated by it. The bracket J is the buffer capacity, and it is the term the influent weak acid controls.

6.3.3 Reagent Valve, Influent Pump and Instruments

The reagent valve is a first-order actuator with a rate limit, and its flow follows the ordinary valve law, as shown in Equations (6.7) and (6.8).

$$\tau_v \frac{\mathrm{d}v}{\mathrm{d}t} = v_a - v, \qquad \left| \frac{\mathrm{d}v_a}{\mathrm{d}t} \right| \le \dot{v}_{\max} \tag{6.7}$$

$$Q_2 = Q_{r,\max} \, f(v) \, \sqrt{p_h}, \qquad f(v) = \frac{v}{100} \tag{6.8}$$

where v is the stem position in percent, v_a is the commanded position v_{cmd} after rate limiting, τ_v is the travel time constant, $\dot{v}_{\max}$ is the slew limit in %/min, $Q_{r,\max}$ is the capacity at full travel and p_h is the header pressure as a fraction of nominal. The linear trim of (6.8) is the default and is what is usually installed. Setting the rangeability R_v above unity replaces f with the equal-percentage characteristic $(R_v^{v/100} - 1)/(R_v - 1)$. The influent pump ramps between throughputs through a first-order lag of time constant τ_Q, and the electrode fouling fraction ϕ grows at a rate the operator sets, as shown in Equation (6.9).

$$\tau_Q \frac{\mathrm{d}Q_1}{\mathrm{d}t} = Q_1^{\mathrm{set}} - Q_1, \qquad \frac{\mathrm{d}\phi}{\mathrm{d}t} = \frac{r_f}{6000} \tag{6.9}$$

where r_f is the fouling rate in percent per hour and the constant 6000 converts to a fraction per minute. The state ϕ is clipped to the interval $[0, 1]$ after every step, v to $[0, 100]$, W_b to non-negative values and Q_1 to at least 0.1 L/min.

The electrode is the only pH instrument the controller reads, and it carries a transport delay, a fouling offset, an electrode lag that grows with fouling and colored noise, as shown in Equation (6.10).

$$\tau_m \, (1 + \kappa \phi) \, \frac{\mathrm{d}\mathrm{pH}_m}{\mathrm{d}t} = \mathrm{pH}(t - \theta) + \delta \phi + n(t) - \mathrm{pH}_m \tag{6.10}$$

Table 6.2 Model Parameters. Every Constant in This Application Is Displayed and Adjustable, so Nothing Is Hidden From the Student.

Symbol	Meaning	Value	Unit
V	working volume	400	L
τ_Q	influent pump ramp	0.25	min
$\mathrm{p}K_a$	acetic acid dissociation	4.76	—
$\mathrm{p}K_w$	ion product of water at 25 °C	14.00	—
C_b	caustic concentration	0.50	mol/L
C_{HCl}	trim acid concentration	0.50	mol/L
$Q_{r,\max}$	reagent valve capacity	4.5	L/min
R_v	valve rangeability	1.0	—
τ_v	valve travel time constant	0.08	min
$\dot{v}_{\max}$	valve slew limit	400	%/min
τ_m	clean electrode time constant	0.20	min
θ	transport delay to electrode	0.15	min
κ	electrode lag multiplier, fouled	9	—
δ	electrode offset, fouled	−0.9	pH
τ_F	reagent flow transmitter lag	0.05	min
—	permit lower limit	6.0	pH
—	permit upper limit	9.0	pH
—	caustic price, 100 % basis	0.60	$/kg
—	cost of discharge off permit	12	$/min

where θ is the transport delay from the injection point to the electrode, τ_m is the clean electrode time constant, κ is the lag multiplier at full fouling, δ is the reading offset at full fouling and n is a first-order-filtered Gaussian sequence scaled so that its stationary standard deviation equals the specified value. The offset is added before the electrode lag rather than after it, because the drift is a property of the glass. Equation (6.10) is advanced by the exact zero-order-hold update $\mathrm{pH}_m^k = \alpha\,\mathrm{pH}_m^{k-1} + (1-\alpha)\,x_k$ with $\alpha = \exp(-\Delta t/\tau)$, which is stable at any sample period even while τ is changing. The reagent flow transmitter is a clean first-order lag of time constant τ_F on Q_2 with no delay and no noise, and it is what makes an inner flow loop worth having.

6.3.4 Initialization, Numerics and Scorecard

The constants are listed in Table 6.2. Every scenario opens at a genuine steady state. The valve is placed at the travel that holds the pH setpoint, found by bisection on (6.4) through the steady form of (6.2) and (6.3), or at the operator's manual position when the scenario starts in manual, and the two invariants are set to the values that position produces. The test suite confirms that the resulting derivative vector is zero to within a relative 10^{-8} on every state.

The engine of Section 1.6.1 integrates the model with its fourth-order Runge-Kutta method on a fixed step of 0.01 min (0.6 s). The controller executes every 0.02 min (1.2 s) and the historian logs 20 samples per simulated minute, one every 3 s. The default seed of the noise generator is 20260118. The scorecard accumulates the caustic consumed as kilograms of 100 % sodium hydroxide, the minutes spent outside the permit on the true pH, the integrated absolute error against the setpoint, the total variation of the valve stem, and a cost that adds caustic purchased to a flat charge for every minute off permit. The exported file holds the true pH beside the electrode reading, both invariants, the local process gain, the controller gain in force and the scheduling multiplier at every logged sample, so the gain map behind Figure 6.2 can be rebuilt from the record of a run.

6.3.5 Assumptions

The model is detailed on the chemistry that sets the gain and thin everywhere else, and each simplification fails somewhere.

The acid-base equilibria are instantaneous. Proton transfer completes in microseconds against a 4.94 min residence time, so treating the ionic species as algebraic functions of two invariants is more accurate than writing differential equations for them, not less [4]. The assumption fails the moment the effluent carries something that reacts slowly, such as metal hydroxides coming out of solution, carbonate solids, or lime slurry in place of caustic, because those are minutes-scale processes that buffer the tank and no algebraic charge balance represents them.

One monoprotic weak acid. Real effluent is a mixture with several pK_a values spread across the range, so its curve has several inflections rather than one cliff. Dissolved carbon dioxide is the important omission, because a tank open to atmosphere buffers itself near pH 6 and again near pH 10 on a timescale of minutes, which is the timescale of the control problem.

Fixed temperature, and with it fixed K_a and K_w. At 25 °C neutrality is pH 7.00 and at 40 °C it is near pH 6.8, so a tank that swings fifteen degrees over a day has a titration curve that moves under the controller. The temperature compensation fitted to a real transmitter corrects the Nernst slope of the glass and not the equilibrium constants of the liquid, so this error survives a properly compensated instrument.

Concentrations are used where activities belong. At the 5 mmol/L of this design point the activity correction is worth a few hundredths of a pH unit. In a brine effluent it is worth several tenths and the equilibrium constants shift with it.

One perfectly mixed volume plus a lumped delay. Real neutralization tanks are badly mixed by this standard, because reagent enters at the surface into a stream orders of magnitude more acidic and the local pH there is nothing like the bulk. A defensible model of a real vessel is two or three compartments in series, so the single delay of 0.15 min understates how much of a pH loop's dead time is mixing rather than piping [1].

The valve is ideal apart from its lag, its rate limit and its pressure term. There is no stiction, no deadband and no hysteresis. That is the most consequential omission, because a sticky reagent valve is one of the commonest causes of a real pH limit cycle and it is routinely misdiagnosed as a tuning problem [2]. Chapter 4 includes a case with valve stiction.

Electrode fouling is linear in one state. Real fouling is nonlinear, coats asymmetrically, and often fails toward a fixed reading rather than a fixed offset. The failure the application cannot show is the worst one, in which the trace goes quiet and looks like perfect control.

6.4 The Gain That Is Not a Number

6.4.1 Process Gain Across One Valve Stroke

Combining (6.6) with the steady form of (6.2) and (6.3) gives the gain from valve travel to pH, as shown in Equation (6.11).

$$\frac{\partial \mathrm{pH}}{\partial v} = \frac{1}{Q_t}\frac{\mathrm{d}Q_2}{\mathrm{d}v}\left[\frac{\partial \mathrm{pH}}{\partial W_a}(-C_b - W_a) + \frac{\partial \mathrm{pH}}{\partial W_b}(-W_b)\right] \tag{6.11}$$

where $\mathrm{d}Q_2/\mathrm{d}v = Q_{r,\max}\sqrt{p_h}/100$ for linear trim. Sweeping the valve from shut to wide open at the design composition and evaluating (6.11) at each steady state gives a gain of 5.191×10^{-3} pH per percent of travel with the valve wide open at pH 12.33 and 101.98 pH per percent at pH 7.77. The ratio is 1.96×10^4, which is 4.29 orders of magnitude on one valve, one tank and one reagent. The tangent gain at the setpoint itself is 33.78 pH per percent. Figure 6.2 draws the curve and both operating points.

6.4.2 Two Correct Step Tests With Incompatible Answers

The identification test is the step test of Section 1.4.4, and it is performed twice. A step of 1 % on the valve from 12 %, out on the flat part of the curve, moves the tank from pH 2.790 to pH 2.872 by $t = 40$ min. The gain is 0.082 pH per percent of valve, and 63 percent of the change appears in the tank pH 5.32 min after the step. The invariant W_a, which is linear in the mixing, reaches its 63 percent

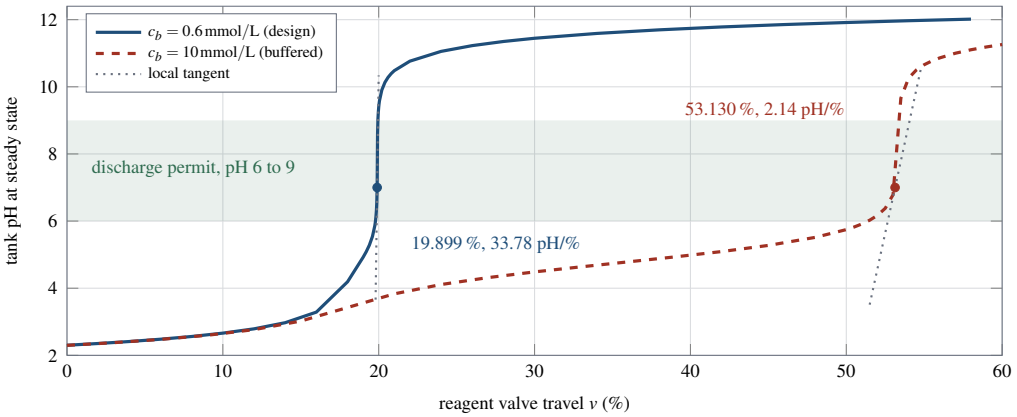


Figure 6.2 Steady-state titration curves of the neutralization tank against reagent valve travel. The design influent (solid curve) reaches pH 7.00 at 19.899 % of travel with a local slope of 33.78 pH per percent, and the whole permit band spans 0.155 % of the same valve. Raising the influent weak acid from 0.6 to 10 mmol/L (dashed curve) moves the same setpoint to 53.130 % and flattens the slope to 2.14 pH per percent, while the influent pH itself moves only from 2.301 to 2.298. The dotted lines are the local tangents at the two operating points, each marked by a filled circle, and the shaded band is the discharge permit.

point at 5.02 min, the 4.96 min residence time at this valve position plus the 0.08 min valve lag, and the curvature of the titration curve, slight on this part, adds the remaining 0.30 min. The electrode sees the response through the 0.15 min transport delay and the 0.20 min lag, and its noise-free reading reaches 63 percent of the change 5.68 min after the step. That is an ordinary first-order-plus-dead-time response with $\tau_p = 4.94$ min (the design residence time) and $\theta_p = 0.35$ min (the transport delay and electrode lag in series). The IMC-PI rule [16] of Equation (6.12), which Section 1.4.5 derives, with a 1.5 min closed-loop target returns $K_c = 33$ percent of valve per pH unit and a reset time equal to τ_p, which the application loads as 5.2 min.

$$K_c = \frac{\tau_p}{K_p\,(\tau_c + \theta_p)}, \qquad \tau_I = \tau_p \tag{6.12}$$

The same test at the setpoint uses a step of 0.03 %, which is thirty-three times smaller, and moves the tank from pH 7.003 to pH 8.701. The secant gain between the two steady states is 56.66 pH per percent, a factor of 691 against the 0.082 the identical procedure measured on the flat part, and a step a thirty-third of the size produced a response twenty times larger. Applying (6.12) to the setpoint tangent of 33.78 with the same 1.5 min closed-loop target returns $K_c = 0.079$, and the application loads 0.092, which corresponds to a 1.24 min target. Both tests were performed correctly. A step test on a pH loop measures the gain at the point where the test was done, and the invariant trend is the honest picture, because W_a traces the same clean first-order response in both tests and moves by the same 0.28 mmol/L per percent of valve (0.279 mmol/L for the 1 % step and 0.0083 mmol/L for the 0.03 % step).

Pause. Try Exercise 6.1 before reading on.

6.4.3 Buffering Moves the Curve Without Moving the Inlet

Acetic acid at pK_a 4.76 is barely dissociated at pH 2.3 and almost completely dissociated at pH 7, so it contributes very little to the influent pH and a great deal to the reagent demand. Raising c_b from 0.6 to 10 mmol/L moves the influent pH from 2.301 to 2.298, which no instrument would report as a change. At the setpoint it divides the tangent gain by 15.8, from 33.78 to 2.14 pH per percent

through the J term of (6.5), raises the reagent demand from 19.899 to 53.130 % of valve, and widens the permit band from 0.155 to 1.973 % of travel. Both curves appear in Figure 6.2. Nobody puts a titrator on an effluent line, so the control system discovers this at the same moment the permit does, which is why the curve has to be estimated from the demand the loop turns out to require [15].

Pause. Try Exercise 6.2 before reading on.

6.5 Controller Design

6.5.1 The Shared PID Algorithm

Both loops use the velocity-form algorithm of the framework, which Section 1.6.2 gives in general. Equations (6.13) and (6.14) restate it in the symbols of this chapter, because the windup analysis and the results cite them.

$$\begin{aligned} e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\ D_k &= -K_c \tau_D\, a \frac{\mathrm{d}\mathrm{PV}_k^{f}}{\mathrm{d}t} \\ s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \end{aligned} \tag{6.13}$$

$$\begin{aligned} c_k &= u_{\mathrm{bias}} + s_k + I_{k-1} \\ u_k &= \mathrm{slew}(\mathrm{clamp}(c_k, u_{\min}, u_{\max}),\, u_{k-1},\, \dot{u}_{\max},\, \Delta t_c) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right] \Delta t_c \end{aligned} \tag{6.14}$$

The accumulator s_k is the running sum of Equation (1.24) and the demand c_k is the output before limiting of Equation (1.25). The slew operator holds $|u_k - u_{k-1}|$ to at most $\dot{u}_{\max}\Delta t_c$, where Δt_c is the control period. Here $a = +1$ for both loops, which is reverse action (Section 1.4.3), because the process gain from reagent to pH is positive and a rising measurement must be met with a falling output. The accumulator equals the positional $P + D$ for the gains in force, so the demand keeps the identity $c_k = u_{\mathrm{bias}} + P_k + D_k + I_{k-1}$ and a gain change acts on the whole error at the next execution. That property matters on this plant, because the scheduled architecture changes the gain at every execution. Anti-windup is back-calculation with tracking time constant T_t, which defaults to τ_I and is set to 10^9 minutes when the student switches anti-windup off [17]. Sections 1.4.9 and 1.4.10 give the reasoning behind the limit, the correction of the integral and the bumpless transfer between Manual and Auto.

The outer loop is bounded to a reagent demand of 0 to 100 % and carries a derivative filter of 0.12 min. The inner flow loop is bounded to a valve travel of 0 to 100 % and has no derivative term. Tuning constants are bounded to $0 \le K_c \le 40$, $0.2\,\mathrm{min} \le \tau_I \le 120\,\mathrm{min}$ and $0 \le \tau_D \le 10\,\mathrm{min}$ on the outer loop. Both execute every 0.02 min, independent of the 0.01 min physics step. The default derivative time is zero, because derivative action on a pH electrode multiplies its noise by whatever the local slope happens to be and that slope moves by four orders of magnitude.

6.5.2 Four Architectures

The architecture switch selects what the controller measures and what its output means, and each step repairs exactly one defect. In the PI architecture the controller acts on the measured pH directly and its output is the valve position, so K_c has units of percent of valve per pH unit. In the GS architecture the same loop multiplies K_c by the ratio of titration-curve slopes between the measurement and the setpoint, as shown in Equation (6.15).

$$K_c^{\mathrm{eff}} = K_c\,\mathrm{clamp}\left(\frac{u^{*\prime}(\mathrm{pH}_m)}{u^{*\prime}(\mathrm{pH}_{\mathrm{SP}})},\ 0.02,\ 3000\right) \tag{6.15}$$

where $u^*(\mathrm{pH})$ is the reagent demand an assumed influent would need to sit at that pH and the slope is taken as a central difference over ± 0.04 pH. The clamp is necessary because the unclamped ratio reaches 873 at pH 2.5 and 6510 with the valve wide open at pH 12.33. Equation (6.15) is the schedule of Equation (1.22) with the slope of an assumed titration curve in place of a measured gain curve. In the LIN architecture the measured pH and the setpoint are both pushed through u^* before the controller sees them, so the controller works on a reagent demand in percent of valve. The inverse is Equation (6.16).

$$u^*(\mathrm{pH}) = \frac{100}{Q_{r,\max}} \cdot \frac{Q_1 \hat{c}_a + Q_3 C_{\mathrm{HCl}} - \hat{W}_a(\mathrm{pH})\,(Q_1 + Q_3)}{C_b + \hat{W}_a(\mathrm{pH})} \tag{6.16}$$

where $\hat{W}_a(\mathrm{pH})$ is (6.4) evaluated at $h = 10^{-\mathrm{pH}}$ with the assumed weak acid, and the dilution of the weak acid by the reagent stream is resolved by eight fixed-point iterations. At steady state the demand that holds the measured pH is the valve position, so the process gain in this variable is 1 by construction, and (6.12) on a unit-gain 4.94 min lag with 0.35 min of dead time and the same 1.5 min target returns $K_c = 2.7$, dimensionless, which the scenarios round to 3. In the CASCADE architecture the same linearized outer loop writes a reagent flow setpoint $Q_2^{\mathrm{SP}} = Q_{r,\max} u/100$ and FIC-401 drives the valve against the transmitter reading.

Three properties of (6.15) and (6.16) are worth stating plainly. Both are built from an assumed influent composition $(\hat{c}_a, \hat{c}_b)$, which the student sets on the advanced panel, because nobody measures buffering on line. The C_b in the denominator of (6.16) is the nameplate caustic concentration and not the strength of the delivery that is actually in the tank, because the strength of a tanker load is not measured either, so a weaker delivery is unmeasured by the whole control system and not only by the inner flow loop. Neither contains the header pressure term of (6.8), so the demand is a percent of valve at nominal header pressure and a pressure change appears to the outer loop as an unmeasured disturbance. The consequence is that these structures are detuned rather than destabilized when the assumption is wrong, which is measured in Section 6.6.

6.5.3 Windup Against a Rate Limit

Section 1.4.9 treats windup at a position limit, and the rate limit of this valve adds a second case. The controller that writes the valve position is told the stem's slew rate through $\dot{u}_{\max}$ in (6.14), which is the outer loop when there is no cascade and the inner flow loop when there is. Its back-calculation therefore unwinds against what the actuator delivered rather than against what was asked for. This is the case that is easy to miss, because a valve limited to 6 %/min takes 16.7 min to travel its full range and is saturated in rate for minutes at a time while sitting nowhere near either stop. The outer loop in the cascade architecture is given an infinite slew rate instead, since its output is a flow setpoint and not a stem position. No external reset signal passes from the inner loop to the outer one, so while the stem is rate-limited the outer integral is unwound only at its own 0 to 100 % bound on the demand.

6.6 Results

Twelve scenarios are arranged as a laboratory sequence, listed in Table 6.3 with the mechanism each exercises and what a learner must do to complete it. Completion rules and the graded challenge work as Section 1.6.4 describes, and Reset keeps the learner's settings (Section 1.6.1). The challenge is judged on the case as given, so changing a model parameter or a disturbance, or firing an extra disturbance button, voids the attempt. Every figure reported below was reproduced by driving the shipped model headlessly at the fixed step of 0.01 min, exactly as the browser steps it.

6.6.1 Dosing by Hand

Scenario 1 opens in manual and steps the valve from 10 to 21 % in six moves of 5, 4, 0.8, 0.15, 0.15 and 0.9 %, holding each for 8 min. The steady-state answers are the whole lesson. The first move, five percent of the valve, carries the tank from pH 2.662 to pH 3.104, which is 0.088 pH per percent. The fourth move is 0.15 % of travel, from 19.80 to 19.95 %, and it carries the tank from pH 6.034 to

Table 6.3 Scenarios and Completion Requirements. The Letter Beside Each Number Gives the Kind: E Is an Experiment, D a Demonstration and C the Challenge.

#	Scenario	What it exercises	Required to complete
1 E	Dose by hand in six steps	the width of the permit in valve travel	Hold Manual through the six scripted demand steps from 10 to 21 % of valve to the end of the last 8 min hold at $t = 54$ min.
2 E	Step test at 12 % of valve	correct FOPDT identification	In Manual, read the settled pH before the 1 % demand step at $t = 10$ min and after it at $t = 40$ min, for the gain and the time to 63 percent of the change.
3 E	Step test at the setpoint	the same procedure, a different plant	In Manual, follow W_a through the 0.03 % step at $t = 10$ min and run to $t = 40$ min, then take the ratio of the gain to the one from case 2.
4 D	$K_c = 33$ run at pH 7	limit cycling from too much loop gain	Run the full hour in Auto on the fixed-gain PI loop with $K_c = 33$, then read the time in permit and the valve travel.
5 D	$K_c = 0.092$, acid load step	the opposite failure, too little	On the PI loop with $K_c = 0.092$, see the tank fall below pH 6 after the acid step at $t = 25$ min and follow the demand to $t = 70$ min, estimating when it would reach the 34.1 % the new load needs.
6 E	Gain scheduled, same step	scheduling from first principles	On the scheduled loop, bring the demand to 33.5 % after the step at $t = 25$ min and run to $t = 70$ min, comparing the low point and the time off permit with case 5.
7 E	Linearized, same step	inverting the measurement	On the linearized loop with $K_c = 3$, run through the step to $t = 45$ min and on to $t = 70$ min.
8 E	Buffering 0.6 to 10 mmol/L	an unmeasured curve change	Four runs, each structure in place before the batch change at $t = 15$ min and held for 75 min, with Reset between them: PI with $K_c = 0.092$, the scheduled loop, the linearized loop with $K_c = 3$, and the linearized loop with the assumed weak acid set to 10 mmol/L.
9 E	Influent flow 80 to 155 L/min	gain and dynamics moving together	On the linearized loop, run to $t = 45$ min with the second line on (influent above 150 L/min) and on to its shutdown at $t = 65$ min.
10 D	Seven hours of electrode fouling	drift inside the loop	On the linearized loop, run to the end of the shift at $t = 420$ min and through the electrode cleaning at $t = 430$ min to $t = 445$ min.
11 E	Header pressure and reagent strength	the exact scope of a cascade	Run the cascade through the header drop to $t = 30$ min and the weaker caustic to $t = 85$ min, then Reset, switch to the linearized loop before $t = 15$ min and run to $t = 30$ min.
12 C	Graded challenge, four upsets	structure before tuning	At $t = 240$ min, more than 45 % of the run inside the permit and the true pH within ± 1.0 of the setpoint, with the model parameters and disturbances as given.[a]

[a] With the shipped configuration (fixed-gain PI, $K_c = 0.092$, anti-windup off) the run is 8.5 % in permit and ends 4.06 pH from the setpoint, failing both parts. With the suggested configuration (cascade on the linearized signal, $K_c = 3$, $\tau_I = 4$ min, back-calculation) it is 49.7 % in permit, 4.7 points above the threshold, and ends 0.001 pH from the setpoint.

pH 9.035, which is 20.0 pH per percent and crosses the permit from just inside the bottom to just outside the top. The same actuator is worth 226 times more in the fourth move than in the first, with nothing about the tank, the valve or the reagent changed. Eight minutes is only 1.6 residence times, so the trend after the fourth hold reads pH 6.29 rather than the settled 9.04, and the scenario is more convincing when each step is left longer. The manual actuator resolves to 0.001 % of valve for this reason, with arrow keys that move it by 0.05 and 0.005 %, and a helper button places it at the travel that holds any setpoint the student types.

> **TRY IT IN THE APP** **1 · Dose It by Hand, and Find the Cliff**
>
> Load the case and press Run. The valve steps from 10 to 21 % of travel in six moves, each held for 8 min. Watch the pH trend after each move. The first move, five percent of the valve, lifts the tank from pH 2.662 to 3.104, while the fourth, only 0.15 % of travel, carries it across the whole permit.

6.6.2 Two Opposite Failures of a Fixed Gain

Scenario 4 closes the loop at pH 7 with the settings the flat-part step test honestly produced, $K_c = 33$ and $\tau_I = 5.2$ min, and applies no disturbance at all. The loop gain at the setpoint is $33 \times 33.78 = 1115$,

Table 6.4 Measured Performance After an Influent Strong Acid Step From 5 to 9 mmol/L at $t = 25$ min, All Values Read at $t = 70$ min. Every Run Uses $\tau_I = 5.2$ min and $\tau_D = 0$.

	PI	Scheduled	Linearized	Cascade
K_c	0.092	0.092	3.0	3.0
Minimum pH	2.54	4.16	3.51	3.53
Return inside permit (min)	never	15.0	21.3	21.4
Time off permit (min)	45.0	15.0	21.3	21.3
IAE (pH·min)	195.3	36.2	54.0	54.1
Valve travel (%)	3.9	19.9	15.1	15.5
Caustic (kg)	1.34	1.82	1.81	1.81
Cost ($)	540	181	256	257

so the loop is unstable by three orders of magnitude and the only thing that bounds the oscillation is the valve reaching its stops. Over one hour the tank cycles between pH 4.50 and pH 10.97, spends 57.5 of 60 min outside the permit for an in-permit score of 4.1 %, and moves the stem through 5231 % of its travel, which is fifty full strokes in an hour. The mean pH over the cycle is 8.57, inside the permit, so a plant that reported an average would look compliant.

> **TRY IT IN THE APP** **4 · A Gain Tuned on the Flat Part, Run at the Setpoint**
>
> Load the case, which closes the loop at pH 7 with $K_c = 33$ and $\tau_I = 5.2$ min and applies no disturbance, and run the full hour. Watch the valve travel trend swing between its stops while the pH cycles between 4.50 and 10.97. Then read the time in permit on the scorecard.

Scenario 5 uses the settings the setpoint deserves, $K_c = 0.092$, and steps the influent strong acid from 5 to 9 mmol/L at $t = 25$ min. The loop is perfectly stable and it never gets back. The tank falls to pH 2.54 and is still at pH 2.62 at $t = 70$ min, off permit for all 45 min since the step. The reagent demand has to reach 34.12 % and is climbing at 0.079 % per minute, so it would need roughly three hours. The mechanism is on the log-gain trend. At pH 2.5 the plant gain is 0.0384 pH per percent against 33.78 at the setpoint, a factor of 880, so the loop gain has fallen from 3.11 to 0.0035. The loop is not unstable. It is asleep.

6.6.3 Scheduling, Inversion and Cascade on One Load Change

Table 6.4 runs the identical acid load change on all four architectures with each one's proper gain, and reports the numbers the application computes for itself. The integrated absolute error and the off-permit time both accrue from the start of the run, so each is quoted at the same observation point of $t = 70$ min, 45 min after the step.

Scheduling and inversion both convert an unrecoverable excursion into a recovery of a quarter of an hour, and the fixed-gain column is worse on every measure at once. The scheduled loop is the faster of the two on this disturbance, reaching the permit 15.0 min after the step against 21.3 min and bottoming out at pH 4.16 against pH 3.51. That ordering is worth being honest about, because the IMC tuning the linearized plant deserves is more conservative than the effective gain the scheduler happens to apply out on the flat part, where K_c^{eff} reaches 3.14 at a measured pH of 5.20 against the 0.092 it holds at the setpoint. What inversion buys is not speed. It is that the integral in (6.14) accumulates percent-of-valve-minutes rather than pH-minutes, and a percent of valve means the same thing at pH 3 and at pH 7 while a pH unit does not. The valve moves 15.1 % of travel over the disturbance where the scheduler moves 19.9 %, and there is no slope estimate to amplify electrode noise. The cascade column is indistinguishable from the linearized one because this disturbance never touches the reagent flow, which is the result Section 6.6.5 depends on. Figure 6.3 draws the same table as bars.

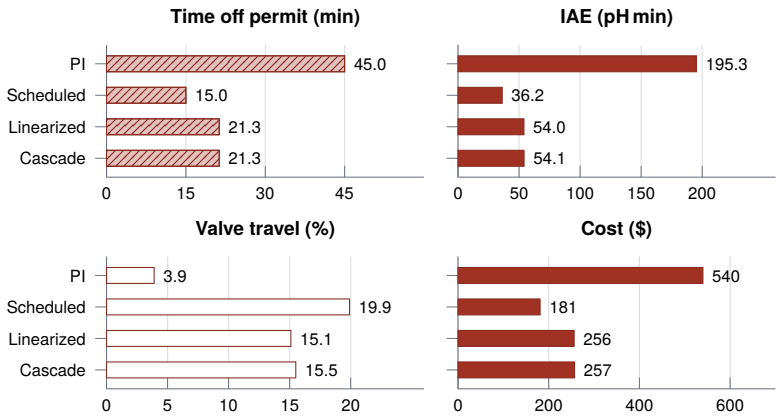


Figure 6.3 The four architectures of Table 6.4 after the influent strong acid step at $t = 25$ min, read at $t = 70$ min: time off permit (hatched bars), integrated absolute error and cost (solid) and valve travel (open). Notice that the fixed-gain PI loop moves the valve least and is worst on every other measure, and that the cascade matches the linearized loop because this disturbance never touches the reagent flow.

Table 6.5 Measured Performance After the Influent Weak Acid Steps From 0.6 to 10 mmol/L at $t = 15$ min, Read at $t = 80$ min. The Demand the New Influent Requires Is 53.130 % of Valve.

	PI	Scheduled	Linearized	Told the truth
K_c	0.092	0.092	3.0	3.0
Assumed $\hat{c}_b$ (mmol/L)	0.6	0.6	0.6	10
Final pH	4.06	6.69	5.71	7.00
Final demand (%)	23.67	52.97	50.36	53.13
Return inside permit (min)	never	31.9	never	10.4
Time off permit (min)	65.0	31.9	65.0	12.6
IAE (pH·min)	198.8	74.0	125.8	35.6
Valve travel (%)	3.8	33.4	30.5	133.5

DID YOU KNOW? The Cheapest Caustic Bill, the Dearest Run

From the simulator: after the acid step, the fixed-gain loop buys the least caustic, 1.34 kg against 1.81 to 1.82 kg for the other three structures, and still costs the most, \$540 against \$181 to \$257, because every minute off permit is charged (Table 6.4).

6.6.4 The Disturbance That Defeats Every Structure

Scenario 8 switches the upstream product at $t = 15$ min. Nothing at the inlet changes and the reagent demand nearly triples. Table 6.5 reports where each configuration has reached 65 min later.

The fixed-gain loop never arrives, reaching 23.67 % of the 53.13 % it needs. The scheduler reaches 52.97 % and the permit 31.9 min after the step. The linearized loop reaches 50.36 % and is still climbing at pH 5.71, which is better than the fixed gain and slower than the schedule, because its inverse is built from the clean curve. On that curve the tank at its low point of pH 4.31 appears to be 1.7 % of valve short of the setpoint when the demand is really 26 % short. Setting the assumed weak acid to 10 mmol/L, which is the same structure told the truth, returns the discharge to the permit 10.4 min after the step and holds pH 7.00 at exactly 53.13 % of valve. The last column is the argument for estimating the curve on line rather than fixing it [10, 15].

Of the 12.6 min off permit in the last column, 2.2 min accrue before the step. A controller told that the influent is buffered when it is not applies about sixteen times the intended loop gain on the clean stream's cliff, so it limit-cycles between pH 6.31 and 9.21 and crosses the caustic side of the permit. That is the same defect in the other direction, and it is worth showing.

The return of the clean stream at $t = 80$ min defeats every architecture equally, because 53 % of valve of caustic is entering a stream that now needs 20 %. Running both events in manual with a perfect feedforward, which changes the demand at the same instant as the composition and which no plant can build, still carries the tank down to pH 5.37 on the buffering step and up to pH 10.16 on the return, and leaves 13.9 min outside the permit over the run. That excursion belongs to the tank and not to the controller. No feedback structure removes an unmeasured composition change, and designing for graceful degradation under one is a legitimate goal in its own right [14].

6.6.5 What the Inner Loop Does and Does Not Reject

Scenario 11 drops the caustic header pressure to 0.6 of nominal at $t = 15$ min, which by (6.8) reduces the reagent flow by 23 % without the stem moving at all. With the inner flow loop closed the transmitter reports the change within seconds and the valve opens to compensate, and the peak pH excursion is 0.91 pH. With the same outer loop and no inner loop the pH controller has to discover the disturbance through five minutes of tank, the 0.15 min transport delay and the 0.20 min electrode, and the peak excursion is 2.25 pH. Over the 90 min run the single loop is off permit for 50.9 min with an IAE of 121.7 pH·min against 15.3 min and 39.1 pH·min for the cascade. A caustic delivery arriving 15 % weaker at $t = 70$ min is invisible to the inner loop, because the flow is exactly what was asked for, and it is equally invisible to the inverse of (6.16), which carries the nameplate concentration. The demand setpoint does not move, the tank falls 1.92 pH below the setpoint at $t = 72.6$ min, and only the outer loop recovers it through the pH. The general rule is that an inner loop rejects everything appearing in its own measurement and nothing else. Volume 3, Chapter 2 develops that rule on a level loop with two disturbance paths.

🕮 **Pause.** Try Exercise 6.3 before reading on.

6.6.6 Fouling and the Graded Challenge

Scenario 10 runs seven hours with nothing wrong except a film building on the glass at 14 % per hour. The controller holds the reading at 7.00 for the whole shift while the true discharge drifts to pH 7.886, because an electrode reading low makes the controller overdose, and the reagent demand moves by 0.014 % of valve over the whole shift. The run never leaves the permit and no permit alarm fires. The only warning, from $t = 235.7$ min when the coating passes 55 %, reports that the electrode needs cleaning and not where the discharge has gone. The second effect is less visible and more dangerous, because at 98 % fouling the electrode lag of (6.10) has grown from 0.20 to 1.96 min, a factor of ten of pure phase lag inside the loop. Cleaning the electrode at $t = 430$ min then steps the measurement by 0.88 pH, which the controller cannot distinguish from the tank moving. The reading jumps to pH 7.76 and the controller cuts the caustic, so a maintenance action moves the valve and the tank as a process upset would, here carrying the true pH from 7.88 to 7.10 within 2.2 min.

The challenge applies four unannounced composition changes over four hours, a buffering step at $t = 20$ min, an acid load rise at $t = 60$ min, an acid load fall at $t = 100$ min and the return of the clean stream at $t = 130$ min, with the valve slowed to 6 %/min and anti-windup switched off in the starting configuration. The case is passed with more than 45 % of the run inside the permit and a final pH within ± 1.0 of the setpoint. Table 6.6 reports three configurations over the full 240 min.

The answer is the whole stack and the scorecard says what each layer is worth. Linearization is required because two of the four upsets change the shape of the curve. Cascade is required because a slew-limited valve is the actuator an inner flow loop was invented for. Anti-windup is not a refinement here, since switching back-calculation off takes the same cascade from 49.7 to 10.2 % in permit, ends the run at pH 2.49 rather than pH 7.00, and multiplies the valve travel by 3.9 while raising the cost

Table 6.6 Graded Challenge Over 240 min With a 6 %/min Valve, Four Unannounced Composition Changes, and Every Figure Read at the End of the Run.

	As loaded	Suggested	Suggested, no anti-windup
Architecture	PI	cascade	cascade
K_c, τ_I (min)	0.092, 5.2	3.0, 4.0	3.0, 4.0
Anti-windup	off	back-calc.	off
In permit (%)	8.5	49.7	10.2
Final pH	11.06	7.00	2.49
IAE (pH·min)	836	291	729
Valve travel (%)	15	134	524
Caustic (kg)	4.9	6.8	7.2
Cost ($)	2639	1453	2591

from $1453 to $2591. Half the afternoon is spent outside the permit even when everything is right, because the four upsets move the demand the setpoint requires through 19.90, 53.13, 67.35, 47.80 and 14.57 % of valve, a total of 100 % of travel on a stem that moves at 6 % a minute into a tank of 400 L. The honest engineering answer to that plant is a bigger tank, a smaller valve and a feedforward from an influent analyzer, and knowing that no control structure buys them is part of the lesson.

Figure 6.4 shows the application after that same buffering scenario has run for three hours on the fixed-gain architecture.

6.7 Technology in Practice

The loops of Section 6.5 read AIT-401 through a delay, a lag and a fouling term, and move FCV-401 through a lag, a rate limit and a header pressure. Figure 6.5 draws them with the devices a neutralizer would carry in place of the schematic's symbols.

What Measures It

AIT-401 is a glass electrode with a reference electrode, wired to a transmitter that converts its millivolts to pH. Dipping the probe into an open vessel of the liquid is a common installation in water treatment [18]. The transport delay θ of Equation (6.10) stands for the pipe between the injection point and the probe, the lag τ_m for the response of the glass, and the fouling terms $\kappa\phi$ and $\delta\phi$ for a coating that slows the electrode and shifts its reading. Calibration uses buffers that bracket the expected range, most often pH 4, 7 and 10 [18]. No accuracy or response time is given here, because none was found in an open source that names no maker.

> **TECHNOLOGY Sensor: The Glass pH Electrode**
>
> A glass electrode ends in a thin bulb that passes hydrogen ions and is filled with a buffer at pH 7. A difference between the process liquid and that buffer sets up a voltage across the glass in proportion to the pH difference, and a reference electrode completes the circuit. The slope is 59.17 millivolts per pH unit at 25 °C and 68.11 at 70 °C, so the probe carries a temperature sensor for compensation. The glass must stay wet, and the same hydration slowly wears it away, so every electrode has a limited life. At high pH, sodium ions add a voltage of their own. In dirty service such as wastewater the bulb needs regular cleaning, with no solid tool touching the fragile glass [18].

What Moves It

In a plant, FCV-401 is a caustic valve whose actuator is driven through a positioner, a closed-loop controller in its own right that holds the stem at the commanded position [18]. The lag τ_v and the slew limit $\dot{v}_{\max}$ of Equation (6.7) stand for that actuator, and the term $\sqrt{p_h}$ of Equation (6.8) for the

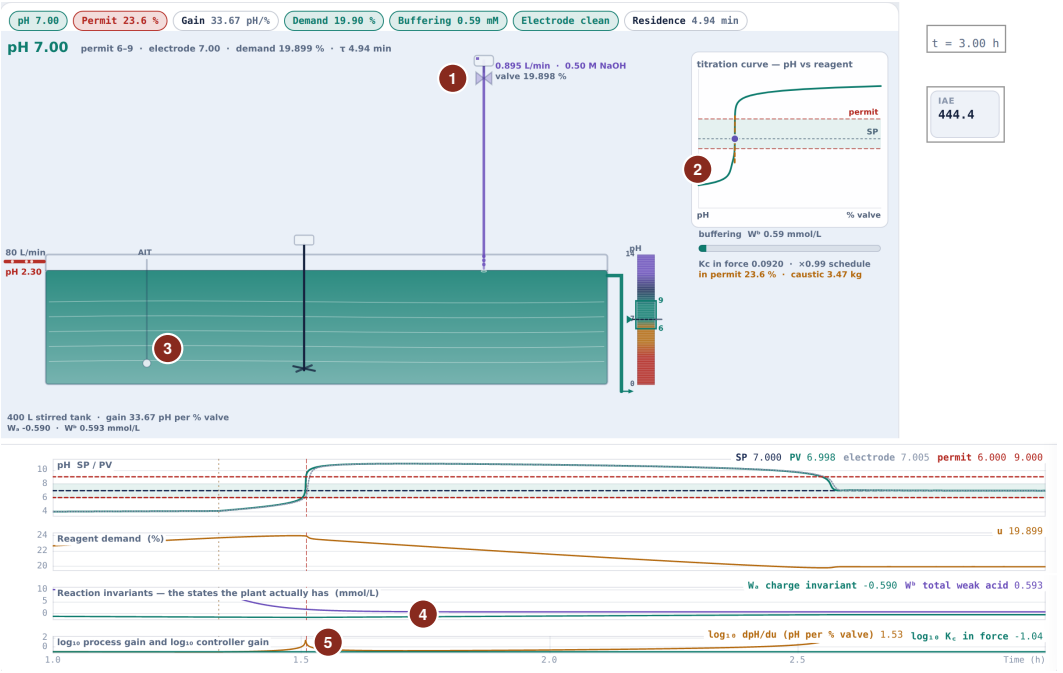


Figure 6.4 Scenario 8 at $t = 3.00$ h (framed readout), after the buffering step at $t = 15$ min and the return of the clean stream at $t = 80$ min, with the fixed-gain PI architecture. The figure is a composite: the plant view above four of the six trend strips, with the elapsed time and the IAE tile moved beside the plant view. The tank has settled back to pH 6.998 at a reagent demand of 19.899 %, and the scorecard records 23.6 % of the run inside the permit, 3.47 kg of caustic and an IAE of 444.4 pH·min. The titration curve beside the tank carries the operating point and the shaded permit band, the invariant strip shows W_b dropping from 10 (Table 6.5) back to 0.593 mmol/L when the clean stream returns, and the log-gain strip shows the process gain swinging while the controller gain holds a straight line. Callouts: (1) the caustic valve FCV-401 and its flow, (2) the operating point on the titration curve, (3) the electrode AIT-401, (4) the invariant strip, (5) the log-gain strip.

supply header behind it. FT-401 measures the caustic flow for the inner loop of the cascade, which is why a header pressure change is rejected there before it reaches the tank.

> **TECHNOLOGY** **Final Element: The Caustic Valve and the Acid Valve**
>
> Where a stream can stray to either side of neutral, one pH controller can drive two reagent valves in exclusive sequence. Both are shut at half output, the acid valve opens as the output rises toward full scale and the caustic valve as it falls toward zero [18]. The chapter's acid trim pump plays the acid side. An air-to-open valve fails closed when its air supply is lost, and the risks of the process, not habit, should decide which way each valve fails [18]. On this tank a caustic valve that fails shut sends the influent to the outfall at the acidity it arrived with.

Where the Controller Runs

A distributed control system runs production at one site in water and wastewater treatment and in chemical manufacturing, and its loops run with cycle times from milliseconds to minutes [19]. AIC-401 and FIC-401 execute every 0.02 min (1.2 s), well inside that range. The inverse titration curve of the linearized architecture fits a multi-segment characterizer, a function block of such a

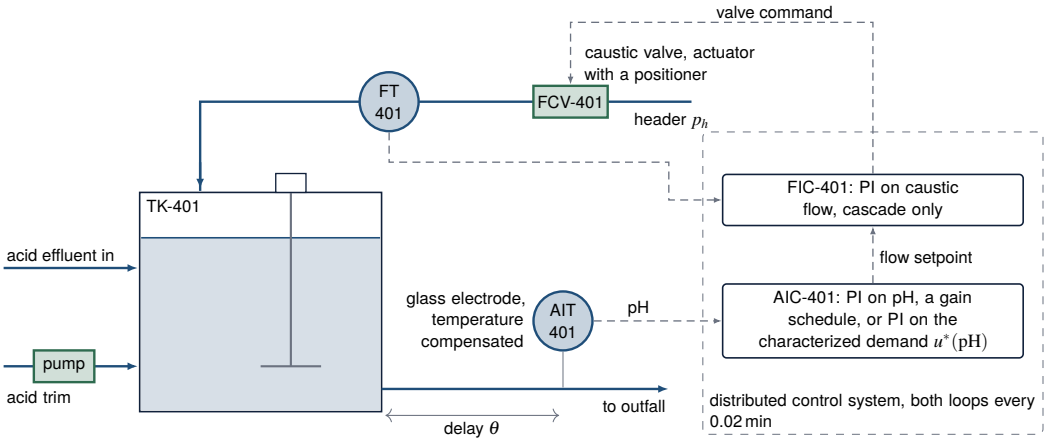


Figure 6.5 The loops of the chapter with plant devices in place of the schematic's symbols. AIT-401, a glass electrode on the discharge, reads the pH after the transport delay θ. AIC-401 on the control platform acts on the pH directly, through a gain schedule, or through the characterized demand, and in the cascade it sets the flow that FIC-401 holds by moving FCV-401 against FT-401. Without the cascade, AIC-401 writes the valve command itself. Notice that the caustic enters at the top of the tank while the pH that the permit judges is read downstream, after mixing and transport.

system that can be programmed to follow almost any continuous curve [18]. The gain schedule is the same PID block with its gain set by another variable, here the slope of the assumed curve at the measured pH.

> **IN PRACTICE** **Mapping the Gain Before Tuning**
>
> An instrument handbook advises exploring a process before tuning it. With the controller in manual, step the output several times across the safe range and compare the size of each response. If the gain grows from one end of the output range to the other, a valve with a different characteristic can counteract it. If the gain follows the measured value instead, the remedy is an adaptive gain controller, whose proportional setting varies as the process changes [18]. A pH loop is the second kind, and its manual steps have to shrink as the tank nears neutral.

What Is Modernizing It

Three developments touch this loop. The first is fieldbus, which lets the PID function block run inside the field devices, in the transmitter or in the valve positioner [18]. The second is nonlinear predictive control. Its industrial models are usually linear dynamic subsystems with an added nonlinearity [20], which is the shape of this plant, two linear invariants and one algebraic map to pH. The third is self-tuning PID, which has done well in the market [20], although the outer loop here would need one that follows a gain moving by four orders of magnitude.

6.8 Limitations

The model is not a design model for a neutralization plant. One perfectly mixed tank cannot support a claim about how many stages a real effluent needs, and plants that hold a tight permit use two or three in series precisely because one cannot [1]. There is no ratio control and no influent analyzer, so the first structural improvement a real plant would make is absent, and scenario 9 would largely disappear with it. Volume 3, Chapter 6 computes a dose from an upstream analyzer and keeps the downstream measurement as a slow trim. Carbonate, temperature and activity effects are all outside the chemistry, and each of them moves the titration curve on a timescale comparable with the control

problem. The valve has no stiction, which matters because a stiction limit cycle looks like the one in scenario 4 and has an entirely different cause. Electrode fouling is one linear state and cannot produce the failure that matters most in service, which is a probe that stops responding and reports a flat and plausible trace. Economics are a caustic price and a flat charge per minute off permit, whereas a real consent order is triggered by a grab sample, so the cost tile compares two runs rather than predicting compliance. Repeatability is itself unrealistic, because freezing the integration step and the controller period reproduces one trajectory on every machine while a distributed control system does no such thing. Where the individual assumptions break is set out in Section 6.3.5.

Two elements of the scenarios are stylized on purpose. A buffering change that arrives as an instantaneous step is not what an upstream product change does, but a step is the disturbance whose response can be read off a trend. A valve slowed to 6 %/min in the challenge is slower than a healthy actuator, although it is the only way to make rate saturation visible on a plant whose dominant time constant is five minutes.

> ⚠ **OPEN CHALLENGE** **A Controller That Learns Its Own Curve**
>
> Both the schedule and the inverse rest on an assumed titration curve, and buffering is not measured on line. A controller that found the curve by itself would remove it, and a simple universal controller that adjusts itself has long been a goal of process engineers, with the relay autotuner of PID control only a partial answer [9]. A survey of industrial predictive control found that only two adaptive versions had reached the market despite a strong incentive, and it did not expect that to change without a theoretical breakthrough. Self-tuning PID controllers, by contrast, have sold well [20]. Progress would be an estimator that tracks the curve from the dosing record and the pH alone, and that an operator could trust.

6.9 Discussion

The three structural remedies in this application are not alternatives of equal standing. Scheduling leaves the controller working in pH, where an error of one unit is a few micromoles per liter near the setpoint and several millimoles per liter out on the flat part, so the proportional term is corrected while the integral accumulates a quantity with no consistent meaning. Inversion corrects the whole error signal and reduces to one characterizer block on the transmitter in a modern control system, which is the same function block that linearizes a valve pointed at a measurement instead. Cascade corrects neither nonlinearity and removes the actuator from the outer loop, which is why Table 6.4 shows it as indistinguishable from inversion on a composition change and Section 6.6.5 shows it worth a factor of 2.5 in peak excursion on a header pressure change. The order in which they are introduced is the order in which a plant learns them, and each is measured against the same scorecard. Volume 7, Chapter 3 sets a gain schedule against feedback linearization on a mechanical plant.

The equipment decisions dominate all three. The reagent valve passes 4.5 L/min against a duty of 0.8955 L/min, which is five times oversized and is normal, because the valve is specified for the worst influent anyone imagined. Halving the caustic concentration from 0.50 to 0.25 mol/L moves the design demand from 19.899 to 39.798 % of travel and widens the permit band from 0.155 to 0.312 %, at the cost of twice the reagent volume. Fitting equal-percentage trim makes the installed gain far more nearly constant across the range without any control system change at all. The application exposes the capacity, the rangeability, the caustic strength and the tank volume as adjustable parameters so that the trade between capital and control structure can be computed rather than argued, and the outcome of that computation is often that the money is better spent on the valve than on the algorithm.

The model also supports methods beyond the implemented loops. The two linear invariant states with an algebraic output map are the natural form for a nonlinear model predictive controller, which enforces the valve rate limit of the challenge explicitly instead of defending against it with back-calculation. Moving horizon estimation on the same equations can infer the unmeasured c_a and c_b from the pH measurement and the reagent demand, which is the on-line titration curve estimate that Table 6.5 shows to be worth 52 min of permit time [10, 15]. Identifying the static map from data

rather than from equilibrium constants gives the Wiener structure that has been applied to the same plant [11]. Nothing in the model has to change for any of them, because the difficult part was already separated out into one algebraic line.

6.10 Conclusion

This chapter studied an effluent neutralization tank as an interactive laboratory for plants whose gain is a function of the operating point rather than a constant. The simulator integrates five states, the two reaction invariants, the valve stem, the influent pump and the electrode fouling fraction, and recovers the pH each step by bisecting a charge balance, so the plant is linear and the measurement is not. The reagent valve holds pH 7.00 at 19.899 % of travel and the whole permit band of pH 6 to 9 occupies 0.155 % of it, while the gain from valve to pH runs from 5.191×10^{-3} to 101.98 pH per percent across one stroke, a factor of 1.96×10^{4}.

Settings obtained correctly from a step test at 12 % of valve produce a limit cycle between pH 4.50 and pH 10.97 at the setpoint with 5231 % of valve travel in an hour, and settings obtained correctly at the setpoint leave the tank at pH 2.62 three quarters of an hour after a load change. Scheduling the gain returns the discharge to the permit 15.0 min after that load change and inversion of the titration curve does so in 21.3 min, against never for the fixed gain. An inner reagent flow loop reduces the excursion from a header pressure change from 2.25 to 0.91 pH and does nothing about a weaker caustic delivery.

A fuller model would add a second tank in series, ratio control from the influent flow meter, valve stiction as an independent source of cycling, and a carbonate species so that the curve has more than one inflection. The pH neutralization tank is an accessible browser platform for learning when a loop needs a different structure rather than a better tuning.

KEY POINTS

- The tank mixes caustic smoothly, and only the pH scale turns that smooth change into a cliff.
- The whole permit of pH 6 to 9 is 0.155 % of the reagent valve's travel, and across one stroke the gain changes by a factor of 1.96×10^{4}.
- A step test measures the gain where it is made: the same procedure returns gains that differ by a factor of 691.
- After the acid step, scheduling returns the discharge to the permit in 15.0 min and inversion in 21.3 min. The fixed gain never does.
- An inner flow loop rejects a header pressure change but not a weaker caustic delivery.

Where else this applies. The gain of Volume 7, Chapter 3 also moves with its operating point, and there a schedule meets feedback linearization.

References

[1] F. G. Shinskey, *pH and pIon Control in Process and Waste Streams*. New York, NY: John Wiley & Sons, 1973.

[2] G. K. McMillan and R. A. Cameron, *Advanced pH Measurement and Control*, 3rd ed. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2004.

[3] U.S. Environmental Protection Agency, "Secondary treatment, 40 CFR 133.102," Code of Federal Regulations, Title 40, Volume 24, U.S. Government Publishing Office, govinfo.gov, 2024.

[4] T. K. Gustafsson and K. V. Waller, "Dynamic modeling and reaction invariant control of pH," *Chemical Engineering Science*, vol. 38, no. 3, pp. 389–398, 1983.

[5] ——, "Myths about pH and pH control," *AIChE Journal*, vol. 32, no. 2, pp. 335–337, 1986.

[6] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

[7] C. Kravaris and C.-B. Chung, "Nonlinear state feedback synthesis by global input/output linearization," *AIChE Journal*, vol. 33, no. 4, pp. 592–603, 1987.

[8] R. A. Wright and C. Kravaris, "Nonlinear control of pH processes using the strong acid equivalent," *Industrial & Engineering Chemistry Research*, vol. 30, no. 7, pp. 1561–1572, 1991.

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

[10] M. A. Henson and D. E. Seborg, "Adaptive nonlinear control of a pH neutralization process," *IEEE Transactions on Control Systems Technology*, vol. 2, no. 3, pp. 169–182, 1994.

[11] A. D. Kalafatis, L. Wang, and W. R. Cluett, "Linearizing feedforward–feedback control of pH processes based on the Wiener model," *Journal of Process Control*, vol. 15, no. 1, pp. 103–112, 2005.

[12] T. J. McAvoy, E. Hsu, and S. Lowenthal, "Dynamics of pH in controlled stirred tank reactor," *Industrial & Engineering Chemistry Process Design and Development*, vol. 11, no. 1, pp. 68–70, 1972.

[13] R. C. Hall and D. E. Seborg, "Modelling and self-tuning control of a multivariable pH neutralization process part I: Modelling and multiloop control," in *1989 American Control Conference*. IEEE, 1989, pp. 1822–1827.

[14] T. K. Gustafsson and K. V. Waller, "Nonlinear and adaptive control of pH," *Industrial & Engineering Chemistry Research*, vol. 31, no. 12, pp. 2681–2693, 1992.

[15] M. Gulaian and J. Lane, "Titration curve estimation for adaptive pH control," in *1990 American Control Conference*. IEEE, 1990, pp. 1414–1419.

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

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

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

[20] 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.

6.11 Exercises

6.1 Concept check · 2 min

The step tests at 12 % of valve and at the setpoint return gains of 0.082 and 56.66 pH per percent, while the invariant W_a moves by the same 0.28 mmol/L per percent of valve in both. Where does the difference between the two gains arise?

(a) In the charge balance, an algebraic map from the invariants to pH that is steepest near neutrality

(b) In the mixing of the tank, whose time constant shortens as the contents approach the neutral point

(c) In the reagent valve, whose linear trim passes more caustic per percent near the design travel

(d) In the electrode, whose lag and transport delay stretch a small change into a large reading

6.2 Concept check · 1 min

The influent weak acid rises from 0.6 to 10 mmol/L while the strong acid stays at 5 mmol/L. What does the model predict for the inlet pH, the reagent demand at pH 7 and the process gain there?

(a) The inlet pH falls by about one unit, and the demand and the gain both stay nearly as they were

(b) The inlet pH barely moves, the demand nearly triples and the gain falls by a factor of about 16

(c) The inlet pH barely moves, the demand stays nearly fixed and the gain falls by a factor of about 16

(d) The inlet pH barely moves, the demand nearly triples and the gain rises by a factor of about 16

6.3 Concept check · 2 min

In scenario 11 the cascade meets a drop of the caustic header pressure to 0.6 of nominal and, later, a caustic delivery that is 15 % weaker. Which of the two does the inner reagent flow loop reject?

(a) Both, because each one changes the amount of base that reaches the tank per percent of valve travel

(b) Neither, because the lag of the flow transmitter hides both from the inner loop until the pH has moved

(c) The pressure drop only, because it changes the measured flow and the weaker caustic leaves that flow as requested

(d) The weaker caustic only, because the inverse of the titration curve carries the true strength of the delivery

6.4 Theory and calculation · 5 min

At the design point the influent is $Q_1 = 80\,\mathrm{L/min}$ with $c_a = 5.0\,\mathrm{mmol/L}$ and $c_b = 0.6\,\mathrm{mmol/L}$, the caustic is $C_b = 0.50\,\mathrm{mol/L}$, the valve passes 4.5 L/min at full travel and the acid trim is off. Take $\mathrm{p}K_a = 4.76$ and $\mathrm{p}K_w = 14$.

(a) Show from Equation (6.4) that $W_a = -f\,W_b$ at pH 7, and find $f = K_a/(K_a + h)$. *Checkpoint:* $f = 0.9943$.

(b) Use the steady forms of Equations (6.2) and (6.3) to find the reagent flow and the valve travel that hold pH 7. *Checkpoint:* 0.8955 L/min and 19.90 %.

(c) Repeat part (b) for $c_b = 10\,\mathrm{mmol/L}$. *Checkpoint:* 2.391 L/min and 53.13 %.

6.5 Theory and calculation · 8 min

The process gain at the setpoint follows from the buffer term. Take pH 7, $K_a = 1.738 \times 10^{-5}$, $K_w = 10^{-14}$, $C_b = 0.50\,\mathrm{mol/L}$ and a valve of 4.5 L/min with linear trim at nominal header pressure.

(a) At the design point the tank holds $W_b = 0.593\,\mathrm{mmol/L}$. Find J from Equation (6.5). *Checkpoint:* $J = 35.7$.

(b) Find $\partial\mathrm{pH}/\partial W_a$ from Equation (6.6). *Checkpoint:* -1.215×10^5 pH per mol/L.

(c) Find the gain from valve travel to pH with Equation (6.11) and $Q_t = 80.90\,\mathrm{L/min}$. Use $W_a = -f\,W_b$ with $f = K_a/(K_a + h)$ to simplify the bracket. *Checkpoint:* 33.8 pH per percent.

(d) With $c_b = 10\,\mathrm{mmol/L}$ the tank holds $W_b = 9.710\,\mathrm{mmol/L}$ and $Q_t = 82.39\,\mathrm{L/min}$. Find J and the gain. *Checkpoint:* $J = 554$ and 2.14 pH per percent.

6.6 Theory and calculation · 5 min

The step test at 12 % of valve gives a process gain of 0.082 pH per percent, $\tau_p = 4.94\,\mathrm{min}$ and $\theta_p = 0.35\,\mathrm{min}$.

(a) Find K_c from Equation (6.12) with a closed-loop target of 1.5 min, and the loop gain $K_c K_p$ that the rule designs for. *Checkpoint:* 32.6 percent of valve per pH unit, and 2.67.

(b) The loop is closed at the setpoint with $K_c = 33$, where the tangent gain is 33.78 pH per percent. Find the loop gain. *Checkpoint:* 1115.

(c) With $K_c = 0.092$ the tank sits at pH 2.5 after a load step, where the gain is 0.0384 pH per percent. Find the loop gain there and at the setpoint. *Checkpoint:* 0.0035 and 3.11.

6.7 App challenge · 8 min

Load scenario 2, A step test on the flat part, and change no setting. The reagent demand steps from 12 to 13 % at 10 min. Run until just past 40 min with the default seed and export the CSV file. The column `ph` holds the true tank pH.

(a) Read `ph` in the rows at 10 and 40 min and compute the gain. *Checkpoint:* 2.790 and 2.872, within 0.002, and 0.082 pH per percent.

(b) Find the time after the step at which `ph` has covered 63.2 % of that change. *Checkpoint:* 5.3 min, within 0.1 min.

(c) Load scenario 3, The same step test at the setpoint, whose step is 0.03 % of valve at 10 min. Run until just past 40 min, read the same two rows and compute the gain and its ratio to the gain of part (a). *Checkpoint:* 7.003 and 8.701, 56.6 pH per percent, a ratio near 690.

Compare the change of the column `Wa_mmol` per percent of valve in the two files.

6.8 App challenge · 8 min

Load scenario 5, A gain tuned at the setpoint, and what it costs elsewhere. The influent strong acid steps from 5 to 9 mmol/L at 25 min. Change no setting, run until just past 70 min with the default seed and export the CSV file.

(a) Find the lowest `ph` after the step. In the row at 70 min read `ph`, the reagent demand `demand_pct` and the time off permit `metric_offspec_min`. *Checkpoint:* 2.54. Then 2.62, 23.73 % and 45.0 min.

(b) From `demand_pct` at 40 and 60 min compute the climb rate. Estimate how long after 70 min the demand would need to reach the 34.12 % that the new load requires. *Checkpoint:* 0.078 % per minute and about 130 min.

(c) Load scenario 6, Gain scheduling, and repeat the run. Find the lowest `ph`, the time off permit at 70 min and the largest value of the gain in force, `kc_eff`. *Checkpoint:* 4.16, 15.0 min and 4.33.

Accept a pH within 0.02, a demand within 0.05 % and a time off permit within 0.2 min.

6.9 App challenge · 9 min

Load scenario 8, The buffering changes without warning, and set the Control panel to Advanced. The influent weak acid steps from 0.6 to 10 mmol/L at 15 min. Use the default seed, run each case until just past 80 min and export the CSV file.

(a) Set Control architecture to Linearized and Controller gain to 3, then run. In the row at 80 min read `ph`, `demand_pct` and `metric_offspec_min`. *Checkpoint:* 5.71, 50.36 % and 65.0 min.

(b) In the group for the controller's titration curve, set the weak acid of the titration model to 10 mmol/L. Press Reset, run again and read the same row. *Checkpoint:* 7.00, 53.13 % and 12.6 min.

(c) In the file of run (b) read `metric_offspec_min` in the row at 15 min, before the batch changes, and explain it. *Checkpoint:* 2.2 min.

Accept a pH within 0.02, a demand within 0.05 % and a time within 0.2 min.

6.10 Application area · 3 min

The reagent valve passes 4.5 L/min at full travel, and the design influent needs 0.8955 L/min of 0.50 mol/L caustic at pH 7. The permit of pH 6 to 9 spans 0.155 % of valve travel.

(a) Find how many times the valve is oversized and the width of the permit as a reagent flow. *Checkpoint:* 5.0 times and 7.0 mL/min.

(b) The plant changes to 0.25 mol/L caustic. Estimate the new demand and the new width of the permit in valve travel, and name the cost. *Checkpoint:* 39.8 % and about 0.31 %.

6.11 Application area · 4 min

Scenario 10 fouls the electrode at 14 % per hour for a shift of 420 min. A fully fouled electrode reads 0.9 pH low, and its lag is the clean 0.20 min multiplied by $1 + 9\phi$.

(a) Find the time at which the fouling fraction ϕ passes 0.55, where the app raises its electrode warning, and ϕ at the end of the shift. *Checkpoint:* 235.7 min and 0.98.

(b) Find the electrode lag and the reading offset at the end of the shift, and the true pH while the controller holds the reading at 7.00. *Checkpoint:* 1.96 min, -0.88 pH and pH 7.88.

(c) Explain why no permit alarm fires and why the pH trend that the controller sees shows nothing wrong.

6.12 Review the agent · 5 min

An AI assistant is asked to move the fixed-gain loop of scenario 5 to the linearized architecture and replies as follows. "The IMC rule at the setpoint gave $K_c = 0.092$ with $\tau_I = 5.2$ min, and the loop is stable there. Linearization changes only the signal that the controller reads. The tank, the valve and the electrode are the same, so the time constant and the dead time are unchanged. Keep $K_c = 0.092$ and $\tau_I = 5.2$ min, select Linearized, and the acid load step will be rejected about as fast as the gain schedule rejects it." Find the flaw in the proposal and say how it could be checked.