INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 6

Municipal Water Treatment Simulator for Feedforward Control of Coagulant Dosing

Pages 125–144

Edited by John D. Hedengren

How to cite this chapter. "Municipal Water Treatment Simulator for Feedforward Control of Coagulant Dosing," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 125–144.

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

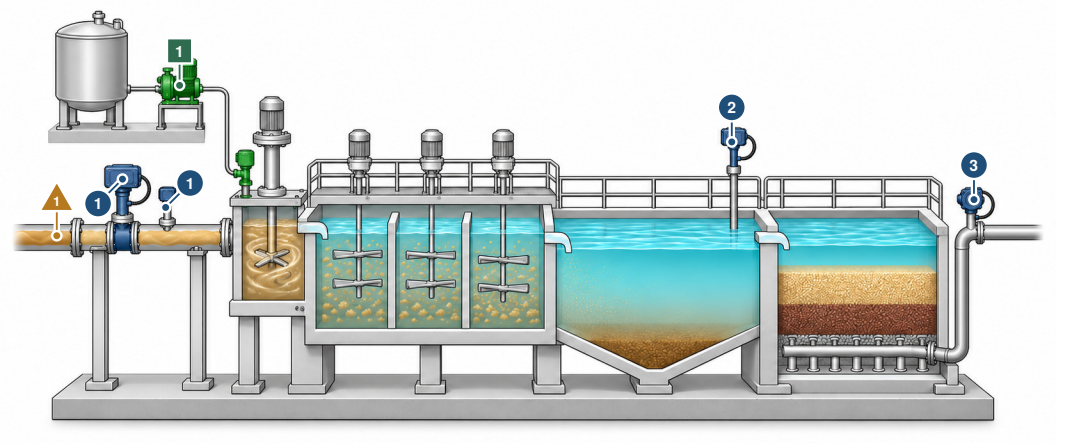
The app for this chapter
https://apmonitor.com/apps/water-treatment/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v03-feedforward-cascade/#water-treatment

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

CHAPTER 6

Municipal Water Treatment Simulator for Feedforward Control of Coagulant Dosing



The controller moves 1 alum metering pump

The controller reads 1 raw-water turbidity and temperature 2 settled-water turbidity 3 filtered-water turbidity

What disturbs it 1 storm turbidity and organic carbon in the raw water

The raw-water meters read a storm before the dosing point, while the settled-water meter waits behind the whole basin. Illustration, not to scale.

IN PLAIN TERMS Dose the Water Before the Storm Arrives

The main idea. A water plant adds a chemical that makes fine mud clump together so it can settle out. The right amount depends on how muddy the river water is as it comes in.

The challenge. The meter that judges the dose sits about three hours downstream, but a storm can muddy the river in half an hour. A controller that waits for that meter acts far too late, and the extra chemical it adds then makes the water worse.

Why it matters. Measuring the river water as it enters lets the plant set the dose in time, and the slow meter only fine-tunes it.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history of the filtered-water rules (Section 6.1), the fact on the by-product limit (Section 6.3), the try-it segments of scenarios 4 and 8 (Section 6.5), the fact on alum use (Section 6.8) and the key points
Study	40 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 6.5), Technology in Practice (Section 6.6), Limitations (Section 6.7) and the open challenge

The control question The settled-water turbidimeter that judges the coagulant dose reads three hours after the dosing pump, and a storm raises the raw turbidity fivefold in half an hour. Where should the dose be computed, and what is left for the slow measurement, the pH loop and the filter sequence to do?

Learning objectives

- Compute the coagulant demand, the dose ratio and the destabilized fraction from the model for a given raw water and dose.
- Explain why a feedback loop on the settled water crosses the dose optimum, and read the sign change of the process gain off the dose curve.
- Compare feedback with feedforward plus a multiplying trim on a storm, and say what the trim corrects that the feedforward cannot see.
- Judge the worth of a fast pH loop and a filter-to-waste timer from measured runs, and the cost of an integrating trim left on a fault.
- Locate the window between the disinfection credit and the by-product limit, and name the variable that moves it.

Prerequisites Sections 1.4.2 (where a disturbance enters), 1.4.3 (static feedforward), 1.4.4 (ratio with a trim), 1.6.2 (velocity form) and 1.6.4 (scenarios and exported files).

The app https://apmonitor.com/apps/water-treatment/

See also Volume 5, Chapter 4 and Volume 5, Chapter 3 for control ahead of a slow analyzer, and Volume 5, Chapter 7 for diverting material judged on a delayed measurement. Volume 8, Chapter 13 for a diversion valve under supervisory logic, and Volume 8, Chapter 12 for discrete units that fall into step on one shared header.

6.1 Introduction

The coagulant dose is the most important decision a conventional surface water plant makes, and the instrument that reports whether the decision was right sits at the far end of a flocculation basin and two sedimentation stages.

The simulator measures the separation. Stepping the alum from 31.3 to 38 mg/L in manual with everything else fixed moves the settled water from 1.200 to 2.435 NTU. Two percent of that change appears after 32.25 min, sixty-three percent after 193.6 min and ninety-five percent after 402.7 min. That is the plant the coagulant controller is asked to close a loop around.

The disturbance is faster by two orders of magnitude. A storm on the intake raises the raw turbidity from 12 to 60 NTU in half an hour, and a plant at 1200 m^3/h produces 600 m^3 of water in that half hour. By the time the settled turbidimeter has reported the first two percent of the problem, the water that caused it is already in the clearwell.

No tuning changes that arithmetic, and Section 6.5.2 shows that tightening the loop makes matters actively worse rather than merely insufficient. What changes it is that the disturbance is *measured*, entry point 2 of Section 1.4.2. A raw-water turbidimeter reads in seconds, costs a few thousand dollars and is installed on essentially every plant of this size. The coagulant demand is a known function of what it reports. Computing the dose from the disturbance and demoting the settled turbidimeter to a slow ratio trim takes the peak settled turbidity through the identical storm from 20.918 NTU to 3.600 NTU and the time above the filtered water specification from 141.9 min to none. Industrial practice has converged on the same structure, with the dose computed from online raw water measurements and corrected slowly [1].

Two further properties make this plant richer than a pure dead-time example. The coagulant loop's process gain reverses sign at the dose optimum, and a quarter of the plant is a discrete sequence rather than a feedback loop.

The physical fundamentals behind the model are the ordinary ones. Each unit is a material balance on a well-mixed volume, so its time constant is the hydraulic residence time and the train is a chain of first-order lags in series, the structure a first course develops as higher-order and apparent-dead-time behavior [2, 3]. The driving force for clarification is the gravitational settling velocity of the floc

against the surface overflow rate, and the driving force for disinfection is the product of free chlorine residual and contact time. The separation of time scales is what organizes the control problem, because the pH loop closes in minutes, the chlorine loop in tens of minutes and the coagulant loop in hours. The experiments built on that structure cover dead time and high-order lag, feedforward from a measured disturbance, ratio control, sign-changing process gain, cascade and trim structure, integral windup against a fault outside the loop's authority, sampled and filtered measurements, and the interaction between discrete sequencing and continuous loops [2, 4]. The behavior matters industrially because turbidity is the regulated surrogate for pathogen removal, so a coagulation excursion is a reportable compliance event rather than a quality inconvenience, and because coagulant is the largest chemical line on a conventional plant's operating budget.

> ⧖ HISTORY **From 0.5 to 0.3 NTU, 1989 to 2002**
>
> A federal rule of 1989 required surface-water systems in the United States that could not qualify to avoid filtration to filter their water as well as disinfect it, by 29 June 1993 at the earliest. For conventional filtration it held the filtered water at or below 0.5 NTU in 95 % of each month's measurements and never above 5 NTU [5, 6]. In 1993 a single waterborne outbreak of cryptosporidiosis in one city caused more than 400,000 cases [6]. A rule published in December 1998 tightened the standard for systems serving 10,000 people or more, from January 2002, to 0.3 NTU in 95 % of the measurements and never above 1 NTU. It added continuous turbidity monitoring of every individual filter [5, 6]. Those two numbers are the filtered-water specification this chapter grades against.

This chapter works with a simulated municipal water treatment plant whose coagulant measurement is three hours downstream of its actuator, so that the case for feedforward is measured rather than asserted, and it separates the three kinds of control problem the same plant carries. Section 6.2 presents the coagulation chemistry that gives the dose curve its interior optimum. Section 6.3 gives the model of the train, from the coagulant demand and the achievable target to the filters and the chlorine contact tanks. Section 6.4 describes the coagulant feedforward with its ratio trim and the pH, chlorine and level loops around it. Section 6.5 lists the scenarios and measures the storm, the unmeasured organic carbon, the soft water, the backwash sequence, the disinfection window and the graded shift. Section 6.6 describes the instruments, pumps and control platforms of a real plant, and Sections 6.7 to 6.9 state the limitations, discuss the three problems and summarize the results.

6.2 Background

Conventional surface water treatment has had the same unit sequence for more than a century, with coagulation and rapid mix, flocculation, sedimentation, granular media filtration and disinfection in series, and the design and operating basis for each unit is standard material [7]. The dose has traditionally been set from jar tests run once a shift or once a day, which is an open-loop procedure calibrated on a sample rather than on the water in the basin. Online raw water instrumentation changed what is possible, and dose control computed from continuous raw water measurements and corrected by a slow feedback trim is now the recommended structure [1].

Coagulation is not monotone in dose. Below the demand the colloid is still charged and does not aggregate. Above it the coagulant reverses the charge and restabilizes the particles, and the mechanism that dominates depends on both dose and pH in a way that has been mapped for aluminum since the nineteen eighties [8, 9]. The simulator represents this as a log-normal in the dose ratio multiplied by a Gaussian in pH,

$$\eta = \eta_{\max} \exp\left[-\frac{\ln^2(C_a/C_a^*)}{2\sigma_a^2}\right] \exp\left[-\frac{(\mathrm{pH}-\mathrm{pH}^*)^2}{2\sigma_p^2}\right] \tag{6.1}$$

with $\sigma_a = 0.45$ on the natural log of the ratio and $\sigma_p = 0.60$ pH units. Table 6.1 is the resulting curve at the design raw water, and Figure 6.1 plots its settled-water column against the dose.

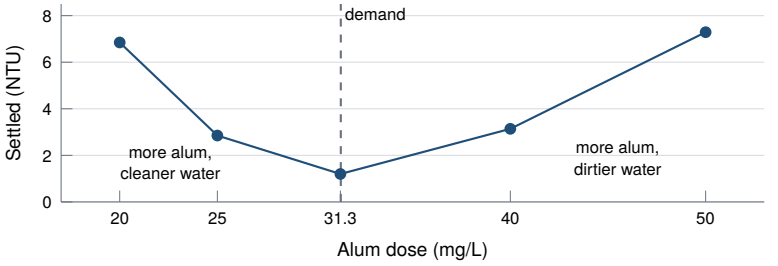


Figure 6.1 The settled-water column of Table 6.1 against the alum dose at the design raw water, with the demand marked by the dashed line. Notice that the slope changes sign at the demand: left of it more alum clears the water, right of it more alum makes it worse, and a settled-water reading alone does not say which side the plant is on.

Table 6.1 Settled and filtered water against alum dose at the design raw water of 12 NTU, 3 mg/L organic carbon and 18 °C. The demand is 31.32 mg/L.

Alum (mg/L)	η	pH	Settled (NTU)	Filtered (NTU)
20	0.598	6.947	6.848	0.5712
25	0.868	6.927	2.851	0.0773
31.3	**0.985**	6.900	**1.200**	**0.0200**
40	0.848	6.862	3.139	0.0925
50	0.568	6.814	7.289	0.6875

A feedback controller responding to high settled turbidity by adding coagulant is correct on the left of Table 6.1 and destructive on the right, and nothing in its measurement distinguishes the two. The same controller is an integrator with a three-hour plant in front of it, so once it crosses the optimum the error and the dose grow together rather than closing. The process gain has changed sign and the loop has become a positive feedback path. Section 6.5.2 measures that behavior and Figure 6.4 shows it in progress at a dose ratio of 3.85 and a destabilization of 1.1 %.

The pH term in (6.1) carries a second nonlinearity. Alum consumes alkalinity, 0.5 mg/L of it as calcium carbonate per mg/L of alum, so the coagulant moves its own pH. On a well-buffered water the effect is small. On a soft water it dominates, and Section 6.5.4 measures a plant that is dosed correctly and produces 9.7 NTU settled water because of it.

6.3 Process Model

Figure 6.2 shows the train. Every unit is a well-mixed volume whose time constant is its hydraulic residence time, and the twenty-seven states span three orders of magnitude in time constant. The dosing pumps and the rapid mix respond in 0.5 min, each flocculator holds 12 min, each clarifier stage holds 80 min, each chlorine contact stage holds 15 min, and a filter run at design lasts about 900 min.

6.3.1 Working Units and the Design Point

The clock is in minutes, flows in m^3/h, turbidity in NTU, doses in mg/L and filter headloss in meters. The design raw water is 12 NTU at 18 °C with 3 mg/L of organic carbon and 80 mg/L of alkalinity. At 1200 m^3/h the plant requires 31.32 mg/L of alum at pH 6.900, which destabilizes 98.5 % of the particles, settles to 1.200 NTU at a surface overflow rate of 1.5 m/h and filters to 0.0200 NTU at 7.5 m/h. A chlorine feed of 2.173 mg/L holds 1.000 mg/L of free residual and earns 15.0 mg · min/L of disinfection credit.

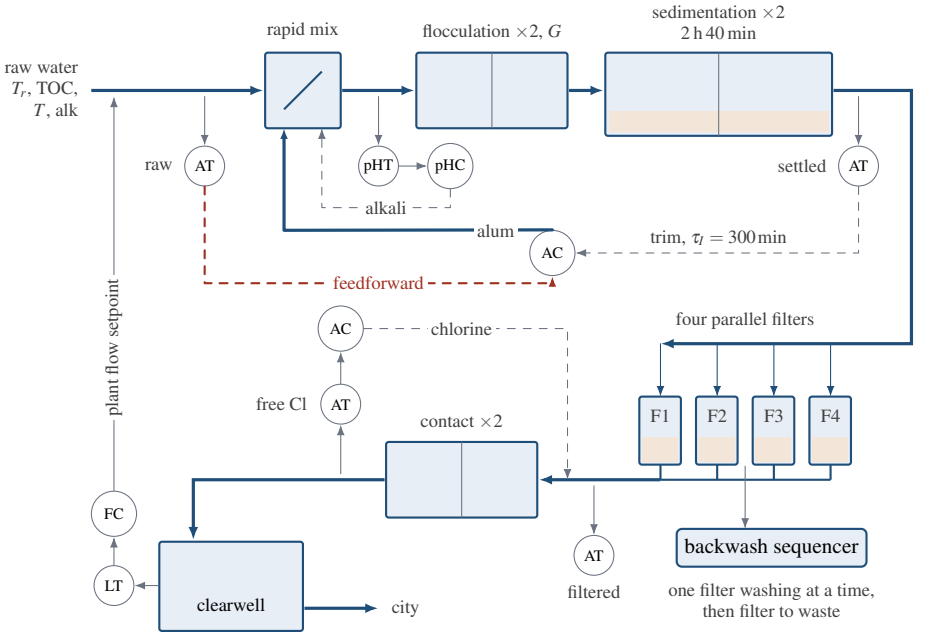


Figure 6.2 The simulated plant. The clarification train runs left to right along the top and the filters, contact tanks and clearwell run right to left below it. The coagulant dose is computed from the raw turbidimeter through the thick dashed path labeled feedforward and only trimmed by the settled turbidimeter, through the thin dashed path labeled trim, which is three hours downstream of the dosing point. The pH controller acts on the same rapid mix through the alkali pump and settles in about two minutes. The chlorine residual controller and the clearwell level controller close on the two ends of the plant. The filters are not a loop at all.

6.3.2 Coagulant Demand and the Achievable Target

The demand rises sub-linearly with turbidity, linearly with organic carbon and sharply in the cold,

$$C_a^* = \left(d_0 + d_1\sqrt{T_r} + d_{\mathrm{TOC}}\,\mathrm{TOC}\right)\left(1 + k_T\frac{(T_c - T)_+}{T_c}\right) \tag{6.2}$$

where C_a^* is the demand in mg/L of alum, T_r is the raw turbidity in NTU, TOC is the raw organic carbon in mg/L, T is the raw water temperature in °C and the subscripted plus sign denotes the positive part. The coefficients are $d_0 = 8.0\,\mathrm{mg/L}$, $d_1 = 5.0\,\mathrm{mg/L}$ per $\mathrm{NTU}^{0.5}$, $d_{\mathrm{TOC}} = 2.0$, $k_T = 0.35$ and $T_c = 15\,^{\circ}\mathrm{C}$, and the cold-water term is an empirical ratio of Celsius temperatures rather than a rate expression. Equation (6.2) gives 31.32 mg/L at design, 36.44 mg/L for the same water at 8 °C and 60.01 mg/L for 55 NTU runoff at 12 °C carrying 5.5 mg/L of organic carbon.

A second scaling governs what the clarifier can reach. No dose makes 60 NTU river water settle to 1.2 NTU, because a clarifier removes a fraction rather than an amount. The achievable target therefore moves with the load,

$$T_s^{\mathrm{SP}} = T_{s,0}^{\mathrm{SP}}\left(\frac{T_r}{T_{\mathrm{ref}}}\right)^{m}\left(\frac{\mu(T)}{\mu(18\,^{\circ}\mathrm{C})}\right)^{1.213}\left(\frac{Q}{Q_d}\right)^{1.85} \tag{6.3}$$

where $T_{s,0}^{\mathrm{SP}}$ is the operator-set target at the reference raw water $T_{\mathrm{ref}} = 12\,\mathrm{NTU}$, $\mu(T)$ is the water viscosity relative to 20 °C, $Q_d = 1200\,\mathrm{m^3/h}$ is the design flow and $m = 0.72$. The operator's 1.2 NTU becomes 2.32 NTU at 30 NTU raw, 3.59 NTU at 55 NTU and 4.27 NTU at 55 NTU and 12 °C. A trim loop held to a fixed setpoint through a storm is given a target the plant cannot reach, and it

integrates across the coagulation optimum in trying to reach it. The simulator shapes (6.3) through two lags of 87.5 min each before presenting it to the trim, so that the setpoint moves at approximately the speed of the plant.

6.3.3 Flocculation and Sedimentation

Each flocculator carries a floc size index whose target has an interior maximum in the mixing intensity,

$$S^{*} = \eta \, \frac{G}{G_{\mathrm{opt}}} \, \mathrm{e}^{\,1 - G/G_{\mathrm{opt}}} \min\!\left(1.35, \left(\frac{T_r}{T_{\mathrm{ref}}}\right)^{0.12}\right) \tag{6.4}$$

where G is the velocity gradient and $G_{\mathrm{opt}} = 45/\mathrm{s}$ is the value at which the floc grows fastest, because too little energy does not bring the particles together and too much tears the floc apart faster than it forms. Each stage relaxes toward S^{*} at $k_g = 0.12/\mathrm{min}$. Settling velocity follows Stokes through the square of the size index and inversely with viscosity [10], $v_s = v_{s0} S^{2} / \mu_r(T)$ with $v_{s0} = 11.24\,\mathrm{m/h}$, which gives 7.21 m/h at design. The basin removes a fraction per stage,

$$\mathrm{pass} = \left(1 + \frac{v_s A_s / n}{Q}\right)^{-n}, \qquad n = 2. \tag{6.5}$$

where $A_s = 800\,\mathrm{m}^2$ is the plan area of the basin. Turbidity is transported through the basin unchanged and (6.5) is applied algebraically at the outlet rather than as a sink inside the tanks. That choice sets the lag the coagulant controller sees. A well-mixed tank containing a removal sink responds at $V/(Q + v_s A_s)$ rather than at V/Q, which would give the train a much shorter lag than the three-hour transport response, and the lag is the subject of this chapter.

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

6.3.4 Filtration and Backwash Sequencing

A filter is a log-removal exponent with three modifiers,

$$\mathrm{pass} = \exp\left[-\frac{k_f S}{v/v_0} \underbrace{\left(1 - 0.95\,\mathrm{e}^{-t_r/2}\right)}_{\text{ripening}} \underbrace{\left(1 - 0.35\,\hat{H}^{4}\right)}_{\text{breakthrough}}\right] \tag{6.6}$$

where k_f is the removal coefficient, S the floc size index, v the filtration rate against its design value v_0, t_r the minutes since the filter returned to service and $\hat{H}$ the fraction of available head consumed. The ripening term matters because the maintenance activity that relieves breakthrough creates it. A freshly washed filter passes almost everything for the first minute, which has been documented on operating plants since 1980 [11], and the resulting spike appears directly in the regulated measurement.

Headloss builds with the product of rate, applied turbidity and time, which gives runs of about 900 min at design and much shorter runs in a storm. The sequencer is discrete. A filter is queued when its headloss reaches 2.20 m or its run reaches 1200 min. Only one filter may be washing at a time: it draws wash water for 10 min and remains isolated for 14 min. An optional 7 min filter-to-waste period follows. A second filter may start washing while the first sends filtrate to waste. Only filters in service contribute to product quality, while the filtration rate includes filters sending water to waste.

6.3.5 Disinfection and By-Product Formation

Two chlorine contact stages carry a first-order decay driven by organic carbon and particles. The credit and the by-products are computed from the same residual and the same contact time,

$$\mathrm{CT} = 60\,C\,\frac{\mathrm{BF}\cdot V_c}{Q}, \qquad \mathrm{THM} = k_{\mathrm{THM}}\,\mathrm{TOC}\,C\,\frac{V_c}{Q} \tag{6.7}$$

where C is the free chlorine residual leaving the second contact stage, V_c is the total contact volume of 600 m^3, Q is the plant flow in m^3/h, and the factor 60 converts hours to minutes for CT. BF is a baffle factor of 0.50 relating the effective contact time T_{10} to the nominal residence time [12]. Both quantities scale with residual and with contact time, so raising the residual to earn the credit raises the by-products in exact proportion, and the only variable that separates them is organic carbon. Enhanced coagulation can remove it in a real plant [13], but the simulated TOC is an unremoved raw-water disturbance.

> ✦ DID YOU KNOW? **The By-Product Limit Is a Federal Number**
>
> The 80 µg/L by-product limit of this chapter is the federal maximum contaminant level for total trihalomethanes, 0.080 mg/L. The same rule names enhanced coagulation, enhanced softening or granular activated carbon, with chlorine as the disinfectant, as the best available technology for meeting it [5].

6.3.6 Initialization and Assumptions

Every stage is a first-order tank, so the steady state marches forward analytically from the raw water rather than iterating, and `init` returns an exact equilibrium. The manual helpers set alum to the true jar-test demand together with its pH correction, calculate chlorine for the selected residual, or request a backwash of the filter with the greatest headloss.

Turbidity is a single lumped particle population split into a settleable and a non-settleable fraction. The coagulation surface (6.1) is an empirical shape rather than a mechanism. Backwash water is scaled by 0.12 because the filter run clock is compressed by about ten times, so that the plant water balance remains correct at the compressed rate. Integration is fixed-step fourth-order Runge-Kutta at 0.05 min, and the controllers execute every 0.2 min, independent of the integration step.

6.4 Controller Design

6.4.1 Velocity Form PID Algorithm

Every continuous loop uses the framework's velocity-form PID of Section 1.6.2, Equations (1.9) to (1.12), with derivative on the measurement [14] and the ISA action convention. The scan is $\Delta t_c = 0.2\,\mathrm{min}$. Back-calculation acts only on the integral, with the tracking time equal to τ_I, and the integral update is disabled when the integral gain is zero. All default settings have $\tau_D = 0$, and the level loop also exposes derivative time and filtering. The nominal settings below follow SIMC rules [15] applied to measured gains and lags.

6.4.2 Feedforward With a Ratio Trim

The coagulant path has two forms selectable at run time. Pure feedback puts the settled turbidimeter directly on the dosing pump with $K_c = 4\,\mathrm{mg/L/NTU}$ and $\tau_I = 120\,\mathrm{min}$. The recommended structure is the static feedforward of Section 1.4.3 with the multiplying trim of Equation (1.6). It computes

$$C_a = r\,\hat{d}\left(T_r, \widehat{\mathrm{TOC}}, T\right), \quad r = \mathrm{PID}\left(T_s^{\mathrm{SP}}, T_s\right) \in [0.80,\ 1.35] \tag{6.8}$$

where $\hat{d}$ is (6.2) evaluated on the information available to the control system. Two inputs represent online raw-water information, turbidity and temperature. In the implementation the dose model reads the disturbance values directly, and the displayed raw turbidimeter lag and noise affect only the shaped target. The other two inputs are not online, because the organic carbon is a laboratory number reported once a shift and the demand coefficients come from a jar test of unknown age. The trim r is dimensionless, is direct acting with $K_c = 0.06/\mathrm{NTU}$ and $\tau_I = 300\,\mathrm{min}$, and exists to correct those two stale inputs rather than to reject the storm. Trimming a ratio rather than adding an offset matters,

because a correction identified at 12 NTU raw water is proportionally correct at 60 NTU while an additive correction is not. The trim on the air-to-fuel ratio of Chapter 4 plays the same part for an unmeasured heating value.

🕮 **Pause.** Try Exercise 6.2 before reading on.

6.4.3 pH Control of the Rapid Mix

The pH controller moves an alkali pump against the coagulation pH with $K_c = 30\,\mathrm{mg/L}$ per pH unit and $\tau_I = 1.5\,\mathrm{min}$. Its actuator spans −15 mg/L to 45 mg/L, and the loop is reverse acting because adding alkali raises the pH. The signed equivalent dose B adds alkalinity when positive and consumes it when negative, with $A = A_r - 0.5C_a + 1.25B$ and $\mathrm{pH} = 4.610 + 0.55\ln(\max(2,A))$. The pump retains the sign through its 0.5 min lag, and reagent consumption uses $|B|$ for both acid and alkali. The loop is faster than the coagulant loop by two orders of magnitude and it is what gives the coagulant loop a stationary target, because (6.1) multiplies the dose term by the pH term rather than adding them. A dose exactly equal to the demand destabilizes 98.5 % of the particles at pH 6.90 and 4.3 % at pH 5.40. Here pH enters through the logarithm of the alkalinity left after the alum, a much milder curve than the titration of a strong acid, where Volume 2, Chapter 6 needs a gain schedule.

6.4.4 Chlorine, Clearwell Level and the Backwash Sequencer

Free chlorine goes to the chlorinator with $K_c = 3$, $\tau_I = 20\,\mathrm{min}$ and 2 min of sample line delay ahead of a 1 min analyzer lag. The clearwell level sets the plant flow with $K_c = 15\,\mathrm{m^3/h}$ per percent of level and $\tau_I = 180\,\mathrm{min}$, detuned deliberately so that the clearwell absorbs the city demand rather than passing it to the coagulation basins. The backwash sequencer is a state machine with a single interlock limiting simultaneous washes. With one of four filters washing the remaining three carry the whole flow at 10 m/h instead of 7.5 m/h, which degrades (6.6) for all of them, so a plant whose filters foul together forms a queue that never clears. Discrete units that fall into step on one shared resource are also the problem of the thermostats of Volume 8, Chapter 12.

6.5 Results

6.5.1 Scenarios of the Application

The simulator carries eleven scenarios, and each is credited by a completion rule that checks the specified actions and measured responses rather than the selection of a case (Section 1.6.4). Table 6.2 lists them. All eleven scripted walkthroughs of these rules pass. Runs use seed 20260928, zero default instrument noise and the shipped model parameters and disturbances, including each scenario's scheduled changes.

Two elements are stylized on purpose. The filter run clock is compressed by about ten times, which brings a full wash cycle inside a single scenario so that the ripening spike and the wash queue can be studied together. The storm arrives as two steps rather than as a ramp, because a step is the disturbance whose response can be read directly off a trend.

6.5.2 Feedback Against a Storm

Scenario 4 raises the raw water from 12 to 35 to 60 NTU between 30 and 60 min with the coagulant loop in pure feedback and everything else in service. Scenario 5 runs the same storm with the coagulant on feedforward plus trim.

Table 6.3 is the central measurement of this chapter, and the mechanism is better followed in time than read as a single ratio. Ten minutes after the raw water reaches 60 NTU the feedforward plant has already moved the alum from 31 to 53 mg/L and its settled water is 1.28 NTU. The feedback plant is at 32 mg/L, because its measurement has barely moved, and its settled water is 1.57 NTU. At 140 min the feedback plant is at 55 mg/L and 6.86 NTU, at 200 min it is at 76 mg/L and 9.64 NTU, and at 300 min it is at its 140 mg/L stop and 20.92 NTU. Figure 6.3 draws the table as bars.

Table 6.2 Scenarios and Completion Requirements

#	Scenario	Kind	Completion requirements
1	Train lag	Experiment	Start the run, then set manual alum to 35.5–37 mg/L after at least 0.2 min and observe a settled-water change of at least 0.45 NTU.
2	Manual storm	Demonstration	Observe a fixed-dose violation after 180 min. Reset, raise the dose before 110 min and compare at 180 min.
3	Dose curve	Experiment	Hold 22, 47 and 63 mg/L manual doses (within 1 mg/L) for 240 min each, with an overdose alarm in the last case.
4	Feedback	Demonstration	Run feedback for 420 min. Reset, change gain or integral time before the storm and compare a 330 min run by 360 min.
5	Feedforward	Experiment	Hold a near-demand dose ratio for 60 min after the second storm step, then run to 420 min.
6	Unmeasured	Experiment	Observe the demand jump and settled offset, then the enabled trim's slow correction through 480 min.
7	Soft water	Experiment	Observe pH loss and restoration. Reset with pH control enabled and compare a 230 min run.
8	Sequencing	Experiment	Compare 355 min runs with Auto and Auto + FTW, enabling FTW within the first 20 min after reset.
9	Disinfection	Demonstration	Observe low CT after the flow increase, the bloom dip, and the CT/THM tradeoff after the residual target rises.
10	Commission	Experiment	Step pH, chlorine and level targets and wait for their response tolerances. Step the turbidity target and observe 180 min.
11	Challenge	Challenge	At 600 min, maximum logged settled error at most 1 NTU, off-specification time at most 60 min, zero turbidity-violation and CT-shortfall times, and maximum logged THM strictly below 80 micrograms/L.

Table 6.3 The same storm under two coagulant structures, measured over 300 min.

	Feedback only	Feedforward plus trim
Peak settled (NTU)	20.918	3.600
Peak filtered (NTU)	2.821	0.126
Above 0.30 NTU (min)	141.9	0
Above 1.00 NTU (min)	47.15	0
Peak alum (mg/L)	140.0	53.5
Alum consumed (kg)	415	306
Settled turbidity IAE (NTU min)	1555.6	28.3

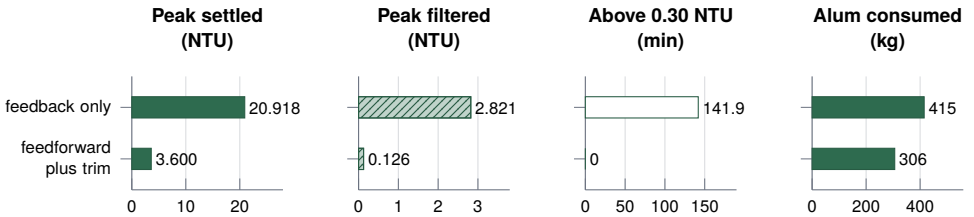


Figure 6.3 Table 6.3 redrawn: the storm of scenarios 4 and 5 under the two coagulant structures, measured over 300 min. Bars give the peak settled and filtered turbidity, the minutes the filtered water spent above 0.30 NTU and the alum consumed. Notice that the structure that doses earlier keeps the filtered water in specification throughout and uses less alum, not more.

The last part of that trajectory is not a slow controller. It is a controller that has crossed the optimum of Table 6.1. When the storm passes at 300 min the raw water falls to 20 NTU and the

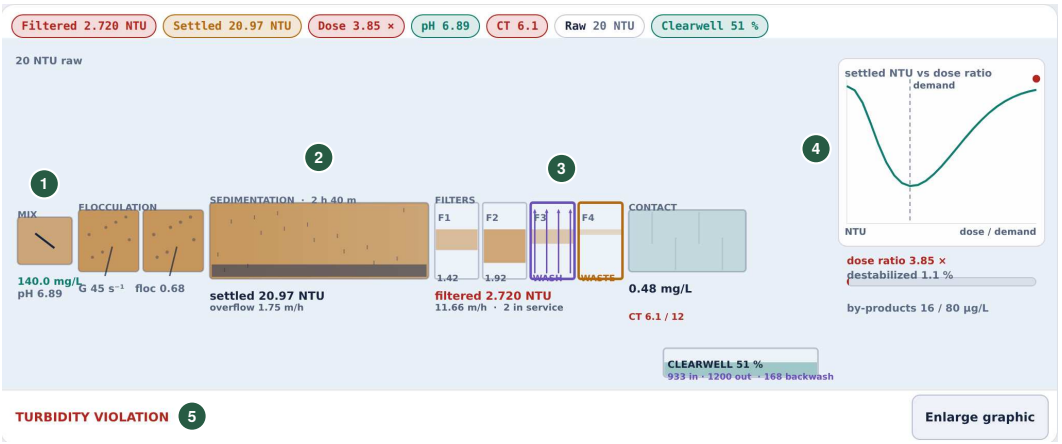


Figure 6.4 The plant view of scenario 4 at 300.2 min, just after the storm passes, cropped from the app window. The crop is a composite: an empty band of the view between the filters and the clearwell is cut out, and the run time is not in the crop. (1) The coagulant controller has integrated to its 140 mg/L stop, read under the rapid mix. (2) The sedimentation basin is drawn with its actual solids loading, and the settled water is 20.97 NTU against a shaped target of 3.61 NTU, which a trend outside the crop shows. (3) Two filters remain in service: one is washing and one is sending filtrate to waste, and the filtered water is 2.720 NTU. (4) The dose ratio is 3.85 and the operating point is drawn on the far side of the coagulation minimum, where destabilization has fallen to 1.1 %. (5) The plant is in turbidity violation.

demand with it to 36.4 mg/L, which leaves the dose ratio at 3.85 and destabilization at 1.1 %. Every additional milligram makes the settled water worse, which the controller registers as a larger error, which produces more coagulant. Figure 6.4 shows that state directly, with the operating point plotted on the far branch of the coagulation curve.

🕮 **Pause.** Try Exercise 6.1 before reading on.

> **TRY IT IN THE APP** **4 · Feedback Against a Storm**
>
> Load the scenario and press Run. The storm takes the raw water to 60 NTU by 60 min. Watch the coagulant trend: the dose climbs slowly, and by 300 min it sits at its 140 mg/L stop with the settled water near 20.92 NTU. Then load 5 · Feedforward against the same storm and run it. The dose moves with the storm, and the settled water peaks at 3.600 NTU.

6.5.3 Unmeasured Disturbances and the Ratio Trim

Scenario 6 runs the same plant with feedforward only and moves a variable the feedforward cannot see. At 30 min, an algal bloom takes the organic carbon from 3 to 6.5 mg/L, which raises the true demand from 31.32 to 38.32 mg/L. The modeled demand does not move, because the dose model is using the laboratory organic carbon number reported that morning. The dose ratio falls to 0.817, destabilization to 89.1 %, and the settled water comes to rest at 2.13 NTU against a 1.21 NTU target.

That offset is what the trim exists to remove, and switching it in at 240 min starts moving the dose upward. The correction is slow. By 600 min the alum has reached 33.4 mg/L and the settled water 2.04 NTU, still short of the 1.23 NTU target. A trim with a 300 min integral time on a 190 min plant does not converge within six hours and is not intended to. Feedforward rejects the fast measured disturbance and the trim corrects the slow model error, and neither substitutes for the other. Volume 5, Chapter 4 trims a residence-time prediction from a late kappa analyzer on the same division of labor.

Table 6.4 Snowmelt with and without the pH controller in service, measured over 300 min. The dose is computed correctly in both columns.

	pH loop out	pH loop in
Minimum pH	5.353	5.918
Time below pH 6.5 (min)	170.35	0.85
Minimum destabilization	2.8 %	25.8 %
Peak settled (NTU)	9.739	1.257
Peak filtered (NTU)	3.831	0.253
Above 0.30 NTU (min)	169.55	0
Above 1.00 NTU (min)	127.35	0

6.5.4 Effect of the pH Loop on Coagulation

Scenario 7 drops the raw alkalinity from 80 to 25 mg/L on snowmelt, with the coagulant feedforward computing the correct dose throughout. The scenario starts with the pH loop gain at zero and puts the loop in service at 200 min. The right-hand column of Table 6.4 is the same run with the pH loop in service from the start.

The two columns of Table 6.4 differ by a factor of 7.7 on the settled water peak, bought by a loop that settles in two minutes and whose reagent costs $0.58/kg. The coagulant controller computes the same dose in both columns and produces unusable water in one of them, because the variable it does not measure has moved. The two minima in the right-hand column are transients of about a minute at the instant the alkalinity steps, before the alkali pump has responded, while the left-hand column remains below pH 6.5 for 170.35 min. The peak filtered turbidity of 0.253 NTU in the right-hand column is a filter ripening spike rather than a coagulation effect, because scenario 7 runs the filters without a filter-to-waste period.

The recovery is far slower than the pH response that starts it. When the pH loop is put in service at 200 min the pH is back at 6.86 within eleven minutes, and the plant does not follow. Seven hundred minutes later the settled water is still 4.09 NTU. Two mechanisms hold it there and neither is the sedimentation lag. The alkali pump is saturated at 45 mg/L because the alum has to be neutralized as well as the raw water, and the coagulant trim wound its ratio to its 1.35 limit, or 42.3 mg/L of alum, during the four hours it spent correcting a fault outside its authority, which places the plant on the over-dosed side of Table 6.1 at 78.7 % destabilization. An integrating trim left running against a disturbance it cannot reach leaves a legacy that outlasts the disturbance. In the same run with the pH loop in service from the beginning, destabilization does not fall below about 90 % once the alkali pump has responded.

6.5.5 Backwash Sequencing and Filter Ripening

Filter runs at design are about 900 min and the four filters start staggered, at 225, 450, 675 and 880 min of run age, so that the wash queue clears. Scenario 8 loads them with a raw water step to 70 NTU at 100 min.

The measurement that matters is the effect of the wash itself. Over 600 min and eight backwashes with immediate return to service, the peak filtered turbidity is 0.804 NTU and the filtered water is above the 0.30 NTU specification for 2.9 min. Adding a 7 min filter-to-waste period after each wash takes the peak to 0.169 NTU and the time off specification to zero, at a cost of 518 m^3 of total water loss against 12 172 m^3 produced. The loss includes 224 m^3 of wash water. No continuous loop is involved in the spike. It is created by the maintenance action, it is predictable to the minute, and it is removed by a timer. Sending material in doubt to waste is also the job of the flow diversion valve of Volume 8, Chapter 13.

Table 6.5 The graded ten-hour shift, measured over 600 min. Each row after the second changes one element of the recommended configuration.

Configuration	Peak filt. (NTU)	Off spec (min)	Violation (min)	Min CT (mg · min/L)	Alum (kg)
As loaded	7.971	493.90	456.00	0.35	379
Suggested	0.279	0	0	13.97	660
feedback only	23.341	454.55	332.60	0.21	1549
no filter-to-waste	0.951	5.55	0	13.97	634

> **TRY IT IN THE APP** **8 · Four Filters and a Queue**
>
> Run the scenario with Filter sequencing on Auto and watch the filtered-water trend: it spikes each time a washed filter returns to service. Press Reset, switch Filter sequencing to Auto + FTW within the first 20 min and run again. Over 600 min the chapter measures a peak of 0.804 NTU without the timer and 0.169 NTU with it.

6.5.6 Disinfection Credit and By-Product Limits

Scenario 9 raises the city demand to 1650 m^3/h, which the clearwell level controller answers by running the plant at its 1700 m^3/h limit. The chlorine residual controller holds its setpoint to within 0.08 mg/L through the flow ramp, and the disinfection credit fails anyway. Contact time is volume divided by flow, so T_{10} falls from 14.88 min to 10.59 min and the credit from 14.60 to 10.59 mg · min/L against a requirement of 12 mg · min/L. The compliance failure is produced by a variable that is not on the chlorine loop at all. An algal bloom at 150 min doubles the organic carbon and the chlorine demand with it, which takes the residual to a transient minimum of 0.745 mg/L and the credit to 8.48 mg · min/L.

Raising the residual target to 1.6 mg/L restores the credit to 17.94 mg · min/L and takes the by-products from 59.0 to 101.1 µg/L against an 80 µg/L limit. Both quantities in (6.7) are proportional to residual and to contact time, so no setting of the chlorine loop widens the window. The only variable that does is the organic carbon. Removing it would require an upstream process change that this model does not implement.

6.5.7 Graded Ten-Hour Shift

The challenge is ten hours against a storm that takes the raw water to 55 NTU, an organic carbon rise, a runoff dilution of the alkalinity to 40 mg/L and a morning demand peak, with the target requiring the settled water to track its shaped setpoint, the filtered water to stay under 0.30 NTU for all but an hour and never to exceed 1.00 NTU, the disinfection credit above 12 mg · min/L and the by-products under 80 µg/L.

The current suggested configuration uses FF + trim and Auto + FTW, with the nominal gains given above except for the stronger trim $K_c = 0.12/\text{NTU}$ and $\tau_I = 150\,\text{min}$. Its residual target is 1.15 mg/L and its level target is 65 %. It passes every clause: maximum logged settled-water error is 0.845 NTU, no CT shortfall is accumulated and the maximum by-products are 70.15 µg/L. Replacing feedforward with feedback fails quality and disinfection, drives the coagulant to its stop and uses 1549 kg of alum against 660 kg. Removing filter-to-waste raises the filtered-water peak to 0.951 NTU and accumulates 5.55 min off specification. It still meets every grading clause on the shipped seed. Filter-to-waste is therefore a recommended quality safeguard rather than a mandatory control setting in the challenge verdict.

The settled turbidity integral against the shaped target is 9549 NTU min as loaded and 341 NTU min recommended.

The quality timers integrate the physical outputs at each three-second step, and the settled-error and by-product clauses read the exported historian. Changing model parameters or disturbance sliders,

or firing additional disturbances, voids challenge credit. Changing controller settings remains part of the task.

6.6 Technology in Practice

Figure 6.5 draws the coagulant feedforward of Section 6.4.2 with its ratio trim, and the pH loop beside it, with their hardware.

What Measures It

Three turbidimeters carry the coagulant structure. The raw-water meter at the intake feeds the dose model together with a raw-water thermometer, the settled-water meter at the basin outlet drives the trim, and the filtered-water meter reads the number the plant is graded on. Guidance for operators names this use of the raw meter: turbidity readings of a surface source serve to adjust the coagulant and disinfectant doses [16]. In the model the dose calculation reads the raw turbidity and temperature directly, and the raw meter's displayed lag and noise act only on the shaped target. The filtered water must be measured at least every four hours, or continuously if the plant validates the online meter on a regular basis [5].

> **TECHNOLOGY** **Sensor: A Nephelometric Turbidimeter**
>
> A nephelometric turbidimeter shines a beam through the sample and reads, with a photodetector, the light scattered at right angles to it. More scattered light means higher turbidity, and the reading is calibrated against a formazin suspension. The standard method covers 0 to 40 NTU, higher values after dilution, and asks for a sensitivity of 0.02 NTU or better below 1 NTU. Fine air bubbles read high, while coarse sediment that settles out and color in the water read low [17]. Online meters can measure below 0.1 NTU reliably if they are carefully calibrated and maintained [16]. In comparison studies of samples between 0.35 and 0.72 NTU, laboratories using different meters agreed to a relative standard deviation just under 20 % [6].

What Moves It

Alum and alkali enter the rapid mix through two dosing pumps, which the model gives a 0.5 min lag. The alum pump stops at 140 mg/L, and the alkali pump spans −15 mg/L to 45 mg/L, a negative value standing for acid. A chlorinator sets the chlorine feed, and the clearwell level loop sets the plant flow. The model works with the dose itself. On a real plant a dose in milligrams per liter becomes a pump rate only in proportion to the measured water flow, which is the job a chemical metering pump is set up for [18].

> **TECHNOLOGY** **Final Element: A Chemical Metering Pump**
>
> A metering pump packs the pump mechanism, its motor and the drive electronics in one unit, and water plants use such pumps to inject small flows of coagulant, disinfectant and pH chemicals. A standard 4 to 20 mA signal sets its output from zero to full flow, and knobs on the pump set what full flow is. Many use a plunger driven by a solenoid at variable intervals, so the chemical arrives in pulses rather than as a steady stream. Each chemical's flow must stay in proportion to the flow of water treated, which is why the pump is set by the control system and not by hand [18]. In this chapter the pump follows its dose after a 0.5 min lag and stops at 140 mg/L.

Where the Controller Runs

A distributed control system runs production within one site in industries that include water and wastewater treatment, with feedback or feedforward loops holding the key conditions [19]. Control cycle times range from milliseconds to minutes [19]. In the app every controller executes each 0.2 min, against a coagulant plant that takes 193.6 min to reach sixty-three percent of a step, so the scan never limits the loop. The structure of Equation (6.8) maps onto standard blocks. A calculation

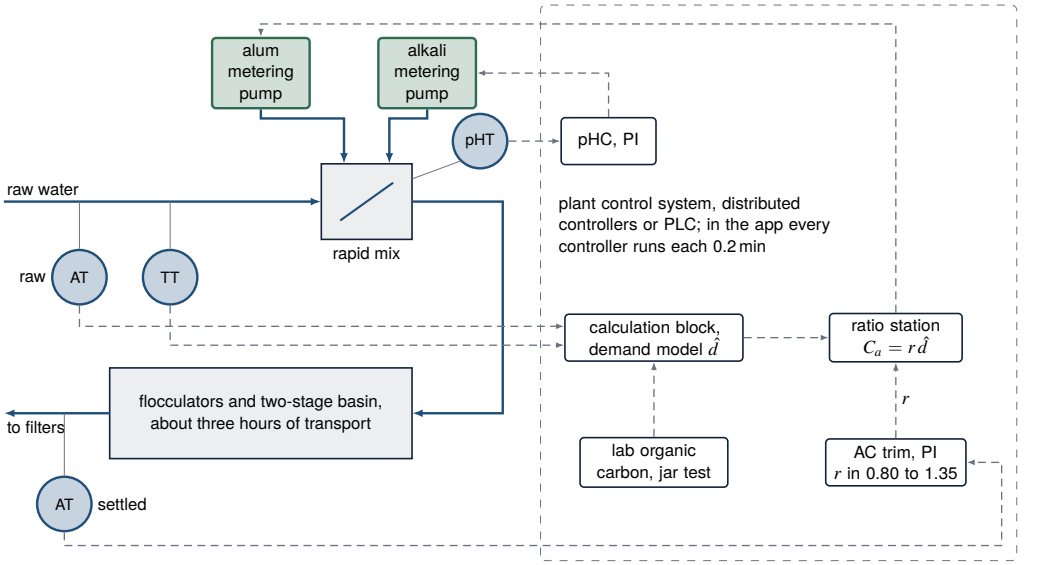


Figure 6.5 The coagulant feedforward with its ratio trim, and the pH loop, on real hardware. The nephelometric turbidimeter AT and the thermometer TT at the intake feed a calculation block that evaluates the demand model of Equation (6.2), with the organic carbon and the jar-test coefficients entered by hand. A ratio station multiplies the demand by the factor r of the trim controller, whose measurement is the settled turbidimeter at the basin outlet, and sends the dose to the alum metering pump. The pH controller moves the alkali metering pump on the same rapid mix. Solid arrows are process streams and dashed gray arrows are signals. Notice that the only signal that waits on the three-hour basin is the slow trim.

block evaluates the demand model, and a signal multiplying relay, or ratio station, multiplies its result by a factor that an operator would otherwise set by hand [18]. Here the trim's PI output sets that factor, between its limits of 0.80 and 1.35.

> **IN PRACTICE** **Every Filter, Every Fifteen Minutes**
>
> Large surface-water plants in the United States must monitor the turbidity of each individual filter continuously and record it every 15 min. A filter above 1.0 NTU in two consecutive readings 15 min apart is reported, and the plant produces a filter profile or names an obvious cause. A filter still above 0.5 NTU in two consecutive readings at the end of its first four hours after a backwash is reported in the same way [5]. Guidance for operators names start-up, filter ripening and the end of a run as the times when spikes occur, asks that they be minimized and ideally eliminated, and lists sending filtered water to waste among the immediate responses to an exceedance [16].

What Is Modernizing It

Three developments bear on this plant. The first is continuous monitoring tied to alarms. Turbidity can be read continuously, with the results linked to automatic alarm systems that allow a rapid response and remedial action when performance deviates [16]. The second is the industrial internet of things: sensors and instruments networked together, with computing at the edge that processes and acts on data where it is collected [19]. The same connectivity raises the need for resilience, safety and security [19]. The third is model-based control and estimation, which the Discussion (Section 6.8) proposes for this train: a predictive controller that evaluates the coagulation surface instead of inferring a gain from the error, and an estimator that turns the trim into a soft sensor for the organic carbon and the jar-test bias.

6.7 Limitations

Turbidity is a single lumped particle population with a settleable and a non-settleable fraction and one floc size index per stage. Real raw water carries a particle size distribution, and the reason a plant that meets its turbidity limit can still fail a pathogen limit is that the organisms of concern sit at the small end of that distribution. Turbidity is a surrogate for pathogen removal rather than a measure of it, and the strength of that surrogacy is itself a research question [20]. The simulator is not a substitute for that literature.

The coagulation surface (6.1) is an empirical shape fitted to give recognizable behavior, not a mechanism. It reproduces the peak, the pH window and the cold-water penalty, and it cannot distinguish charge neutralization from sweep floc, cannot represent polymer aids, and has no memory of the coagulant already in the basin. Data-driven models fitted to individual plants capture considerably more of the real surface [21], and none of them transfer between plants either.

Organic carbon is a single number that affects coagulant demand, chlorine decay and by-product formation, and that is removed by nothing. A real plant's enhanced coagulation removes a large fraction of it, which changes both the disinfection window and the coagulant demand downstream, so the model omits the upstream removal effect discussed in Section 6.5.6.

The filters share one settled water quality and differ only in headloss and run age. Real filters differ in media, in condition and in how well they were washed, and the diagnosis of which filter is the problem is a routine operating task the simulator cannot stage. Backwash water is scaled by 0.12 to keep the plant water balance right against a compressed run clock, so the wash water fraction of production is representative and any single wash is not. Sludge is not modeled, and sludge handling is where a real plant's storm response is often actually limited [22].

Finally, the plant has one train. Real plants have several and can take one out of service, which is the response an operator would reach for in scenario 9 before touching the chlorine setpoint.

> **OPEN CHALLENGE** **Turbidity Stands In for Pathogens**
>
> Low turbidity is a proven indicator that pathogens have been removed, yet no direct proportional relation between removing turbidity and removing pathogens has been shown, and studies that link turbidity to illness in communities have given mixed results. A durable, low-cost meter sensitive below 1 NTU is still missing [16]. Progress would be a measurement closer to the organisms themselves, cheap enough for every filter, so that the trim and the filter sequence could close on the quantity that matters.

6.8 Discussion

The plant carries three control problems with different answers, and separating them is the purpose of this chapter.

The first is a measured fast disturbance in front of a slow plant, and it has a structural answer rather than a tuning answer. The dose belongs on the disturbance measurement and the slow instrument belongs on the model. Table 6.3 gives a factor of 5.8 on the peak settled turbidity and a factor of 55 on the integral, and no tuning of the feedback structure approaches either. The general result is that when a disturbance is measured earlier than its effect, feedback on the effect is the wrong place to spend effort [2].

> **DID YOU KNOW?** **The Earlier Dose Uses Less Alum**
>
> From the simulator: through the same storm, feedforward plus trim peaks at 53.5 mg/L of alum and uses 306 kg, while feedback alone climbs to its 140.0 mg/L stop and uses 415 kg (Table 6.3). Acting early takes less coagulant.

The second is a fast inexpensive loop whose value is indirect. Table 6.4 shows the pH controller holding a variable that carries no specification of its own and deciding whether the plant meets the

one that does. Multiplicative interactions of the form of (6.1) are common in process chemistry and uncommon in introductory block diagrams, and they carry a consistent implication for commissioning order, because the inexpensive fast loop is placed in service first and removed last. The legacy the coagulant trim leaves behind in the same scenario is a second result, which is that an integrating controller acting on a fault outside its authority does not simply fail to correct it, it accumulates an error that outlasts the fault.

The third problem is not a loop. Four filters, an interlock and a timer produce a predictable turbidity spike at a predictable moment, and the answer is a sequence rather than a controller. Continuous control theory does not address it and the regulated measurement records the result. Discrete sequencing interacting with continuous control accounts for a large part of what a plant control system does, and it is underrepresented in control education [4, 23].

The same equations also suit methods the simulator does not implement. A model predictive controller on the twenty-seven state train could hold the filtered water and the disinfection credit as explicit constraints and coordinate the alum, alkali and chlorine doses across the three-hour lag [24]. Because it would evaluate (6.1) rather than infer a gain from the error, it would handle the sign reversal of Table 6.1 directly. An estimator on the raw, settled and filtered turbidimeters could track the organic carbon and the jar-test bias, which would turn the trim of (6.8) into a soft sensor for the two inputs the dose model does not measure.

6.9 Conclusion

This chapter studied a municipal water treatment plant that separates a coagulant decision from its evidence by three hours, and the simulator measures the separation as 32.25 min of apparent dead time and 193.6 min to sixty-three percent against a storm that arrives in thirty. Feedback on that plant does not merely respond slowly. It integrates past the coagulation optimum, reaches its actuator stop at 140 mg/L, destabilizes 1.1 % of the particles and puts the plant in turbidity violation for 47.15 min, while a feedforward computed from a raw turbidimeter and trimmed by the same slow instrument holds the same storm with no time off specification and 306 kg of alum against 415 kg. Three further results follow from the same model, a pH loop worth a factor of 7.7 on the settled water peak and carrying no specification of its own, a filter-to-waste timer that removes a violation no continuous loop reaches, and a disinfection window whose two limits move together and cannot be widened by the chlorine residual setting alone. The model also carries the constraints and the estimation problem that model predictive control and moving horizon estimation would address. The plant measures what feedforward is worth where feedback cannot arrive in time.

KEY POINTS

- The meter that judges the dose reads about three hours after the pump, while a storm arrives in half an hour.
- Feedback on the settled water crosses the coagulation optimum to its 140 mg/L stop, while feedforward with a slow ratio trim holds the same storm with no time off specification.
- A pH loop worth a factor of 7.7 on the settled-water peak carries no specification of its own.
- A filter-to-waste timer removes a violation that no continuous loop reaches.

Where else this applies. Chapter 4 and Volume 5, Chapter 4 trim a ratio or a prediction from a late measurement.

References

[1] 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.

[2] 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.

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

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

[5] U.S. Environmental Protection Agency, "National primary drinking water regulations, 40 CFR part 141, sections 141.64, 141.73, 141.74 and 141.173 to 141.175," Code of Federal Regulations, Title 40, edition of July 1, 2025, govinfo.gov, 2025.

[6] ——, "National primary drinking water regulations: Interim enhanced surface water treatment; final rule," *Federal Register*, vol. 63, no. 241, pp. 69 478–69 521, 1998, december 16, 1998.

[7] J. C. Crittenden, R. R. Trussell, D. W. Hand, K. J. Howe, and G. Tchobanoglous, *MWH's Water Treatment: Principles and Design*, 3rd ed. Hoboken, NJ: John Wiley & Sons, 2012.

[8] A. Amirtharajah and K. M. Mills, "Rapid-mix design for mechanisms of alum coagulation," *Journal AWWA*, vol. 74, no. 4, pp. 210–216, 1982.

[9] J. K. Edzwald, "Coagulation in drinking water treatment: Particles, organics and coagulants," *Water Science and Technology*, vol. 27, no. 11, pp. 21–35, 1993.

[10] Y. Adachi and Y. Tanaka, "Settling velocity of an aluminium-kaolinite floc," *Water Research*, vol. 31, no. 3, pp. 449–454, 1997.

[11] A. Amirtharajah and D. P. Wetstein, "Initial degradation of effluent quality during filtration," *Journal AWWA*, vol. 72, no. 9, pp. 518–524, 1980.

[12] J. D. Paccione, D. M. Follansbee, P. L. Young, and D. M. Dziewulski, "Analytical determination of the baffle factor for disinfection contact systems based on hydraulic analysis, disinfection kinetics, and Ct tables," *Journal of Environmental Engineering*, vol. 142, no. 7, p. 04016026, 2016.

[13] G. Hua and D. A. Reckhow, "Comparison of disinfection byproduct formation from chlorine and alternative disinfectants," *Water Research*, vol. 41, no. 8, pp. 1667–1678, 2007.

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

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

[16] World Health Organization, "Water quality and health: Review of turbidity. information for regulators and water suppliers," World Health Organization, Geneva, Technical brief WHO/FWC/WSH/17.01, 2017.

[17] U.S. Environmental Protection Agency, "Method 180.1: Determination of turbidity by nephelometry, revision 2.0," Environmental Monitoring Systems Laboratory, Office of Research and Development, Cincinnati, OH, Tech. Rep., 1993.

[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] R. K. X. Bastos, D. B. Viana, and P. D. Bevilacqua, "Turbidity as a surrogate for Cryptosporidium removal by filtration in drinking-water QMRA models," *Water Supply*, vol. 13, no. 5, pp. 1209–1219, 2013.

[21] C. W. Baxter, S. J. Stanley, and Q. Zhang, "Development of a full-scale artificial neural network model for the removal of natural organic matter by enhanced coagulation," *Journal of Water Supply: Research and Technology—AQUA*, vol. 48, no. 4, pp. 129–136, 1999.

[22] N. N. Reza, S. Peldszus, and P. M. Huck, "Quantifying drinking water treatment plant robustness with respect to turbidity: Examining raw water and treatment," *Water Research*, vol. 295, p. 125550, 2026.

[23] B. W. Bequette, "Process control practice and education: Past, present and future," *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[24] 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.

6.10 Exercises

6.1 Concept check · 2 min

In Scenario 4 the storm passes at 300 min, the raw water falls to 20 NTU and the feedback controller holds the alum at its 140 mg/L stop while the pH loop holds pH 6.90. Why does the settled water get worse, not better, with every further milligram of alum?

(a) The pH loop cannot hold its target against so much alum, so the pH leaves its window
(b) The flocculators shear the larger floc apart once the dose reaches the pump limit
(c) Filters are washing, so the remaining beds carry more of the settled turbidity
(d) The dose is about four times the demand, so the extra coagulant restabilizes the particles

6.2 Concept check · 1 min

The trim r of Equation (6.8) multiplies the computed dose instead of adding a correction to it. Why?

(a) A correction found at 12 NTU raw water stays proportionally right at 60 NTU
(b) A multiplying trim follows the storm faster than an additive trim of equal tuning
(c) A ratio held between 0.80 and 1.35 cannot wind up against a fault
(d) A ratio needs no settled turbidimeter, so the three-hour lag drops out of it

6.3 Concept check · 2 min

The model applies the basin removal of Equation (6.5) algebraically at the outlet instead of as a sink inside each well-mixed tank. What would a sink inside the tanks change?

(a) The settled turbidity at steady state would be lower for every dose
(b) Each basin stage would respond at $V/(Q+v_sA_s)$, so the lag would be far shorter
(c) The coagulation optimum would move to a higher dose ratio
(d) Headloss would stop building on the filters between washes

6.4 Theory and calculation · 6 min

Scenario 4 starts at the design raw water of 12 NTU, 3 mg/L organic carbon and 18 °C, takes the turbidity to 60 NTU, and later lets it fall to 20 NTU. The feedback controller ends at its 140 mg/L stop, and the pH loop holds pH 6.90.

(a) Find the coagulant demand of Equation (6.2) at 12 NTU and at 60 NTU. *Checkpoint:* 31.32 mg/L and 52.73 mg/L.

(b) Find the dose ratio of the 140 mg/L dose at 60 NTU and at 20 NTU. *Checkpoint:* 2.66 and 3.85.

(c) Find the destabilized fraction η of Equation (6.1) for each ratio, and compare it with a dose equal to the demand. *Checkpoint:* 9.4 % and 1.1 %, against 98.5 %.

6.5 Theory and calculation · 6 min

The step of the Introduction moves the alum from 31.3 to 38 mg/L and the settled water from 1.200 to 2.435 NTU, with 2 % of the change at 32.25 min and 63 % at 193.6 min.

(a) Add the time constants of the train from Section 6.3 (dosing 0.5 min, two flocculators, two clarifier stages). Then read a first-order-plus-dead-time model off the step, taking the 2 % time as the dead time. *Checkpoint:* 184.5 min. $\theta \approx 32.25$ min, $\tau \approx 161.35$ min.

(b) Find the steady-state gain of the step, and the gain between 25 and 31.3 mg/L from Table 6.1. *Checkpoint:* $+0.184$ and -0.262 NTU per mg/L.

(c) The feedback loop raises the dose by $K_c = 4\,\mathrm{mg/L}$ per NTU of settled error. Find K_cK_p on each side of the optimum and say on which side the loop is positive feedback. *Checkpoint:* -1.05 below the optimum (negative feedback), $+0.74$ above it (positive feedback).

6.6 Theory and calculation · 6 min

Take Equation (6.7) with $V_c = 600\,\text{m}^3$, BF $= 0.50$ and the by-product coefficient $k_{\text{THM}} = 26$ (µg/L per mg/L of TOC, per mg/L of residual, per hour of contact), the value in Show Model.

(a) Find T_{10}, the CT credit and the by-products at a residual of 1.00 mg/L with 3 mg/L of organic carbon, at 1200 m^3/h and at 1700 m^3/h. *Checkpoint:* 15.0 min, 15.0 mg · min/L, 39 µg/L. At 1700 m^3/h, 10.59 min, 10.59 mg · min/L, 27.5 µg/L.

(b) Show that THM/CT does not depend on the residual or the flow. With 6.5 mg/L of organic carbon, find the by-products when the credit just meets 12 mg · min/L, and the credit when the by-products reach 80 µg/L. *Checkpoint:* 67.6 µg/L and 14.2 mg · min/L.

(c) Find the organic carbon above which no residual meets both limits. *Checkpoint:* 7.69 mg/L.

6.7 App challenge · 8 min

Run Scenario 4, *Feedback against a storm*, with its default settings and the default seed to 300 min, and export the CSV file. Then do the same with Scenario 5, *Feedforward against the same storm*.

(a) For Scenario 4, read `metric_offspec_min` and `turb_set_ntu` in the row at 300 min, and `alum_mgL` and `alum_need_mgL` at 70 min. *Checkpoint:* 141.9 min and 20.92 NTU. 32.2 mg/L dosed against 52.73 mg/L needed.

(b) Read the same four values for Scenario 5. *Checkpoint:* 0.0 min and 3.60 NTU. 52.6 mg/L against 52.73 mg/L.

(c) Read `metric_alum_kg` at 300 min in both files, and explain why the structure with the worse water uses more coagulant. *Checkpoint:* 414.5 kg and 306.0 kg.

Tolerances: 0.1 min, 0.02 NTU, 0.2 mg/L and 1 kg.

6.8 App challenge · 8 min

Run Scenario 7, *The loop that makes the other loop possible*, as loaded, with the default seed, to 300 min. The scenario starts with pH loop gain 0 and sets it to 30 at 200 min.

(a) Read the Off-spec time tile and the largest `turb_set_ntu` in the exported file. *Checkpoint:* About 169.6 min and 9.74 NTU.

(b) Press Reset, set pH loop gain to 30 before Run, run to 300 min and read the same two values. Find the ratio of the settled peaks. *Checkpoint:* 0.0 min and 1.26 NTU, a ratio of about 7.7.

(c) Reload the scenario as shipped, run to 900 min, and read `turb_set_ntu`, `base_mgL` and `ratio_r` in the last row. Explain why the settled water has not recovered. *Checkpoint:* 4.09 NTU, 45 mg/L and 1.35.

Tolerances: 0.2 min, 0.02 NTU and 0.1 mg/L.

6.9 App challenge · 8 min

Run Scenario 8, *Four filters and a queue*, as loaded (Filter sequencing Auto), with the default seed, to 600 min, and export the CSV file.

(a) Read the Off-spec time and Backwashes tiles. Then read the columns `metric_wasted_m3` and `metric_produced_m3` in the last row of the file. *Checkpoint:* 2.9 min, 8, 215 m^3 and 12 193 m^3.

(b) Press Reset, set Filter sequencing to Auto + FTW before Run, and repeat. Find the share of production sent to waste in each run. *Checkpoint:* 0.0 min, 8, 518 m^3 and 12 172 m^3. 1.8 % against 4.3 %.

(c) Find the largest `turb_filt_ntu` in each file. The first run logs a lower peak than the 0.804 NTU of Section 6.5.5. Explain why. *Checkpoint:* About 0.59 NTU and 0.17 NTU.

Tolerances: 0.1 min, 2 m^3 and 0.01 NTU.

6.10 Application area · 3 min

An operator sees the Filtered water tile stay below 0.30 NTU through the storm of Scenario 5 and concludes that the plant met its pathogen-removal obligation.

(a) What does turbidity measure, and why can a plant meet its turbidity limit and still fail a pathogen limit?

(b) Name two features of real filters that the simulator leaves out, and what each would add to an operator's work.

6.11 Application area · 4 min

A manager asks whether a raw-water turbidimeter pays for itself. Use Table 6.3 and the coagulant price of \$0.32/kg in Show Model.

(a) Find the alum cost of the storm under each structure and the saving per storm. *Checkpoint:* \$132.8 and \$97.9, a saving of about \$35.

(b) Taking \$3000 as an example price for the instrument, how many such storms does the chemical saving alone need to repay it? What does the case for the instrument rest on instead? *Checkpoint:* About 86 storms. Compliance: 141.9 min off specification and 47.15 min in violation avoided.

6.12 Review the agent · 5 min

An AI assistant reviews Scenario 4 and writes: "The turbidity loop is too sluggish for a three-hour plant. Double the gain to 8 mg/L per NTU and halve the integral time to 60 min. The dose will then follow the storm, and the off-spec time over 300 min should fall by about half, to near 70 min." Find the flaw, and say how you would check the claim in the app.