INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 2 Practical PID Control

CHAPTER 5

Cooling Tower Makeup Simulator for Identification and Tuning of Integrating Loops

Pages 119–146

Edited by John D. Hedengren

How to cite this chapter. "Cooling Tower Makeup Simulator for Identification and Tuning of Integrating Loops," in *Practical PID Control*, John D. Hedengren, Ed., vol. 2 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 5, pp. 119–146.

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

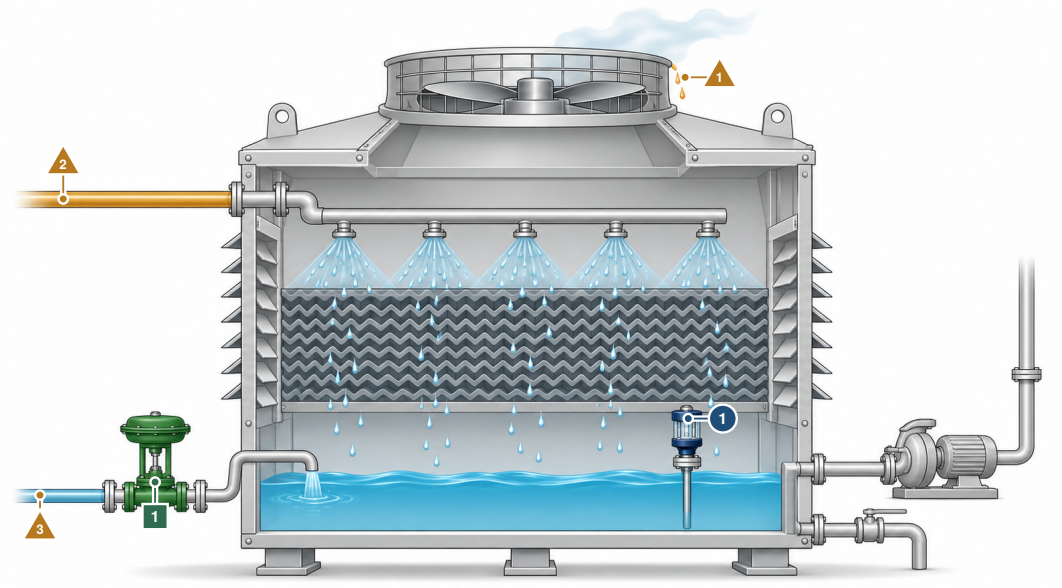
The app for this chapter
https://apmonitor.com/apps/cooling-tower/

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

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

CHAPTER 5

Cooling Tower Makeup Simulator for Identification and Tuning of Integrating Loops



The controller moves 1 the makeup valve

The controller reads 1 the basin level

What disturbs it 1 water carried off as vapor and drift 2 more heat in the return water 3 a sag in the makeup supply pressure

Notice that nothing that drains the basin depends on its level, so the makeup valve alone must match the losses or the level keeps moving. Illustration, not to scale.

IN PLAIN TERMS **Keeping a Cooling Tower's Pool Topped Up**

The main idea. A cooling tower cools water by letting some of it evaporate, and more leaves as spray and drained water. A controller opens a refill valve to keep the pool under the tower at a steady depth.

The challenge. None of these losses changes with the depth, so any mismatch makes the pool rise or fall at a steady rate. One quick test of the valve sets up the controller, and a delay read too short makes the valve swing from shut to wide open.

Why it matters. The refill water is treated and paid for, and a low pool can stop the pumps that cool the plant.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on water use and the history of the float regulator (Section 5.2), the try-it segment of case 6 (Section 5.5.2), the fact on reset time (Section 5.5.3), the try-it segment of case 1 (Section 5.6) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 5.6), Technology in Practice (Section 5.7), Limitations (Section 5.8) and the open challenge

The control question A cooling tower basin loses water by evaporation, blowdown and drift, and none of the three depends on the level, so the level ramps where a drained tank would settle. What does a step test measure on such a plant, and how do its two numbers become a controller that neither cycles nor wears out the makeup valve?

Learning objectives

- Explain from the water balance why the basin level integrates, and compute the valve position that balances the losses.
- Identify an integrating gain and an apparent dead time from the slope of a step response.
- Compute SIMC and Ziegler–Nichols settings from the identified model, and predict the result of a dead time read too short.
- Predict the proportional offset after a load change, and explain why the reset time of an integrating loop must be long.
- Distinguish a makeup capacity limit from a tuning problem.

Prerequisites Sections 1.4.4, 1.4.5 and 1.4.6 (step test, tuning rules, integrating process), 1.4.7 (ultimate gain) and 1.4.9 (windup). Chapter 4 runs the same workflow on a plant that settles. Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/cooling-tower/

See also Volume 1, Chapter 2 for the tank that regulates itself and Volume 1, Chapter 9 for another inventory that integrates. Volume 9, Chapter 2 for predictive control of a delayed integrating plant and Volume 9, Chapter 4 for offset-free estimation.



5.1 Introduction

A gravity-drained tank, the plant of Volume 1, Chapter 2, is the first level loop most process control courses present, and it is atypical. Its outlet flow rises with the head above the nozzle, so an inflow that is slightly too large raises the level until the outflow grows to match, and the tank reaches a steady state without control. Such a plant is self-regulating. A cooling tower basin has no such mechanism. Water leaves the basin as vapor carried off by the fan, as a bleed stream through the blowdown valve, and as droplets entrained past the drift eliminators, and the rate of each is fixed by the process heat load, by the position of the blowdown valve and by the airflow. None of the three is a function of the basin level. The consequence is a plant with a pole at the origin, a level that ramps rather than settles, and an operator who can hold the basin steady only by matching two flows exactly and indefinitely.

The same structure appears throughout a process plant. Surge drums, storage tanks, boiler drums, hoppers and reactor holdups behave the same way whenever the discharge is set by a pump or a downstream controller rather than by hydrostatic head, and integrating inventories are the ordinary case in a process plant rather than the exception [1, 2]. The water level of the pool in Volume 1, Chapter 9 is a pure integrator for the same reason. The physical origin is conservation of mass with outflows independent of the stored inventory, the structure textbooks treat as the integrating or pure capacity process [3, 4]. What changes for the engineer is the whole identification and tuning sequence. An open-loop step on a self-regulating plant is read for a steady-state gain, a time constant and a dead time. An open-loop step on an integrating plant never reaches a steady state, so the reading is a change of slope and a delay, the fitted model is integrating plus dead time (IPDT), and the tuning rules that follow place the dead time alone in the denominator of the gain [5, 6]. Stability of the closed loop then rests entirely on the ratio of the closed-loop speed to that dead time [7].

Simulators of cooling towers are usually built to answer thermal questions. The Merkel enthalpy-difference method and the more complete Poppe and effectiveness-NTU formulations predict outlet water temperature against wet-bulb temperature and air flow [8, 9], and design packages extend that to fill selection and fan power [10]. The water side receives less attention even though it is where the control loop lives [11]. However, a student who is handed an IPDT transfer function and asked for PI settings never performs the experiment that produces the two numbers, and the arithmetic of the

tuning rule is the easy part. This chapter works with an interactive cooling tower basin whose holdup, valve capacity, transport delay and actuator lag are hidden, so that the model must be identified from a measured slope and an identification error appears as sustained oscillation rather than as a small correction. The scenarios cover integrating dynamics, dead time in series with an actuator lag, open-loop identification, model-based PID tuning, anti-windup and the diagnosis of an actuator capacity limit, principles a standard curriculum develops with self-regulating examples [3, 7]. Section 1.4.6 introduces the integrating process with dead time and its tuning rule, and this chapter carries both through on one plant.

The chapter is organized as follows. Section 5.2 reviews evaporative cooling, blowdown and makeup level control. Section 5.3 gives the three-state water and solids balance with its instrument chain and its actuator model, written out in the detail needed to rebuild the simulator. Section 5.4 derives the transfer function from valve to level, identifies it from the slope of a step response and shows why the apparent dead time exceeds the transport delay. Section 5.5 sets out the level controller, the two tuning rules computed from the identified model, the proportional offset, the role of the reset time and a linear model predictive controller (MPC) built on the same identified integrating model. Section 5.6 measures these designs on a scenario sequence that begins with manual operation and ends at a makeup capacity limit that no tuning removes. Section 5.7 sets the loop in its plant hardware, Sections 5.8 and 5.9 state what the model omits and what level control costs in water, and Section 5.11 holds the exercises.

5.2 Background

Evaporative cooling is a mature technology used at large scale. A tower rejects heat by evaporating roughly one percent of the circulating flow for every 10 °F of cooling range, which is why the water side of a tower is a consumption problem and not merely a recirculation problem. Thermoelectric generation, refining, chemicals and data centers all depend on it, and a review of operational water factors for electricity generation puts recirculating cooling in a class where consumption per unit of output is far higher than for once-through cooling [12]. Every gallon that evaporates leaves its dissolved solids behind in the basin, so a second stream has to be discarded to keep the concentration bounded, and a third is lost as fine droplets past the eliminators.

> **DID YOU KNOW? Where a Building's Water Goes**
>
> In commercial and institutional buildings, cooling towers are often the largest single use of water, at 20 to 50 percent or more of a facility's total. Drift, the mist carried off in the air, is 0.05 to 0.2 percent of the flow through a tower, and modern drift eliminators can cut it below 0.005 percent [13].

Analysis of the thermal side matured first. Merkel's enthalpy-driving-force treatment of the 1920s was reduced to the tower characteristic and the demand curves that designers still use [8], and later work quantified how far that treatment departs from a full heat and mass balance and how the evaporated fraction should be corrected for it [9, 11]. The water balance followed the chemistry rather than the thermodynamics. The ratio of basin conductivity to makeup conductivity defines the cycles of concentration, and the blowdown flow that holds a target ratio is fixed once evaporation and drift are known. Raising the target reduces water use and moves the basin closer to scaling and corrosion limits, which is why blowdown moved from a fixed manual bleed to a conductivity switch that opens a valve only when the basin passes its limit. Blowdown streams are concentrated enough that their treatment and reuse are studied in their own right [14].

Makeup level control followed a different path. The earliest arrangement is a mechanical float valve in the basin, which is a proportional controller whose gain is set by the float geometry and whose offset is rarely measured. Pneumatic and then electronic level transmitters replaced the float, and the loop became one line in a distributed control system. The difficulties characteristic of level loops survived the change of hardware. A basin level offers no natural damping, its measurement is

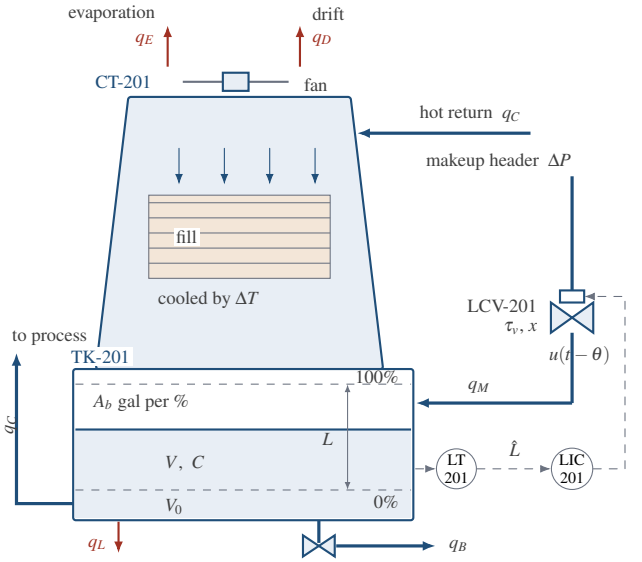


Figure 5.1 Cooling tower CT-201 and basin TK-201 with the makeup, blowdown, drift and evaporation streams. Circulating water q_C returns hot, is cooled through the fill by ΔT and falls into the basin, which holds volume V above the transmitter zero at conductivity C and volume V_0 below it. Makeup q_M enters through valve LCV-201, whose stem x follows the delayed controller output $u(t-\theta)$ with lag τ_v against header differential ΔP. LT-201 reports the level $\hat{L}$ over a span holding A_b gallons per percent, and LIC-201 drives the valve.

noisy because the surface is disturbed by the returning water, and the loops are frequently detuned on purpose so that the inventory absorbs upsets rather than passing them downstream [2]. Tuning rules written specifically for integrator-plus-dead-time plants appeared for the same reason [1, 6], and identification methods for plants with a pole at the origin are still an active subject [15], because the settled step response that a self-regulating plant provides is not available and the model must be constructed from a slope.

> **HISTORY** **The Float Regulator, About 270 BC to 1791**
>
> Around 270 BC the Greek inventor Ktesibios built a float regulator for a water clock. It held the water in a tank at a constant depth, so that a steady flow from its bottom filled a second tank at a constant rate, and the level in the second tank told the time. The float worked an inflow valve: as the level fell, the valve opened and refilled the tank, as the ball and cock of a flush toilet does today. The device returned in the late 1700s, when the boilers of steam engines and household water reservoirs both needed a held level. A patent of 1758 describes a float valve on a steam boiler, and by 1791 the float regulator was in common use in steam engines [16]. The first makeup controller of a cooling tower basin was the same device.

5.3 Process Model

Tower CT-201 is an induced-draft counterflow unit that circulates 3000 gpm of water through a process returning it 10 °F hotter at design load. Cooled water falls into basin TK-201, which feeds the circulating pump. Makeup comes from a treated-water header through control valve LCV-201, blowdown leaves through a valve of its own, and level transmitter LT-201 measures a 1.2 ft band on a 250 ft^2 basin. Figure 5.1 shows the layout, the instrument tags and each symbol used in the equations that follow. LT-201 supplies the controlled variable in percent of that span, and LCV-201 position is the only manipulated variable in the application.

Table 5.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
L (LT-201)	CV, measured	60 %	30–85 % SP
u (LCV-201)	MV	42.5 %	0–100 %
heat load	DV, measured	100 %	50–150 %
ΔP	DV, measured	100 %	40–120 %
blowdown lift	DV, measured	100 %	0–300 %
m_D drift	DV, unmeasured	1×	1–40×
q_L leak	DV, unmeasured	0 gpm	0–15 gpm
V	state	1350 gal	0–2250 gal
x	state	42.5 %	0–100 %
C	state	2000 µS/cm	$\geq C_M$

Table 5.1 classifies what the loop interacts with. Three alarm limits bound normal operation, a high level at 85 %, a low level at 40 % and a pump-suction trip at 20 %. None of the three stops the run. Each one latches an annunciation, so a test that carried the basin outside its limits leaves a record. The plant is expressed in the units of utility operation, flows in gpm, volumes in gal, level and valve position in percent, conductivity in µS/cm and time in minutes. Three states are integrated.

5.3.1 Water Balance on the Basin

The basin is a single well-mixed inventory whose level is linear in volume over the transmitter span, and the balance that governs it is Equation (5.1).

$$A_b \frac{\mathrm{d}L}{\mathrm{d}t} = q_M - q_E - q_B - q_D - q_L \tag{5.1}$$

where A_b is the basin holdup per percent of span in gal/%, L is the level in percent, q_M is makeup, q_E is evaporation, q_B is blowdown, q_D is drift and q_L is a basin leak, all in gpm. Both sides carry units of gal/min. The state actually integrated is the volume above the transmitter zero, $V = A_b L$, in gallons.

Nothing on the right-hand side of (5.1) is a function of L, and that absence is the subject of this chapter. Substituting a head-driven outlet $q_{\text{out}} = k\sqrt{L}$ would add a term that grows as the tank fills, would linearize to a first-order lag, and would give the tank a steady state of its own. The cooling tower basin has no such term, so a constant imbalance produces a constant slope and the level travels until it reaches the overflow weir at 100 % or the empty basin at 0 %. Both walls are enforced twice. The derivative in (5.1) is set to zero when the state sits at a bound and the imbalance points outward, so the fixed-step integrator cannot pass through the wall and return, and the state is clipped again after every step. Volume clipped off at the top latches a basin-overflow annunciation. While the level sits on the weir, the surplus that the blocked derivative discards is recomputed after each step and metered as a spill rate and a totalized spill volume, and a deficit at the empty basin is reported the same way, so the water sent to drain is accounted for during a sustained overflow.

🕮 **Pause.** Try Exercise 5.1 before reading on.

5.3.2 Evaporation, Blowdown and Drift

Evaporation is written in the form a tower vendor uses, as a fraction of the circulating flow proportional to the cooling range, as shown in Equation (5.2).

$$q_E = 0.001\, f_E\, q_C\, \Delta T, \qquad \Delta T = \Delta T_{\text{nom}} \frac{\text{heat load}}{100\,\%} \tag{5.2}$$

where q_C is the circulation rate in gpm, ΔT is the cooling range in °F, f_E is the fraction of the duty rejected as latent heat, and the coefficient 0.001/°F carries the units. The heat load is an operator-set

input that moves ΔT and therefore q_E with no dynamics between them. Drift is the entrained liquid that escapes the eliminators, taken as a fixed fraction of circulation in Equation (5.3).

$$q_D = f_D \, q_C \, m_D \tag{5.3}$$

where f_D is the design drift fraction and m_D is a multiplier that represents damaged or fouled eliminators. Blowdown is sized from the cycles of concentration $N = C/C_M$ that the water treatment program holds, as shown in Equation (5.4).

$$q_{B,0} = \max\left(0,\ \frac{q_E}{N-1} - f_D q_C\right), \qquad q_B = q_{B,0}\,\frac{\text{lift}}{100\,\%} \tag{5.4}$$

where N is floored at 1.05 and the design drift $f_D q_C$ is credited against the required bleed because drift removes solids as effectively as blowdown does. The multiplier m_D deliberately does not appear in (5.4), so raising the drift on a fixed bleed dilutes the basin and pulls the achieved cycles below target, which is what a real fixed bleed does. In the alternative blowdown mode the lift term is replaced by an on-off valve that passes $3q_{B,0}$ while open and nothing while shut, switched on basin conductivity with a hysteresis band of $\pm 1.5\,\%$ about NC_M.

At the design point the three losses are $q_E = 25.50\,\text{gpm}$, $q_B = 6.225\,\text{gpm}$ and $q_D = 0.15\,\text{gpm}$, which total 31.875 gpm.

5.3.3 Makeup Valve, Transport Delay and Level Transmitter

The makeup valve is linear in stem position with a square-root pressure term, and its stem follows a delayed command through a first-order actuator lag, as shown in Equations (5.5) and (5.6).

$$\tau_v \frac{\mathrm{d}x}{\mathrm{d}t} = u(t-\theta) - x \tag{5.5}$$

$$q_M = q_{M,\max}\,\frac{x}{100}\sqrt{\frac{\Delta P}{\Delta P_{\text{nom}}}} \tag{5.6}$$

where x is the stem position in percent, u is the controller output in percent, θ is the combined transport and measurement dead time in minutes, τ_v is the actuator lag in minutes, $q_{M,\max}$ is the valve capacity at full lift and nominal differential, and $\Delta P/\Delta P_{\text{nom}}$ is the header pressure ratio. Delaying the command in (5.5) rather than delaying the flow is equivalent and keeps the delay line out of the right-hand side, which the fixed-step integrator evaluates four times per step. The delay line stores each controller output and is read at every controller execution, so a command issued at one execution, the first of a run included, reaches the stem exactly 0.50 min later and is held over the following 0.02 min scan. The stem is clipped to 0–100 % after each step. Equation (5.6) makes the valve gain a function of header pressure, so a model identified at 100 % pressure overstates the makeup available at 40 % by a factor of $1/\sqrt{0.4} = 1.58$.

The level transmitter is the only instrument the controller reads, and it is sampled at each controller execution. It adds a constant bias, then noise made by passing a Gaussian sequence through a first-order filter of 0.15 min and rescaling it so that the stationary standard deviation matches the specified value, then a fixed transmitter filter of time constant τ_S, as shown in Equation (5.7).

$$\tau_S \frac{\mathrm{d}\hat{L}}{\mathrm{d}t} = L + b + n(t) - \hat{L} \tag{5.7}$$

where $\hat{L}$ is the measured level, b is the transmitter bias and n is the noise sequence. The filter is advanced by the exact zero-order-hold update $\hat{L}_k = \alpha \hat{L}_{k-1} + (1-\alpha)(L + b + n)$ with $\alpha = \exp(-\Delta t/\tau_S)$, which is stable at any sample period. A separate first-order filter of time constant 1.0 min produces a smoothed rate of change of $\hat{L}$, reported in %/min as the instrument for reading a step test. The makeup, evaporation and blowdown flow indications carry no delay and no filter.

Table 5.2 Model Parameters. Hidden Constants Are Never Displayed and Are the Quantities the Step Test Recovers.

Symbol	Meaning	Value	Unit	Visibility
q_C	circulation rate	3000	gpm	adjustable
ΔT_{nom}	design cooling range	10	°F	adjustable
f_E	evaporation factor	0.85	—	adjustable
f_D	drift fraction	5×10^{-5}	—	adjustable
N	target cycles	5	—	adjustable
C_M	makeup conductivity	400	µS/cm	adjustable
V_0	volume below zero	6000	gal	adjustable
τ_S	transmitter filter	0.05	min	adjustable
—	high-level alarm	85	%	adjustable
—	low-level alarm	40	%	adjustable
—	pump-suction trip	20	%	adjustable
A_b	holdup per percent	22.5	gal/%	hidden
$q_{M,\max}$	makeup valve capacity	75	gpm	hidden
θ	makeup dead time	0.50	min	hidden
τ_v	valve actuator lag	0.33	min	hidden

5.3.4 Dissolved Solids and the Conductivity Switch

Evaporation removes water and leaves its solids behind, so the basin concentrates until the streams that carry solids out balance the solids arriving in the makeup. The solids balance is Equation (5.8).

$$\frac{\mathrm{d}(CV_t)}{\mathrm{d}t} = C_M q_M - C\,(q_B + q_D + q_L) \tag{5.8}$$

where C is the basin conductivity in µS/cm, C_M is the makeup conductivity and $V_t = \max(100\,\text{gal},\ V_0 + V)$ is the total wetted volume with V_0 the basin volume below the transmitter zero. Evaporation appears nowhere on the right of (5.8), which is exactly why the basin concentrates. Expanding the product and substituting (5.1) gives the form that is integrated, as shown in Equation (5.9).

$$\frac{\mathrm{d}C}{\mathrm{d}t} = \frac{q_M\,(C_M - C) + q_E\,C}{V_t} \tag{5.9}$$

The state is floored at C_M, since the basin cannot be cleaner than its makeup. At the design point (5.9) is stationary at $C = NC_M = 2000\,\mu\text{S/cm}$, which is one of the tests in the verification suite shipped with the application. In the on-off blowdown mode this second loop becomes a square-wave load on the level loop, and the period of the square wave is set by V_t, which is why the cycling scenario reduces V_0 from 6000 gal to 1500 gal.

5.3.5 Parameters, Initialization and Numerics

The constants are listed in Table 5.2. Eleven of them are adjustable. Four are hidden and never displayed, and those four set the dynamics that the step test of Section 5.4.2 recovers.

Every scenario opens at a genuine steady state. The volume is initialized to $V(0) = A_b L_{\text{SP}}$, the conductivity to $C(0) = NC_M$, and the valve position to the opening that balances the present losses exactly, as shown in Equation (5.10).

$$x(0) = \frac{100\,(q_E + q_B + q_D + q_L)}{q_{M,\max}\sqrt{\Delta P / \Delta P_{\text{nom}}}} \tag{5.10}$$

At the design point (5.10) returns 42.5 %, and the same expression supplies the flow-balancing aid available in manual mode. Two scenarios override the initial state, one starting the level 8 % below setpoint at $V = 1170$ gal and one starting the conductivity high at 2100 µS/cm.

The integrator is classical fourth-order Runge-Kutta at a fixed 0.005 min (0.3 s) step. LIC-201 executes on a 0.02 min (1.2 s) period, and the historian records one sample at the same 0.02 min interval, 50 samples for every simulated minute. The default seed of the noise generator is 20260825. Section 1.6.1 describes the three clocks, the seeded noise and the speed setting, which this app shares with the others.

5.3.6 Assumptions

The model is deliberately thin on the thermal side and detailed on the water side, and each simplification fails somewhere.

One well-mixed basin, linear in volume. The transmitter span is treated as a prism of constant area, and the water held up in the fill, the distribution deck and the return piping is ignored except through the discrete second-pump event, which removes 3 % of span (67.5 gal) instantly. Surface waves from the returning water and wind are represented only by the measurement noise term. The assumption fails on a basin with a sloped floor or a sump, where the gallons per percent change with level.

Evaporation is algebraic in the heat load. Equation (5.2) has no thermal capacitance, no wet-bulb dependence and no transport lag through the fill, so a heat-load change reaches the water balance at once. A real tower takes a minute or two to respond, and f_E itself falls in cold or humid weather because more of the duty leaves as sensible heat [9, 11]. The assumption overstates how quickly a load upset appears as a level ramp.

Cycles of concentration are the only chemistry. Scaling indices, corrosion, biological control, pH and chemical feed are outside the model, and conductivity stands in for total dissolved solids. The model therefore cannot be used to choose a target N, only to show what holding one costs in water.

The valve is ideal apart from lag and pressure. There is no stiction, no hysteresis, no equal-percentage trim and no finite stroking rate, and the pressure term in (5.6) assumes the valve does not itself destroy the differential it works against. Any of those would change the identified gain and would add a limit cycle that tuning cannot remove.

Dead time is one lumped constant. The delay in (5.5) covers piping transport, valve positioner response and measurement, and it is held fixed. In a plant the transport component falls as makeup flow rises.

The numerics are more repeatable than a control system. An identical trajectory on every machine is convenient and false, because a real level input arrives with jitter, a coarser resolution and a scan order that moves.

5.4 Dynamic Analysis

5.4.1 Transfer Function From Valve to Level

Linearizing (5.1), (5.5) and (5.6) about any operating point gives the plant in Equation (5.11), because the losses contribute nothing to the derivative with respect to L.

$$\frac{L(s)}{U(s)} = \frac{K_I\, \mathrm{e}^{-\theta s}}{s\,(\tau_v s + 1)}, \qquad K_I = \frac{q_{M,\max}/100}{A_b}\sqrt{\frac{\Delta P}{\Delta P_{\mathrm{nom}}}} \tag{5.11}$$

where K_I is the integrating gain in %/min of level per percent of valve. At the design condition $K_I = 0.75/22.5 = 0.0333\,\%/(\mathrm{min}\,\%)$, and the verification suite confirms that value directly by evaluating the derivative at a valve position 10 % above balance. The linearization in (5.11) does not depend on the level, which distinguishes the basin from a drained tank. A gravity-drained tank of area A and outlet coefficient k linearizes to $K_p/(\tau_p s + 1)$ with $\tau_p = 2A\sqrt{L_s}/k$ and a gain that both depend on the operating level. The basin has one gain, one delay and no process time constant, and none of the three moves with level.

Two consequences follow immediately. The steady-state gain of (5.11) is infinite, so there is no process gain to read from a settled step response and no self-regulation. The open-loop phase begins

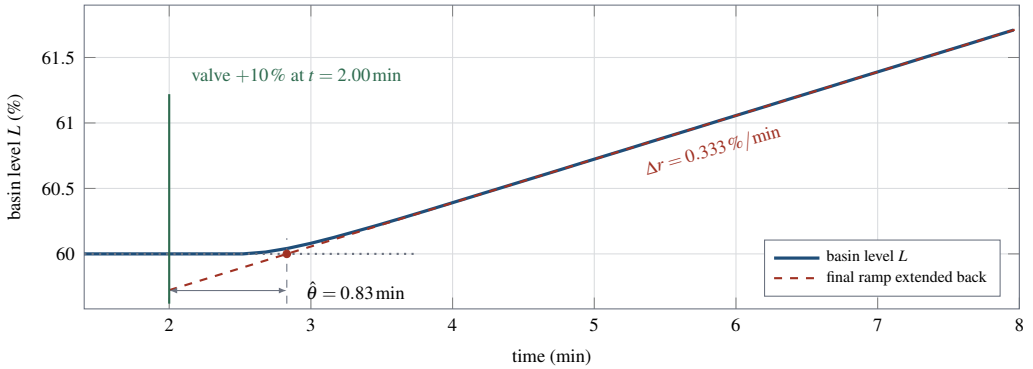


Figure 5.2 Identification of the basin from one open-loop step, recorded by the application. The makeup valve is commanded 10 % open at $t = 2.00$ min from the balanced position. The level departs from its pre-step value by 0.01 % 0.65 min later and reaches a straight ramp of 0.3333 %/min, giving $\hat{K}_I = 0.0333\,\%/(\text{min}\,\%)$. Extending that ramp backward to the pre-step level places the intersection 0.83 min after the step, against a true transport delay of 0.50 min.

at -90° and falls monotonically, so the loop has phase margin only because the dead time is small compared with the speed the tuning asks for.

5.4.2 Identification From the Step-Response Slope

The identification is a 10 % open step on the makeup valve from the balanced position, applied as a calibrated command so that the step size is exact. It is the integrating counterpart of the step test of Section 1.4.4. What is read from the trend is a change of slope and a delay, not a gain and a time constant, as shown in Equation (5.12).

$$\hat{K}_I = \frac{\Delta r}{\Delta u}, \qquad \hat{\theta} = t_\times - t_{\text{step}} \tag{5.12}$$

where Δr is the change in level rate in %/min, Δu is the step in valve percent, and $t_\times$ is the time at which the final ramp, extended backward, crosses the pre-step level. The gain $\hat{K}_I$ is the K' of Equation (1.13). The construction is shown against data recorded by the application in Figure 5.2, which plots the basin level L that the application trends beside the transmitter reading $\hat{L}$. Figure 5.5 shows the application during the same test.

The slope fitted by least squares from 6 to 12 min is 0.3333 %/min against a commanded step of 10 %, which gives $\hat{K}_I = 0.0333\,\%/(\text{min}\,\%)$ and matches the true K_I of (5.11) to four significant figures. The procedure does not require the level to settle, and waiting longer adds no information, because the ramp continues at the same slope until the basin overflows. The application grades the entered $\hat{K}_I$ and $\hat{\theta}$ against the plant and reports each error as a percentage.

5.4.3 Why the Apparent Dead Time Exceeds the Transport Delay

The plant in (5.11) is not IPDT. It carries a pure delay of 0.50 min and, in series with it, a valve lag of 0.33 min. A step test cannot separate the two, because a first-order lag on the input of an integrator produces exactly the same asymptotic ramp as a pure delay of the same magnitude, offset by τ_v. Integrating (5.5) and (5.1) for a step of size Δu at $t = 0$ gives Equation (5.13).

$$L(t) = L_0 + K_I \Delta u \left[t - \theta - \tau_v \left(1 - \mathrm{e}^{-(t-\theta)/\tau_v} \right) \right], \quad t \geq \theta \tag{5.13}$$

The exponential decays, and the surviving asymptote is a straight line of slope $K_I \Delta u$ that crosses L_0 at $t = \theta + \tau_v$. The effective dead time an IPDT fit recovers is therefore the sum, as shown in

Equation (5.14).

$$\theta_{\text{eff}} = \theta + \tau_v = 0.50 + 0.33 = 0.83\,\text{min} \tag{5.14}$$

The value read from Figure 5.2 is 0.83 min, in agreement with (5.14). The same construction applied to the transmitter reading $\hat{L}$ places the intersection at 0.89 min, because the 0.05 min transmitter filter of (5.7) and the 0.02 min scan at which $\hat{L}$ is refreshed add their own lag to the apparent dead time. An alternative and easier reading is the moment the trend visibly departs from the flat line, which occurs 0.65 min after the step at a resolution of 0.01 % and tends toward the true transport delay of 0.50 min as the resolution of the trend improves. That reading is wrong for tuning purposes even though it is correct as a statement about the plant, and Section 5.6 measures what the difference costs. The grading compares $\hat{\theta}$ against θ_{eff} of (5.14) rather than against θ, for that reason. Section 1.4.4 reports the same reading on a self-regulating plant, where a small lag is read almost entirely as dead time.

🕮 **Pause.** Try Exercise 5.2 before reading on.

5.5 Controller Design

5.5.1 Velocity Form PID Algorithm

LIC-201 runs the velocity-form algorithm that the framework apps share (Section 1.6.2), written for this loop in Equations (5.15) and (5.16). The derivative acts on the measurement, so moving the setpoint alone cannot kick the valve (Section 1.4.8).

$$\begin{aligned} e_k &= L_{\text{SP},k} - \hat{L}_k \\ \Delta u_k^P &= K_c\,(e_k - e_{k-1}) \\ \Delta u_k^I &= \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \\ \Delta u_k^D &= -\,K_c\,\tau_D\,\frac{\hat{L}_k^f - 2\hat{L}_{k-1}^f + \hat{L}_{k-2}^f}{\Delta t_c} \end{aligned} \tag{5.15, 5.16}$$

The controller is reverse acting, because the process gain is positive and a rising level calls for less makeup (Section 1.4.3). Derivative comes from a second difference on a filtered measurement, with $\hat{L}_k^f = \alpha \hat{L}_{k-1}^f + (1-\alpha)\hat{L}_k$ and $\alpha = \exp(-\Delta t_c/\tau_f)$, where the filter constant τ_f defaults to 0.05 min. The proportional and derivative increments are formed as differences of the complete terms $K_c e_k$ and $-K_c\tau_D(\hat{L}_k^f - \hat{L}_{k-1}^f)/\Delta t_c$, which equals (5.16) while the gains are constant and makes a retuned gain act at once. The increments are combined, the demand is clipped to the valve span, and the integral is then corrected by back-calculation [17], as shown in Equation (5.17).

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= s_k + I_{k-1} \\ u_k &= \min(100,\ \max(0,\ c_k)) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{\tau_t}\,(u_k - c_k) \end{aligned} \tag{5.17}$$

Here s_k accumulates the proportional and derivative increments without limit, c_k is the position the algorithm demands with the integral accumulated through the previous scan, u_k is that demand clipped to the valve span, and $\tau_t = \tau_I$ is the tracking time of the back-calculation. These are Equations (1.24) to (1.26) with s_k for S_k, c_k for v_k and no bias. Section 1.4.9 explains why the integral alone is corrected and the running sum is never truncated. With $\tau_t = \tau_I$ the correction cancels Δu_k^I exactly when the clipped excess $c_k - u_k$ equals the proportional term $K_c e_k$, so during a sustained saturation without derivative action the integral holds at the applied output instead of growing, as Section 5.6 measures. The bias is zero and I is seeded at start-up with the initial valve command, so the first

scan contributes the whole positional proportional term. The output carries no rate limit. Transfer between manual and automatic is bumpless in both directions through the tracking of Section 1.4.10. Execution is every 0.02 min, independent of the 0.005 min physics step. Tuning constants are bounded to $0 \le K_c \le 60\,\%/\%$, $0.2\,\text{min} \le \tau_I \le 60\,\text{min}$ and $0 \le \tau_D \le 2\,\text{min}$. The top of the τ_I range is read as an infinite reset time, which removes both the integral increment and the back-calculation and leaves a proportional-plus-derivative loop.

5.5.2 SIMC and Ziegler-Nichols Settings

Both rules read $\hat{K}_I$ and $\hat{\theta}$ exactly as they were entered and never consult the hidden constants, so an identification error propagates directly into the tuning. Equation (5.18) is the SIMC setting for an integrator with dead time [5], which is Equation (1.14) with $\hat{K}_I$ for K'.

$$K_c = \frac{1}{\hat{K}_I(\tau_c + \hat{\theta})}, \qquad \tau_I = 4(\tau_c + \hat{\theta}), \qquad \tau_D = 0 \tag{5.18}$$

where τ_c is the closed-loop time constant and the application uses $\tau_c = \hat{\theta}$, the usual fast-but-robust choice. The reaction-curve PID settings [18] follow from the ultimate gain and period of an IPDT plant under proportional control (Section 1.4.7). The phase of $K_I e^{-\theta s}/s$ reaches -180° at $\omega_u = \pi/(2\theta)$, at which the magnitude is $2K_I\theta/\pi$, giving Equation (5.19).

$$K_u = \frac{\pi}{2K_I\theta}, \qquad P_u = \frac{2\pi}{\omega_u} = 4\theta \tag{5.19}$$

Substituting (5.19) into the classical $0.6K_u$, $P_u/2$, $P_u/8$ settings of Table 1.2 gives the rule the application implements, as shown in Equation (5.20).

$$K_c = \frac{0.94}{\hat{K}_I\hat{\theta}}, \qquad \tau_I = 2\hat{\theta}, \qquad \tau_D = 0.5\hat{\theta} \tag{5.20}$$

Results are rounded to a resolution of 0.1 on K_c and τ_I and 0.01 on τ_D, and clamped to the tuning bounds before they are loaded. With $\hat{K}_I = 0.0333$ and $\hat{\theta} = 0.83$, the effective dead time of (5.14) that the identification of Figure 5.2 recovers to within 0.01 min, (5.18) returns $K_c = 18.1\,\%/\%$ and $\tau_I = 6.6\,\text{min}$, while (5.20) returns $K_c = 34.0\,\%/\%$, $\tau_I = 1.7\,\text{min}$ and $\tau_D = 0.42\,\text{min}$. These are the settings the scenarios preload, and the identified gain is entered to four decimals, so the step-test reading reproduces both.

Equation (5.20) has $\hat{\theta}$ alone in a denominator, with no additive term to soften a bad reading. Reading the dead time as the visible departure of the trend, 0.50 min instead of 0.83 min, raises the gain to 56.4 %/% with $\tau_I = 1.0\,\text{min}$ and $\tau_D = 0.25\,\text{min}$, the tuning scenario 6 loads (from the exact $K_I = 1/30$, while an entry of 0.0333 gives 56.5). Equation (5.19) evaluated with the effective dead time of (5.14) gives $K_u = \pi/(2 \cdot 0.0333 \cdot 0.83) = 56.8\,\%/\%$, so within the IPDT model the rule assumes, the optimistic reading places the proportional gain at 99 % of the ultimate gain. The simulated loop has more margin than that, because part of its apparent dead time is the 0.33 min lag rather than a pure delay. A proportional-only test of the loop, transmitter filter included, reaches sustained oscillation at $K_c = 65.5\,\%/\%$ (beyond the tuning bound of 60) with a period of 3.39 min, so 56.4 is 86 % of the actual ultimate gain. The cycling reported in Section 5.6 comes from pairing that gain with the reset time $\tau_I = 2\hat{\theta} = 1.0\,\text{min}$ that the same reading error sets, since the same gain and derivative settle with $\tau_I = 1.7\,\text{min}$ or with no integral action at all. Equation (5.18) is far less sensitive, because its denominator is $\tau_c + \hat{\theta} = 2\hat{\theta}$ and the same error moves K_c by the same factor from a much lower starting value.

Table 5.3 Effect of Reset Time Alone on a 5 % Setpoint Move at $t = 1$ min, $K_c = 18.1\,\%/\%$ and $\tau_D = 0$ Throughout, Read at $t = 26$ min. Settling Is Into the $\pm 2\,\%$ Band.

τ_I (min)	6.6	2.0	1.0
Overshoot (%)	26.2	49.7	53.7
Settling time (min)	2.38	6.44	13.48
IAE (% min)	18.00	22.76	38.37
Valve travel (% of span)	132.3	249.3	625.2
Time saturated (%)	8.9	12.1	33.8

TRY IT IN THE APP 6 · Aggressive Tuning, Optimistic Dead Time

Load the case, which preloads $K_c = 56.4$, $\tau_I = 1.0$ min and $\tau_D = 0.25$ min from a dead time read as 0.50 min, and run it to 30 min. After the heat load rises at 5 min, watch the makeup valve trend: from 20 min it swings between 4.7 and 100 % while the level stays between 59.09 and 60.84 %. Then Reset, enter $\hat{\theta} = 0.83$ min and choose the ZN PID rule before 5 min.

5.5.3 Proportional Offset and Reset Time

The pole at the origin also changes what proportional action alone can achieve. With $\tau_I \to \infty$ and the delay neglected, the loop transfer from setpoint to level is $K_c K_I/(s + K_c K_I)$, which is unity at $s = 0$, so a proportional controller on an integrating process tracks a setpoint change with no offset at all. A self-regulating plant cannot do that (Equation (1.3)). The offset reappears for load changes, and its size is worth computing because it explains the peak deviations reported later. Writing a sustained load as an equivalent valve demand d in percent, the steady level satisfies Equation (5.21).

$$L_\infty - L_{\mathrm{SP}} = -\frac{d}{K_c} \tag{5.21}$$

A 30 % rise in heat load raises evaporation from 25.50 to 33.15 gpm and, through (5.4), blowdown from 6.225 to 8.14 gpm, a total of 9.5625 gpm that the valve must supply. Dividing by the valve capacity of 0.75 gpm per percent gives $d = 12.75\,\%$, so (5.21) predicts a residual deviation of 0.70 % at $K_c = 18.1$ and 0.375 % at $K_c = 34$ before integral action removes it. With the reset time at the top of its range, so that the loop is proportional only, the level settles 0.704 % and 0.375 % below setpoint at those two gains after the load step, as (5.21) predicts.

Integral action is therefore needed only to erase (5.21), and it carries a cost. A PI controller on (5.11) places two open-loop poles at the origin, so the open-loop phase starts at -180° and the loop is stable only because the PI zero at $1/\tau_I$ supplies phase lead below the crossover frequency. Shortening τ_I moves that zero above the crossover, the phase lead is lost, and the loop oscillates slowly about the setpoint instead of settling. This is the mechanism behind the long reset times that tuning rules written for integrating plants recommend [1, 6], and behind the factor of four in (5.18). Section 1.4.6 finds the same trend on the running example of Chapter 1. Table 5.3 measures it on a 5 % setpoint move with the proportional gain held at 18.1 %/%.

Cutting τ_I from the SIMC value of 6.6 min to 1.0 min multiplies the settling time by 5.7, from 2.38 to 13.48 min, and the valve travel by 4.7, from 132.3 to 625.2 % of span, on a loop whose proportional gain never changed. At $\tau_I = 1.0$ min the level stays inside the band after 13.48 min but never stops rolling, with an amplitude of 1.9 % about the setpoint and the valve swinging between 0 and 94 % from 16 to 26 min, so its travel keeps growing with the length of the run. This failure mode is particular to integrating processes, and it reverses the intuition a first-order plant teaches, where reducing τ_I shortens the response over a wide range.

> ✦ DID YOU KNOW? **A Shorter Reset Time That Slows the Loop**
>
> From the simulator: with the gain held at 18.1 %/% on a 5 % setpoint move, cutting the reset time from 6.6 to 1.0 min lengthens the settling time from 2.38 to 13.48 min and raises the valve travel from 132.3 to 625.2 % of span (Table 5.3).

5.5.4 Model Predictive Control on the Identified Integrator

The same two identified numbers also serve as the prediction model of a third design, an MPC that moves LCV-201 directly. Only the comparison case, scenario 10, starts under the MPC. It is the receding-horizon controller of Section 1.6.3 [19, 20], and it makes a decision every $T_s = 0.1$ min, five controller executions. Its model is Equation (5.11) with the actuator lag folded into the dead time, the IPDT form the step test identifies, plus an unmeasured input disturbance, as shown in Equation (5.22).

$$\hat{L} - \hat{L}_0 = z, \qquad \frac{\mathrm{d}z}{\mathrm{d}t} = \hat{K}_I \left[u(t - \hat{\theta}) - u_0 + \delta \right] \tag{5.22}$$

where u_0 and $\hat{L}_0$ are the command and the LT-201 reading when the estimator starts, z is the predicted level deviation in percent and δ is a load the model cannot measure, expressed in percent of valve. On an integrator a constant δ is a ramp rather than an offset, so a heat-load change, a leak, a sagging header and an error in $\hat{K}_I$ all register as makeup that is not arriving, and the load d of Equation (5.21) appears as $\delta = -d$. The MPC reads $\hat{K}_I$ and $\hat{\theta}$ as they were entered for the rules, so a misread step test misleads it as it misleads them, and with either entry at zero the PID keeps the valve. A least-squares IPDT fit to ± 10 % valve steps on the simulator, without noise, returns $\hat{K}_I = 0.03324$ and $\hat{\theta} = 0.873$ min on the transmitter with a root-mean-square residual of 0.19 % of the response, the transient of the lag that the IPDT form cannot follow, and the comparison case uses the 0.0333 and 0.83 min read from Figure 5.2. The dead time is a record of the last $D = \mathrm{round}(\hat{\theta}/\Delta t_c) = 42$ commands sent.

A steady-state Kalman filter on (z, δ), the integrator sampled at Δt_c and augmented with δ as a random walk, which is the usual route to offset-free predictive control [21] and the subject of Volume 9, Chapter 4, runs at every execution whichever design holds the valve, as shown in Equation (5.23).

$$\begin{aligned} \bar{z}_k &= \hat{z}_{k-1} + \hat{K}_I \Delta t_c \left(u_{k-1-D} - u_0 + \hat{\delta}_{k-1} \right) \\ \nu_k &= \hat{L}_k - \hat{L}_0 - \bar{z}_k \\ \hat{z}_k &= \bar{z}_k + \ell_z \nu_k, \qquad \hat{\delta}_k = \hat{\delta}_{k-1} + \ell_\delta \nu_k \end{aligned} \tag{5.23}$$

where ν_k is the innovation and the gains $\ell_z = 0.1064$ and $\ell_\delta = 8.977$ percent of valve per percent of level follow from noise variances r, $10^{-6} r$ and $\rho\, r / \hat{K}_I^2$ on the measurement, z and δ, with $\rho = 0.1$, which gives the estimation error a time constant of 0.36 min. At rest the plan holds the valve at $u_0 - \hat{\delta}$, so $\hat{\delta}$ supplies the integral action, and after the heat load of scenario 5 it settles at −12.750 %, the $-d$ computed for Equation (5.21).

Each decision starts from the deviation z_0 that the D commands in transit will leave, as a Smith predictor would, and solves the quadratic program of Equation (5.24) on the model resampled at T_s.

$$\begin{aligned} \min_{v,\varepsilon} \; & \frac{T_s}{1\,\mathrm{min}} \sum_{j=1}^{N_p} \left(\frac{\hat{L}_0 + z_j - L_{\mathrm{SP}}}{1\,\%} \right)^2 \\ & + \lambda \sum_{i=1}^{12} \left(\frac{v_i - v_{i-1}}{5\,\%} \right)^2 + c_\varepsilon \varepsilon^2 \\ \text{s.t. } & z_j = z_{j-1} + \hat{K}_I T_s \left(v_{i(j)} - u_0 + \hat{\delta}_k \right) \\ & 0 \le v_i \le 100\,\%, \qquad \varepsilon \ge 0 \\ & L_{\mathrm{lo}} - \varepsilon \le \hat{L}_0 + z_j \le L_{\mathrm{hi}} + \varepsilon, \quad j = 1, 4, \ldots, 148 \end{aligned} \tag{5.24}$$

Table 5.4 Scenario Sequence and the Mechanism Each Exercises

#	Scenario	What it exercises
1	Operate by hand, $u = 30\,\%$	no self-regulation, exact flow matching
2	Step test from balance	IPDT identification from a slope
3	First automatic loop, $K_c = 3$	proportional recovery from 8 % low
4	SIMC PI installed	model-based tuning, reset time
5	$+30\,\%$ heat load at $t = 5$	load rejection on an integrator
6	Ziegler-Nichols from $\hat{\theta} = 0.50$	identification error, not rule error
7	On-off conductivity blowdown	square-wave load from a second loop
8	Level noise $\sigma = 0.4\,\%$	derivative action on a ramping PV
9	130 % load, 40 % pressure, 12 gpm leak	makeup capacity limit, anti-windup
10	Three controllers from one model	SIMC, ZN and MPC from one fit, move suppression
11	Graded challenge, four upsets	the whole workflow under time

where z_j is the deviation j steps beyond the dead time, $v_1, \ldots, v_{12}$ are valve positions held over blocks of $(1,1,1,1,2,2,3,4,6,9,15,105)$ steps that cover $N_p = 150$ steps or $T_p = 15\,\text{min}$, $v_{i(j)}$ is the block containing step j, v_0 is the command applied now, $\lambda = 0.3$ is the move suppression, ε is one slack priced at $c_\varepsilon = 10^3$ per $\%^2$, and $L_{\text{lo}} = \min(40\,\%, L_{\text{SP}} - 2\,\%)$ and $L_{\text{hi}} = \max(85\,\%, L_{\text{SP}} + 2\,\%)$ are the level alarms kept 2 % from the setpoint. With these scales $\lambda = 1$ charges a 5 % valve move like a 1 % level error held for one minute. Only the valve range is a hard constraint. There is no terminal weight, because the last block holds the valve for 10.5 min and a plan that leaves the basin ramping pays for the ramp over that whole block. The dual active-set solver [22] that the apps share handles the resulting 13 variables and 125 inequality rows.

The period, the horizon, λ and ρ were set once, on the comparison case, by the tuning rule common to the predictive controllers of the apps. Of 75 settings (λ from 0 to 10, ρ from 0.01 to 100), 53 finish the case without a rejected plan and with at most the 150.2 % of valve travel of the SIMC PI, and among those within 10 % of the lowest IAE (13.06 % min) the largest λ and then the smallest ρ give $\lambda = 0.3$ and $\rho = 0.1$ (14.06 % min). A solution is discarded if the solver does not certify it as optimal, if its first move is not finite, or if that move leaves the 0–100 % range, and the PID, held on the applied command at every execution as in manual mode, takes the valve without a bump (0.034 % when a rejection is forced) until the next acceptable plan. The MPC reads only LT-201, the applied command, the setpoint, the two alarm limits and its own tuning, so the plant of Equation (5.11), the identified pair of Equation (5.12) and the model (5.22) stay separate, and a change to the header or to a hidden constant is model mismatch for the MPC.

5.6 Results

Eleven scenarios, listed in Table 5.4, are arranged as a laboratory sequence, and Table 5.5 gives the requirement for completing each. Section 1.6.4 describes what a scenario fixes and how its completion is judged. Reset restarts a case from $t = 0$ with the reader's own settings kept, while reloading the case restores its shipped settings, so the comparisons in Table 5.5 are made by changing one setting between restarts. The challenge is judged on the case as given, and changing a model parameter or a disturbance voids the attempt. A run exports as a CSV record at one sample every 0.02 min (1.2 s) of simulated time. The record holds the true level beside the measured one and the commanded valve beside the actual stem, so the slope fit behind Figure 5.2 can be reproduced in a spreadsheet or in Python.

Scenario 1 makes the point with arithmetic before any controller is involved. A valve at 30 % delivers 22.5 gpm against 31.875 gpm of losses, a deficit of 9.375 gpm that by (5.1) is a level ramp of –0.417 %/min. From 60 % the basin reaches the pump-suction trip at 20 % in 96.8 min and empties in 144.8 min, the ramp times plus the 0.83 min effective dead time, and no valve position other than 42.5 % stops the motion. Scenario 3 shows the same physics under proportional-plus-reset

Table 5.5 Scenarios and Completion Requirements

#	Scenario	Kind	Required to complete
1	Operate by hand	Experiment	In Manual, hold the level rate within ±0.05 %/min of zero for 1 min, change the heat load, rebalance the valve, and hold the rate within ±0.01 %/min for 0.5 min at the new load.
2	Step test	Experiment	In Manual, step the valve from balance to at least 52 % after the first minute and hold it for 3 min, then enter $\hat{K}_I$ within 25 % of K_I (0.025 to 0.0417) and $\hat{\theta}$ within 25 % of $\theta + \tau_v$ (0.62 to 1.04 min).
3	First automatic loop	Experiment	Recover from 8 % low to 59.8 % with $K_c = 3$ in Auto, then Reset, apply the SIMC rule to the identified model before $t = 0.5$ min, and recover again with $K_c \geq 9$.
4	SIMC tuning, reset time	Experiment	Raise the setpoint to 65 % in Auto after 0.5 min and run 5 min with $\tau_I \geq 6$ min, then repeat the step in runs of the reloaded case with $\tau_I = 2$ min for 8 min and 1 min for 15 min.
5	Heat-load disturbance	Experiment	Run through the +30 % heat-load step at 5 min until the level has stayed within 0.1 % of setpoint for 2 min.
6	Optimistic dead time	Experiment	Run the loaded tuning ($K_c \geq 50$, $\tau_I \leq 1.05$ min) to $t = 25$ min, then Reset, apply the ZN rule from $\hat{\theta} \approx 0.83$ min before the step at 5 min, and run to $t = 25$ min with $25 \leq K_c \leq 45$ and $\tau_D > 0$.
7	On-off blowdown	Experiment	With conductivity-switched blowdown, run past 18 min until the bleed has been shut for 5 min, then switch to continuous blowdown and run 10 min more.
8	Noisy level	Experiment	With $\sigma = 0.4$ %, run 10 min at $\tau_D = 0.4$ min and $\tau_f \leq 0.06$ min, 10 min after Reset with $\tau_D = 0$, and 10 min from the reloaded case with $\tau_f \geq 0.3$ min, then raise the noise above 0.45 %.
9	Valve saturation	Demonstration	Hold the valve at 100 % for 1 min while the leak is on, run past the repair at 18 min until the level is back above 59.5 %, then Reset with the header pressure restored to at least 95 % before $t = 0.5$ min and, from the repair at 18 min, hold the level within 2 % of setpoint, unsaturated, for 10 min.
10	Three controllers	Experiment	Run the setpoint step at 1 min and the heat-load step at 10 min to $t = 20$ min under the MPC with $\lambda \geq 0.25$, again after a Reset with $\lambda \leq 0.05$, then under the SIMC PI after a Reset with $K_c \geq 9$.
11	Challenge	Challenge	At $t = 45$ min, a worst level deviation over the whole run of at most 2 % from the 60 % setpoint, no overflow, no time saturated and less than 250 % of valve travel, with the model parameters and disturbances as given.[a]

[a] With the shipped tuning ($K_c = 3$, $\tau_I = 30$ min) the worst deviation is 5.73 %, at $t = 25$ min, and the case fails although the valve never saturates and travels only 29 %. With the suggested SIMC tuning ($K_c = 18.1$, $\tau_I = 6.6$ min) the worst deviation is 0.70 % with 66 % of travel and no saturation, inside the band with 1.30 % to spare.

control. Starting 8 % low with $K_c = 3$ and $\tau_I = 30$ min, the loop enters the ±2 % band for good after 10.58 min, first reaches 60 % after 17.2 min and accumulates an IAE of 58.37 % min by $t = 20$ min, then overshoots to 61.41 % at 34.4 min, 17.6 % of the initial error. The SIMC tuning ($K_c = 18.1$, $\tau_I = 6.6$ min) reaches 60 % after 5.0 min and settles within 3.98 min after a peak of 61.75 % at 7.1 min, with an IAE of 33.27 % min at $t = 20$ min, while resting against a valve limit for 20.9 % of those 20 min. The saturation figure counts time at either end of the travel, because a makeup valve pinned shut costs the loop as much authority as one pinned open.

Table 5.6 Measured Performance After a +30 % Heat-Load Step at $t = 5$ min, Read at $t = 20$ min. All Four Controllers Act on the Same Plant With $K_I = 0.0333\,\%/(\text{min}\,\%)$ and $\theta_{\text{eff}} = 0.83$ min.

	Default	SIMC PI $\hat{\theta} = 0.83$	ZN PID $\hat{\theta} = 0.83$	ZN PID $\hat{\theta} = 0.50$
K_c (%/%)	3	18.1	34	56.4
τ_I (min)	30	6.6	1.7	1.0
τ_D (min)	0	0	0.42	0.25
IAE (% min)	30.90	4.42	1.04	4.72
ISE (%^{2} min)	75.30	2.01	0.25	1.89
Max. deviation (%)	3.06	0.70	0.40	0.64
Valve travel (% of span)	12.23	20.47	31.68	474.7
Time saturated (%)	0	0	0	0

TRY IT IN THE APP 1 · Operate the Basin by Hand

Load the case and press Run without touching the valve. It stays at 30 % while evaporation, blowdown and drift carry away more water than it supplies. Watch the basin level trend and the Level rate tile: the level falls on a straight line at –0.417 %/min and never levels off. Only a valve at 42.5 % would stop it.

5.6.1 Load Rejection With Four Tunings

Four controllers meet the identical upset in Table 5.6, which reports the performance measures the application computes for itself. The heat load steps from 100 to 130 % at $t = 5$ min from a settled start and every figure is read at $t = 20$ min. Two of the four follow from applying a rule to a correct identification, one is the cautious default tuning, and one follows from supplying Equation (5.20) with the dead time a trend appears to show.

Reading Table 5.6 across the first three columns gives a clean and unsurprising trade. More gain buys a smaller deviation and pays for it in valve travel, with the maximum deviation falling from 3.06 to 0.40 % while the travel rises from 12.23 to 31.68 % of span. The measured peaks also confirm Equation (5.21), since the proportional offset predicted for a 12.75 % load is 0.70 % at $K_c = 18.1$ against a measured 0.70 %, and 0.375 % at $K_c = 34$ against a measured 0.40 %.

The fourth column is a different kind of result. Its IAE of 4.72 % min is indistinguishable from the SIMC value of 4.42 % min and its peak deviation is smaller, yet its valve travel is 474.7 % of span against 20.47 %, a factor of 23. The loop is not rejecting the disturbance. It is cycling. Continuing both runs to $t = 40$ min separates them completely, because the SIMC IAE grows from 4.42 to 4.65 % min and its travel from 20.47 to 20.73 % while the aggressive loop grows from 4.72 to 15.64 % min and from 474.7 to 1485 %, and begins to hit the valve limits for 7.2 % of the time. Over $20 \le t \le 30$ min the makeup valve swings between 4.7 and 100.0 % with a period of 3.75 min, against the 3.32 min that (5.19) predicts for P_u at $\theta_{\text{eff}} = 0.83$ min and the 3.39 min of the proportional-only test. The level itself only moves between 59.09 and 60.84 %, which is the reason the error measures look respectable and the reason a cycling level loop can persist unnoticed in operation.

The two Ziegler-Nichols columns separate identification error from rule error. Both were produced by the same rule and the same arithmetic. One is stable and the fastest controller in the table, and the other carries 99 % of the ultimate gain the IPDT model predicts together with a reset time cut by the same factor, so the degradation comes from the identification and not from the rule. Rules that place $\hat{\theta}$ in an additive denominator, such as (5.18), are tolerant of the same reading error by construction. Figure 5.3 draws the four columns of Table 5.6 side by side.

5.6.2 Derivative Action on a Noisy Level

Scenario 8 raises the level noise to $\sigma = 0.4$ % and switches derivative action on at $\tau_D = 0.4$ min. Read at $t = 30$ min, the valve travel is 11 791 % of span with an IAE of 2.27 % min, and the valve touches

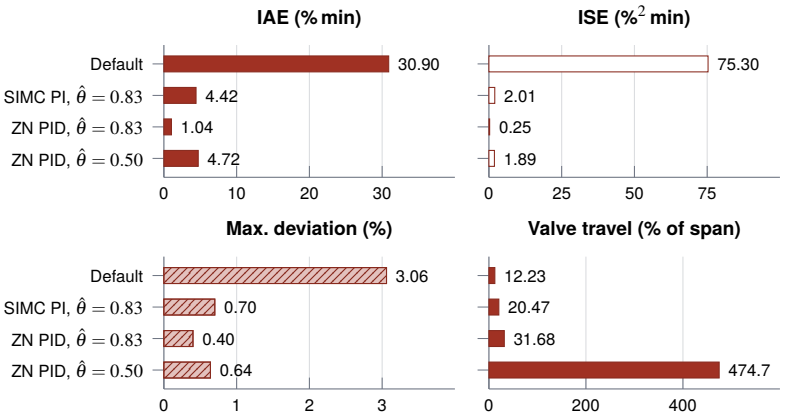


Figure 5.3 The four tunings of Table 5.6 after the +30 % heat-load step, read at $t = 20$ min: integrated absolute error and valve travel (solid bars), integrated squared error (open) and largest deviation (hatched). Notice that the tuning from the misread dead time keeps its error measures close to the SIMC values while its valve travel is 23 times larger, the signature of a cycling loop.

a limit for 1.6 % of the run. Removing derivative action entirely gives 1669 % of travel and an IAE of 2.43 % min, so derivative action obtained a 6.6 % reduction in accumulated error at a cost of 7.1 times the valve motion. Raising the derivative filter from its default 0.05 min to the maximum 0.5 min recovers most but not all of the difference, at 2982 % of travel. The mechanism is in (5.11). Derivative action anticipates, and the level of an integrating plant is already a ramp whose slope changes only when a flow changes, so the derivative term amplifies the measurement noise without adding useful anticipation. The same term adds useful phase lead on the temperature loop of Section 1.4.8, and the scenario supplies the contrast.

Pause. Try Exercise 5.3 before reading on.

Scenario 7 adds the second loop. With the blowdown valve switched by conductivity, an on-off valve passing three times the design bleed lands on the level loop as a repeating load step of 18.7 gpm at every switching, with the bleed open for 9.2 min and shut for 17.6 min once the first cycle has passed. Read at $t = 30$ min the level loop holds a sawtooth with a maximum deviation of 1.35 % and an IAE of 16.83 % min, using 79.05 % of valve travel and never saturating. The interaction is one-way and the basin absorbs it, which is what a buffer volume exists to do.

5.6.3 Makeup Capacity Limit

Scenario 9 applies a 130 % heat load, a 12 gpm basin leak and a makeup header at 40 % of nominal pressure at the same time. By (5.6) the valve can deliver at most $75\sqrt{0.40} = 47.4$ gpm, and the losses total 53.4 gpm, so the level falls at −0.267 %/min with the valve wide open and no controller action can arrest it. Read at $t = 30$ min the measured valve travel is 0.00 % with 100 % of the time saturated, an IAE of 81.68 % min and a maximum deviation of 4.80 %. A valve travel of exactly zero on a loop with a large error is the diagnostic signature of a capacity problem rather than a tuning problem, and separating the two is a routine diagnostic task in practice. Chapter 3 draws the same line between an error that feedback removes and one that an actuator bound imposes. The leak is repaired at $t = 18$ min, after which the same fully open valve delivers a surplus of 6.0 gpm and the level climbs back into the 2 % band 28.52 min after the run started. Throughout the saturated period the back-calculation term of (5.17) holds the integral at the applied 100 %, its fixed point with $\tau_t = \tau_I$, so the valve stays fully open until the measured level reaches the setpoint at 36.1 min, and the level peaks only 0.58 % above setpoint at 39.8 min. A counterfactual run with the back-calculation removed winds the integral to 324 % by $t = 30$ min, holds the valve open until 55.8 min and overshoots to

Table 5.7 Measured Comparison of SIMC PI, ZN PID and MPC Designed From One Identified Integrating Model

Case	Design	IAE (% min)	OS (%)	t_s (min)	e_{max} (%)	Travel (%)	Sat (%)	Plans
10, SP, load	SIMC PI	18.02	26.2	2.38	0.54	150.2	11.6	—
	ZN PID	16.91	39.0	2.38	0.58	227.4	16.3	—
	MPC	14.06	0.5	2.68	0.57	144.1	1.5	200
	MPC, $\lambda = 0$	12.95	6.1	2.38	0.50	2659.4	79.0	200
5, load	SIMC PI	4.42	—	—	0.70	20.5	0	—
	ZN PID	1.04	—	—	0.40	31.7	0	—
	MPC	1.58	—	—	0.56	27.1	0	200
6, $\hat{\theta} = 0.5$	ZN PID	15.64	—	—	0.96	1485.3	7.2	—
	SIMC PI	1.70	—	—	0.55	37.7	0	—
	MPC	2.04	—	—	0.58	61.5	0	400
	MPC, $\lambda = 0$	17.83	—	—	0.97	2456.4	45.5	400
3, 8 % low	SIMC PI	33.27	21.9	—	—	78.9	20.9	—
	MPC	25.09	0.0	—	—	94.5	11.5	200
8, noise	PID, $\tau_D = 0.4$	2.27	—	—	0.21	11790.9	1.6	—
	PI, $\tau_D = 0$	2.43	—	—	0.20	1669.1	0	—
	MPC	4.92	—	—	0.47	1412.4	0	300
	MPC, $\lambda = 3$	3.25	—	—	0.30	741.6	0	300
7, blowdown	SIMC PI	16.83	—	—	1.35	79.0	0	—
	MPC	6.15	—	—	1.10	104.5	0	300
9, capacity	SIMC PI	90.31	—	—	4.80	16.9	80.2	—
	MPC	87.11	—	—	4.80	16.4	74.7	450

IAE, travel and Sat are scorecard values over 20, 20, 40, 20, 30, 30 and 45 min in case order. OS is the overshoot in percent of the step (case 3: of the initial error), t_s the time from the step until the level stays within ±2 %, e_{max} the largest deviation after the load, Sat the percentage of time at either valve stop, and Plans the number of MPC decisions, all accepted. In case 6 every design comes from $\hat{\theta} = 0.50$ min.

65.56 %. A recursion that absorbed the whole clipped excess at every scan would instead unwind the integral to 14 % while saturated, release the valve at 29.5 min and approach the setpoint from below with no overshoot, reaching 59.5 % 0.8 min later than the shipped controller. With the header restored before the run, the same tuning rides out the repair with a 0.88 % peak and no saturation.

The graded challenge applies four upsets over 25 min, a +30 % heat load at $t = 4$, a 5 gpm leak at $t = 12$, a 25 % loss of header pressure at $t = 20$ and a return to normal load at $t = 25$, and is judged at $t = 45$ min on the whole run. The level must stay within ±2 % of the 60 % setpoint with no overflow, no time saturated and less than 250 % of valve travel, with the model parameters and disturbances as given. Starting from the cautious default the challenge is not met, with a worst deviation of 5.73 % at $t = 25$ min, nearly three times the band, an IAE of 134.62 % min and 29 % of valve travel. The SIMC tuning obtained from the step test meets every part of the target on the same disturbance sequence, with a worst deviation of 0.70 %, an IAE of 12.62 % min, 66 % of valve travel and no saturation, which leaves 1.30 % of margin on the band and 184 points on the travel limit.

5.6.4 Three Designs From One Identified Model

The SIMC PI, the ZN PID and the MPC all follow from $\hat{K}_I = 0.0333$ and $\hat{\theta} = 0.83$ min and are compared in Table 5.7. Every case starts from its own seed and instruments, the design is selected before the run, the PI rows keep the tuning of the case or the gains a rule installs, the MPC keeps the settings of Section 5.5.4 everywhere, and no design receives a measured disturbance. Scenario 10 raises the setpoint from 60 to 65 % at 1 min and the heat load by 30 % at 10 min, and every row is reproduced exactly after a restart from $t = 0$.

A 5 % rise of an integrating level needs a fixed volume of makeup, $5/\hat{K}_I = 150$ % min of valve above the balance, and the designs differ in when they deliver it. The SIMC PI kicks the valve by $18.1 \times 5 \approx 90$ %, holds it at 100 % from 1.02 to 3.32 min and overshoots to 66.31 %. The MPC counts the water already in transit, holds 100 % only from 1.72 to 2.00 min and starts closing at 2.02 min with the level still at 60.26 %, so the level peaks at 65.03 %, as shown in Figure 5.4. Its IAE is 14.06 against 18.02 and 16.91 % min for SIMC and ZN, with less travel and 1.5 against 11.6 and 16.3 %

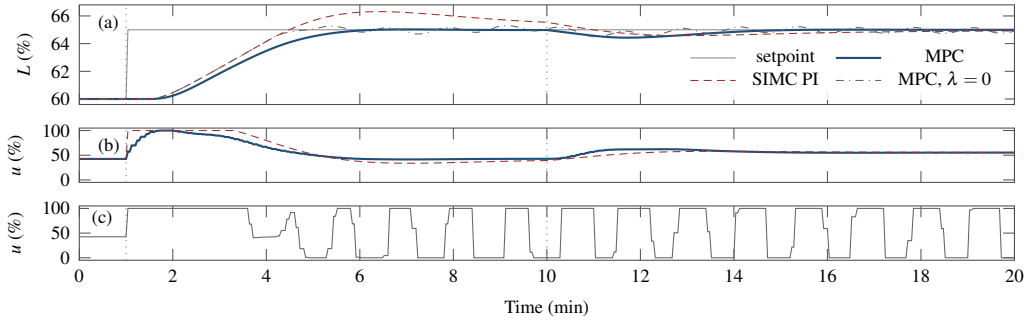


Figure 5.4 Comparison case under the MPC, the SIMC PI and the MPC with $\lambda = 0$, all three designed from $\hat{K}_I = 0.0333\,\%/(\text{min}\,\%)$ and $\hat{\theta} = 0.83\,\text{min}$. The setpoint rises from 60 to 65 % at 1 min and the heat load by 30 % at 10 min (dotted vertical lines). (a) True level under the MPC (thick solid), the SIMC PI (dashed) and the MPC with $\lambda = 0$ (dash-dotted), with the setpoint (thin solid). (b) Valve command under the MPC (thick solid) and the SIMC PI (dashed). (c) Valve command under the MPC with $\lambda = 0$. At $t = 20\,\text{min}$ the scorecard IAE is 14.06, 18.02 and 12.95 % min and the travel 144.1, 150.2 and 2659.4 %.

of the time at a limit, but the approach without overshoot enters the ±2 % band 0.30 min later. The setpoint step of scenario 4 alone gives the same order (12.49 against 18.00 and 16.08 % min).

Move suppression is what keeps the plan usable on an integrator. With $\lambda = 0$, 158 of the 200 plans put the first move on a bound and the valve strokes from one limit to the other every 1.23 min, at a limit for 79.0 % of the run with 2659 % of travel for an IAE only 8 % lower, as shown in Figure 5.4(c). The strokes come from the actuator lag folded into $\hat{\theta}$. By Equation (5.13) a valve move Δu carries the level ahead of the IPDT prediction by up to $K_I \Delta u\, \tau_v/\text{e}$, 0.41 % for a full stroke, the filter books that excursion as a change of δ within its 0.36 min time constant, and a plan with free moves answers with a full stroke back, which a hundredfold smaller ρ also prevents. Without the lag (0.01 min, with the delay raised to keep $\theta + \tau_v$) the same plan travels 298 %, $\lambda = 0.05$ alone brings the shipped plant to 165 %, and on the misread fit of scenario 6, $\lambda = 0$ turns a settled MPC into a cycle between the stops. Larger prices trade error for travel in the usual way. Scenario 10 repeated with a move suppression of 0.1, 1 and 3 gives an IAE of 12.86, 18.11 and 25.29 % min with 154.9, 110.4 and 78.6 % of travel. The model sets whether the plan is right, and the price on valve moves sets whether the valve can live with it.

Three cases are lost outright. On the unmeasured load of scenario 5 the ZN PID keeps the lower IAE and dip (1.04 against 1.58 % min, 0.40 against 0.56 %). With $K_c = 34$ and derivative action it has the valve past the new balance of 55.25 % half a minute after the load, while the move suppression holds the MPC at 46.1 % although $\hat{\delta}$ has already covered 63 % of the load. With $\hat{\theta}$ read as 0.50 min the SIMC rule returns a faster loop ($K_c = 30$) that rejects the load best of all (1.70 against 2.04 % min), while the MPC settles with its travel more than doubled against the correct fit (61.5 against 27.1 %). On the noisy level of scenario 8 the filter passes the noise into $\hat{\delta}$ (standard deviation 17.7 % of valve) and the held moves swing the valve slowly (standard deviation of one-minute averages 7.75 against 2.73 % for the PI without derivative), which the basin integrates into a root-mean-square level error of 0.21 against 0.10 % and an IAE of 4.92 against 2.43 % min.

The MPC wins where the plan can use what it knows about the integrator. From 8 % low it refills without overshoot (peak 60.00 against 61.75 %, IAE 25.09 against 33.27 % min). Under the on-off blowdown $\hat{\delta}$ swings over 27.3 % of valve, about the 24.9 % that one switching of the 18.7 gpm bleed represents, so each switching is canceled as a step of δ (IAE 6.15 against 16.83 % min for 32 % more travel). At the capacity limit the MPC has no integrator to unwind and leaves 100 % at 33.62 min with the level at 59.36 %, 2.48 min before the PI, which lowers the recovery peak from 60.58 to

60.13 %, with 335 of its 450 plans on the upper bound and at most 12 active-set iterations. On the challenge, reported for information, both designs meet the target, the MPC with an IAE of 5.63 against 12.62 % min for the suggested SIMC tuning.

A dead time read too long is tolerated far less well by the MPC than by the SIMC rule. On scenario 10 run to 40 min at $\lambda = 0.3$ the MPC settles for $\hat{\theta}$ up to 1.1 min, hunts at 1.15 and 1.2 min and cycles between the stops from 1.25 min, and at $\lambda = 3$ it cycles from 1.7 min, while the SIMC PI, whose gain falls as $1/(2\hat{K}_I\hat{\theta})$, settles up to 3 min. The valve travel over those 40 min is 1228 % with $\hat{\theta}$ of 1.2 min and 3307 % with 1.5 min, while the SIMC PI from the same readings settles. With $\hat{\theta}$ longer than the plant's delay the level answers before the model expects it and the filter books the early water as a load (after a 10 % valve step with $\hat{\theta} = 3$ min, $\hat{\delta}$ reaches 10.0 % of valve by 3.5 min), so the plan counts the same water twice. The cycle disappears when the plant's delay is lengthened to match $\hat{\theta}$ and, over the setpoint step, when the estimate is replaced by the bare measurement, and the onset lies where $\hat{\theta}$ exceeds 0.83 min by about the 0.36 min time constant of the filter. Read short, the order reverses, and at $\hat{\theta} = 0.2$ min the MPC hunts while the SIMC PI cycles at its gain bound of 60. All 3920 plans of these runs, scenario 4 and the challenge included, were optimal and the soft band never became active, at about 0.1 ms per plan in Node.js and 0.28 to 0.34 ms in Chrome, 0.3 % of one processor core at the default speed (wall-clock values on one desktop computer).

5.7 Technology in Practice

The loop of Section 5.3 reads LT-201 through a noise term and a transmitter filter and moves LCV-201 through a delay and a lag. Figure 5.6 draws it with the devices a plant basin would carry in the place of the schematic's symbols.

What Measures It

A basin level is read by a hydrostatic transmitter near the floor of the basin or by an echo transmitter mounted above the water. No range, accuracy or response time is given here, because none was found in an open source that names no maker. The noise $n(t)$ of Equation (5.7) stands for the waves that the falling return water raises on the surface, the filter τ_S for the damping that the transmitter or a stilling well adds, and the bias b for a calibration error that no controller can see.

> **TECHNOLOGY** **Sensor: A Level Transmitter Over a Splashing Basin**
>
> A hydrostatic transmitter infers level from the pressure at the bottom of the liquid column, which grows with the height and the density of the liquid, so its calibration rests on that density. An echo transmitter times a sound or radio wave from the top of the vessel to the surface and back. It is immune to density, but foam can fool it, and liquid pouring in through the space above, as the return water rains into a basin, may also cause problems. A turbulent surface is especially hard on echo sensors. A stilling well, a vertical pipe around the sensor, shields it from the splashing and the waves in the rest of the vessel [23].

What Moves It

LCV-201 of Equation (5.6) is a linear valve whose stem follows the delayed command through a lag, and whose flow falls with the square root of the header pressure. That pressure term is the plant's dependence of flow on supply pressure, and the clipped stem is a valve that cannot reach the position the controller asks for. A controller with reset action then winds up in a futile attempt to drive the valve where it cannot go, unless the final element feeds its real position back to the PID block by back-calculation [23], the role of Equation (5.17). The blowdown valve of case 7 is a second final element. A conductivity controller measures the tower water continuously and opens it only when the conductivity passes its set point [13].

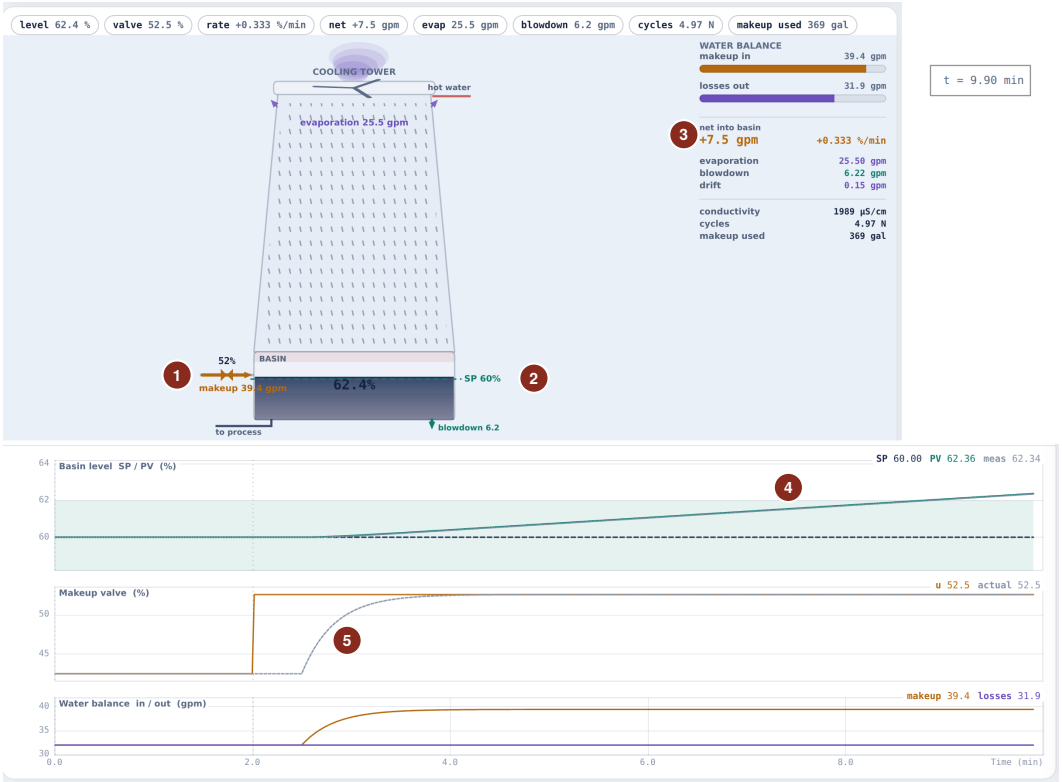


Figure 5.5 Step-test identification in progress in scenario 2 at $t = 9.9$ min, a composite of the application window: the plant view with its readouts above the three trends, and the elapsed-time readout of the header in the gray frame at the upper right. The basin starts in balance at 42.5 % valve (Table 5.1), the makeup valve is commanded 10 % open at $t = 2.00$ min, and the stem follows 0.50 min later with the 0.33 min actuator lag (Table 5.2). The level climbs on a straight ramp of 0.333 %/min to 62.36 % with no sign of settling, while makeup at 39.4 gpm exceeds the 31.9 gpm of losses by 7.5 gpm. Callouts: (1) makeup valve LCV-201, (2) the basin level against its setpoint, (3) the water balance, makeup in against losses out, (4) the level ramp on the trend and (5) the valve command step with the delayed and lagged stem.

> **TECHNOLOGY** **Final Element: The Makeup Valve and Its Failure Position**
>
> A pneumatic valve with a spring goes to a known position, open or closed, when its air fails, while an electric actuator usually stays where it was. An air-to-open valve fails closed, and the risks of the process, not habit, should decide which way it fails [23]. On a basin, a valve that fails shut lets the level fall toward the pump suction, and one that fails open fills the basin to the weir. The flow at a given stem position also depends on the supply pressure upstream, so a sagging header calls for a new valve position [23]. Water guidance asks that the fill valve cut off cleanly when the basin is full [13].

Where the Controller Runs

A distributed control system runs production at one site in power generation, refining and chemical manufacturing, and a programmable logic controller with three-mode PID can serve as the main controller of a smaller system. A tower that cools a building is more often on a building automation system, which runs heating, ventilation, air conditioning and other utility systems. Loops on these

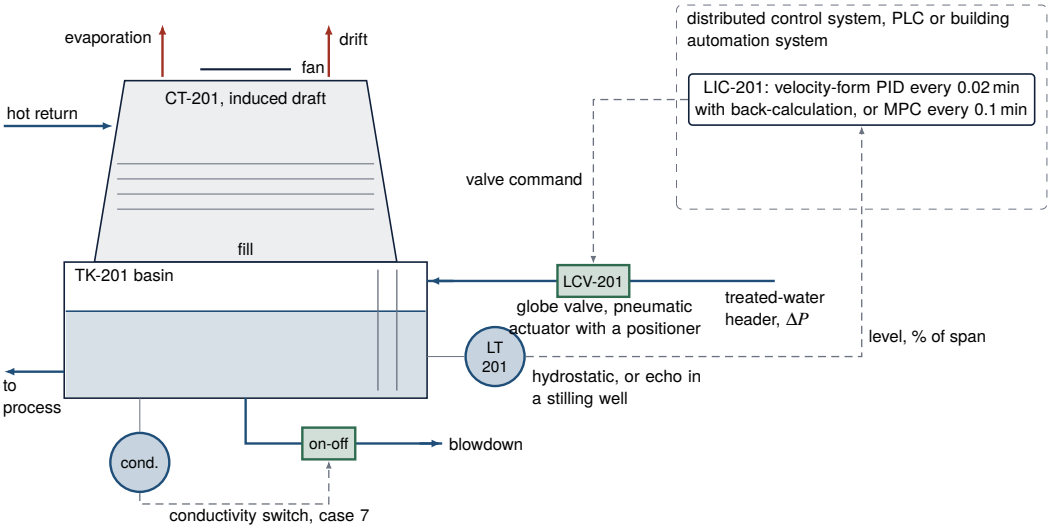


Figure 5.6 The makeup loop of the chapter with plant devices in place of the schematic's symbols. LT-201 reads the basin level, either by hydrostatic pressure or by an echo instrument in a stilling well, and LIC-201 on the control platform moves LCV-201 through its positioner. A conductivity switch opens the blowdown valve in case 7. Notice that nothing that leaves the basin is set by the level, so the makeup valve is the only device that can hold it.

platforms run with cycle times from milliseconds to minutes [24]. LIC-201 executes every 0.02 min (1.2 s) and the predictive controller decides every 0.1 min, both well inside that range. Industrial practice places model predictive control one level above the PID loops that move the valves [20], while the MPC of Section 5.5.4 moves LCV-201 directly and hands it to the PID when a plan is rejected.

> **IN PRACTICE** **Tuning a Level Loop in Service**
>
> An instrument handbook sets out the practice in a few rules. Explore the process with the controller in manual, because in automatic the response mixes the process with the controller's own action. If a step of the output makes the level ramp, the process integrates, and it may be controlled mainly by proportional action, with minimal integral action to remove offset. The lags in the loop and the measurement noise set how much gain it tolerates, since proportional action copies noise onto the output and keeps the valve stem jumping. Integral action is still needed when an outflow or the supply pressure changes [23].

What Is Modernizing It

Three developments touch this loop. The first is automation of the water side. Controllers now handle conductivity and blowdown, pH and chemical dosing, many of them with water meter inputs and alarm relays, and they hold the conductivity much closer to its set point than manual or timed blowdown [25]. The second is fieldbus, which lets the PID function block run inside the field devices, in the transmitter or in the valve positioner [23]. The third is the echo level instrument. In some industries guided-wave radar transmitters have been replacing older displacement level instruments, even at high pressure [23].

5.8 Limitations

The model is not intended for equipment selection. There is no thermal model of the tower at all, so a cold-water temperature, an approach to wet bulb or a fan power cannot be predicted, and the

Merkel and effectiveness-NTU formulations a design calculation requires [8–10] are deliberately omitted. Evaporation is one algebraic expression in the heat load with the weather frozen, which understates how much a real makeup loop moves across a day. Chemistry is one conductivity state, which puts scaling, corrosion and biological control out of reach. The valve is ideal apart from its lag and its pressure term, and stiction would add a limit cycle that looks like the one in Table 5.6 while having a different cause. Chapter 4 includes a case with valve stiction. Upsets arrive as commanded steps and timed events and never as the drifting behavior of a plant. 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. The four features that survived the simplification, the pure integrator, a transport delay in series with an actuator lag, the square-root valve characteristic and the on-off second loop, all survived because a controller has to be designed around each of them. Where the individual assumptions break is set out in Section 5.3.6.

Three elements of the scenarios are stylized on purpose. A basin leak that appears instantly at 12 gpm and is repaired 18 min later is not something a physical basin does, but it isolates a demand that exceeds the actuator and then stops exceeding it. A drift multiplier reaching 40× design is far beyond what a damaged eliminator produces, although it does separate a loss the blowdown calculation credits from one it does not. The heat load arrives as a step because a step is the disturbance whose response can be read off a trend, while a real load ramps.

The predictive controller plans on a fixed IPDT model from the learner's identification. It omits the actuator lag as a lag, which is the source of the strokes at $\lambda = 0$, the transmitter filter, the pressure term of Equation (5.6), which lowers the true gain to 0.0289 %/(min %) at the 75 % header of the challenge, the weir and the empty basin, and the on-off blowdown, which it sees only through $\hat{\delta}$. Nothing re-identifies the model on line, and a dead time read about 0.4 min too long makes it cycle.

5.9 Discussion

Basin level control carries a direct water cost. Recirculating evaporative cooling consumes water at a rate that dominates the water footprint of the plant it serves [12], and at the design point of this chapter the tower evaporates 25.5 gpm, bleeds 6.2 gpm and drifts 0.15 gpm, so the makeup meter runs at 31.9 gpm continuously. All of the makeup stream is softened, dosed and purchased. Overshooting the level costs twice, first because the surplus is eventually blown down or spilled over the weir along with the chemistry dissolved in it, and second because the treated blowdown stream is itself concentrated enough to be a disposal problem [14].

The cycles of concentration set the trade. Raising N from 5 to 7 reduces the required bleed in (5.4) from 6.225 to 4.10 gpm at the same duty, a saving of 6.7 % on total makeup, and moves the basin conductivity from 2000 to 2800 µS/cm, closer to the scaling limit the treatment program allows. The application exposes both N and C_M as adjustable parameters so that the trade can be computed rather than asserted, and the on-off blowdown mode shows what holding a target conductivity does to the neighboring level loop. A level controller tuned for a fast setpoint response is poorly suited to water-efficient operation. A loop slow enough that the basin absorbs the blowdown cycles and the load steps without moving the makeup valve reduces water and chemical use, which is the reasoning behind detuning level loops deliberately [2], and the reason Table 5.6 reports valve travel beside IAE rather than underneath it.

> **OPEN CHALLENGE** **Running More Cycles of Concentration**
>
> Many cooling towers run at two to four cycles of concentration, while six or more might be possible, and going from three to six cuts the makeup by 20 percent and the blowdown by half [13]. The limit is the water itself. As minerals concentrate, the potential for scale and corrosion rises, and treatment pushes those thresholds back without removing them. Reverse osmosis can recover much of the blowdown, but it is difficult to operate correctly and expensive to buy and run [25]. Progress would be a treatment and control scheme that holds high cycles without scale, so that less treated water is bought and less is sent to drain.

The model also supports control methods beyond the implemented loops. Beyond the linear MPC of Section 5.5.4, the three-state balance with a hard makeup limit is suited to a nonlinear MPC, which would carry the header pressure into the valve capacity of scenario 9 and coordinate makeup and blowdown as one two-input problem instead of two independent loops [26]. Moving horizon estimation or Kalman filtering on the same equations can estimate the unmeasured leak and drift disturbances from the level and conductivity measurements, so the capacity diagnosis of scenario 9 becomes a soft sensor for water loss. An economic formulation can trade makeup and treatment cost against the conductivity limit directly rather than through a fixed cycles target. The model equations transfer to the Gekko Optimization Suite in Python with the four hidden constants recast as estimated parameters [27].

5.10 Conclusion

This chapter studied a cooling tower basin as a laboratory for the class of plants that have no steady state of their own. The simulator integrates three states, the basin volume, the makeup valve stem and the circulating-water conductivity, with evaporation set by the cooling range, blowdown sized for a target cycles of concentration, drift proportional to circulation, and a makeup valve carrying a 0.50 min transport delay, a 0.33 min actuator lag and a square-root pressure characteristic. Four constants that fix those dynamics are hidden. A 10 % open-loop valve step identifies $\hat{K}_I = 0.0333\,\%/(\mathrm{min}\,\%)$ from a slope of 0.3333 %/min and, from the backward extension of that ramp, $\hat{\theta} = 0.83\,\mathrm{min}$, the $\theta + \tau_v$ that an integrating fit is bound to report.

SIMC settings from those two numbers reject a +30 % heat load with an IAE of 4.42 % min and 20.5 % of valve travel, while the same reaction-curve rule fed a dead time read as 0.50 min produces a gain of 56.4, 99 % of the ultimate gain the IPDT model predicts, paired with a reset time of 1.0 min, and the loop falls into a sustained 3.75 min cycle costing 475 % of valve travel for no reduction in IAE. Holding the gain fixed and cutting the reset time from 6.6 to 1.0 min multiplies the settling time by 5.7 and the valve travel by 4.7, which is the failure mode particular to a controller with an integrator in front of a plant that already has one.

A linear MPC planned on the same two numbers lowers the IAE of the setpoint and load case of scenario 10 from 18.02 to 14.06 % min with less valve travel, provided its move suppression is not zero, and it loses the unmeasured load step to the ZN PID (1.58 against 1.04 % min). A fuller model would add a wet-bulb-dependent evaporation model, valve stiction as a separate source of cycling, a scaling index alongside conductivity, and a variable-speed fan so that the thermal and water loops can be studied together. The basin remains a compact case for learning to identify and tune a process that never settles.

KEY POINTS

- A cooling tower basin loses water in ways that do not depend on its level, so the level ramps whenever makeup and losses differ.
- One valve step gives the integrating gain from a slope and an apparent dead time of 0.83 min, the transport delay plus the valve lag.
- Fed a dead time read as 0.50 min, the Ziegler-Nichols rule sets a gain of 56.4, and the loop cycles at 474.7 % of valve travel.
- With the gain fixed, cutting the reset time from 6.6 to 1.0 min multiplies the settling time by 5.7.
- A valve held wide open with zero travel marks a capacity limit, not a tuning problem.

Where else this applies. The pool of Volume 1, Chapter 9 integrates for the same reason.

References

[1] B. D. Tyreus and W. L. Luyben, “Tuning PI controllers for integrator/dead time processes,” *Industrial & Engineering Chemistry Research*, vol. 31, no. 11, pp. 2625–2628, 1992.

[2] E. C. Horton, M. W. Foley, and K. E. Kwok, “Performance assessment of level controllers,” *International Journal of Adaptive Control and Signal Processing*, vol. 17, no. 7-9, pp. 663–684, 2003.

[3] 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.

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

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

[6] A. Visioli, “Optimal tuning of PID controllers for integral and unstable processes,” *IEE Proceedings – Control Theory and Applications*, vol. 148, no. 2, pp. 180–184, 2001.

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

[8] D. R. Baker and H. A. Shryock, “A comprehensive approach to the analysis of cooling tower performance,” *Journal of Heat Transfer*, vol. 83, no. 3, pp. 339–349, 1961.

[9] J. C. Kloppers and D. G. Kröger, “A critical investigation into the heat and mass transfer analysis of counterflow wet-cooling towers,” *International Journal of Heat and Mass Transfer*, vol. 48, no. 3-4, pp. 765–777, 2005.

[10] N. Milosavljevic and P. Heikkilä, “A comprehensive approach to cooling tower design,” *Applied Thermal Engineering*, vol. 21, no. 9, pp. 899–915, 2001.

[11] B. A. Qureshi and S. M. Zubair, “Prediction of evaporation losses in wet cooling towers,” *Heat Transfer Engineering*, vol. 27, no. 9, pp. 86–92, 2006.

[12] J. Macknick, R. Newmark, G. Heath, and K. C. Hallett, “Operational water consumption and withdrawal factors for electricity generating technologies: a review of existing literature,” *Environmental Research Letters*, vol. 7, no. 4, p. 045802, 2012.

[13] U.S. Environmental Protection Agency, “WaterSense at work: Best management practices for commercial and institutional facilities, section 6.3: Cooling towers,” U.S. Environmental Protection Agency, WaterSense, Washington, DC, Tech. Rep., Oct. 2012.

[14] N. E. Koeman-Stein, R. J. M. Creusen, M. Zijlstra, C. K. Groot, and W. B. P. van den Broek, “Membrane distillation of industrial cooling tower blowdown water,” *Water Resources and Industry*, vol. 14, pp. 11–17, 2016.

[15] T. Liu and F. Gao, “Identification of integrating and unstable processes from relay feedback,” *Computers & Chemical Engineering*, vol. 32, no. 12, pp. 3038–3056, 2008.

[16] F. L. Lewis, “Introduction to modern control theory,” in *Applied Optimal Control and Estimation*. Englewood Cliffs, NJ: Prentice-Hall, 1992, ch. 1, section “A Brief History of Automatic Control,” reprinted by permission at lewisgroup.uta.edu/history.htm.

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

[18] J. G. Ziegler and N. B. Nichols, “Optimum settings for automatic controllers,” *Transactions of the ASME*, vol. 64, no. 8, pp. 759–765, 1942.

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

[20] 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.

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

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

[23] 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.

[24] 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.

[25] Federal Energy Management Program, "Cooling towers: Understanding key components of cooling towers and how to improve water efficiency," U.S. Department of Energy, prepared by Pacific Northwest National Laboratory, Fact sheet DOE/PNNL-SA-75820, Feb. 2011.

[26] 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.

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

5.11 Exercises

5.1 Concept check · 2 min

The basin water balance is $A_b \frac{\mathrm{d}L}{\mathrm{d}t} = q_M - q_E - q_B - q_D - q_L$. The makeup valve is left in Manual at an opening that supplies slightly less than the losses. What does the level do, and why?

(a) It falls at a constant rate until it meets a limit, because no term on the right depends on the level

(b) It settles at a lower level, because the evaporation rate falls as the basin holds less water

(c) It settles at a lower level, because the blowdown is sized in proportion to the basin inventory

(d) It falls at a rate that fades with time, because the valve passes more water as the level drops

5.2 Concept check · 2 min

After the 10 % valve step of scenario 2, the level trend visibly leaves its flat line 0.65 min after the step. The final ramp, extended backward, meets the old level 0.83 min after the step. Which value belongs in the tuning rules, and why?

(a) 0.65 min, because the first movement marks the moment the makeup starts to reach the basin

(b) 0.83 min, because the valve lag adds its full length to the delay that an integrating fit carries

(c) 0.65 min, because the ramp extension also counts the transmitter filter, which acts outside the loop

(d) 0.83 min, because the transport delay of the makeup line grows while the valve is opening

5.3 Concept check · 1 min

Which controller does the evidence of this chapter support for the noisy basin level of scenario 8?

(a) Proportional only, because the integrator in the basin removes the offset after any load change

(b) PID with a derivative time of 0.4 min, because derivative action anticipates the level ramp

(c) PI with a long reset time and no derivative, because derivative mostly turns noise into valve travel

(d) PI with a short reset time, because faster integral action removes the load offset much sooner

5.4 Theory and calculation · 5 min

The tower circulates 3000 gpm over a design cooling range of 10 °F, with an evaporation factor of 0.85, a drift fraction of 5×10^{-5} and a target of 5 cycles of concentration. The makeup valve passes 75 gpm wide open, and the basin holds 22.5 gal per percent of level.

(a) Find the evaporation, drift and blowdown flows at design load and the valve position that balances them. *Checkpoint:* 25.50, 0.15 and 6.225 gpm, balanced at 42.5 %.

(b) Find the integrating gain K_I from the valve capacity and the holdup. *Checkpoint:* 0.0333 %/min per percent of valve.

(c) Scenario 1 holds the valve at 30 %. Find the level rate, and the time for the level to fall from 60 % to the pump-suction trip at 20 %. Neglect the dead time. *Checkpoint:* −0.417 %/min and 96 min.

5.5 Theory and calculation · 7 min

Equation (5.13) gives the level after a valve step of $\Delta u = 10\,\%$ at $t = 0$, with $K_I = 0.0333\,\%/\text{min}$ per percent, $\theta = 0.50\,\text{min}$ and $\tau_v = 0.33\,\text{min}$.

(a) Show that the final ramp, extended backward, crosses the starting level at $t = \theta + \tau_v$. *Checkpoint:* 0.83 min.

(b) Find how far the level has in fact risen at that time. *Checkpoint:* 0.040 %.

(c) A trend resolves 0.01 % of level. Use the expansion $e^{-x} \simeq 1 - x + x^2/2$ for the time just after θ to estimate when the trend first shows movement. *Checkpoint:* 0.64 min after the step.

5.6 Theory and calculation · 6 min

Scenario 9 starts with the heat load at 130 %, a basin leak of 12 gpm and the makeup header at 40 % of nominal pressure. At 130 % load the evaporation, blowdown and drift total 41.44 gpm. The valve passes 75 gpm wide open at nominal pressure, and the basin holds 22.5 gal per percent.

(a) Find the largest makeup flow the valve can deliver and the total loss. *Checkpoint:* 47.4 and 53.4 gpm.

(b) Find the level rate with the valve wide open and the level at 18 min, starting from 60 %. *Checkpoint:* −0.267 %/min and 55.2 %.

(c) The leak is repaired at 18 min. Find the time at which the level is back within 2 % of the 60 % setpoint. *Checkpoint:* 28.5 min.

State what a higher controller gain would change before the repair.

5.7 App challenge · 8 min

Load scenario 2, Step test (identify the process), which starts in Manual with the makeup valve balanced at 42.5 %. Press Run, wait about one minute, press Step valve +10 %, run for at least 9 min more and export the CSV file. The scenario has no noise, so the default seed plays no part. Take the step time as the time of the last row in which `valve_cmd_pct` still reads 42.5.

(a) Read the true level, column `level_pct`, at 4 and at 8 min after the step. Find the slope of the ramp and the integrating gain. *Checkpoint:* 0.333 %/min and 0.0333 %/min per percent.

(b) Extend the line through those two rows back to 60.00 % and find the apparent dead time. *Checkpoint:* 0.83 min, within 0.02 min.

(c) Repeat part (b) with the transmitter reading, which is the column `level_measured_pct` of the same file. *Checkpoint:* 0.89 min, within 0.02 min.

(d) Under Advanced, enter 0.0333 as the Identified integrating gain and 0.83 as the Identified dead time, and set Tune from my model to SIMC PI. Read the installed tuning and the Model check tile. *Checkpoint:* $K_c = 18.1$ and $\tau_I = 6.6\,\text{min}$. The tile reads excellent, with both errors at 0 %.

5.8 App challenge · 9 min

Load scenario 5, Heat-load disturbance (+30 %). The loop is in Auto with the SIMC tuning, $K_c = 18.1$ and $\tau_I = 6.6\,\text{min}$, and the heat load steps from 100 to 130 % at 5 min. Use the default seed. In each run go a little past 20 min, pause, export the CSV file and read the row at 20.00 min.

(a) Run the scenario as shipped. Read the columns `level_iae` and `level_max_error_pct` and the valve command, `valve_cmd_pct`. *Checkpoint:* 4.42 % min, 0.70 % and 55.5 %.

(b) Set the Integral time to its top, 60 min, which switches integral action off. Press Reset and run again. Read the true level, `level_pct`, and the valve command. *Checkpoint:* 59.30 % and 55.25 %.

(c) Compute the load d in percent of valve and the offset that Equation (5.21) predicts for part (b). *Checkpoint:* 12.75 % and 0.70 %.

Accept an IAE within 0.05 % min, a level within 0.02 % and a valve within 0.1 %. State where each controller stores the extra 12.75 % of valve.

5.9 App challenge · 8 min

Load scenario 10, Three controllers from one model. The setpoint steps from 60 to 65 % at 1 min and the heat load rises by 30 % at 10 min. Use the default seed. In each run go a little past 20 min, pause, read the MV travel tile of the scorecard and export the CSV file. Read `level_iae` in the row at 20.00 min and the highest `level_pct` between 1 and 10 min.

(a) Run the scenario as shipped, with the Controller at MPC and a Move suppression of 0.3. *Checkpoint:* 14.06 % min, 65.03 % and 144 % of travel.

(b) Set the Move suppression to 0 under Advanced, press Reset and run again. *Checkpoint:* 12.95 % min, 65.31 % and about 2700 % of travel.

(c) Set the Controller to PI, press Reset and run again. *Checkpoint:* 18.02 % min, 66.31 % and 150 % of travel.

Accept an IAE within 0.05 % min, a peak within 0.02 %, and a travel within 5 % in parts (a) and (c) and within 200 % in part (b). State which design you would install and what the IAE alone hides.

5.10 Application area · 3 min

The water treatment program proposes to raise the target cycles of concentration from 5 to 7 at the design duty, where the evaporation is 25.50 gpm, the design drift is 0.15 gpm and the makeup conductivity is 400 µS/cm.

(a) Find the blowdown at each target and the saving in total makeup, in gpm, in percent and in gallons per day. *Checkpoint:* 6.225 and 4.10 gpm. The saving is 2.125 gpm, 6.7 % or 3060 gallons per day.

(b) Find the basin conductivity at each target and name what limits the choice. *Checkpoint:* 2000 and 2800 µS/cm.

5.11 Application area · 4 min

A plant historian shows a basin level that stayed within 1 % of its setpoint for a week. Over the same week the makeup valve swung across most of its range about every four minutes.

(a) Name the scorecard measure of this chapter that exposes the problem, and explain why the level error does not. *Checkpoint:* The valve travel.

(b) The chapter gives two causes of such a cycle. Name them, and give one test on the plant that separates them. *Checkpoint:* A dead time read too short, or valve stiction.

5.12 Review the agent · 5 min

An AI assistant reviews the scenario 5 run under the SIMC PI and replies as follows. “The level is still 0.05 % low fifteen minutes after the heat-load step, because the integral time of 6.6 min is long. Cut the integral time to 1.0 min and keep $K_c = 18.1$. The proportional gain sets the stability margin against the dead time, so with the gain unchanged the margin is unchanged, and the level will return about six times faster.” Find the flaw in the proposal and say how it could be checked.