INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 2

Matching Loop Structure to Disturbance Entry in a Tank Level Simulator

Pages 22–45

Edited by John D. Hedengren

How to cite this chapter. "Matching Loop Structure to Disturbance Entry in a Tank Level Simulator," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 22–45.

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

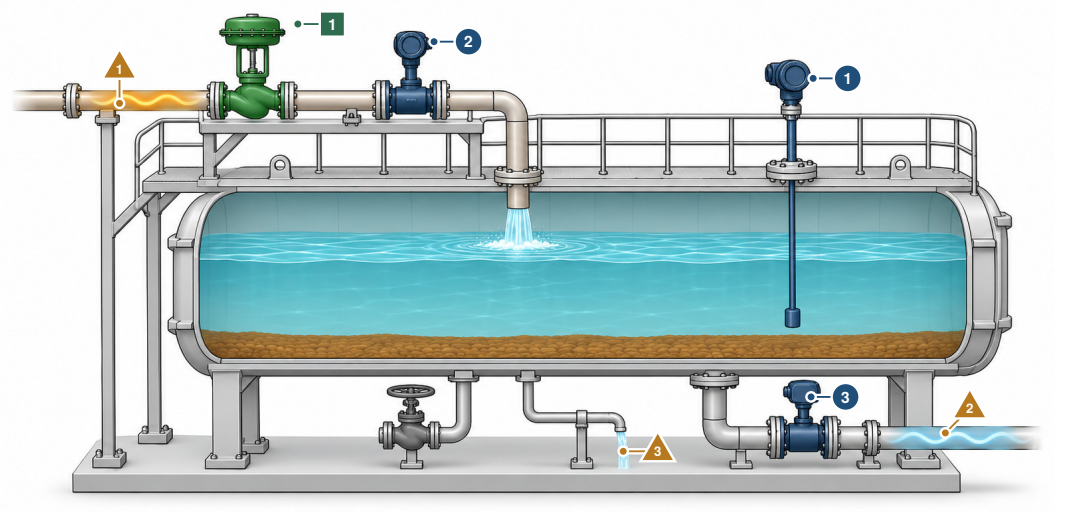
The app for this chapter
https://apmonitor.com/apps/cascade/

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

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

CHAPTER 2

Matching Loop Structure to Disturbance Entry in a Tank Level Simulator



The controller moves 1 the inlet valve on the filling pipe

The controller reads 1 the level in the tank 2 the flow arriving through the valve 3 the flow drawn out at the outlet

What disturbs it 1 the supply pressure ahead of the valve, swinging up and down

2 the outlet demand, stepping up and back 3 a small leak that grows with the level

Notice that the supply upset enters at the valve and the demand upset at the outlet, two different points on one tank, and that the model's tank is a vertical one of uniform area, as in Figure 2.1. Illustration, not to scale.

IN PLAIN TERMS **Fix the Upset Where It Comes In**

The main idea. A controller, a small computer, holds a tank's level by opening and closing the valve that fills it. Two upsets reach the tank: the push of the water at the valve rises and falls, and the draw at the bottom jumps.

The challenge. A second controller that watches the inflow cuts the first upset within seconds but misses the second. Watching the outflow fixes the second, while watching the wrong flow makes the valve jump back and forth.

Why it matters. What the controllers read decides what they can fix, and adjusting them cannot make up for a wrong choice.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the fact on the steam drum (Section 2.1), the history of cascade control (Section 2.2), the try-it segments of scenarios 6 and 4 (Sections 2.4 and 2.5), the fact on the inner gain (Section 2.6) and the key points
🕮	**Study**	45 min	the whole chapter with its three pause points and the twelve exercises at its end
⛰	**Practice**	15 min	the first page, the results (Section 2.6), Technology in Practice (Section 2.7), Limitations (Section 2.8) and the open challenge

The control question A tank held on level by an inlet valve meets two upsets that enter at different points, a supply pressure swing at the valve and a step in the outlet demand. Which control structure rejects which upset, and what does a structure chosen for the wrong one cost?

Learning objectives

- Derive the valve gain, the linearized level model and the ideal outlet feedforward gain from the mass balance, and explain why that compensator is a pure gain.
- Compute the roots and the Jury limit of the inner flow loop, and state how the limit moves with the supply pressure.
- Explain why feedforward from the inlet flow transmitter oscillates and halves the authority of the level controller.
- Measure each structure against each upset separately, and read the level error together with the valve travel.
- Judge a claim about a structure by where the upset enters, not by one aggregate error.

Prerequisites Sections 1.4.2, 1.4.3 and 1.4.5 (entry points, static feedforward, inner-loop speed), 1.6.2 (velocity form) and 1.6.1 (clocks and noise).

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

See also Volume 1, Chapter 2 for a gravity-drained tank whose gain and time constant move with the level, and Volume 1, Chapter 3 for step tests on two tanks whose order changes with a splitter valve. Volume 2, Chapter 2 for the velocity-form PID with back-calculation, and Volume 2, Chapter 5 for the integrating level that this tank becomes with the leak closed.

2.1 Introduction

Tuning a controller changes how quickly a loop rejects a disturbance. Restructuring the loop changes whether it can reject that disturbance at all. The distinction is routine in operating plants, where the fastest route to a steadier column or a steadier furnace is often a secondary measurement and a second controller rather than a better gain, and it is the distinction that control structure design addresses before any tuning rule is applied [1, 2]. A single loop measures its controlled variable, so every disturbance reaches it through the process dynamics that separate the disturbance entry point from the measurement. When those dynamics are slow, the loop responds late. Cascade and feedforward are the two classical remedies, and they are not interchangeable. Cascade inserts a faster measurement between the actuator and the slow process, so upsets that enter upstream of that measurement are caught inside the inner loop [3, 4]. Feedforward measures the disturbance itself and moves the actuator before the controlled variable has responded [5, 6].

Textbook treatments of both structures present a worked transfer function and a simulation in which the improvement is demonstrated [7, 8]. The negative half of the result is harder to convey, that a structure chosen for the wrong disturbance gives almost no benefit and can perform worse than the single loop it replaced. A student who has applied cascade to a supply-pressure disturbance may apply cascade again to a load disturbance, where the inner loop is outside the disturbance path and contributes nothing. An experiment that applies both disturbances together hides this point, because the aggregate error improves for reasons the experiment cannot separate.

Tank level control is well suited to the demonstration. The vessel supplies a capacity large enough that the controlled variable is slow, the inlet valve supplies a manipulated path with a square-root pressure dependence that makes a supply disturbance multiplicative rather than additive, and the outlet draw supplies a load disturbance that enters the same mass balance as the manipulated flow. The model rests on a liquid mass balance, the conservation law that produces first-order level dynamics on a vessel with a level-dependent outflow [7, 8], and the separation between the fast valve and the slow inventory is the time-scale separation that cascade control requires [9]. Two flow transmitters on the two sides of the vessel then make all four structures physically available on one apparatus, so the plant is unchanged when the structure changes. The experiments on this plant cover process order,

disturbance entry location, cascade structure, static feedforward, integral anti-windup and sampled measurement. They apply the sorting of disturbances by entry point that Section 1.4.2 sets out, and this tank is the example that section names for it.

> ✦ DID YOU KNOW? **The Same Two Structures in a Steam Drum**
>
> A steam drum is held by three-element control, which reads the feedwater flow, the steam flow and the level. A study of practice at solar power plants calls it the common industry-standard scheme for drum level [10]. Volume 18, Chapter 4 builds it on a drum as a cascade plus a feedforward, the pairing this chapter measures on a tank.

This chapter places four control structures on one tank in a single browser app and applies the two disturbances separately, so that the pairing between a structure and the disturbance it rejects is measured rather than asserted. The structures themselves are standard. The level loop, both feedforward paths and the level-to-flow cascade share one plant, one actuator and one sample period, so a comparison between them is free of every confounding difference except the structure. The separated disturbance protocol, and the reporting that goes with it, is the method of the chapter. Section 2.2 traces the two structures through the literature. Section 2.3 gives the model in enough detail to rebuild the simulator from this chapter alone. Section 2.4 linearizes it, derives the ideal feedforward compensator and compares the two loops at the frequency of the pressure swing. Section 2.5 states the control algorithm, the four structures, the stability limit of the inner loop and the algebraic feedback that the inlet feedforward closes. Section 2.6 reports a step test, the combined and separated scenarios, an inner-loop gain sweep and the scenarios of the app. Section 2.7 places the loop on real devices and platforms. Sections 2.8 to 2.10 give the limitations, a discussion and the conclusion, and the exercises follow in Section 2.11.

2.2 Background

How cascade control came from pneumatic practice to its first quantitative treatment [3] is told in the history box below, and the design rule still in use is that the secondary loop must be substantially faster than the primary one for the arrangement to give a benefit. Later work sharpened the rule into an explicit comparison against single-loop control across ratios of primary to secondary dynamics, and identified the region in which cascade gives no worthwhile benefit at all [4]. Chapter 3 measures that region on a fermenter, where a detuned inner loop performs worse than no cascade. Variants followed for cases where the secondary measurement responds to the manipulated variable but not to the primary disturbance, which requires the parallel rather than the series arrangement [11], and analytic tuning rules were derived for both controllers of a series cascade from the two identified models [12].

> ⧖ HISTORY **First Results on Cascade Control, 1956**
>
> In the pneumatic form of a boiler feedwater system, the level controller's air signal of 3 to 15 psi is piped to a flow controller as its setpoint, and only the flow controller moves the valve. A swing in feedwater pressure is then countered by the flow controller before the drum level has to react [13]. Much of the development of PID control took place in the power and process industries, where structures such as cascade, selector and midrange control also appeared [9]. The first published results on cascade control date from the mid-1950s, in a paper of 1956 by Franks and Worley. A later view counts cascade as a precursor of hierarchical control, because the outer loop is designed as if the closed inner loop were a fairly simple system [14].

Feedforward control developed along a separate route. Ratio control on blending and combustion services is the oldest form, where a measured wild stream sets the demand for a manipulated stream through a fixed factor [6]. The general form replaces the fixed factor by the compensator of Section 1.4.3, which in the usual case is a lead-lag with a dead time and cannot be realized exactly.

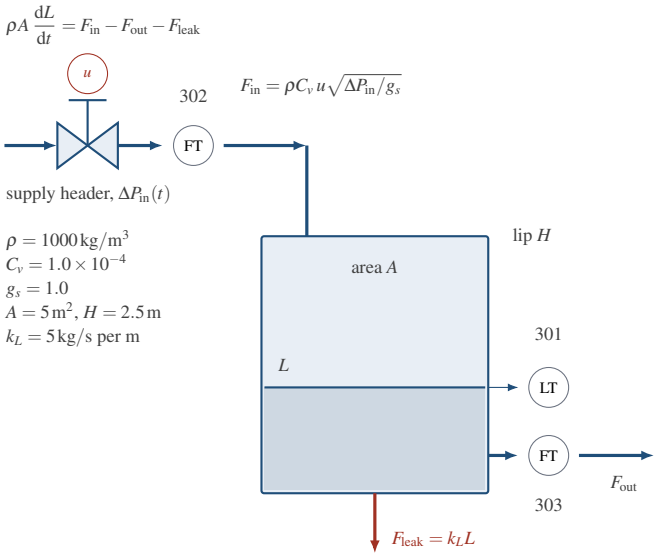


Figure 2.1 Instrumented tank with an inlet control valve. The valve lift u is the manipulated variable in all four structures, the outlet demand F_{out} and the supply pressure drop ΔP_{in} are the two upsets, and LT-301, FT-302 and FT-303 are the three transmitters the controllers are allowed to read.

The design literature is largely about what to do when that inverse is improper or unstable, how to tune the lead and the lag when the two paths carry different delays, and how much robustness the compensator gives up when the plant moves [5, 15, 16]. Advanced PID practice treats both structures as the standard extensions of the single loop [17], and process control texts introduce them together for the same reason, because the two answer complementary questions [7].

The industrial evidence is that structural choices persist and tuning choices do not. A loop is retuned many times over its life and rewired rarely, so a missing secondary measurement is a permanent limitation while a poor gain is a temporary one, as Section 1.1 also argues. Surveys of industrial practice report that the proportional-integral-derivative algorithm remains the dominant regulatory controller and that cascade and feedforward are its most common structural extensions [18]. Practitioner assessments of what plants actually run report the gap between the structures taught and the structures installed, and place the regulatory layer, not the advanced layer, at the center of delivered benefit [2]. Plantwide design methodology reaches the same conclusion from the other direction, treating the selection of controlled and secondary variables as the decision that bounds everything downstream [1].

2.3 Process Model

The simulated plant is a vertical tank of uniform cross-section fed through a control valve from a supply header and drawn down by a metered outlet flow, with a small level-proportional leak that makes the vessel self-regulating. The arrangement and the symbols used throughout are shown in Figure 2.1.

2.3.1 Variable Classification

The classification of every signal the application exposes is given in Table 2.1. One differential state carries the whole plant. The inlet flow is not a state, because the valve law is algebraic, but it becomes a secondary controlled variable when the cascade is selected, and that change of role is the only effect the structure selection has on the process side of the simulation.

Table 2.1 Variable Classification

Symbol	Role	Unit	Range
L	state, primary CV	m	0 to H
L^{SP}	primary setpoint	m	0–2.5
u	MV, valve lift	%	0–100
F_{in}	secondary CV, cascade	kg/s	0–46.9
$F_{\mathrm{in}}^{\mathrm{SP}}$	secondary setpoint	kg/s	0–40
F_{out}	DV, load	kg/s	0–20
$\overline{\Delta P}_{\mathrm{in}}$	DV, mean drop	bar	2–22
a	DV, swing amplitude	bar	0–10
k_L	DV, leak coefficient	kg/s/m	0–10
F_{leak}	internal flow	kg/s	$k_L L$

2.3.2 Mass Balance, Valve Law and Leak

The single state equation is the liquid mass balance on the tank, written for constant density and uniform cross-section, as shown in Equation (2.1).

$$\rho A \frac{\mathrm{d}L}{\mathrm{d}t} = F_{\mathrm{in}} - F_{\mathrm{out}} - F_{\mathrm{leak}} - F_{\mathrm{spill}}, \qquad 0 \le L \le H \tag{2.1}$$

where ρ is the liquid density, A is the tank cross-sectional area, L is the level, F_{in} is the inlet mass flow through the control valve, F_{out} is the metered outlet demand and F_{leak} is the leak. Each Euler substep floors the level at zero and caps it at the lip height H. Any excess inventory leaves as spill: if an unconstrained step would give $L^{*} > H$, the model records $F_{\mathrm{spill}} = \rho A (L^{*} - H)/h$ over substep h, accumulates the spilled mass, and sets $L = H$. All reported operating comparisons remain below the lip, so spill is zero in those runs.

The valve delivers the flow of Equation (2.2), which is the standard incompressible sizing relation written with lift in percent.

$$F_{\mathrm{in}} = \rho\, C_v\, u \sqrt{\Delta P_{\mathrm{in}} / g_s} \tag{2.2}$$

where C_v is the valve coefficient, u is the lift in percent of travel, ΔP_{in} is the pressure drop across the valve in bar and g_s is the specific gravity of the liquid. The constant C_v carries units of m^3/s per percent of lift per $\sqrt{\mathrm{bar}}$, so the product is a mass flow in kg/s. Flow is linear in the lift, which means the valve has no installed characteristic of its own, and it is proportional to the square root of the drop, which is the only nonlinearity in the plant. The inlet flow gain of Equation (2.3) is the quantity every structure in this chapter depends on.

$$K_v(\Delta P_{\mathrm{in}}) = \frac{\partial F_{\mathrm{in}}}{\partial u} = \rho\, C_v \sqrt{\Delta P_{\mathrm{in}} / g_s} \tag{2.3}$$

At the design drop of 12 bar the gain is 0.3464 kg/s per percent, at 2 bar it is 0.1414 kg/s per percent and at 22 bar it is 0.4690 kg/s per percent, a factor of 3.32 across the range the supply sinusoid sweeps.

The leak is taken proportional to level, as shown in Equation (2.4).

$$F_{\mathrm{leak}} = k_L\, L \tag{2.4}$$

where k_L is the leak coefficient. This term is what makes the tank self-regulating instead of integrating. Holding the valve still and waiting gives the steady level of Equation (2.5), and holding the level still

requires the valve position of Equation (2.6).

$$L_\infty = \frac{K_v(\Delta P_{\text{in}})\,u - F_{\text{out}}}{k_L} \tag{2.5}$$

$$u_\infty = \frac{F_{\text{out}} + k_L L}{K_v(\Delta P_{\text{in}})} \tag{2.6}$$

The right-hand side of Equation (2.6) is what the positional control law carries as its bias. At the design point of $L = 1.0\,\text{m}$, $F_{\text{out}} = 2.0\,\text{kg/s}$ and $\Delta P_{\text{in}} = 12\,\text{bar}$ it evaluates to 20.21 %, passing 7.0 kg/s of which 2.0 kg/s leaves through the outlet and 5.0 kg/s through the leak.

2.3.3 Disturbance Inputs

The outlet demand is a metered draw set directly by the operator, so it enters Equation (2.1) additively and is completely independent of the level and of the valve. The disturbance protocol steps it from 2 kg/s to 12 kg/s and back, in steps of exactly $\pm 10\,\text{kg/s}$.

The supply pressure drop is the second upset and it enters multiplicatively through Equation (2.2), as shown in Equation (2.7).

$$\Delta P_{\text{in}}(t) = \max\bigl(0,\ \overline{\Delta P}_{\text{in}} + a\,\sin(t/10)\bigr) \tag{2.7}$$

where $\overline{\Delta P}_{\text{in}}$ is the operator-set mean, a is the amplitude and t is the simulated time in seconds. The default $\overline{\Delta P}_{\text{in}} = 12\,\text{bar}$ and $a = 10\,\text{bar}$ sweep the drop between 2 bar and 22 bar with a period of $2\pi \times 10 = 62.8\,\text{sec}$, and setting the amplitude to zero removes the swing for step tests. The sinusoid runs on the simulated clock and its phase resets with the tank, so a run is repeatable at any speed setting.

The square root has two consequences that precede any controller. The swing is not symmetric in flow, because at a fixed lift of 20.21 % the inlet flow ranges from 2.86 kg/s to 9.48 kg/s rather than symmetrically about 7.0 kg/s. The mean is also displaced, because the average of $\sqrt{\Delta P_{\text{in}}}$ over one cycle is 3.276 against $\sqrt{12} = 3.464$, so a zero-mean pressure swing lowers the mean delivered flow by 5.4 % and the steady lift needed to hold 1.0 m rises from 20.21 % to 21.37 %. Integral action must supply that 1.16 % of lift even though the disturbance itself has zero mean.

2.3.4 Instrumentation and Numerics

Three transmitters are simulated and the controllers are permitted to read nothing else. Each adds an independent zero-mean Gaussian deviate drawn by the Box-Muller transform and each reading is floored at zero, as shown in Equation (2.8).

$$\begin{aligned} L^m &= \max(0,\ L + \sigma_L\,\varepsilon_1) \\ F_{\text{in}}^m &= \max(0,\ F_{\text{in}} + \sigma_F\,\varepsilon_2) \\ F_{\text{out}}^m &= \max(0,\ F_{\text{out}} + \sigma_F\,\varepsilon_3) \end{aligned} \tag{2.8}$$

where $\sigma_L = 0.004\,\text{m}$ is the level transmitter noise, $\sigma_F = 0.05\,\text{kg/s}$ is the noise on both flow transmitters and ε_i are independent standard normal deviates. There is no transmitter lag, no span error, no bias and no filtering. The app draws its noise from the browser's unseeded generator, as the standalone apps of Section 1.6.1 do, so repeated visits differ. The reference experiments below inject three declared random streams into the shipped model for reproducibility.

In manual mode the operator sets the valve lift directly, and the app reports the flow that lift passes at the current mean pressure and the steady level it produces from Equation (2.5). Stepping the lift with the pressure swing at zero amplitude is the step test of Section 2.6.2, from which the tuning is identified.

The plant is advanced by explicit Euler in substeps of at most 0.05 sec, and both controllers execute once per 1.0 sec of simulated time, which is twenty substeps. Simulated time advances at

Table 2.2 Simulation Parameters and Default Tuning

Symbol	Meaning	Value	Group
ρ	liquid density	1000 kg/m^3	fixed
C_v	valve coefficient	1.0×10^{-4}	model
A	tank area	5.0 m^2	model
g_s	specific gravity	1.0	model
H	lip height	2.5 m	model
σ_L	LT-301 noise	0.004 m	instrument
σ_F	FT-302, FT-303 noise	0.05 kg/s	instrument
k_L	leak coefficient	5.0 kg/s per m	disturbance
F_{out}	outlet demand	2.0 kg/s	disturbance
$\overline{\Delta P}_{\text{in}}$	mean supply drop	12 bar	disturbance
a	pressure swing amplitude	10 bar	disturbance
Δt	controller sample period	1.0 sec	fixed
h	integration substep	0.05 sec	fixed
K_{c1}	level gain, direct to valve	20 %/m	tuning
K_{I1}	level reset, direct to valve	0.4	tuning
K_{c1}	level gain, cascade	6.25	tuning
K_{I1}	level reset, cascade	0.125	tuning
K_{D1}	level derivative gain	0	tuning
u_{bias}	level loop bias, direct	20.2 %	tuning
u_{bias}	level loop bias, cascade	7.00 kg/s	tuning
K_{ff}	static feedforward gain	2.887	tuning
K_{c2}	flow loop gain	1.0	tuning
K_{I2}	flow loop reset	1.0	tuning
u_{bias2}	flow loop bias	20.2 %	tuning

$1\times$ to $10\times$ wall time, and as in the framework apps the speed changes the pace and not the result. Within one control execution the order is fixed, the three transmitters are read first, then the selected structure computes a lift, then the lift is applied, then the sample is recorded. Because the transmitters are read before the new lift is applied, the measured inlet flow at sample k reflects the lift commanded at sample $k-1$, which puts one sample of delay in the inner cascade loop and matters for its stability limit.

2.3.5 Parameters

Every parameter of the simulation is listed in Table 2.2 with its default value. No parameter is hidden, and only the liquid density, the controller sample period and the integration substep are fixed in code.

2.3.6 Assumptions

The plant is deliberately simple so that the structural comparison is not contaminated by effects that would need their own explanation. Each assumption below is stated with the condition under which it stops being reasonable.

The tank contents are perfectly mixed and the cross-section is uniform, so one lumped level describes the inventory. This fails for a horizontal drum or a dished head, where the area varies with level and the time constant varies with it.

The valve is algebraic. There is no stroking time, no stiction, no deadband and no positioner dynamics, so a commanded lift appears as a flow within the same sample. The inner cascade loop is the most sensitive to this assumption. A real valve with a stroking time of two or three seconds would put a lag inside the inner loop comparable with the 1 sec sample period, which lowers the inner gain that can be used, and with it the pressure rejection the cascade delivers.

The supply pressure drop is exogenous. Opening the valve does not pull the header down, which fails when the tank inlet takes a large fraction of header capacity and the disturbance becomes partly a consequence of the manipulated flow.

The outlet is a perfectly metered draw independent of level, which is the behavior of a downstream pump on flow control. A gravity outlet would make F_{out} a function of L and would add self-regulation the level loop does not currently have.

The leak is linear in level rather than proportional to $\sqrt{L}$ as an orifice would be. Linearity keeps the plant exactly first order at every operating level, which is convenient for the structural argument but means the gain and the time constant do not move with the level as they do on a gravity-drained vessel. Volume 1, Chapter 2 treats such a vessel, whose gain and time constant both follow the square root of the level.

The transmitters are instantaneous and their noise is white. Real differential-pressure level transmitters carry several seconds of lag and real flow measurements are noisier at the bottom of their span, both of which limit how hard the inner loop and the feedforward gain can be pushed.

There is no transport delay anywhere in the plant. This is the assumption that most favors feedforward, because the ideal compensator for this plant is a constant and remains realizable. A delay in either the manipulated path or the disturbance path would make the ideal compensator improper and would force the approximations that occupy the feedforward design literature [15].

Explicit Euler at 0.05 sec is accurate here because the substep is four orders of magnitude below the 1000 sec level time constant and the only other dynamics in the plant are algebraic.

2.4 Dynamic Analysis

Three models must be kept apart in what follows. The simulation model is Equations (2.1) to (2.7) and is what the application integrates. The control model is the linearization below and is what the tuning and the structural argument are built on. The controller's internal model is nothing at all, because both controllers are fixed-gain proportional-integral algorithms that carry no plant description.

2.4.1 Linearization About the Design Point

Perturbing Equations (2.1) to (2.4) about $L_0 = 1.0\,\text{m}$, $u_0 = 20.21\,\%$, $\Delta P_0 = 12\,\text{bar}$ and $F_{\text{out},0} = 2.0\,\text{kg/s}$ gives the first-order form of Equation (2.9).

$$\tau \frac{\mathrm{d}L'}{\mathrm{d}t} + L' = K_p\, u' + K_d\, F'_{\text{out}} + K_q\, \Delta P'_{\text{in}} \tag{2.9}$$

where the primes denote deviations from the design point and Equation (2.10) defines the four constants.

$$\begin{aligned} \tau &= \rho A / k_L, \quad & K_p &= K_v / k_L, \\ K_d &= -1/k_L, \quad & K_q &= F_{\text{in},0} / (2\,\Delta P_0\, k_L) \end{aligned} \tag{2.10}$$

At the design point these evaluate to $\tau = 1000\,\text{sec}$, $K_p = 0.0693\,\text{m}$ per percent, $K_d = -0.200\,\text{m}$ per kg/s and $K_q = 0.0583\,\text{m/bar}$. The three transfer functions from Equation (2.9) share the same single pole, which is the fact the rest of this section exploits.

2.4.2 Ideal Feedforward Compensator

The ideal feedforward compensator for a measured disturbance is the negative ratio of the disturbance transfer function to the manipulated transfer function, whose general form with a lead-lag and a delay is Equation (1.4). For the outlet demand that ratio is given by Equation (2.11).

$$G_{\text{ff}}(s) = -\frac{G_d(s)}{G_p(s)} = -\frac{K_d/(\tau s + 1)}{K_p/(\tau s + 1)} = \frac{1}{K_v(\Delta P_{\text{in}})} \tag{2.11}$$

The pole cancels exactly, so the compensator is a pure gain with no lead, no lag and no delay. This is not a modeling approximation. It is a structural property of the plant, because the outlet demand and

the manipulated inlet flow enter the same integrator at the same point, so the dynamics between the disturbance and the level are identical to the dynamics between the valve and the level, and inverting the ratio leaves the algebraic valve gain behind. At 12 bar Equation (2.11) evaluates to 2.8868 % per kg/s, which is the default value of 2.887. At 4 bar the exact value is 5.000 and at 22 bar it is 2.132, and the application reports the exact value for the current mean pressure.

🕮 **Pause.** Try Exercise 2.2 before reading on.

The same construction applied to the pressure upset gives $K_q/K_p = u_0/(2\Delta P_0) = 0.842\,\%/\text{bar}$, also a constant. A pressure feedforward is realizable in principle, but it would be exact only to first order because the true requirement is $u \propto 1/\sqrt{\Delta P_{\text{in}}}$, and it would need a pressure transmitter and an explicit model of the valve. The application does not provide that path. The cascade reaches the same upset without any valve model at all, which is the practical reason cascade is preferred for supply-side disturbances.

2.4.3 Sensitivity at the Disturbance Frequency

The pressure swing has a single frequency, $\omega_d = 0.1\,\text{rad/s}$, and the two loops meet it with very different authority. The level loop sees the process through the 1000 sec lag, so at ω_d the plant contributes $|G_p(j\omega_d)| = K_p/\sqrt{1+(\omega_d\tau)^2} = 6.93\times10^{-4}$ m per percent, a factor of 100 below its static gain. The velocity-form proportional-integral controller contributes $|C_1| = 20.6\,\%/\text{m}$ at the same frequency, so with the one-sample measurement delay the loop gain is $|L_1(j\omega_d)| = 0.0143$ and the sensitivity is $|S_1(j\omega_d)| = 1.005$. The level loop does not attenuate the pressure swing. It amplifies it by half a percent.

The flow loop sees an algebraic plant. Its loop transfer function is $C_2(z)K_v z^{-1}$, with the sample delay identified above, so at ω_d the loop gain is $|L_2| = 3.50$ and the sensitivity is $|S_2| = 0.271$. The inner loop removes 72.9 % of the flow upset at the frequency at which the outer loop removes none of it. That single pair of numbers is the whole argument for cascade on this plant, and it follows from the position of the measurement rather than from any tuning constant.

> **TRY IT IN THE APP 6 · Cascade, Pressure Swing**
>
> Press PID auto, choose Cascade with its outlet trim off, turn the sin(t/10) swing on and leave the demand alone for at least 200 s. Watch the valve lift: it sweeps widely every cycle while the level hardly moves. In the chapter's run the lift spans 14.80 to 42.17 % and the level ripples only 0.0047 m peak to peak, against 0.0214 m under feedback.

The ratio of speeds is far beyond the usual design guidance. Cascade is normally recommended when the inner loop is three to ten times faster than the outer one [4], and here the inner loop settles within a few samples against a 1000 sec outer lag. Section 1.4.5 states the rule as three to five times faster and measures on its heater that this range keeps most of the benefit against an upset inside the inner loop and costs a little against one outside it, while a factor of ten costs nothing measurable there. The tank exceeds both statements of the rule by a wide margin.

2.5 Controller Design

2.5.1 Control Algorithm

Both controllers are the same velocity-form proportional-integral-derivative algorithm and differ only in their gains, their limits and what they measure. The increment applied each sample is shown in Equation (2.12).

$$\begin{aligned} e_k &= \mathrm{SP} - \mathrm{PV}_k \\ s_k &= s_{k-1} + \Delta P_k + \Delta D_k \\ c_k &= s_k + I_{k-1} + K_I e_k \Delta t \\ u_k &= \mathrm{clamp}\big(c_k,\, u_{\min}, u_{\max}\big) \\ I_k &= I_{k-1} + K_I e_k \Delta t + (u_k - c_k) \end{aligned} \tag{2.12}$$

where e_k is the error, s_k is the running accumulator, I_k is the integral term and the increments ΔP_k and ΔD_k are selected by the algorithm variant of Equation (2.13), written with the second difference $\delta^2 x_k = x_k - 2x_{k-1} + x_{k-2}$.

$$\begin{aligned} &\text{A:} \quad \Delta P_k = K_c(e_k - e_{k-1}), && \Delta D_k = \tfrac{K_D}{\Delta t}\,\delta^2 e_k \\ &\text{B:} \quad \Delta P_k = K_c(e_k - e_{k-1}), && \Delta D_k = -\tfrac{K_D}{\Delta t}\,\delta^2 \mathrm{PV}_k \\ &\text{C:} \quad \Delta P_k = -K_c(\mathrm{PV}_k - \mathrm{PV}_{k-1}), && \Delta D_k = -\tfrac{K_D}{\Delta t}\,\delta^2 \mathrm{PV}_k \end{aligned} \tag{2.13}$$

where K_c is the proportional gain, $K_I = K_c/\tau_I$ is the integral gain in output units per unit of error per second and $K_D = K_c\tau_D$ is the derivative gain. Type B is the default on the level loop and the only variant available on the flow loop. Opening the valve raises both the inlet flow and the level, so the process gain is positive in both loops and both controllers are reverse acting, with the error computed as setpoint minus measurement and positive gains. With $K_D = 0$ and a setpoint that does not move, all three variants produce the same increment, so none of the results in this chapter depends on which one is selected.

Anti-windup is part of the recursion rather than a stage added after it. The accumulator s_k is never clipped, and the integral term I_k absorbs the whole of the clipping through the back-calculation term $(u_k - c_k)$, which is applied only when $K_I \neq 0$. This is the arrangement of Section 1.6.2, with s_k in the role of S_k there and c_k in that of v_k, except that c_k already holds the current integral increment and the whole clipping is absorbed at once, as with a tracking time constant equal to the sample period. Substituting the fourth line of Equation (2.12) into the fifth gives $c_k = u_{k-1} + \Delta P_k + K_I e_k \Delta t + \Delta D_k$, so with integral action the recursion is algebraically identical to clipping the running sum and none of the results below is affected by the distinction. It matters only at $K_I = 0$, for the reason Section 1.6.2 gives. An output parked at a limit stores only the integral contribution consistent with the bounded command, and can leave the limit when the net increment points into range. Transfer into automatic

and every edit of a gain, a setpoint or a bias re-seed the recursion at the positional law $u = u_{\text{bias}} + K_c e$ with the measurement history filled from the current reading. This prevents a derivative kick, but does not guarantee continuity of valve position: the positional bias and error, plus the next integral increment, may differ from the previous command. It makes a proportional-only setting follow that positional law.

The default level loop tuning is $K_{c1} = 20\,\%/\text{m}$ and $K_{I1} = 0.4$, which is $\tau_{I1} = 50\,\text{sec}$. Internal model control on the identified first-order plant would give $K_c = \tau/(K_p(\tau_c + \theta))$ with $\tau_I = \tau$, so $\tau_c = 720\,\text{sec}$ reproduces the default gain but calls for $\tau_I = 1000\,\text{sec}$ [19], and the SIMC modification, which caps the reset at $4(\tau_c + \theta)$ to improve load rejection on lag-dominant plants, still calls for the full 1000 sec at this τ_c [20]. The default reset is twenty times faster than that, which is a deliberate departure. The tank is slow enough that the internal model control reset would need roughly 3000 sec to clear a load step, and the loop tolerates the aggressive setting because the actuator has no lag of its own to destabilize it.

2.5.2 Structure Definitions

The four structures are shown in Figure 2.2, and the calculations that produce the valve signal are collected in Equations (2.14) to (2.17). In all four the manipulated variable is the same valve lift, the sample period is the same 1 sec and the plant is untouched.

Structure 1 is the plain level loop of Equation (2.14), whose output is the valve lift and whose limits are the valve limits.

$$u_k = \text{PI}_1(L_k^m), \qquad 0 \le u_k \le 100 \tag{2.14}$$

Structure 2 subtracts a trim proportional to the deviation of the measured inlet flow from its design value, as shown in Equation (2.15).

$$u_k = \text{clamp}\left(\text{PI}_1(L_k^m) - K_{\text{ff}}\left(F_{\text{in},k}^m - F_{\text{in},0}\right), 0, 100\right) \tag{2.15}$$

where $F_{\text{in},0} = 2.0\,\text{kg/s} + k_L L^{\text{SP}}$, which is 7.0 kg/s at the default setpoint. The sign is negative because a rise in inlet flow calls for the valve to close.

Structure 3 adds a trim proportional to the deviation of the measured outlet demand from its design value, as shown in Equation (2.16).

$$u_k = \text{clamp}\left(v_k + K_{\text{ff}}\left(F_{\text{out},k}^m - F_{\text{out},0}\right), 0, 100\right) \tag{2.16}$$

where $v_k = \text{PI}_1(L_k^m)$ is the level controller output and $F_{\text{out},0} = 2.0\,\text{kg/s}$. The reference is the fixed 2.0 kg/s of the design point rather than the current demand, so a demand already displaced when the structure is selected leaves a standing trim that the level loop absorbs through its bias.

Structure 4 replaces the direct connection with a cascade, as shown in Equation (2.17).

$$\begin{aligned} F_{\text{in},k}^{\text{SP}} &= \text{clamp}\left(v_k + r\left(F_{\text{out},k}^m - F_{\text{out},0}\right), 0, 40\right) \\ u_k &= \text{PI}_2\left(F_{\text{in},k}^m\right), \qquad 0 \le u_k \le 100 \end{aligned} \tag{2.17}$$

where $r = K_{\text{ff}}/K_{\text{ff}}^0$, $K_{\text{ff}}^0 = 1/K_v(12\,\text{bar}) = 2.8868$ and the trim term is present only when the outlet feedforward option is enabled. Dividing by K_{ff}^0 converts the gain from percent of lift per kg/s into kg/s of flow setpoint per kg/s of demand, where complete feedforward is unity, so the same number in the same field means full compensation in either place. Selecting the cascade also changes the units of the level loop's output, so the application switches K_{c1} from 20 %/m to 6.25 kg/s per m and the bias from 20.2 % to 7.00 kg/s. The two settings are close to the same loop gain expressed in different units, because $20 \times 0.3464 = 6.93$ kg/s per m against the 6.25 the cascade uses.

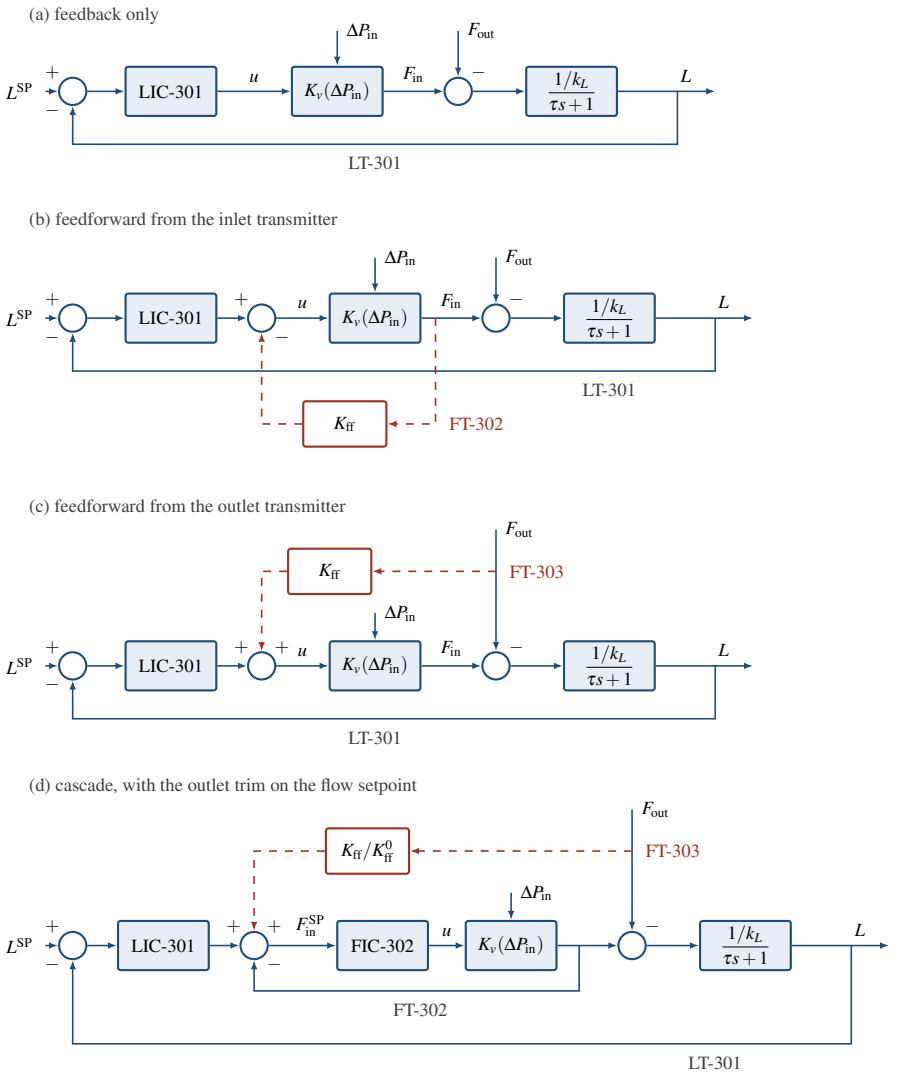


Figure 2.2 The four structures available on the same tank. Solid paths are always active and dashed paths carry the feedforward trim. The supply pressure enters at the valve gain block K_v and the outlet demand enters at the summing junction after it, so (b) taps a signal on the manipulated path while (c) and (d) tap the disturbance.

2.5.3 Inner-Loop Stability Limit

The inner loop is unusual in having no plant dynamics at all, only the one sample of measurement delay identified above. Section 1.4.5 closes an inner loop around a valve lag, with the time constant of Equation (1.7). Here that lag is absent, and the limit comes from the sample delay alone. Writing $F_k^m = K_v u_{k-1}$ and substituting the velocity form gives the characteristic polynomial of Equation (2.18).

$$z^2 + \left[\left(K_{c2} + K_{I2}\Delta t\right)K_v - 1\right]z - K_{c2}K_v = 0 \tag{2.18}$$

The Jury conditions for both roots inside the unit circle reduce to the pair of inequalities in Equation (2.19).

$$K_{c2}K_v < 1 \qquad \text{and} \qquad \left(2K_{c2} + K_{I2}\Delta t\right)K_v < 2 \tag{2.19}$$

For the default $K_{c2} = K_{I2} = 1.0$ at 12 bar the roots are 0.762 and -0.455, so the loop settles in a few samples with a mild alternating component. Holding $K_{c2} = K_{I2} = K$ and $\Delta t = 1$ sec, the second inequality binds and the limit is $K < 2/(3K_v)$, which is 1.925 at 12 bar. The limit moves with the supply pressure, falling to 1.421 at 22 bar and rising to 4.714 at 2 bar, so the sinusoid carries the inner loop in and out of the stable region during every cycle whenever K lies between those two values.

2.5.4 Algebraic Feedback in the Inlet Feedforward

Structure 2 degrades performance, and the mechanism is more specific than the observation that the inlet flow is not the disturbance. Substituting $F^m_{\text{in},k} = K_v u_{k-1}$ into Equation (2.15) shows that the trim closes an algebraic feedback path from the valve to itself, as shown in Equation (2.20).

$$u_k = v_k + K_{\text{ff}} F_{\text{in},0} - K_{\text{ff}} K_v u_{k-1} \tag{2.20}$$

where v_k is the level controller's output. The pole of that recursion sits at $z = -K_{\text{ff}} K_v$. Choosing K_{ff} by Equation (2.11), which is the correct value for the disturbance path, places the pole at exactly $z = -1$ at the design pressure. The default configuration therefore sits precisely on the boundary of a period-two oscillation, and the pressure sinusoid pushes $|K_{\text{ff}} K_v|$ above unity for the half of every cycle in which the drop exceeds 12 bar. The static gain of the same recursion is $1/(1 + K_{\text{ff}} K_v) = 0.5$, so the trim also halves the authority the level controller has over the valve, which is why this structure is worse than plain feedback on a load upset and not merely no better. Section 1.4.7 lists this case among the ways a structure applied to the wrong path does harm.

> **TRY IT IN THE APP** **4 · The Wrong Feedforward**
>
> Press PID auto, choose FF inlet and leave the sin(t/10) swing off. Press Step +10, later Step −10, and run 2 min past the last step. Watch the valve trend move up and down on almost every sample. The chapter reports reversals on 98.0 % of increments with the swing off and a worst level error of 0.3510 m, against 0.2777 m for plain feedback.

2.6 Results

2.6.1 Protocol

Reference trajectories use the current shipped model and controller functions, with twenty Euler steps of 0.05 sec per 1 sec control execution. The ordering reproduces the application: read all transmitters, calculate control, apply the lift, record, then integrate. Selection of automatic control first seeds both recursions from their positional biases. Model time is t. The application labels its first pre-integration sample as elapsed second 1, so exported elapsed time is $t + 1$. Errors use true level, IAE sums absolute error times the sample period, travel sums successive lift changes, and saturation is the fraction of commands at either lift bound.

Each comparison runs 900 sec from the design point, with default noise and LCG seeds 1, 2 and 3. The recurrence is $x_{n+1} = (1664525 x_n + 1013904223) \bmod 2^{32}$, and $x_n / 2^{32}$ supplies the model's Box–Muller draws. Each run draws once for the automatic-control seed and then once per transmitter per execution. Means and half-ranges describe these three runs, not confidence intervals. A noise-free cascade-with-feedforward run is also operated in the built browser application and checked by CSV export. All 900 samples agree to export rounding: 0.00001 m for true level and 0.0001 for lift, flow and pressure.

The steps raise demand from 2 to 12 kg/s and return it to 2 kg/s. Declaring the half-sample phase matters when feedforward nearly cancels the load.

2.6.2 Step Test and Valve Gain

With the swing off, a manual step of the lift from 20.2 % to 30 % identifies $K_p = 0.069$ m per percent and $\tau = 1000$ sec from the export, the values of Equation (2.10). The same step repeated at a mean drop of 4 bar shows the gain fall by $\sqrt{4/12} = 0.58$, the factor by which Equation (2.3) lowers

Table 2.3 Disturbance Scenarios: Demand Steps Halfway Between Executions

Scenario	Pressure swing	Outlet demand
Combined	$10\sin(t/10)$ bar	steps at 100.5, 350.5 s
Pressure only	$10\sin(t/10)$ bar	held at 2 kg/s
Demand only	off	steps at 100.5, 350.5 s

Table 2.4 Combined Scenario: Mean of Three Seeded Runs, With the Worst-Error Half-Range

Structure	max $\lvert e\rvert$ (m)	IAE (m s)	MV travel (%)	Saturated (%)
Feedback only	0.2989 ± 0.0002	145.9	103	0.0
FF from FT-302	0.3590 ± 0.0021	142.0	16295	8.1
FF from FT-303	0.0444 ± 0.0002	12.5	232	0.0
Cascade	0.2929 ± 0.0002	144.8	1219	0.0
Cascade with FF	0.0136 ± 0.0001	3.5	1122	1.1

Table 2.5 Mean Worst True-Level Error for Separate Upsets

Structure	Pressure only (m)	Demand only (m)
Feedback only	0.0184	0.2777
FF from FT-302	0.0162	0.3510
FF from FT-303	0.0183	0.0011
Cascade	0.0039	0.2927
Cascade with FF	0.0039	0.0046

K_v. This is the gain variation the cascade absorbs by measuring the delivered flow, and the static feedforward gain is exact only at the pressure it was computed for.

2.6.3 Combined Scenario

Table 2.4 gives the combined result. The best and worst structures differ by a factor of 26 in worst error with the shipped tunings.

Plain feedback reaches its worst error at model second 316 (exported second 317), while the valve is still ramping against the larger demand. The first seeded run finishes at 1.209 m, outside the ± 0.02 m band. Inlet feedforward has larger worst error and about 158 times the valve travel. It reverses valve direction on 70.0% of consecutive command increments and saturates for 8.1% of samples. With the swing off, its reversal fraction is 98.0%. This is the alternating mode of Equation (2.20), rather than useful anticipation of outlet demand. Its slightly smaller IAE would mislead a comparison that ignored peak error and actuator movement.

Pause. Try Exercise 2.3 before reading on.

Outlet feedforward cuts worst error by a factor of 6.7 and IAE by 11.7. The trim moves the valve by about 28.9% when demand changes, raising travel from 103 to 232%. Cascade alone gives almost the same worst error as feedback, while using 1219% travel. Adding the outlet trim reduces worst error to 0.0136 m, about 3.3 times below outlet feedforward alone.

2.6.4 Separated Scenarios

The separate upsets explain that ranking, as shown in Table 2.5. Cascade divides the pressure error by 4.7, but does not improve the demand error. Outlet feedforward divides demand error by about 250, while leaving pressure error almost unchanged. The two structures combined address both entry points. Inlet feedforward is worse against demand, and its small pressure-error reduction comes with oscillatory valve movement. Figure 2.3 draws the same table as bars, one panel for each upset.

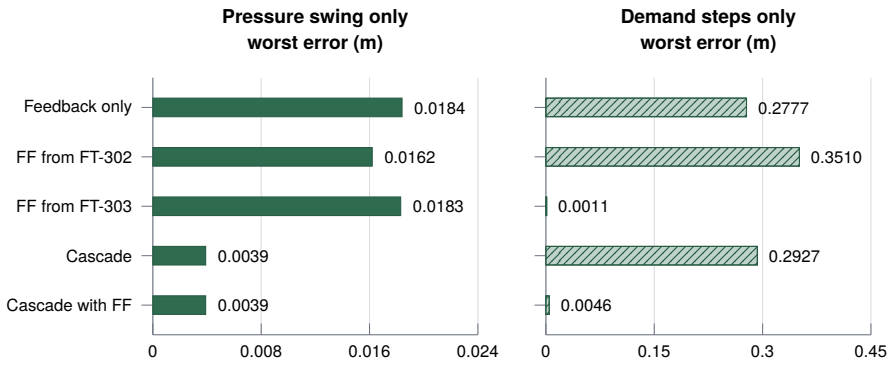


Figure 2.3 The worst level error of Table 2.5 for each structure under one upset at a time, the pressure swing alone (solid bars) and the demand steps alone (hatched bars), each panel on its own scale. Notice that the cascade shortens only the pressure bar and the outlet feedforward only the demand bar, that feedforward from the inlet transmitter lengthens the demand bar, and that only the cascade with the outlet trim is short in both panels.

Table 2.6 Inner Flow Loop Gain Sweep: Means of Three Seeded Runs

K	Roots at 12 bar	Stable below	max $\lvert e \rvert$ (m)	MV travel (%)	Saturated (%)
1.0	0.762, −0.455	44.4 bar	0.0039	254	0.0
1.5	0.701, −0.741	19.8 bar	0.0028	391	0.0
2.0	0.662, −1.047	11.1 bar	0.0744	7599	17.8
3.0	0.614, −1.692	4.9 bar	0.2798	17661	54.3

🕮 **Pause.** Try Exercise 2.1 before reading on.

For the first seeded pressure-only run, over the final 300 s feedback keeps lift between 21.51 and 22.21%, while inlet flow spans 3.04–10.39 kg/s and true level ripples 0.0214 m peak to peak. Cascade sweeps lift through 14.80–42.17%, inlet flow through 5.17–8.90 kg/s, and level through only 0.0047 m. It converts level excursion into actuator travel.

With constant pressure, the ideal trim cancels demand exactly in the noise-free continuous relation. The declared half-sample delay loses $10(0.5)/(\rho A) = 0.001$ m before compensation, and transmitter noise adds to the observed 0.0011 m residual. With swinging pressure the fixed trim is exact only at the mean pressure: a 10 kg/s step needs 21.3% lift at 22 bar but 70.7% at 2 bar, whereas the trim always requests 28.9%. This operating-point mismatch explains the 0.0444 m combined residual [16]. The cascade measures delivered flow rather than predicting it from one valve gain.

2.6.5 Inner-Loop Gain Sweep

The pressure-only test runs for 300 sec, with the same three seeds and $K_{c2} = K_{I2} = K$, not just a change in proportional gain. The pressure boundary solves $(2K_{c2} + K_{I2}\Delta t)K_v = 2$. Raising K from 1.5 to 2.0 increases travel by a factor of 19.4 and worst error by 26.6. The nominal boundary $K = 1.925$ lies between those values. At $K = 2$, mean successive movement is 38.4% above 20 bar but only 1.5% below 4 bar. Stability therefore changes within each pressure cycle. At $K = 1.5$, the frozen-pressure boundary is exceeded briefly near the crest, yet the finite-time response stays small. This is a local stability analysis of a moving operating point, not a guarantee based on the mean pressure.

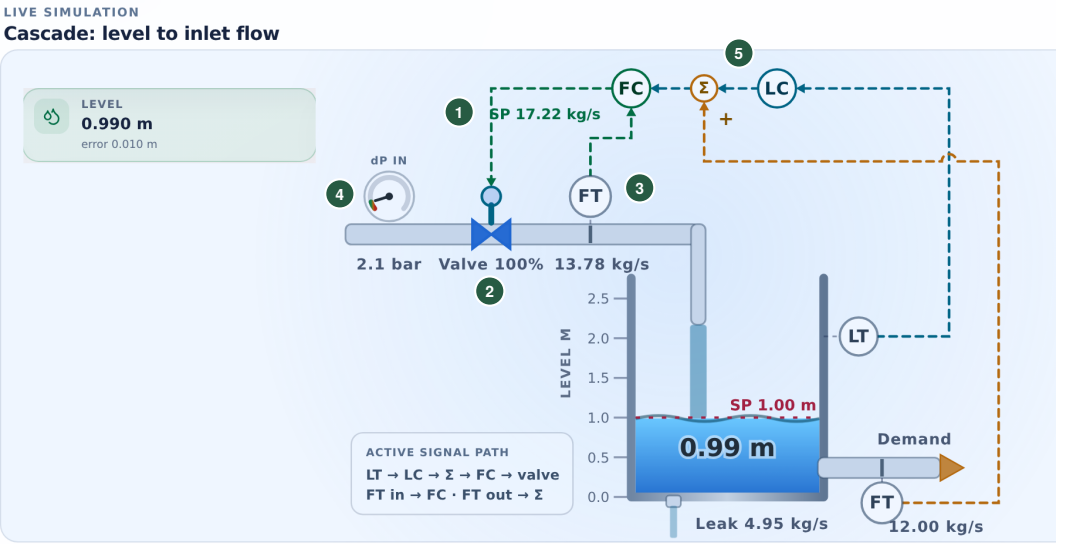


Figure 2.4 Plant view of the current application in a noise-free cascade-with-outlet-feedforward run, with the level tile of the status column set into the view's empty upper left corner, a composite of two parts of one screenshot. Demand rises from 2 to 12 kg/s at model second 100.5 (Table 2.3) while the pressure swing remains on, and the capture is taken at 300 s. The crop shows pressure at about 2.1 bar (4), lift at 100% (2) and, in the tile, level at 0.990 m. The commanded 17.22 kg/s flow setpoint (1) of the flow controller exceeds what the valve can deliver near the trough, as the inlet flow transmitter (3) shows, exposing the actuator constraint while the tank inventory limits the level excursion. The outlet trim joins the level controller's output at the summing junction (5).

> ✦ **DID YOU KNOW?** **Half a Unit Too Much Inner Gain**
>
> From the simulator: on the pressure-only test, raising both inner gains from 1.5 to 2.0 lifts the worst level error from 0.0028 to 0.0744 m, the valve travel from 391 to 7599 % and the time at a valve stop from 0.0 to 17.8 % (Table 2.6). The nominal limit of 1.925 lies between the two.

2.6.6 Application

A cascade run with the outlet trim engaged is shown in Figure 2.4.

2.6.7 Scenarios of the Application

The app has no scenario selector. Its thirteen scenarios are set up by hand from the controls, and the app tracks each as a required unit that combines a history predicate with an explicit learner confirmation. Table 2.7 summarizes the automatic predicates. History is segmented when mode, structure, feedforward state, K_{ff} or either inner gain changes. A demand-step segment needs both an upward and a downward step, 120 s after the last step, and no pressure swing after the first step. A pressure segment needs 200 s of swing since its last demand step. These checks establish activity, not a graded numerical explanation or independently measured learning outcome.

The inner-boundary unit checks proportional gain, whereas the stability sweep in Table 2.6 changes both proportional and integral gains. The noise unit combines a current setting with stored history and does not require that history to have been collected under the current noise value. The metric unit requires confirmation of a comparison but does not calculate whether level error improved. Progress and certificates should be interpreted with those limits.

Table 2.7 Scenarios and Completion Requirements

#	Scenario	Kind	Completion requirements
1	Design point	Experiment	A learner-started session and a manual quiet tail of at least 120 s, with swing off and true level within 0.02 m of SP.
2	Feedback steps	Experiment	Automatic feedback and the demand-step segment defined above.
3	Feedback swing	Experiment	Automatic feedback and the pressure segment defined above.
4	Inlet feedforward	Demonstration	Automatic inlet feedforward and a demand-step segment.
5	Outlet feedforward	Experiment	Automatic outlet feedforward and a demand-step segment.
6	Cascade swing	Experiment	Automatic cascade without feedforward and a pressure segment.
7	Cascade steps	Demonstration	Automatic cascade without feedforward and a demand-step segment.
8	Both structures	Experiment	Cascade with feedforward: one pressure segment and one demand-step segment, which can persist across segments.
9	Feedforward gain	Exploration	Outlet feedforward at gains 1.44, 2.887 and 5.77 (specified tolerances), each with a demand step and 120 s afterward.
10	Inner boundary	Demonstration	Cascade at inner proportional gains 1, 1.5 and 3 (tolerance 0.05), each with at least 200 s of swing.
11	Misleading metric	Demonstration	Cascade with both inner gains within 0.05 of 3, in one segment lasting at least 900 s.
12	Square-root gain	Exploration	One manual segment spanning mean pressure at most 4.5 and at least 21.5 bar, with lift variation at most 0.05%.
13	Noise limit	Demonstration	Current flow-noise setting at least 0.3 kg/s plus a cascade segment with inner proportional gain at least 5 and duration at least 120 s.

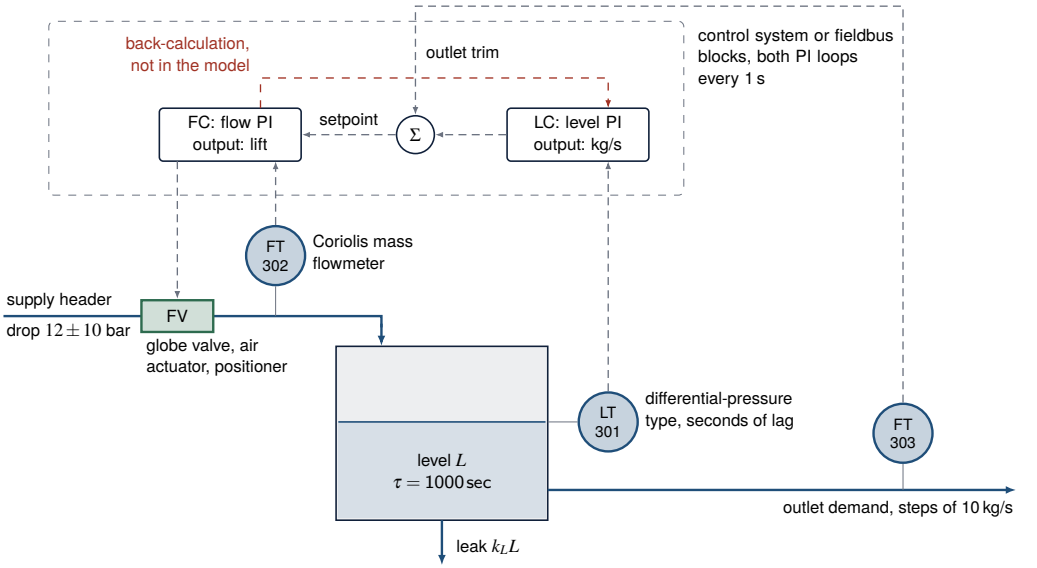


Figure 2.5 The level cascade with its outlet trim on real hardware. The level transmitter LT-301 feeds the level PI, whose output in kg/s is summed with the trim from FT-303 and sent as the setpoint of the flow PI, which reads FT-302 and moves the valve FV. Solid arrows are process streams and dashed gray arrows are signals. Notice the dashed path from the flow controller back to the level controller: a fieldbus sends this back-calculation signal so that the level controller can be told to stop integrating when the valve is at a stop, and the model of this chapter has no such path.

2.7 Technology in Practice

Figure 2.5 draws the loop on real hardware, which adds the lags, the noise and the limits the model leaves out.

What Measures It

A real level transmitter, such as the differential-pressure type that the Assumptions of Section 2.3 name, carries several seconds of lag, little against the 1000 sec level time constant. The flow transmitters matter more, because the inner loop and the outlet trim act on every sample they report. A mass flowmeter on the inlet line would read FT-302 in kg/s, as the model does, and the model's instantaneous reading with white noise of 0.05 kg/s stands for a meter whose own lag is short against the 1 sec sample.

> **TECHNOLOGY** **Sensor: A Coriolis Meter Under the Inner Loop**
>
> A Coriolis meter shakes one or more tubes that carry the liquid at their resonant frequency. Moving mass makes the tubes twist, the phase shift between two points on a tube measures the mass flow, and the frequency gives the density, so the meter reads kilograms per second directly, as FT-302 does in the model [13]. It is linear and accurate and needs no straight pipe ahead of it. Vibration from nearby machines and gas bubbles trapped in bent tubes cause errors in the flow reading, straight tubes are less sensitive at low flow, and the meter is costly in large pipe sizes [13]. An inner loop trusts its meter, so every such error moves the valve.

What Moves It

The final element is the inlet control valve, algebraic in the model: flow follows lift within the sample, linear in the lift and proportional to the square root of the supply drop. A real valve adds its stroking lag inside the inner loop, which the Assumptions name as the effect that lowers the usable inner gain. It is customary in process plants to close yet another loop on the stem with a valve positioner, at the lowest level of the hierarchy, to reduce the effects of friction and nonlinearity [9], so a level cascade on such a valve is three loops deep.

> **TECHNOLOGY** **Final Element: A Pneumatic Valve That Needs Time to Stroke**
>
> A pneumatic actuator moves the stem only as fast as air can fill or vent it. A large actuator takes time to change pressure, so the valve lags its signal. A volume booster, a relay that copies the signal one to one with far more air capacity, shortens the lag but cannot remove it while the air flow is finite [13]. When the pressure drop across a valve changes, the valve also responds more near closed than near open. A flow loop wrapped around the valve confines that distortion to the flow controller, and the level controller sees a more predictable response [13].

Where the Controller Runs

Both controllers would run in a distributed control system or in fieldbus function blocks. Loops there are sometimes nested or cascaded, with cycle times from milliseconds to minutes [21], and the app's 1 sec period sits inside that range. The standard PID function block of a fieldbus carries the chapter's structures as inputs: a cascade input that takes a remote setpoint, a feedforward input with its own gain, and a back-calculation input and output. The flow PID may run inside the transmitter or inside the valve positioner, with no separate controller [13]. In Equation (2.17) the flow setpoint is clamped at 40 kg/s and the lift at 100 %, but no signal passes from the lift clamp back to the level controller. Figure 2.4 shows the valve at that clamp with a setpoint the valve cannot deliver.

What Is Modernizing It

Two developments change where and how a cascade runs. First, control moves into the field: a function block chain propagates the status of every measurement, so a block fed a bad reading can hold its last good value or fail to a set end of its range instead of acting on bad data [13]. Second, the loops of a cascade can close over communication networks, where delay and lost packets eat into the inner loop's margin and schemes that transmit only on events save network load [22].

IN PRACTICE **Keep the Secondary in Cascade, and Tell the Primary**

A secondary controller has three modes: manual, automatic with a local setpoint, and cascade, in which the primary sets its setpoint. Left in automatic, it defeats the cascade as surely as manual defeats any loop [13]. Windup also crosses between the loops. The secondary's anti-windup works in the usual way because it drives the actuator, but the primary must be told when the secondary goes into anti-windup [9]. Fieldbus function blocks pass a back-calculation signal up the chain for this, so a valve that cannot follow, or a selector that drops the primary's output, stops the integral action above it [13].

2.8 Limitations

The plant is a single first-order state and its manipulated path is algebraic, which is the condition under which static feedforward is exactly right and cascade is maximally effective. Both results weaken on a plant with transport delay or valve dynamics, and this application cannot show that weakening because it contains neither.

The performance numbers depend on the default tuning of the application, and that tuning is not optimal for any of the four structures. The level loop's reset of 50 sec against a 1000 sec process is aggressive, and a slower reset would lower the valve travel of every row of Table 2.4 while raising the recovery time. The comparison is between four structures at one tuning, not between four structures each at its own best tuning.

The pressure sinusoid is a single frequency of large amplitude. A single frequency makes the sensitivity argument clean, but it also means the numbers in Table 2.5 say nothing about a broadband supply upset, where the inner loop would attenuate the low-frequency content and pass the content above its own bandwidth. The amplitude is large enough that the square-root nonlinearity contributes harmonics, which is why the measured ripple reduction of 4.6 is not exactly the 3.7 that the ratio of the two linear sensitivities predicts.

Neither the valve position nor the flow setpoint is rate limited, so the cascade is free to sweep the valve 27 % every cycle without penalty. A real installation would either accept the wear or detune the inner loop, and the application offers no way to quantify that trade-off.

The app's random stream is unseeded, so visits are not bit-identical. The reference runs use explicit seeds. Their half-ranges describe only those three realizations, and do not establish the spread over all possible sensor noise.

The model now caps level at the 2.5 m lip and meters spilled mass. Its zero-level floor still clips an infeasible withdrawal without separately metering unmet demand. High-demand dryout therefore needs care in mass-accounting studies, and the results here concern operation near the design level.

OPEN CHALLENGE **An Inner Loop Closed Over a Network**

A recent survey treats cascades whose loops close through a real-time communication network. Delays are then unavoidable in both the inner and the outer loop, and with lost packets they can reduce performance or even cause instability. Schemes that send data only when an event calls for it lighten the network, and cyber-attacks that change the loop's dynamics are part of the problem [22]. On this tank one sample of delay already caps equal inner gains at 1.925 at 12 bar. Progress would be an inner loop whose stability holds with a network's varying delay and losses inside it.

2.9 Discussion

The result that a structure rejects one disturbance and not the other is not a property of this tank. It is a property of where measurements sit relative to where energy and material enter, and it recurs wherever a slow inventory or a slow temperature is held by a fast actuator. A furnace outlet temperature loop that does not cascade to fuel flow will pass every fuel-header pressure swing into the process fluid, and no retuning of the temperature controller will change that, because the temperature controller cannot see the header. Chapter 4 puts flow loops under the fuel and the air of such a heater. A blending

loop without a ratio on the wild stream will chase composition after the fact for the same reason [6], and Chapter 7 blends syrup into water on such a ratio. The measurement is the constraint and the gain is not.

The economic asymmetry follows from the same observation. Adding a flow transmitter and a second controller is a capital project, while changing a gain is a configuration change, so plants under-invest in structure and over-invest in tuning [2]. The measurements in this chapter quantify the return on the added transmitter. The cascade with the outlet trim reaches 0.0136 m of worst error where the default single-loop tuning reaches 0.2989 m, at the reported settings. The comparison does not prove an optimum over all possible tunings, and the default level loop has a loop gain of only 0.0143 at the disturbance frequency. Deciding which variables to measure and control is the decision that bounds every later one [1].

The model of Section 2.3 also supports control methods beyond the four structures implemented. The single state, the algebraic valve law and the two measured disturbances translate directly into a model predictive control formulation, in which one controller coordinates the valve against its 0 % to 100 % travel limits and the level bounds while the square-root pressure dependence enters the prediction model instead of being absorbed as gain variation [23]. Moving horizon estimation or Kalman filtering on the same model would reconstruct the supply pressure drop and the leak coefficient from the three transmitters, replacing the fixed feedforward gain with one that tracks the operating point and removing the 0.0444 m residual the static compensator leaves when the pressure swings. Both methods run on differential and algebraic models of exactly this form in the APMonitor framework [23] and the Gekko Optimization Suite in Python [24], and the exported runs supply the data such an extension would be identified against.

2.10 Conclusion

This chapter placed four control structures on one simulated tank and applied two upsets that enter the plant at different points. The plant is a single mass balance with a square-root valve law, a level-proportional leak, a 1000 sec time constant and three noisy transmitters, and every structure drives the same valve through the same velocity-form proportional-integral algorithm on the same 1 sec sample period. Feedforward from the outlet transmitter is exact rather than approximate on this plant, because the demand and the manipulated flow enter the same integrator and the ideal compensator collapses to the constant 2.887 % per kg/s. Cascade to the inlet flow reaches the supply pressure upset instead, attenuating it by 72.9 % at the 0.1 rad/s swing frequency at which the level loop attenuates none of it. Applying the two upsets separately measures the pairing, because the cascade divides the pressure error by 4.7 and leaves the demand error at 0.2927 m, while the outlet feedforward divides the demand error by about 250 and leaves the pressure error at 0.0183 m. Running both together reaches 0.0136 m on the combined scenario against 0.2989 m for the single loop. This comparison is made at one tuning on a plant with no transport delay and no valve dynamics, and extending it to a delayed manipulated path, a broadband supply upset and a rate-limited valve would show how far the pairing survives. The app lets a reader choose a control structure by measurement rather than by rule of thumb.

KEY POINTS

- Choose a loop structure by where the upset enters the tank, not by tuning.
- The cascade divides the pressure error by 4.7, and outlet feedforward divides the demand error by about 250.
- The inner loop removes 72.9 % of the flow upset at 0.1 rad/s, where the level loop removes none.
- Feedforward from the inlet transmitter puts a pole at $z = -1$ and halves the level controller's authority.

Where else this applies. Flow loops under a heater's fuel and air return in Chapter 4, and inner-loop speed on a fermenter in Chapter 3.

References

[1] S. Skogestad, "Control structure design for complete chemical plants," *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

[2] F. G. Shinskey, "Process control: As taught vs as practiced," *Industrial & Engineering Chemistry Research*, vol. 41, no. 16, pp. 3745–3750, 2002.

[3] R. G. Franks and C. W. Worley, "Quantitative analysis of cascade control," *Industrial & Engineering Chemistry*, vol. 48, no. 6, pp. 1074–1079, 1956.

[4] P. R. Krishnaswamy, G. P. Rangaiah, R. K. Jha, and P. B. Deshpande, "When to use cascade control," *Industrial & Engineering Chemistry Research*, vol. 29, no. 10, pp. 2163–2166, 1990.

[5] J. L. Guzmán and T. Hägglund, "Simple tuning rules for feedforward compensators," *Journal of Process Control*, vol. 21, no. 1, pp. 92–102, 2011.

[6] T. Hägglund, "The Blend station – a new ratio control structure," *Control Engineering Practice*, vol. 9, no. 11, pp. 1215–1220, 2001.

[7] 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.

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

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

[10] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

[11] W. L. Luyben, "Parallel cascade control," *Industrial & Engineering Chemistry Fundamentals*, vol. 12, no. 4, pp. 463–467, 1973.

[12] Y. Lee, S. Park, and M. Lee, "PID controller tuning to obtain desired closed loop responses for cascade control systems," *Industrial & Engineering Chemistry Research*, vol. 37, no. 5, pp. 1859–1865, 1998.

[13] 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.

[14] D. Gromov, "Analysis of hierarchical structures for hybrid control systems," Ph.D. dissertation, Technische Universität Berlin, Berlin, 2010.

[15] M. Hast and T. Hägglund, "Low-order feedforward controllers: Optimal performance and practical considerations," *Journal of Process Control*, vol. 24, no. 9, pp. 1462–1471, 2014.

[16] E. J. Adam and J. L. Marchetti, "Designing and tuning robust feedforward controllers," *Computers & Chemical Engineering*, vol. 28, no. 9, pp. 1899–1911, 2004.

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

[18] ——, "The future of PID control," *Control Engineering Practice*, vol. 9, no. 11, pp. 1163–1175, 2001.

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

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

[21] 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.

[22] M. Fallahnejad, "A survey of recent breakthroughs in networked cascade control systems," *International Journal of Robust and Nonlinear Control*, vol. 36, no. 5, pp. 2531–2541, 2026.

[23] 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.

[24] 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.

2.11 Exercises

2.1 Concept check · 2 min

With the pressure swing off, the outlet demand steps from 2 to 12 kg/s and back. Plain feedback leaves a worst level error of 0.2777 m and the level-to-flow cascade leaves 0.2927 m (Table 2.5). Why does the cascade not help against this upset?

(a) The inner flow loop is too slow, so it adds a lag in front of the level loop
(b) The inner flow loop runs close to its stability limit at the default gains
(c) The demand enters the balance downstream of FT-302, so the inner loop never sees it
(d) The noise on FT-302 hides the demand step from the inner flow controller

2.2 Concept check · 2 min

Why is the ideal feedforward compensator from the outlet demand to the valve a pure gain on this tank, with no lead, lag or delay?

(a) Because the plant has no transport delay, so every compensator for it is static
(b) Because the integral action of the level loop supplies the missing dynamics
(c) Because the leak makes the tank self-regulating, which removes the pole from the ratio
(d) Because the demand and the inlet flow enter one balance, so both paths share one pole

2.3 Concept check · 1 min

Structure 2 trims the lift by $K_{\mathrm{ff}}(F_{\mathrm{in}}^{m} - F_{\mathrm{in},0})$ with the default $K_{\mathrm{ff}} = 2.887$. What does Equation (2.20) predict at the design pressure?

(a) A pole at $z = -1$, an alternating valve and half the level controller's authority
(b) Exact cancellation of the pressure swing, since 2.887 is the ideal gain $1/K_v$
(c) A pole at $z = +1$, so the lift drifts slowly away from its design value
(d) No change from plain feedback, since FT-302 cannot see the outlet demand

2.4 Theory and calculation · 6 min

The tank of Table 2.2 holds $L = 1.0\,\mathrm{m}$ with $F_{\mathrm{out}} = 2.0\,\mathrm{kg/s}$, $k_L = 5.0\,\mathrm{kg/s}$ per m, $\rho = 1000\,\mathrm{kg/m^3}$, $A = 5.0\,\mathrm{m^2}$, $C_v = 1.0 \times 10^{-4}$ and $g_s = 1.0$.

(a) Find the valve gain K_v at 12 bar and the lift that holds the level. *Checkpoint:* 0.3464 kg/s per percent and 20.21 %.

(b) Find K_p and τ of the linearized model, and the steady level change when the lift steps from 20.2 to 30 % with the swing off. *Checkpoint:* 0.0693 m per percent, 1000 s and 0.679 m.

(c) Find the ideal outlet feedforward gain at 12 bar and at 4 bar, and the factor by which K_p falls at 4 bar. *Checkpoint:* 2.887 and 5.000 % per kg/s, and 0.577.

2.5 Theory and calculation · 8 min

The inner flow loop has $K_{c2} = K_{I2} = K$, $\Delta t = 1\,\mathrm{s}$ and the characteristic polynomial of Equation (2.18). The supply drop follows Equation (2.7) with a mean of 12 bar and an amplitude of 10 bar.

(a) Find the roots at the default $K = 1.0$ and 12 bar. *Checkpoint:* 0.762 and -0.455.

(b) Find the largest stable K at 12 bar, and the drop above which $K = 1.5$ is unstable. *Checkpoint:* 1.925, and 19.75 bar.

(c) For how many seconds of each swing cycle does the drop exceed that boundary for $K = 1.5$? *Checkpoint:* 13.7 s of every 62.8 s.

2.6 Theory and calculation · 6 min

The static outlet feedforward of structure 3 uses the fixed gain $K_{\mathrm{ff}} = 2.887\,\%$ per kg/s while the supply drop swings between 2 and 22 bar.

(a) Find the lift change that passes an extra 10 kg/s at 22, 12 and 2 bar, and compare it with what the trim requests. *Checkpoint:* 21.3, 28.9 and 70.7 %, against 28.9 %.

(b) With the swing off, the demand steps by 10 kg/s halfway between two executions of the 1 s control period. Find the level lost before the trim acts, and the most that can be lost. *Checkpoint:* 0.001 m, and 0.002 m.

(c) Find the gain a feedforward from the supply drop would need at the design point. *Checkpoint:* 0.842 % per bar.

2.7 App challenge · 8 min

In Settings set Level transmitter noise and Flow transmitter noise to 0. Leave the pressure swing on at its default 12 ± 10 bar and the demand at 2 kg/s. For scenario 3, Feedback only, pressure swing: at speed 1× press Reset, select Feedback and press PID auto within the first 5 s, then raise the speed to 10×, press Stop once the clock passes 900 s and download the Results file. Repeat for scenario 6, Cascade, pressure swing, selecting Cascade with + outlet FF off. Use the rows from 10 to 15 min of each file.

(a) Read the range, highest minus lowest, of `level_true_m` in each run, and the ratio of the two. *Checkpoint:* 0.0214 m and 0.0047 m, within 0.0005 m, a ratio near 4.6.

(b) Read the lowest and highest `valve_pct` and `inlet_flow_kgs` in each run. *Checkpoint:* Feedback 21.6 to 22.1 % and 3.06 to 10.33 kg/s. Cascade 15.1 to 42.0 % and 5.26 to 8.86 kg/s. Within 0.2 on each.

Compare the ratio of part (a) with the ratio of the sensitivities $|S_1|/|S_2|$ at the swing frequency.

2.8 App challenge · 8 min

In Settings set both transmitter noises to 0, and switch the pressure swing off. For scenario 2, Feedback only, demand steps: at speed 1× press Reset, select Feedback and press PID auto within the first 5 s, then raise the speed to 10×, press Step +10 when the clock reads about 100 s and Step −10 at about 350 s, run to 470 s and download the Results file. Repeat with FF outlet (scenario 5, The right feedforward) and with Cascade, + outlet FF off (scenario 7, Cascade, demand steps). For each run find the largest $|1 - \texttt{level_true_m}|$. Accept 0.002 m on each.

2.9 App challenge · 8 min

Scenario 9, Check Kff exactly. In Settings set both transmitter noises to 0, and switch the pressure swing off. At speed 1× press Reset, select FF outlet, set Kff to 1.44 and press PID auto within the first 5 s. Raise the speed to 10×, press Step +10 when the clock reads about 100 s, run to 350 s and download the Results file. Repeat with Kff at 2.887 and at 5.77.

(a) For each gain read the lowest and the highest `level_true_m` after the step. *Checkpoint:* 1.44: lowest 0.861 m. 2.887: within 0.002 m of 1 m. 5.77: highest 1.280 m. Within 0.003 m.

(b) Predict the two nonzero errors from the feedback error of 0.2775 m for the same step, taking the error to scale with the share of the load the trim leaves uncompensated. *Checkpoint:* 0.139 m below and 0.277 m above the setpoint.

2.10 Application area · 3 min

In the pressure-only run the cascade sweeps the valve between 14.80 and 42.17 % on every 62.8 s cycle of the swing (Section 2.6).

(a) Estimate the valve travel per day, counting two sweeps per cycle, in full strokes of 100 %. *Checkpoint:* About 750 full strokes per day.

(b) Name the two responses to this wear that the Limitations section gives, and say what the app cannot tell you about them.

2.11 Application area · 4 min

The model takes the leak as linear in level. Suppose the tank drained instead through an orifice, $F_{\text{leak}} = c\sqrt{L}$, with $c = 5\,\text{kg/s}$ per $\sqrt{\text{m}}$, so that the same 5 kg/s leaves at 1 m.

(a) Linearize the leak and find τ and K_p at 1 m and τ at 2 m, with $K_v = 0.3464\,\text{kg/s}$ per percent. *Checkpoint:* 2000 s and 0.139 m per percent at 1 m, and 2828 s at 2 m.

(b) Say whether the ideal outlet feedforward gain changes with the level.

2.12 Review the agent · 5 min

An AI assistant reviews the cascade and replies as follows. “In the 900 s combined scenario, raising both inner gains from 1.0 to 3.0 lowers the worst level error from about 0.29 to about 0.22 m. The tighter flow loop rejects the supply swing faster, so set $K_{c2} = K_{I2} = 3.0$.” Find the flaw in the recommendation and say how you would check it.