INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 1 Process Dynamics, Modeling, and Identification

CHAPTER 3

Two-Tank Level Control With a Splitter Valve That Sets the Plant Order

Pages 53–83

Edited by John D. Hedengren

How to cite this chapter. "Two-Tank Level Control With a Splitter Valve That Sets the Plant Order," in *Process Dynamics, Modeling, and Identification*, John D. Hedengren, Ed., vol. 1 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 53–83.

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

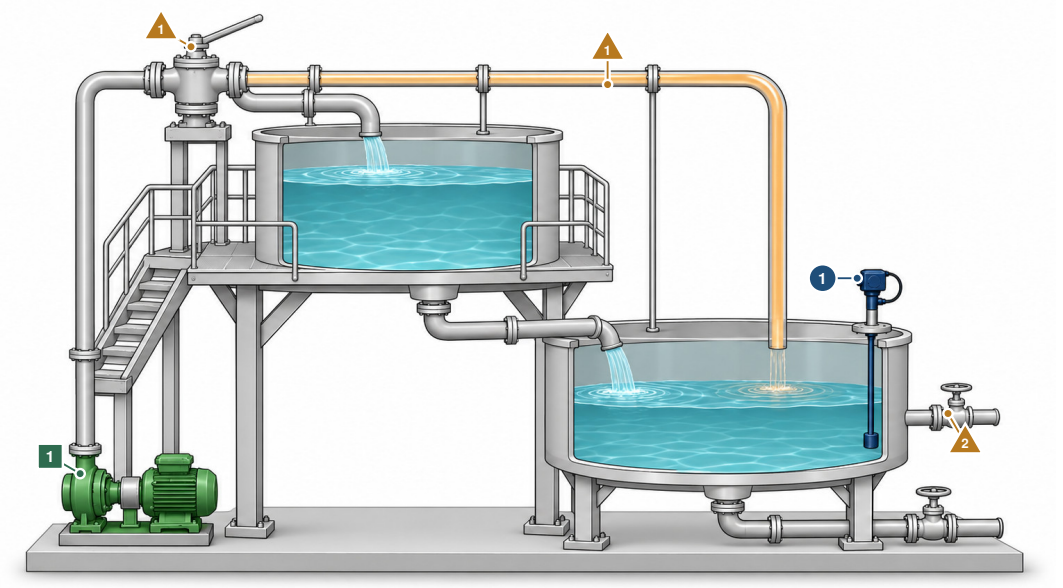
The app for this chapter
https://apmonitor.com/apps/2tank/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v01-dynamics-modeling/#2tank

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

CHAPTER 3

Two-Tank Level Control With a Splitter Valve That Sets the Plant Order



The controller moves 1 the pump that feeds both paths

The controller reads 1 the level in the lower tank

What disturbs it 1 the splitter, set by hand, and the bypass it feeds 2 a draw-off from the lower tank

Notice that the bypass reaches the lower tank at once, while the other path must first fill the upper tank, whose own level indicator is not drawn. Illustration, not to scale.

IN PLAIN TERMS **One Pump, Two Tanks and a Choice of Path**

The main idea. A pump fills a top tank that drains into a bottom tank, and a controller holds the bottom tank at its target depth. A hand valve can send part of the water straight to the bottom tank.

The challenge. Whichever path the water takes, the bottom tank ends at the same depth, so a slow look shows no change. But the path through the top tank is slower, and settings tuned for one path can make the level swing on the other.

Why it matters. A plant that reroutes its flow for repairs can keep a controller tuned for a path that is gone.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on five tanks in series (Section 3.2), the history of the draining tank (Section 3.3), the try-it segments of cases 2 and 9 (Sections 3.4 and 3.8), the fact on the draw-off (Section 3.7) and the key points
Study	70 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Sections 3.6 and 3.7), Technology in Practice (Section 3.9), Limitations (Section 3.10) and the open challenge

The control question A splitter valve that sends pump flow past the top tank leaves the bottom level and the process gain where they were, yet turns a second-order plant into a first-order one. What does a step test read at each splitter setting, and what does a tuning or a controller model carried from one lineup to the other do?

Learning objectives

- Derive the steady levels, the pump bias and the linearized time constants of two non-interacting tanks from their mass balances.
- Explain why the splitter changes the poles and the zero of the pump transfer function but not its gain.
- Fit a gain, time constant and apparent dead time to a step test, and tell a secant gain from a tangent gain.
- Predict the proportional offset and the reachable level under a draw-off, and the top-tank level that rejecting it requires.
- Judge a tuning or a predictive controller's model against the plant structure it was identified on.

Prerequisites Sections 1.4.2 and 1.4.3 (balances and linearization), 1.4.5 (two tanks in series), 1.4.6 (half rule) and 1.4.8 (two-point fit), and Chapter 2 for one gravity-drained tank. Sections 1.6.2 and 1.6.3 are needed for the controllers.

The app https://apmonitor.com/apps/2tank/

See also Volume 2, Chapter 4 for step tests on a plant whose model moves with load, Volume 2, Chapter 6 for gain scheduling and Volume 2, Chapter 3 for a bound set by the supply. Volume 9, Chapter 3 for predictive controllers that differ only in their step test, and Volume 9, Chapter 4 for offset-free estimation.



3.1 Introduction

A controller is commissioned against a plant, and the identified model reflects the piping configuration that was in service on the day of the step test. Operating-point dependence is taught carefully, because a nonlinear gain moves with throughput and every process control course shows how to recognize it. Structural change receives less attention. A bypass opens, a vessel is taken out of service, a recycle is lined up, and the loop that used to see two lags in series now sees one. The steady-state map may not move at all when that happens, so the operator has no reason to expect a change, the trend looks the same in the long run, and the only symptom is that a controller tuned against the previous structure exhibits oscillation.

Two tanks in series with a splitter valve are the smallest apparatus in which that structural change can be produced deliberately. One pump feeds a three-way valve. Part of the flow goes to a top tank that drains into a bottom tank, and the rest bypasses the top tank and enters the bottom tank directly. Whatever the split, all of the pump flow leaves through the same bottom drain, so the resting level in the bottom tank depends on the pump position alone. The split changes only the dynamic path from the pump to the bottom tank. With the bypass shut the pump must fill the top tank before the bottom tank moves at all, and with the bypass wide open the top tank empties and stays empty. The same pump step therefore produces the same final level through a plant that is second order in one case and first order in the other.

The model behind the apparatus is a mass balance on each vessel with gravity discharge through a fixed restriction as the driving force, which is the standard construction for tank dynamics in process control texts [1, 2]. The experiments built on it cover process order, apparent dead time from lags in series, graphical identification, internal model control tuning, proportional offset, integral action, anti-windup and gain scheduling, and the limits they expose follow from the bound that effective delay places on the bandwidth of any linear feedback loop [3, 4]. The behavior matters industrially because bypasses, spared vessels and seasonal lineups change plant structure without changing the material balance, so a tuning carried across such a change degrades without any steady-state symptom.

Coupled tanks are already the standard apparatus for teaching multivariable control, and the quadruple-tank rig makes exactly this kind of point by moving a transmission zero across the

imaginary axis with a pair of manual splits [5]. However, that experiment needs two pumps and two controlled levels before the effect appears, and the structural change it demonstrates is a property of the multivariable zero rather than of the plant order that a single loop has to be tuned against. This chapter works with a two-tank simulation in which the splitter fraction converts a second-order overdamped level plant into a first-order one without disturbing the operating point. It quantifies what a controller tuned at either extreme of that fraction does at the other.

Section 3.2 places two tanks in the plantwide and tuning literature. Section 3.3 writes down the simulator in enough detail (two states, two outflow laws, three disturbances, four hard limits, an instrument model and a numerical scheme) for an independent build to reproduce its trajectories. Section 3.4 shows that the pair is strictly lower triangular and non-interacting, that its upper time constant is the lower one scaled by the share of flow still passing the top tank, and that the splitter reaches the bottom level through a transfer function with a zero at the origin. It then reads one gain and three dynamic models from step tests at three splitter positions. Section 3.5 describes the velocity-form PID and the linear model predictive controller that can replace it in automatic, with its moving horizon estimator, its table of models selected by the splitter setting and a relay test that proposes PI constants from a limit cycle. Section 3.6 runs the tunings identified at the two ends of the splitter against both structures. Section 3.7 treats the draw-off and the leak, which move the steady state, and shows that the top tank caps the load a fixed lineup can reject. Section 3.8 turns the splitter into a test of a controller model as well as of a tuning. Section 3.9 names the devices a plant pair of tanks would carry. Sections 3.10 to 3.12 state the limitations, discuss the industrial setting and conclude, and Section 3.13 holds the exercises. One sequence of experiments thus produces identification, tuning, tuning mismatch and a controllability bound from the travel of a single valve.

3.2 Background

The reason two tanks appear in a control course at all is that process plants are built out of units that must be decoupled from one another. A vessel placed between two units is there so that a flow disturbance arriving from upstream does not have to be passed downstream at the moment it arrives, and the inventory in that vessel supplies the surge capacity that absorbs the disturbance. That requirement shapes the whole control structure of a plant rather than the tuning of one loop, because every holdup has to be regulated by some flow, the direction in which each holdup is regulated has to be consistent with how throughput is set, and the loops that do the regulating are chosen before any of them is tuned [6]. Level is where the structural question is most visible, because a level is an integral of the imbalance between two flows and cannot be left uncontrolled.

Placing vessels in series rather than in parallel has a specific dynamic cost. Each additional holdup contributes a lag, and lags in series accumulate into an apparent dead time that no controller can undo. The half rule of Equation (1.16) states that consequence quantitatively for tuning, and it is exactly the trade a second tank imposes on a level loop [4]. The limit is fundamental rather than a matter of tuning skill, because effective delay bounds the achievable bandwidth of any linear controller regardless of its structure [7]. The practical answer, when the plant structure itself is the thing that moves, is to schedule the controller against whatever variable indexes the change, and the gain-scheduling literature records both the appeal and the pitfalls of doing so [8, 9].

> **✦ DID YOU KNOW? Five Lags in Series in a Gold Circuit**
>
> A gold leaching circuit strings five stirred tanks in series, each a first-order lag on the one before it (Volume 15, Chapter 7). For tuning, such a chain is reduced to one lag and a delay by the half rule: the largest neglected time constant is split evenly between the effective delay and the smallest time constant kept [4].

Tank rigs became the laboratory vehicle for all of this because they are inexpensive, slow enough to observe in real time, and nonlinear. The quadruple-tank process uses two pumps and four vessels so that a pair of manual splits moves a multivariable zero between the left and right half planes [5], and

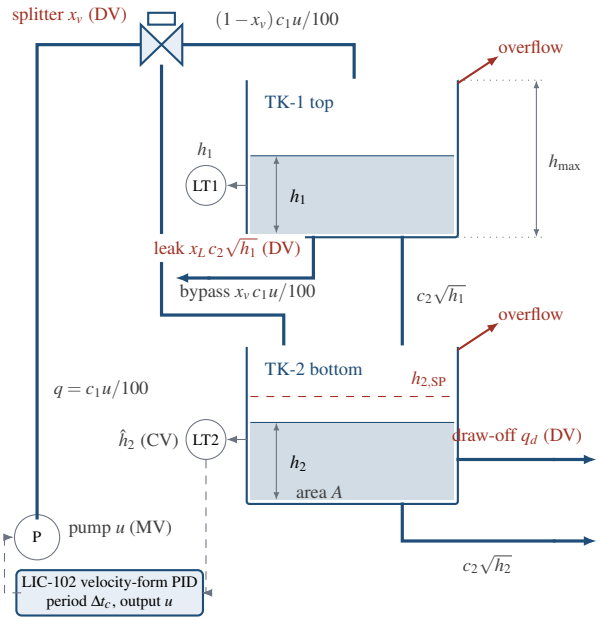


Figure 3.1 Two gravity-drained tanks in series behind one pump and a three-way splitter valve. The splitter sends the fraction x_v of the pump flow past TK-1 and straight into TK-2, TK-1 drains into TK-2 through $c_2\sqrt{h_1}$, and TK-2 drains to waste through $c_2\sqrt{h_2}$. LIC-102 controls the bottom level $\hat{h}_2$ with the pump. The splitter fraction, the bottom-tank draw-off q_d and the top-tank leak x_L are set by the operator and are never controlled.

it has served since as a plant for constrained and model-based control studies [10]. Two-tank rigs are equally standard, coupled through an orifice so that a backstepping controller has to be written for the pair [11], or driven over a network as a remote laboratory [12]. The textbook chapter that introduces tanks in series also separates interacting from non-interacting pairs [1], which decides whether the plant transfer matrix is full or triangular, and the apparatus of this chapter is the non-interacting case with a variable structure, a combination the existing rigs do not offer.

3.3 Process Model

The simulated plant is a pair of open, vertical, constant-area tanks. A single pump draws from an unlimited supply and discharges into a three-way splitter valve. The fraction $1 - x_v$ of the pump flow enters the top tank TK-1, the fraction x_v bypasses TK-1 and enters the bottom tank TK-2 directly, TK-1 drains by gravity into TK-2, and TK-2 drains by gravity to waste. Transmitter LT2 on the bottom tank reports the controlled variable, transmitter LT1 on the top tank gives an indication only, and the pump position is the manipulated variable, while the splitter fraction, an extra draw-off on TK-2 and an extra leak on TK-1 stay under operator control in manual and in automatic alike. Every symbol used below appears on the flowsheet of Figure 3.1.

Table 3.1 sorts the signals into their roles. There are two states and one actuator, so both tanks are driven by the same pump and no cascade exists. Four hard limits bound the states, an overflow limit and a dry limit on each tank. The plant is written with both levels in meters, the pump, the splitter and the leak in percent of travel, the draw-off as a level velocity in m/s and time in seconds. The splitter and leak settings enter the balances as the fractions x_v and x_L, each the setting divided by 100.

Table 3.1 Signal Roles, Ranges and Nominal Values at the 0.5 m Bottom-Tank Operating Point With the Splitter Shut and No Load

Symbol	Role	Nominal	Range	Unit
h_2	state, bottom level	0.5013	0–$h_{\max}$	m
h_1	state, top level	0.5013	0–$h_{\max}$	m
$\hat{h}_2$ (LT2)	CV, bottom level	0.5013	0–$h_{\max}$	m
$\hat{h}_1$ (LT1)	indication, top level	0.5013	0–$h_{\max}$	m
$h_{2,\mathrm{SP}}$	setpoint	0.50	0–$h_{\max}$, step 0.01	m
u	MV, pump position	35.4	0–100, step 0.1	%
$100\,x_v$	DV, splitter setting	0	0–100, step 1	%
q_d	DV, bottom draw-off	0	0–0.030, step 0.001	m/s
$100\,x_L$	DV, top-tank leak	0	0–100, step 1	%
σ_n	DV, LT2 noise 1σ	0.004	0–0.020	m
$\sigma_{n,1}$	DV, LT1 noise 1σ	0.004	0–0.020	m
overflow 1, 2	outputs, two flags	false	—	—

3.3.1 Level Balances

Both tanks have the same constant cross-sectional area A, and A is divided out of every term before it reaches the source, so the two balances are written directly as level velocities. Each is the mass balance of Equation (1.2) for one gravity-drained tank [1, 2], as shown in Equations (3.1) and (3.2).

$$\frac{\mathrm{d}h_1}{\mathrm{d}t} = c_1\,(1-x_v)\,\frac{u}{100} - c_2\,(1+x_L)\,\sqrt{h_1} \tag{3.1}$$

$$\frac{\mathrm{d}h_2}{\mathrm{d}t} = c_1\,x_v\,\frac{u}{100} + c_2\sqrt{h_1} - c_2\sqrt{h_2} - q_d \tag{3.2}$$

where c_1 is the pump coefficient in m/s at full pump, c_2 is the gravity-drain coefficient in $\mathrm{m}^{1/2}/\mathrm{s}$, u is the pump position in percent, $x_v \in [0,1]$ is the splitter setting over 100, $x_L \in [0,1]$ is the leak setting over 100 and $q_d \geq 0$ is the draw-off in m/s of bottom-tank level. Every term on both right-hand sides carries units of m/s. The outflow terms are the Torricelli form with the discharge coefficient, the orifice area, the tank area and $\sqrt{2g}$ all lumped into the single constant c_2, which is the only combination a controller can distinguish. Substituting the default constants gives $\mathrm{d}h_1/\mathrm{d}t = 0.0008\,(1-x_v)\,u - 0.04\,(1+x_L)\sqrt{h_1}$ and $\mathrm{d}h_2/\mathrm{d}t = 0.0008\,x_v\,u + 0.04\sqrt{h_1} - 0.04\sqrt{h_2} - q_d$ in m/s.

> ⧖ HISTORY **The Law of the Draining Tank, 1644**
>
> Evangelista Torricelli (1608–1648), the last student of Galileo, stated the law that drains both tanks of this chapter. Liquid leaving a small hole at the bottom of a vessel flows out as fast as a body that has fallen freely from the height of the liquid surface [13]. His work on water jets appeared in 1644, in a section on the motion of waters within a treatise on falling bodies, and it showed that every jet leaving a hole in a filled tube traces a parabola. It counts among the founding results of fluid mechanics [14]. Torricelli had no calculus. Written as a differential equation for the level, his law is now a standard exercise of a first course in differential equations [13].

The structure of Equations (3.1) and (3.2) establishes the property on which the rest of the analysis depends. The tanks are *non-interacting*. The discharge from TK-1 is $c_2\sqrt{h_1}$ and depends on the top level alone, because TK-1 drains into the vapor space of TK-2 rather than below its liquid surface, so nothing that happens in the bottom tank propagates back into the top one. The state matrix is lower triangular for that reason, and it stays lower triangular at every splitter setting. An arrangement in which a connecting orifice sits below both surfaces makes the discharge proportional to $\sqrt{h_1 - h_2}$ and produces a genuinely coupled pair with a full state matrix and a different control problem [11]. Neither vessel in this application is coupled in that sense.

3.3.2 Steady State and the Bias Term

Setting both derivatives to zero with $x_L = 0$ gives the resting levels of Equations (3.3) and (3.4).

$$\bar{h}_2 = \left[\frac{c_1\,\bar{u}/100 - q_d}{c_2}\right]^2 \tag{3.3}$$

$$\bar{h}_1 = (1 - x_v)^2 \left[\frac{c_1\,\bar{u}/100}{c_2}\right]^2 \tag{3.4}$$

The splitter fraction is absent from Equation (3.3) and appears squared in Equation (3.4). At steady state every unit of pump flow has to leave through the bottom drain no matter which path it took to get there, so the resting bottom level is a function of the pump position and the draw-off alone, while the resting top level falls as $(1 - x_v)^2$ and reaches zero when the bypass is wide open. Inverting Equation (3.3) at $q_d = 0$ gives the pump position that holds a given level, which is the bias of a positional control law and the value the integral term of Section 3.5 has to carry at rest, as shown in Equation (3.5).

$$\bar{u} = \frac{100\,c_2\sqrt{h_{2,\mathrm{SP}}}}{c_1} = 50\sqrt{h_{2,\mathrm{SP}}} \tag{3.5}$$

where the second form uses the default coefficients and returns a pump position in percent. Holding 0.50 m requires 35.36 % of pump travel whether the splitter is shut, half open or wide open, which states the splitter invariance of the operating point directly.

3.3.3 Overflow, Dry Tanks and the Instrument

Each tank stops at the rim. When a level has reached $h_{\max}$ and its balance is still positive, the balance is replaced by zero for that step and an overflow flag is raised for that tank, as shown in Equation (3.6).

$$\frac{\mathrm{d}h_i}{\mathrm{d}t} \leftarrow 0 \;\text{ if }\; h_i \geq h_{\max} \text{ and } \frac{\mathrm{d}h_i}{\mathrm{d}t} > 0, \quad i \in \{1,2\} \tag{3.6}$$

The spilled liquid leaves the model. It is not routed into the tank below and its rate is not reported. Both levels are additionally clipped to $[0, h_{\max}]$ after every integration substep, which is what enforces the dry limit, and at $h_i = 0$ the corresponding outflow term vanishes on its own because $\sqrt{0} = 0$. The draw-off does not vanish at $h_2 = 0$, so a dry bottom tank under draw-off is held at zero by the clip rather than by the physics.

Both transmitters add independent Gaussian noise to the true level and clip the result to the tank, as shown in Equation (3.7).

$$\hat{h}_i = \min(h_{\max},\ \max(0,\ h_i + \sigma_{n,i}\,\eta_i)), \qquad \sigma_{n,2} \equiv \sigma_n \tag{3.7}$$

where σ_n is the noise standard deviation of LT2, the transmitter the loop reads, $\sigma_{n,1}$ is that of LT1, each in m and set separately in the application, and η_i is an independent unit normal deviate obtained through a Box-Muller pair from the seeded generator of the simulation engine, redrawn for each of the two instruments on every controller execution. Nothing else is added. Neither instrument carries a lag, an offset or a drift, and its noise is white at the 2 Hz rate. The noise enters the readings only. The states h_1 and h_2, and with them the drains of Equations (3.1) and (3.2), are computed from the true levels. Only $\hat{h}_2$ is used, and the controller acts on it. The deviate η_1 is drawn first whatever $\sigma_{n,1}$ is, so $\hat{h}_2$, and every closed-loop result in this chapter, is independent of $\sigma_{n,1}$. Both readings are displayed and recorded beside the true levels, the top-tank trend above the bottom-tank trend as the tanks stand in the plant, so the error the loop reacts to and the indication LT1 gives can each be read against the state underneath it.

Table 3.2 Model, Controller and Numerical Parameters. Values in the First Two Blocks Are Adjustable in the Application.

Symbol	Meaning	Default	Range	Unit
Adjustable model parameters				
c_1	pump coefficient	0.08	0.02–0.20	m/s
c_2	drain coefficient	0.04	0.01–0.10	$m^{1/2}/s$
h_{max}	overflow height	1.0	0.5–2.0	m
σ_n	LT2 transmitter noise	0.004	0–0.020	m
$\sigma_{n,1}$	LT1 transmitter noise	0.004	0–0.020	m
Adjustable controller parameters				
K_c	proportional gain	300	0–1000	%/m
τ_I	integral time	60	0.05–500[a]	s
τ_D	derivative time	0	0–50	s
$h_{2,SP}$	level setpoint	0.50	0–h_{max}	m
T_p	MPC prediction horizon	300	40–600	s
λ	MPC move suppression	0.3	0–3	—
—	MPC estimator	Kalman	or MHE	—
—	speed multiplier	1	0.5–10	—
Fixed in the source				
u_{bias}	controller bias	$u(0)$	—	%
T_t	tracking time	τ_I	—	s
Δt	integration substep	0.05	—	s
Δt_c	controller sample period	0.5	—	s
T_s	MPC decision period	2	—	s
N_e	MHE window	40	—	executions
ε_r	relay hysteresis band	0.012	—	m
—	substeps per execution	10	—	—
—	historian rate	2	—	samples/s
—	noise seed	20260824	—	—
$h_i(0)$	initial levels	steady	—	m
A	tank area	—	—	divided out

[a] The top of the range, 500 s, switches integral action off.

3.3.4 Parameters and Numerics

Table 3.2 collects every constant the simulator needs, and the proportional gain moves in steps of 3 %/m. Every parameter edit is guarded before it reaches the model, c_1 and c_2 taking a lower bound of machine epsilon, h_{max} a lower bound of 0.1 m and both noise levels a lower bound of zero. A change of tank height re-clips both levels into the new geometry and moves the top of the setpoint range with it, pulling a setpoint above the new rim down to the rim. The three tuning parameters are bounded to their table ranges, and the top of the integral-time range switches integral action off.

Every case starts at steady operation rather than from empty tanks. In manual the two levels are set to the steady state of Equations (3.1) and (3.2) at the pump in force, and in automatic the pump that holds the setpoint is found by bisection on that steady state and the levels are set there, in both cases with the splitter, the draw-off, the leak and the parameters in force. A brim-full top tank passes on at most $c_2\sqrt{h_{max}}$, and a bottom tank whose inflow cannot meet the draw-off starts empty. A setting chosen before a run starts therefore changes the starting state instead of producing a step, and it gives the same run as the same setting carried through a restart. A step test is made after the run has started.

Equations (3.1) and (3.2) are integrated with explicit Euler on a fixed 0.05 s substep, with the overflow test of Equation (3.6) applied inside every substep and the $[0, h_{max}]$ clip applied after it. Nothing is gained from a higher-order integrator on this right-hand side, because the four clipping bounds contribute more error than the smooth dynamics do. This explicit Euler step replaces the framework's Runge–Kutta integrator. Every tenth substep, which is 0.5 s of simulated time, the

controller executes and the historian writes one row, on the simulated clocks of Section 1.6.1, so the trajectory at 10× is the trajectory at 1× and a noisy run repeats exactly. The export advances its time column in exact 0.500 s increments. A change to the pump, the splitter, the draw-off or the leak reaches the plant at the next controller execution, so every such step lands on the 0.5 s grid. One detail of the export matters for identification. A row stamped t carries the true levels at t and the measured level read at t, which is the reading the controller execution at t acts on, while its pump is the value held over $[t - \Delta t_c, t]$, the output of the execution at $t - \Delta t_c$. A fit of either level against the pump column therefore has to move the pump column back by Δt_c. Without that shift the shut-splitter test of Section 3.4 returns an apparent dead time of 18.0 s instead of 18.5 s.

3.3.5 Assumptions

The model keeps the mechanisms that change how the loop has to be tuned and discards the rest. Each entry below names what the simplification buys and the situation that breaks it.

1. *Non-interacting discharge.* The flow from TK-1 to TK-2 is $c_2\sqrt{h_1}$ and is independent of h_2, which requires the TK-1 outlet to discharge above the TK-2 liquid surface at all times. It fails as soon as TK-2 is flooded to the level of that nozzle, at which point the driving head becomes $h_1 - h_2$ and the pair becomes genuinely interacting [11]. Nothing in the source detects that condition, so the non-interacting form is used even when TK-2 is at the rim.

2. *Equal and constant tank areas.* The two balances share one c_1 and one c_2 and carry no area at all, which forces both tanks to have the same cross-section and the same drain geometry. Two vessels of different size would need separate coefficients, and a conical or horizontal cylindrical vessel would add a geometric nonlinearity on top of the hydraulic one.

3. *Quasi-steady outflow.* Each drain passes the steady Torricelli flow for whatever level stands above it at that instant, which gives the liquid in the outlet lines no momentum of its own. The approximation survives as long as a tank takes far longer to empty than its outlet column takes to accelerate, and the default constants satisfy that by a wide margin. A long discharge line, a fast valve stroke or any water-hammer question breaks it.

4. *Ideal splitter.* The split is a pure fraction applied to the pump flow, with no pressure drop, no installed characteristic, no stroking time, no hysteresis and no interaction between the two branches. The two paths of a real three-way valve have different resistances, so the split of a real valve depends on the levels it is feeding. The splitter also moves instantly, which makes every splitter experiment a true step.

5. *Ideal pump.* Flow is strictly proportional to the pump signal with no head dependence, no minimum flow, no slew limit and no deadband. A real centrifugal pump delivers less as the discharge head rises, which would couple the pump flow to h_1 and to h_2 and add a second interaction path that this model does not have.

6. *Overflow discards liquid.* Above $h_{\max}$ the excess leaves the model rather than cascading into the tank below or into a bund. A spill therefore breaks the mass balance, and its rate appears nowhere. What the flag does establish is that an inventory has escaped the loop, and that is all it establishes. Containment volume and spill duration lie outside what the model carries.

7. *Ideal transmitters apart from white noise.* Neither LT1 nor LT2 has lag, bias, drift, a resolution limit or a rangeability limit, and their noise is white, Gaussian and independent. The dominant errors in a real level measurement are slow ones, from a moving surface and from density, and they are correlated rather than white.

8. *Omitted entirely.* Inlet flow disturbances upstream of the pump, downstream pressure on either drain, vapor space pressure, temperature, composition, agitation, a cascade to a flow controller and any third vessel are absent. Each omission was made for the same reason, because none of these mechanisms changes the structural behavior the splitter produces.

3.4 Effect of the Splitter Fraction on Plant Order

3.4.1 Linearized State-Space Model

Taking first-order terms of Equations (3.1) and (3.2) around a resting point $(\bar{h}_1, \bar{h}_2, \bar{u})$, as Equation (1.4) does for one tank, produces the two-state deviation model of Equation (3.8), the state-space form of Equation (1.7).

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} h_1' \\ h_2' \end{bmatrix} = \begin{bmatrix} -1/\tau_1 & 0 \\ 1/\tau_1 & -1/\tau_2 \end{bmatrix}\begin{bmatrix} h_1' \\ h_2' \end{bmatrix} + \begin{bmatrix} c_1(1-x_v)/100 \\ c_1 x_v/100 \end{bmatrix} u' \tag{3.8}$$

in which a prime marks a departure from the resting point and the two time constants follow Equation (3.9).

$$\tau_1 = \frac{2\sqrt{\bar{h}_1}}{(1+x_L)\,c_2}, \qquad \tau_2 = \frac{2\sqrt{\bar{h}_2}}{c_2} \tag{3.9}$$

The zero in the upper right corner of the state matrix in Equation (3.8) is the whole content of the non-interacting claim, and it is a zero at every operating point rather than an approximation. The two poles of the plant are therefore $-1/\tau_1$ and $-1/\tau_2$ and are read straight off the diagonal. Each time constant is that of Equation (1.5) for one tank. Substituting Equation (3.4) into Equation (3.9) at $x_L = 0$ gives the central relation of the analysis, as shown in Equation (3.10).

$$\tau_1 = (1-x_v)\,\tau_2 \tag{3.10}$$

The upper time constant is the lower one scaled by the share of flow that still goes through the top tank. At the 0.5013 m operating point $\tau_2 = 35.40$ s, so shutting the splitter gives a repeated pole at –0.0282/s and opening it to 50 % moves one of the two poles to –0.0565/s while leaving the other exactly where it was. Opening it fully sends the upper pole to infinity, because an empty tank contributes no lag.

3.4.2 Pump and Splitter Transfer Functions

Eliminating h_1' from Equation (3.8) gives the transfer function from the pump to the controlled level, as shown in Equation (3.11).

$$G_p(s) = \frac{h_2'(s)}{u'(s)} = \frac{K_p\,(\tau_z s + 1)}{(\tau_1 s + 1)(\tau_2 s + 1)}, \qquad \tau_z = x_v\,\tau_1 \tag{3.11}$$

where $K_p = c_1\tau_2/100$ is the process gain in m/% of pump travel. With $x_v = 0$ it is the two non-interacting lags of Equation (1.10). Three properties of Equation (3.11) organize the results. The gain K_p contains no x_v at all, so the steady-state map is untouched by the splitter, as Equation (3.3) already showed. The numerator zero at $-1/\tau_z$ and the pole at $-1/\tau_1$ both travel outward together as the splitter opens, and at $x_v = 1$ they coincide and cancel exactly, leaving the first-order plant $K_p/(\tau_2 s + 1)$. Between the ends the zero is faster than the pole it partially cancels by the factor $1/x_v$, and combining Equations (3.10) and (3.11) gives $\tau_z = x_v(1-x_v)\tau_2$, which is largest at a half open splitter and reaches 8.85 s there. At the default operating point $K_p = 0.0283$ m/%, and that number is the same at $x_v = 0$, at $x_v = 0.5$ and at $x_v = 1$.

The splitter itself is a disturbance, and the transfer function it enters through is the reason the operating point does not move, as shown in Equation (3.12).

$$G_v(s) = \frac{h_2'(s)}{x_v'(s)} = \frac{c_1\bar{u}}{100}\,\frac{s}{(s+1/\tau_1)(s+1/\tau_2)} \tag{3.12}$$

Table 3.3 Manual Step Tests From 35.4 % to 45.4 % Pump at Three Splitter Positions, $\sigma_n = 0$, $q_d = 0$, $x_L = 0$. The Predicted Columns Come From Equations (3.4) to (3.11) at the Starting Point.

Splitter x_v	0 %	50 %	100 %
$\bar{h}_1$ before (m)	0.5013	0.1253	0.0000
$\bar{h}_1$ predicted (m)	0.5013	0.1253	0.0000
$\bar{h}_2$ before (m)	0.5013	0.5013	0.5013
$\bar{h}_2$ after (m)	0.8245	0.8245	0.8245
$\hat{K}_p$ (m/%)	0.0323	0.0323	0.0323
t_{28} (s)	42.4	21.1	13.7
t_{63} (s)	90.1	55.3	42.3
$\hat{\tau}_p$ (s)	71.6	51.3	42.9
$\hat{\theta}_p$ (s)	18.5	4.0	−0.6
τ_1 predicted (s)	35.40	17.70	0
τ_2 predicted (s)	35.40	35.40	35.40
τ_z predicted (s)	0	8.85	0
IMC K_c (%/m)	66.6	174.0	309.4
IMC τ_I (s)	71.6	51.3	42.9

The numerator of Equation (3.12) is a bare s, so the disturbance transfer function has a zero at the origin and its steady gain is exactly zero. A splitter step at a fixed pump position produces a transient in the bottom level and returns it to the level it started from with no offset and no control action of any kind. Starting from a shut splitter, where $\tau_1 = \tau_2 = \tau$, the response to a step of size Δx_v is $(c_1\bar{u}/100)\,\Delta x_v\, t\, e^{-t/\tau}$, which peaks at $t = \tau$ with a height of $(c_1\bar{u}/100)\,\Delta x_v\,\tau/e$. For the half open step at the default operating point that estimate is 0.184 m at 35.4 s, and the running application produces a peak of 0.176 m at 31.8 s before returning to 0.5013 m, which is the level it left. Two further disturbances act differently and are treated in Section 3.7, because the draw-off and the leak both move the steady state that the splitter leaves alone.

Pause. Try Exercise 3.1 before reading on.

3.4.3 Step Tests at Three Splitter Positions

Table 3.3 reports three step tests carried out in manual mode with the instrument noise set to zero, each raising the pump by 10 % from 0.5013 m at rest. The splitter position and a pump of 35.4 % are chosen before each run starts, so both tanks start at rest at that splitter position, and the pump is raised 10 s into the run. Times are counted from the controller execution at which the pump moves. The graphical fit takes the gain as the secant slope over the step, the time constant as $\hat{\tau}_p = 1.5\,(t_{63} - t_{28})$ and the apparent dead time as $\hat{\theta}_p = t_{63} - \hat{\tau}_p$, the two-point fit of Equation (1.18).

The three tests return the same gain to three decimal places and three quite different dynamic descriptions. The apparent dead time falls from 18.5 s with the splitter shut to 4.0 s half open and to −0.6 s wide open. The same construction applied to the exact first-order response gives −0.03 s, so the small negative value comes from the gain rising along the step rather than from any lead in the plant. The fitted time constant falls from 71.6 s to 42.9 s. The steady levels reached in the top tank, 0.5013, 0.1253 and 0.0000 m, match the $(1 - x_v)^2$ scaling of Equation (3.4) exactly. The vessels are unchanged across the three tests, and only the flow path differs.

Pause. Try Exercise 3.2 before reading on.

Two differences between the fitted and the predicted columns are worth naming, because both are informative rather than errors. The fitted gain 0.0323 m/% exceeds the linearized 0.0283 m/% by 14 %, because the step raises the bottom level from 0.5013 to 0.8245 m and the local gain $2c_1^2\bar{u}/(100\,c_2)^2$ grows with level, so the secant over a large step lies above the tangent at its start. The fitted time constant at the shut splitter, 71.6 s, exceeds the 58.1 s that the same graphical construction returns

when applied to the exact linear second-order response, for the same reason. Chapter 2 follows the same growth of gain and time constant with level in a single tank. The apparent dead time, on the other hand, is a structural quantity rather than an operating-point one, and the 18.5 s measured against the 17.8 s predicted from Equation (3.11) agrees to 4 %.

Pause. Try Exercise 3.3 before reading on.

Three distinct models have now appeared and none substitutes for another. Equations (3.1) and (3.2) are the simulation model, exact inside the assumptions of Section 3.3.5. Each column of Table 3.3 is an identified control model, good near one level at one splitter setting only. The PID of Section 3.5 holds no model, only a gain, two time constants, a tracking time and a sample period. The predictive controller of Section 3.5.4 carries one identified control model, fitted near 0.50 m with the splitter shut, and keeps it at every splitter setting unless it is given the table of Section 3.5.5, five such models fitted at fixed splitter settings.

> **TRY IT IN THE APP 2 · Open-Loop Actuator Step**
>
> Load the case and press Run. At 10 s the pump steps from 35.36 to 45.36 % with the controller in manual. Watch the two level trends. The upper tank rises at once, while the lower tank, fed only through it, starts slowly and has covered 63 percent of its rise only about 90 s after the step (Table 3.3).

3.5 Controller Design

3.5.1 Velocity Form

LIC-102 runs the velocity-form PID of Section 1.6.2 [15], written here in the loop's own symbols. Error is taken against the measured bottom level in Equation (3.13), and the three increments follow in Equations (3.14) to (3.16).

$$e_k = h_{2,\mathrm{SP}} - \hat{h}_{2,k} \tag{3.13}$$

$$\Delta u_k^P = P_k - P_{k-1}, \quad P_k = K_c\, e_k \tag{3.14}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\, e_k\, \Delta t_c \tag{3.15}$$

$$\Delta u_k^D = D_k - D_{k-1}, \quad D_k = -K_c \tau_D\, (\hat{h}_{2,k} - \hat{h}_{2,k-1})/\Delta t_c \tag{3.16}$$

with K_c in %/m, τ_I and τ_D in s and $\Delta t_c = 0.5$ s. Both P_k and D_k are formed with the gains in force at scan k. Equation (3.17) combines the three increments and clips the sum to the actuator span.

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= u_{\mathrm{bias}} + s_k + I_{k-1} \\ u_k &= \min(100,\ \max(0,\ c_k)) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{3.17}$$

In the symbols of Equations (1.24) to (1.27), $s_k = S_k$, $c_k = v_k$, $u_{\mathrm{bias}} = u_b$ and the action is $a = +1$. Here s_k accumulates the proportional and derivative increments without limit and telescopes to $P_k + D_k$ exactly, c_k is the output the algorithm demands, I_k is the integral state and $T_t = \tau_I$ is the tracking time of the back-calculation. The integral used at scan k was accumulated through scan $k-1$. The bias u_{bias} is the pump in force when the run starts, which is the pump that holds the starting steady state, and the integral starts from zero, so u_{bias} carries the starting operating point and I_k every change from it. Away from the limits $u_{k-1} = c_{k-1}$, and differencing the second line gives $c_k = u_{k-1} + \Delta u_k^P + \Delta u_k^D + \Delta u_{k-1}^I$, the velocity law with the integral increment of the previous scan. The recursion is therefore algebraically identical to the positional law $u = P + I + D$ only while the pump stays inside its range. The back-calculation runs whenever $K_c/\tau_I \neq 0$. Setting τ_I to the top of

its range, 500 s, switches integral action off, so neither Δu_k^I nor the back-calculation is applied, I_k keeps its last value and the law becomes proportional, or proportional-derivative, control about the fixed bias $u_{\text{bias}} + I$.

As in every framework app, the derivative acts on the measurement and a gain edit acts at once on the whole proportional term. A setpoint move therefore steps the pump by $K_c \Delta h_{2,\text{SP}}$ with no derivative kick, and a gain edit steps it by $\Delta K_c \, e_k$ and leaves the integral untouched. The controller is reverse acting, because raising the pump raises the level, so the pump has to fall as the measured level rises, and no action switch is offered. The gains are entered as K_c, τ_I and τ_D, so the integral time of a tuning correlation is typed in directly and the integral gain K_c/τ_I follows from it.

The gain units require attention, because K_c is expressed in percent of pump travel per meter of level and is therefore a number in the hundreds. A tuning quoted as $K_c = 2.0$ and $K_I = 0.03$ for a pump signal on a zero to one scale is $K_c = 200\,\%/\text{m}$ and $\tau_I = 66.7\,\text{s}$ in the units of Table 3.2. The two are the same controller, although the proportional gain moves in steps of 3 %/m, so 201 %/m is the nearest setting.

3.5.2 Bias, Transfer and Setpoint Edits

The velocity form carries its operating point in the bias and the integral state, so where the recursion starts decides them. In manual the integral is overwritten so that the demand equals the manual pump, as shown in Equation (3.18).

$$I_k = u_{\text{man},k} - u_{\text{bias}} - s_k \qquad \text{(manual)} \tag{3.18}$$

Both transfers between manual and automatic are therefore bumpless, as Section 1.6.2 explains. A run that starts in automatic starts from the steady state of its setpoint with u_{bias} at the pump that holds it and the integral at zero, so nothing moves before the first scripted event except the response to transmitter noise.

Editing the setpoint, the integral time or the proportional gain during a run changes the next increments and nothing else, so integral action built up against a standing load survives the edit. A loop with $K_c = 66\,\%/\text{m}$ and $\tau_I = 71.65\,\text{s}$ holding 0.500 m against a 0.010 m/s draw-off sits with the pump at 47.86 %, the 35.36 % bias of its starting steady state plus 12.50 % carried by the integral. Raising the setpoint by 0.01 m moves the pump to 48.52 % in one sample, which is the proportional step $K_c \Delta h_{2,\text{SP}} = 0.66\,\%$, and the level climbs to the new setpoint without first falling below 0.500 m.

3.5.3 Saturation and Windup

Anti-windup is the back-calculation term of Equation (3.17) [15]. While the pump is clipped, $(\Delta t_c/T_t)(u_k - c_k)$ drives I toward the value that reproduces the applied pump with a tracking time of τ_I, so the integral alone is corrected and the running sum s_k is never truncated. Setting the integral update to zero on a limit gives $c_k = u_k + K_c e_k$, so with $\tau_D = 0$ the sum $u_{\text{bias}} + I$ settles at the limit value itself and the demand returns inside the range as soon as the error changes sign. Held on the 100 % limit for 700 s by the 0.020 m/s draw-off of Section 3.7, the bias and the integral together reach 99.98 % while the demand stands at 116.5 %. Removing the draw-off releases the pump 16.5 s later, at the first execution after the level crosses the setpoint. That sum then still has to unwind from 100 % to the 35.4 % the unloaded plant needs, the pump stays above 32.5 %, and TK-1 stays at its rim for 104 s, so the bottom level climbs to 0.940 m before it returns. Back-calculation bounds the integral by the actuator range, and a load that holds the pump on a limit for longer than τ_I still leaves a full limit's worth of integral to unwind.

3.5.4 Model Predictive Control

A model predictive controller (MPC), the option of Section 1.6.3, can take the pump in automatic mode instead of the PID. Every case except cases 9 and 10 starts with the PID, and while the MPC acts the PID follows the applied pump so that either can hand over to the other. Every $T_s = 2\,\text{s}$, four controller executions, the MPC computes a pump trajectory over a horizon, sends its first move

and keeps it until the next computation, the receding-horizon arrangement of industrial predictive control [16] whose constrained form is analyzed in [17]. It plans on two lags in series about the design point $\bar{u}_m = 35.36\,\%$, $\bar{h}_m = 0.50\,\mathrm{m}$, as shown in Equation (3.19).

$$\begin{aligned} \tau_a \frac{\mathrm{d}z_1}{\mathrm{d}t} &= -z_1 + K_m\,(u - \bar{u}_m + d) \\ \tau_b \frac{\mathrm{d}z_2}{\mathrm{d}t} &= z_1 - z_2, \qquad h_m = \bar{h}_m + z_2 \end{aligned} \tag{3.19}$$

Here z_1 and z_2 are model states in m, h_m is the predicted bottom level, d is an unmeasured disturbance in percent of pump travel, $K_m = 0.028\,28\,\mathrm{m}/\%$, $\tau_a = 39.7\,\mathrm{s}$ and $\tau_b = 31.1\,\mathrm{s}$. The constants are a least-squares fit to the mean response to pump steps of ± 3 percentage points from rest at the design point, with the splitter shut, no load and the noise off, over 600 s. The residual is $7.8 \times 10^{-6}\,\mathrm{m}$ root mean square, under 0.01 % of the response, the fitted dead time of 0.14 s is dropped, and a first-order-plus-dead-time fit leaves 143 times that residual. At $x_v = 0$ Equation (3.11) gives the same gain and two lags of 35.36 s, whose 70.71 s sum the fit reproduces to 0.02 s. The model is the shut-splitter plant only, with no splitter, no zero, no top-tank level, no rim and no square-root drain.

Offset-free tracking comes from a steady-state Kalman filter on Equation (3.19), sampled with a zero-order hold at Δt_c and augmented with d as a random walk, the disturbance model of offset-free predictive control [18]. At every execution, whichever controller holds the pump, it updates as in Equation (3.20).

$$\begin{aligned} \bar{z}_k &= \Phi\,\hat{z}_{k-1} + \Gamma\left(u_{k-1} - \bar{u}_m + \hat{d}_{k-1}\right) \\ \varepsilon_k &= \hat{h}_{2,k} - \bar{h}_m - \bar{z}_{2,k} \\ \hat{z}_k &= \bar{z}_k + \ell_z\,\varepsilon_k, \qquad \hat{d}_k = \hat{d}_{k-1} + \ell_d\,\varepsilon_k \end{aligned} \tag{3.20}$$

where Φ and Γ are the sampled matrices of the two lags, u_{k-1} is the pump applied over the period just ended, $\hat{h}_{2,k}$ is the reading the PID also receives, and $\ell_z = [1.615,\ 0.202]^{\mathsf{T}}$ and $\ell_d = 315.9\,\%/\mathrm{m}$ are the converged filter gains obtained with a measurement variance $r = \sigma_n^2$, a state noise variance $10^{-6} r$ and a disturbance noise variance $q_d = 100\,r/K_m^2 = 2.00\,\%^2$. A load the model lacks, including an error in its own dynamics, appears in $\hat{d}$ as an equivalent pump offset that the plan holds constant, so a constant load leaves no offset. A 1 mm prediction error moves $\hat{d}$ by 0.32 %.

Each decision minimizes a quadratic cost over predictions of Equation (3.19) discretized at T_s, across $T_p = 300\,\mathrm{s}$, or $N = 150$ steps, about four times the sum of the lags. Twelve pump positions v_j, held over blocks of $1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15$ and 105 steps, minimize Equation (3.21) subject to the hard bounds $0 \le v_j \le 100$.

$$J = \frac{T_s}{\tau_n} \sum_{k=1}^{N} \left(\frac{h_{m,k} - h_{2,\mathrm{SP}}}{y_s} \right)^2 + \lambda \sum_{j=1}^{12} \left(\frac{v_j - v_{j-1}}{u_s} \right)^2 \tag{3.21}$$

In it $h_{m,k}$ is the level predicted k steps ahead with $\hat{d}$ held, v_0 is the pump applied now, the factor T_s/τ_n with $\tau_n = 70.7\,\mathrm{s}$ makes the first sum a time integral, $y_s = 0.03\,\mathrm{m}$ is the scorecard settling band, $u_s = 100\,\%$ and $\lambda = 0.3$. There is no level band, rim, rate limit, terminal weight or preview, because the controller is told neither the draw-off nor the leak, and the splitter fraction only through the schedule of Section 3.5.5. The decisions are departures from $\bar{u}_m$ with a ridge of 10^{-9} on the scale of the weights, solved by the dual active-set algorithm of [19]. The values of λ and q_d follow one rule for the seven single-loop studio apps, the largest λ within 10 % of the best true-level IAE on cases 3 and 4 among settings that move the pump no more than the shipped PID, and for this plant the rule takes the largest q_d of its grid.

The answer of the solver is discarded when it is not optimal, when its first move is not a finite number, or when that move lies outside 0 to 100 %. The PID, which tracks the applied pump as it tracks the operator in manual, then takes over without a bump until a valid plan returns. A rejection

Table 3.4 The Scheduled Model's Entries, Fitted at 0.50 m and $K_m = 0.028\,28\,\mathrm{m}/\%$, Against Two Lags Without the Parallel Path and the Hand Linearization

x_v (%)	τ_a (s)	τ_b (s)	σ	RMS (µm)	RMS_2 (µm)	$(1-x_v)\tau_2$ (s)
0	39.67	31.05	0	7.8	—	35.36
25	25.01	36.77	0.257	5.4	443	26.52
50	16.91	35.90	0.504	5.4	411	17.68
75	8.19	35.55	0.748	9.2	162	8.84
100	8.19[a]	35.38	1	17.5	—	0

RMS_2 is the residual of two lags in series without the parallel path. The 0 % entry is Equation (3.19). [a]Borrowed from the 75 % entry.

forced at 36 s of case 9 without noise changes the pump by 1.9 %, the size of the next PID moves, and the fresh plan that takes it back moves it by 49 %. A plan is already computed while a case loads, so an MPC run is logged from $t = 0$. The controller reads only the LT2 reading, the applied pump, the setpoint, the pump limits and, when scheduled, the splitter setting, which keeps the plant of Equations (3.1) and (3.2), the instrument of Equation (3.7) and the model of Equation (3.19) apart.

3.5.5 Model Scheduled on the Splitter Setting

The splitter is a hand valve that the operator sets and no loop moves, so the controller may be told its position, as a lineup entry would tell it, without reading a process state. With the scheduled model chosen, the plan and the filter use the entry of a five-model table nearest the told setting, a multiple-model predictive controller [20]. Each entry passes a share σ of the input around the first lag of Equation (3.19), as shown in Equation (3.22).

$$\begin{aligned} \tau_a \frac{\mathrm{d}z_1}{\mathrm{d}t} &= -z_1 + K_m(1-\sigma)\,v, \qquad v = u - \bar{u}_m + d \\ \tau_b \frac{\mathrm{d}z_2}{\mathrm{d}t} &= z_1 - z_2 + K_m \sigma\, v \end{aligned} \tag{3.22}$$

Its transfer function is Equation (3.11) with $\tau_z = \sigma\tau_a$, so the balances linearized at 0.50 m give $\sigma = x_v$, $\tau_a = (1-x_v)\tau_2$ and $\tau_b = \tau_2 = 35.36\,\mathrm{s}$, with K_m unchanged because the splitter moves the path and not the steady level. Each entry was fitted as Equation (3.19) was, with its splitter set before the run, as listed in Table 3.4. The fits return σ within 0.007 of x_v, $K_m = 0.028\,28\,\mathrm{m}/\%$ throughout and lag sums within 1.1 % of the hand values. At 100 % a pump step cannot see τ_a, so the 75 % entry's value stands in, and replacing it by 0.5 or 17.7 s moves the IAE of five runs that enter that entry by at most 0.09 m s.

Weights, horizon, blocks and filter variances are those of Section 3.5.4, each entry with its own filter gain. The update over the period just ended uses the entry then in force, the new entry takes over before the plan, and $\hat{z}$, $\hat{d}$ and the input record carry over, so neither the estimated level nor the pump moves at a change. The carried $\hat{z}_1$ changes meaning, because each entry counts its upper path from the outflow of its own split, a shift of about $K_m\bar{u}_m\Delta\sigma$ (0.50 m at the 0 to 50 % opening) that the filter books as load, $\hat{d}$ reaching 26.0 % 5 s after the opening of case 10 without noise and 2.3 % after 35 s. Re-mapping $\hat{z}_1$ by script cuts the IAE of case 10 from 1.53 to 1.35 m s but takes TK-1 to its rim when the splitter goes to 50, 100 and 0 %, so it was not adopted. Scheduled plans must be optimal, with finite moves, predictions and cost and every move inside 0 to 100 %, a rejection is followed by a three-decision return, and an MHE restarts its window on the new entry from the filter's prediction.

3.5.6 Moving Horizon Estimation

The state from which each plan starts may instead come from a moving horizon estimator (MHE), while the filter of Equation (3.20) keeps running and logging under either choice. The MHE works on the augmented state $\xi = [z_1\ z_2\ d]^{\mathsf{T}}$ of that filter and, at execution k, fits the readings $j = s, \ldots, k$ with $s = k - N_e$ through the bounded least squares of Equation (3.23), the constrained estimator analyzed

in [21].

$$\begin{aligned} \min \; & \|\xi_s - \bar{\xi}_s\|^2_{\bar{P}^{-1}} + \sum_{j=s}^{k-1} \omega_j^{\mathsf{T}} Q^{-1} \omega_j \\ & + \sum_{j=s}^{k} \frac{\left(\hat{h}_{2,j} - \bar{h}_m - z_{2,j}\right)^2}{r} \\ \text{s.t. } & \xi_{j+1} = \Phi_e \, \xi_j + \Gamma_e \, (u_j - \bar{u}_m) + \omega_j \\ & 0 \le \bar{h}_m + z_{i,j} \le h_{\max,0}, \qquad i = 1,2 \end{aligned} \tag{3.23}$$

The decisions are ξ_s and the noise ω_j on all three states, $\Phi_e = \left[\begin{smallmatrix} \Phi & \Gamma \\ 0 & 1 \end{smallmatrix}\right]$ and $\Gamma_e = [\Gamma^{\mathsf{T}} \; 0]^{\mathsf{T}}$ repeat the prediction of Equation (3.20), u_j is the pump applied after reading j, $Q = \mathrm{diag}(10^{-6} r, 10^{-6} r, q_d)$ and r are the variances that designed the filter, $\bar{\xi}_s$ is the estimate the MHE made at execution $s-1$ advanced one step, and the steady prior covariance $\bar{P}$ of the filter is held fixed as the arrival cost. With $N_e = 40$ the window spans 20 s and 41 readings, which gives the active-set solver of Section 3.5.4 123 variables and 164 bound rows. Both model levels stay between the floor and $h_{\max,0} = 1\,\mathrm{m}$, the declared rim rather than the one set for a run, and d is free because no declared value limits it. Bounding z_1 treats the slower fitted lag as the top tank, whereas the plant has two equal lags of 35.36 s. With no bound active the problem is the filter's own, and a first window starts from the filter's prediction. A window left without an optimal, finite answer inside the bounds gives way to a Kalman measurement update of the MHE's one-step prediction, clipped into the bounds, and the history marks the substitution.

3.5.7 Relay Autotuning

PI constants for the PID can also come from a limit cycle that the loop produces under the relay feedback test of [22], an action that no shipped case starts with. Taken from the PID in automatic, a relay sets the pump to $u_0 \pm \delta_r$, with u_0 the mean pump of the last 30 s and δ_r equal to 5 % of the span, or to 0.9 of the gap between u_0 and the nearer limit when that is smaller. It switches side whenever the LT2 reading passes the setpoint by more than $\varepsilon_r = 3\sigma_{n,0} = 0.012\,\mathrm{m}$, three times the declared noise $\sigma_{n,0}$, and a cycle with a half under four executions is discarded. Once the period and the swing $\pm\alpha$ of two successive accepted cycles agree to within 5 %, the first harmonic of the relay output, corrected for the hysteresis, gives the ultimate gain K_{cu}, and the Tyreus–Luyben PI constants follow from K_{cu} and the period P_u as in Equation (3.24).

$$K_{cu} = \frac{4\delta_r \sin(\pi D)}{\pi \sqrt{\alpha^2 - \varepsilon_r^2}}, \qquad K_c = \frac{K_{cu}}{3.2}, \qquad \tau_I = 2.2\,P_u \tag{3.24}$$

Here D is the fraction of each cycle that the relay spends on its high side, near one half only when u_0 is the pump that holds the setpoint. The constants are written to three figures with $\tau_D = 0$ and stay editable. A test stops unconverged after 30 cycles or 600 s or on a second request, and the PID, which follows the applied pump, resumes without a bump. The angle $\psi = \arcsin(\varepsilon_r/\alpha)$ separates the cycle from -180° because of the hysteresis, and from $\psi = 10^\circ$ upward the result is flagged as a cycle placed by the band and not by the tanks.

3.6 Results

The app ships eleven cases, and every one starts at steady operation with the bottom level at 0.50 m, so the first scripted event of each, at 10 s, acts on a settled plant. Case 1 holds the reference operating point in manual with the pump at 35.36 %, and case 2 steps the pump from 35.36 to 45.36 %. Case 3 steps the setpoint to 0.60 m under the shipped tuning, case 4 applies a draw-off of 0.006 m/s in automatic and case 5 the same draw-off in manual. Cases 6 and 7 step the setpoint to 0.58 m at $K_c = 120$ and 600 %/m. Case 8 steps it to 0.90 m on a smaller pump ($c_1 = 0.035\,\mathrm{m/s}$, 80.8 % at

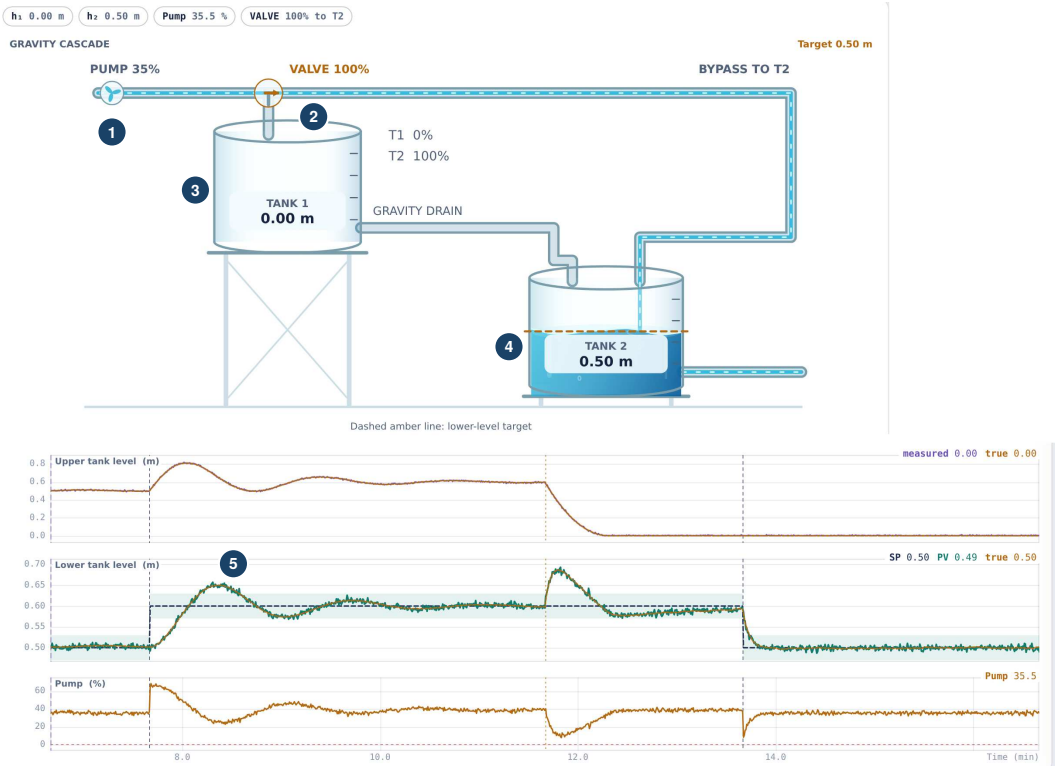


Figure 3.2 One controller meeting both plant structures in a single run built from case 1, shown over the last 10 min of 1000 s with the shipped $K_c = 300\,\%/\mathrm{m}$, $\tau_I = 60\,\mathrm{s}$ and $\tau_D = 0$ and both transmitter noises at their 0.004 m default (levels quoted are true levels). With the pump set to 35.4 % before the run starts, the bottom level rests at 0.501 m until automatic is engaged at 400 s. With the splitter shut the setpoint step from 0.500 to 0.600 m at 460 s rings through peaks of 0.650 and 0.613 m and troughs of 0.574 and 0.593 m and drives the top tank to 0.810 m. Opening the splitter to 100 % at 700 s drains the top tank and lifts the bottom level 88 mm in 7.95 s before the loop recovers it, and the step back to 0.500 m at 820 s on the first-order plant settles inside 0.005 m in 9.75 s, its lowest point 3.8 mm below the setpoint. The trends draw these levels without printing them, and the values are those of the same run. The picture is a composite of the application's window, the plant view above its three trend strips, with the header and the side panels cut. Callouts: (1) pump, (2) splitter valve, (3) top tank TK-1 and its gravity drain, (4) bottom tank TK-2 with the dashed target line, (5) the bottom level ringing after the setpoint step with the splitter shut.

0.50 m). Cases 9 and 10 run the predictive controller of Section 3.5.4, and case 11 is the graded challenge of Section 3.8. The experiments of this section and the next are built from case 1 with the settings chosen before Run, and Table 3.7 collects them.

The two extreme rows of Table 3.3 were converted into proportional-integral tunings with the internal model control correlation [23], with $\tau_c = \max(0.1\hat{\tau}_p, 0.8\hat{\theta}_p)$, $K_c = \hat{\tau}_p/[\hat{K}_p(\tau_c + \hat{\theta}_p)]$ and $\tau_I = \hat{\tau}_p$. The shut-splitter fit gives $\tau_c = 14.8\,\mathrm{s}$ and a slow tuning of $K_c = 66.6\,\%/\mathrm{m}$ and $\tau_I = 71.6\,\mathrm{s}$. The open-splitter fit has no dead time to protect against, so with the small negative fit taken as zero τ_c collapses to $0.1\hat{\tau}_p = 4.3\,\mathrm{s}$ and the fast tuning is $K_c = 309.4\,\%/\mathrm{m}$ and $\tau_I = 42.9\,\mathrm{s}$. The proportional gains differ by a factor of 4.6 and the integral gains K_c/τ_I by a factor of 7.7, and the process gain they were both derived from is identical. On the 3 %/m grid of the gain field the two tunings are entered as $K_c = 66$ with $\tau_I = 71.65\,\mathrm{s}$ and $K_c = 309$ with $\tau_I = 42.95\,\mathrm{s}$, and those are the values run below.

Table 3.5 The Same Setpoint Step From 0.500 to 0.600 m Under Two Tunings and Two Splitter Positions, $\sigma_n = 0$. Settling Is the Last Crossing of a ±0.005 m Band and Travel Is $\sum |u_k - u_{k-1}|$ Over the 400 s After the Step.

	slow tuning $K_c = 66$, $\tau_I = 71.65$		fast tuning $K_c = 309$, $\tau_I = 42.95$	
Splitter x_v	0 %	100 %	0 %	100 %
Peak h_2 (m)	0.6078	0.5996	0.6587	0.6000
Overshoot (%)	7.8	none	58.7	none
Peak h_1 (m)	0.6347	0	0.8354	0
IAE (m s)	4.255	3.580	4.950	0.419
Pump travel (%)	12.4	10.3	144.0	58.4
Settling (s)	117.9	159.7	210.8	12.3
Damped period (s)	—	—	79.8	—
Predicted ζ	0.532	real poles	0.199	real poles

Table 3.5 reports each tuning driving each structure through the same setpoint step from 0.500 to 0.600 m, with the transmitter noise set to zero so that the manipulated-variable travel measures control action rather than noise. Each run is case 3 with the tuning, the splitter and the noise changed before it starts, so the loop starts in automatic at rest at 0.500 m and the setpoint moves at 10 s. The predicted damping ratio is that of the dominant closed-loop pole pair of the linear model of Equation (3.11) in feedback with $K_c[1 + 1/(\tau_I s)]$.

The diagonal of Table 3.5 pairs each tuning with the structure it was identified on, and the off-diagonal entries quantify the cost of the mismatch. The fast tuning on the plant it was derived from settles the step inside 0.005 m in 12.3 s with an IAE of 0.419 m s and no overshoot, because the proportional-integral zero at −0.0233/s very nearly cancels the closed-loop pole at −0.0228/s and leaves a single closed-loop time constant of 3.96 s. Moving the same controller to the shut-splitter plant without changing a gain produces 58.7 % overshoot, five peaks above the setpoint inside the window, three of them outside the ±0.005 m band, on a damped period of 79.8 s, and 144.0 % of pump travel against the 58.4 % it spent on the fast plant. The linear model predicts that outcome closely, with a dominant pole pair at $\zeta = 0.199$ and a damped period of 76.6 s. The added lag has cost the loop essentially all of its phase margin, which is the result the effective-delay literature predicts for any controller of any structure once a dead time appears [4, 7].

The reverse mismatch produces no visible symptom and a larger performance loss. The slow tuning on the fast plant never overshoots and never rings, and it takes 159.7 s to settle against the 12.3 s the fast tuning needs, an IAE of 3.580 m s against 0.419 m s. A factor of 8.5 in integrated error is invisible on a trend unless the two runs are put side by side, which is why the slow tuning survives on a plant that could be controlled eight times more tightly. Figure 3.2 shows the first mismatch and the well-matched case produced by one controller in one run. Figure 3.3 draws four rows of Table 3.5 as bars.

One further measurement does not appear in the table. The fast tuning on the shut-splitter plant drives the top tank to 0.8354 m, which is 84 % of the rim, while moving the bottom tank by 0.100 m. The overshoot that a level trend shows is in the controlled vessel, but the inventory that comes closest to spilling is in the uncontrolled vessel.

The remedy that the structure itself suggests is to schedule the tuning on the splitter position, which is a measured variable and moves slowly. Volume 2, Chapter 6 schedules a PI loop on a gain that moves four orders of magnitude. Both entries in the IMC row of Table 3.3 come from the same correlation applied to the same plant at two settings of one valve, and interpolating them on x_v is the textbook shape of a scheduled controller. The scheduling literature is equally clear about what makes this case easy and the general case hard, because the scheduling variable here is an exogenous input

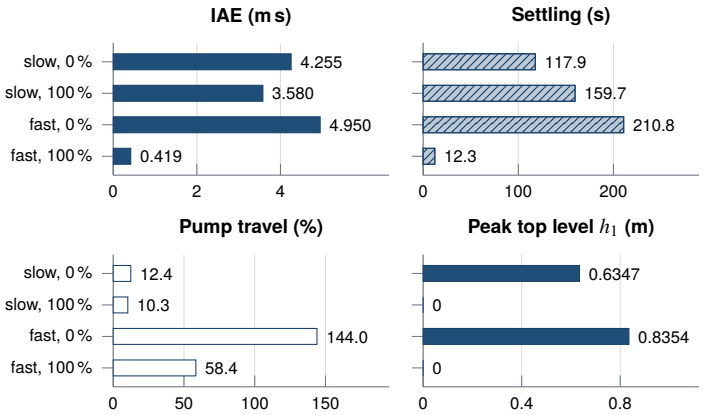


Figure 3.3 The setpoint step of Table 3.5 under the slow tuning (identified with the splitter shut) and the fast tuning (identified with it wide open), each run at splitter settings of 0 and 100 %: integrated absolute error and peak top-tank level (solid bars), settling time (hatched) and pump travel (open). Notice that each tuning settles fastest on the structure it was identified on, and that the fast tuning on the shut splitter spends the most pump travel and lifts the top tank highest, a level no loop controls.

Table 3.6 Load Disturbance Responses From a Settled 0.500 m With $K_c = 66\,\%/\text{m}$, $\sigma_n = 0$, Measured Over the 700 s After the Upset

Upset	x_v (%)	τ_I (s)	Minimum h_2 (m)	Final h_2 (m)	Final u (%)
$q_d = 0.010$	0	71.65	0.3345	0.5000	47.86
$q_d = 0.010$	0	off	0.3273	0.3796	43.30
$x_L = 100\%$	0	71.65	0.2824	0.4988	70.66
$q_d = 0.020$	0	71.65	0.1818	0.2500	100.0
$q_d = 0.020$	100	71.65	0.3020	0.4998	60.35

rather than a state, and the plant is linear in every respect except the square-root drains [8, 9]. The application schedules the model of its predictive controller on the splitter rather than the PID's tuning (Section 3.5.5), and Section 3.8 measures both against the shipped loop.

3.7 Load Disturbance Rejection

The other two disturbances behave in the way the splitter does not, because both of them move the steady state. Table 3.6 reports the responses, all with the splitter shut unless stated, the transmitter noise at zero and the slow tuning of Section 3.6 in automatic. Each run starts at rest at 0.500 m with automatic chosen before it starts, and the upset is applied at 10 s. The draw-off rows use a larger upset than case 4, whose 0.006 m/s keeps the top tank below its rim (Section 3.7.1).

3.7.1 Draw-Off Response and Proportional Offset

A draw-off of 0.010 m/s is 35 % of the 0.0283 m/s that the bottom drain passes at 0.5 m. Left uncontrolled with the pump held at the 35.36 % that held 0.500 m, Equation (3.3) puts the resting level at 0.2089 m, a fall of 0.291 m. Under proportional-only control with the bias left at that value, the loop would settle where the positional law and Equation (3.3) agree, which is the root of $c_1(\bar{u} + K_c e)/100 = c_2\sqrt{h_{2,\text{SP}} - e} + q_d$. At $K_c = 66\,\%/\text{m}$ that root is $e = 0.1204\,\text{m}$ at a pump position of 43.30 %, so proportional feedback alone would recover 59 % of the uncontrolled excursion and leave the rest as a standing offset, which is the classical result. The application runs that case directly, because the top of the integral-time range switches integral action off and a run started in

automatic keeps the 35.36 % bias of its starting steady state. The level falls to 0.3273 m and settles at 0.3796 m with the pump at 43.30 %, the predicted root to four decimal places, and the 0.120 m offset is unchanged 1500 s after the upset. The slow tuning's $\tau_I = 71.65$ s removes the offset entirely and takes the pump to 47.86 %, the value that Equation (3.3) requires.

What the successful rejection costs is not visible in the controlled vessel. Holding 0.500 m against that draw-off requires the top tank to sit at 0.9160 m by Equation (3.4), which is 92 % of the rim, and the transient carries it to 0.9319 m. With the shipped tuning, $K_c = 300\,\%/\text{m}$ and $\tau_I = 60$ s, the same draw-off carries TK-1 to the rim, where it spills for 31 s. The disturbance was applied to the bottom tank and the vessel it fills is the top one. Case 4 therefore applies 0.006 m/s, which needs TK-1 at 0.735 m, and the shipped tuning's transient peaks at 0.849 m (0.900 m even at $K_c = 1000\,\%/\text{m}$). The bottom level dips to 0.440 m at 28.5 s and is back within ±0.03 m from 44.8 s. Recovered by hand in case 5, the same draw-off is met by a pump of 42.9 % set at 10 s, and that setting completes the case at 123 s.

3.7.2 Leak Response

A 100 % leak doubles the discharge coefficient of the top tank and throws half of its drain away, so at a shut splitter the pump has to deliver twice as much to pass the same flow to TK-2. Equation (3.1) at steady state gives $u = 200 c_2 \sqrt{h_{2,\text{SP}}}/c_1 = 70.71\,\%$, exactly twice the undisturbed bias, and the loop reaches 70.66 % and 0.4988 m within 700 s while still converging. The leak also shortens τ_1 by the factor $1 + x_L$ through Equation (3.9), so a leak shortens the loop time constant while raising the pump position the same level requires. At a fully open splitter the top tank is empty and the leak has no effect at all, which is a compact demonstration that a disturbance path can be removed by a piping change rather than by a controller.

> **+ DID YOU KNOW? A Bypass That Lifts the Reachable Level**
>
> From the simulator: against a draw-off of 0.020 m/s, the shut splitter leaves the pump at 100.0 % and the bottom level at 0.2500 m. With the splitter wide open, the same tuning holds 0.4998 m with the pump at 60.35 % (Table 3.6).

3.7.3 Reachable Level Under Draw-Off

The last two rows of Table 3.6 are the same 0.020 m/s draw-off applied at the two ends of the splitter, and they do not differ in speed. They differ in whether the setpoint is reachable at all. With the splitter shut, all flow entering TK-2 must pass through TK-1, whose discharge is bounded by $c_2\sqrt{h_{\max}}$ once the top tank is brim-full. Substituting that bound into Equation (3.2) gives the highest bottom level a shut splitter can sustain against a draw-off, as shown in Equation (3.25).

$$h_2^{\max}\big|_{x_v=0} = \left[\frac{c_2\sqrt{h_{\max}} - q_d}{c_2}\right]^2 \tag{3.25}$$

At $q_d = 0.020\,\text{m/s}$ that ceiling is 0.2500 m, and the application settles at 0.2500 m with the pump parked on 100 % for 71 % of the 700 s and the top tank at the rim and overflowing for 93 % of it. The limit belongs to the plant rather than to the tuning. Opening the splitter to 100 % routes the whole pump flow around the bottleneck, and the identical disturbance is then rejected completely, to 0.4998 m after 700 s and 0.5000 m after 900 s, with the pump at 60.35 % against the 60.36 % of Equation (3.3) and the top tank empty. The IAE over the 900 s after the upset falls from 226.0 m s to 27.1 m s. The splitter therefore changes two things that are usually thought about separately, the order of the plant and the size of the load that the plant can reject, and it changes neither the process gain nor the nominal operating point. Volume 2, Chapter 3 meets a bound of the same kind in a composition that the upstream supply limits.

Table 3.7 Experiment Sequence Built From Case 1 and the Mechanism Each One Exposes

Experiment	Setup	What it exposes
Splitter bump	manual, pump fixed, x_v 0 to 50 %	zero steady gain of Equation (3.12), 0.176 m transient
Step test ×3	manual, 10 % pump step at $x_v = 0$, 50, 100 %	one $\hat{K}_p$, three $(\hat{\tau}_p, \hat{\theta}_p)$ pairs
Cross tuning	IMC from each end, run at both ends	58.7 % overshoot one way, 8.5× the IAE the other
Draw-off, P only	$q_d = 0.010$, integral off	the predicted 0.120 m standing offset
Draw-off, PI	same upset, $\tau_I = 71.65$ s	offset removed, TK-1 at 92 % of its rim
Leak	$x_L = 100$ % at $x_v = 0$, then at $x_v = 1$	a disturbance path removed by piping
Draw-off ceiling	$q_d = 0.020$ at both splitter ends	0.25 m ceiling and pump saturation
Fill and drain	from empty, pump 100 % then 0 %	overflow, dry limit and $2\sqrt{h}/c_2$

The pump sets a second limit that no tuning can move. In case 8 the smaller pump sits at 100 % from the execution that sees the step to 0.90 m, and the level climbs only to 0.763 m in 300 s, toward the full-pump level $(c_1/c_2)^2 = 0.766$ m of Equation (3.3).

3.7.4 Filling and Draining Against the Limits

Two manual runs bound the apparatus. A pump of 0 % chosen before the run starts puts both tanks at rest empty. With the pump raised to 100 % 10 s into the run and the splitter shut, TK-1 reaches the 1.0 m rim 19.3 s after the pump starts, matching the 19.3 s obtained by integrating Equation (3.1) in closed form, and it spills from then on. TK-2 rises toward the rim but does not reach it, holding 0.9979 m after 300 s, because with the splitter shut its only inflow is the discharge of a tank that is itself pinned at $h_{\max}$, so h_2 approaches h_1 asymptotically from below and can never exceed it. Any splitter opening at all removes that protection, because the bypass delivers $c_1 x_v \bar{u}/100$ into TK-2 on top of the TK-1 discharge. Shutting the pump from that state empties TK-1 in 49.8 s against the $2\sqrt{h_{\max}}/c_2 = 50.0$ s of the analytic drain expression, and empties TK-2 27.3 s later, because TK-2 keeps receiving what TK-1 is losing until TK-1 is dry.

3.8 Predictive Control on Both Plant Structures

Table 3.8 runs cases 3, 4, 8, 9 and 11 under each controller for 300 s, or 120 s for the challenge. Each run selects its controller before it starts, which gives both controllers identical events and an identical noise realization. The PID runs each compared case at its shipped tuning, $K_c = 300\,\%/\text{m}$ and $\tau_I = 60$ s, except the challenge, which ships $K_c = 60\,\%/\text{m}$ and suggests 300 %/m, while the MPC runs everywhere with the settings given in Section 3.5.4. The comparison is a second cross-mismatch experiment beside Table 3.5, because the shipped PID gain is close to the open-splitter IMC gain of Table 3.3 while the MPC model belongs to the shut splitter. Transmitter noise accounts for most of the pump travel. Over the quiet first 10 s of case 3 alone the PID moves the pump by 31.7 % and the MPC by 28.5 %.

With the splitter shut the model matches the plant and the MPC wins cases 3 and 4. At the setpoint step of case 3 it holds the pump at 100 % for 6 s, peaks at 0.618 m against 0.649 m and settles 101 s sooner. Under the draw-off of case 4 the level falls to 0.461 m against 0.440 m, and $\hat{d}$ averages −7.38 % over the last 100 s against the −7.50 % that Equation (3.3) implies. The improvement is paid for in the uncontrolled vessel, as with the fast PID tuning of Section 3.6. The model has no top tank, and its larger pump moves carry TK-1 to 0.898 and 0.983 m in the two cases against 0.807 and 0.849 m, the second only 17 mm below the rim. Case 8 is close to a tie, both controllers holding the pump at 100 % from 10.05 s and ending at 0.763 m. The IAE of the MPC is 1.0 % higher because of the first 10 s, where its plans answer the noisy readings with pump positions of 68.4 to 84.7 %

Table 3.8 Shipped PID and Predictive Controller on Five Cases With the Same Noise Sequence, and the Predictive Controller With Its Model Scheduled on the Splitter or Frozen at the Half-Open Entry

Case	Controller	IAE (m s)	IAE_h	t_s (s)	$h_{1,max}$ (m)	Travel (%)	Limit (%)
3, setpoint	PID	4.27	4.00	132.0	0.807	852	0
	MPC	2.18	1.65	31.0	0.898	750	2.0
	MPC^a_{50}	14.96	14.96	—	1.000	1164	58.7
4, draw-off	PID	3.28	2.86	114.0	0.849	825	0
	MPC	1.74	1.25	37.5	0.983	728	2.0
8, limit	PID	60.82	60.67	—	0.765	52	96.7
	MPC	61.44	61.30	—	0.765	48	96.7
9, splitter	PID	3.14	2.65	19.0	0.502	882	0
	MPC	8.68	8.68	—	0.506	5376	70.7
	MPC^c_s	1.53	1.00	5.5	0.506	574	0.7
	MPC^a_{50}	1.54	1.00	5.5	0.506	557	2.0
11, challenge	PID, $K_c = 60$	7.21	7.11	—	0.755	66	0
	PID, $K_c = 300^b$	2.59	2.51	—	0.849	331	0
	MPC	1.12	0.99	37.5	0.983	333	5.0

IAE uses the measured bottom level, as the scorecard does, and IAE_h the true one. t_s is the time after the last event from which the true bottom level stays within ±0.01 m of the setpoint, with a dash where it never does. $h_{1,max}$ is the highest true top level, Travel is $\sum |u_k - u_{k-1}|$ and Limit is the share of the run with the pump at 0 or 100 %, all on the 2 Hz log. [a]The scheduled model frozen by script at its 50 % entry, not a setting of the application. [b]Suggested, $\tau_I = 60$ s. [c]Scheduled model (case 10), which repeats the MPC rows exactly in the other four cases.

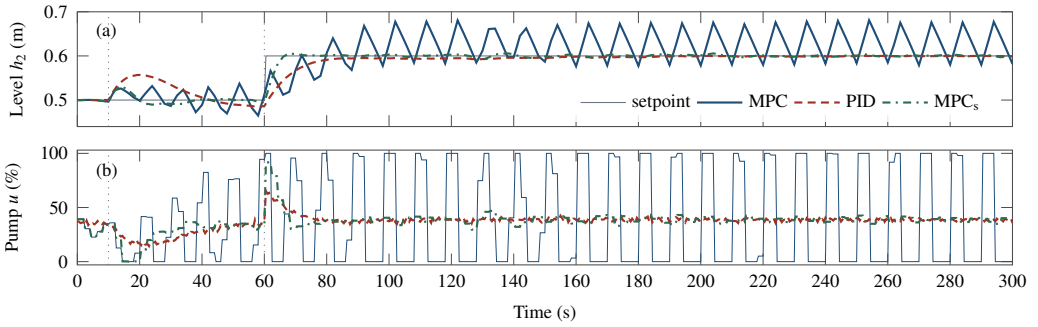


Figure 3.4 The predictive controller with its fixed model (MPC) and with its scheduled model (MPC_s, case 10) and the shipped PID on the events of case 9, each run with an identical noise realization. The splitter opens from 0 to 50 % at 10 s and the setpoint steps to 0.60 m at 60 s (dotted lines). The MPC is solid (thin in panel b), the PID dashed, MPC_s dash-dotted and the setpoint a thin gray line. (a) True bottom level and setpoint. (b) Pump position, logged at 2 Hz.

and leave TK-1 at 0.486 m against 0.500 m, so the climb at full pump starts from a smaller upstream inventory.

Case 9 opens the splitter to 50 % at 10 s, which leaves the operating point unchanged, and steps the setpoint to 0.60 m at 60 s, as shown in Figure 3.4. The MPC loses it. After the step its pump sits at 0 or 100 % for 81 % of the time, the level cycles between 0.512 and 0.681 m with a period near 10 s and never settles, and the true-level IAE is 8.68 m s, 3.3 times the 2.65 m s of the PID, whose pump stays between 14.3 and 65.2 %. The cause is the fast path of the bypass. By Equation (3.11) at $x_v = 0.5$ a 1 % pump step raises the level by 0.82 mm within one decision period, where the model predicts 0.044 mm, and by 4.34 against 0.95 mm after 10 s, though both end at 28.3 mm. The filter books the unexplained response as load, $\hat{d}$ rising from 1.5 to 28.0 % within 5 s of the opening and swinging from −52.5 to 58.5 % after the step, so each plan answers its own earlier moves with a full-range move and 106 of the 150 first moves sit on a limit. Over the first 60 s the cycle costs little, 0.78 against 1.26 m s, and the setpoint step turns it into the loss, 7.91 against 1.39 m s over the last 240 s.

TRY IT IN THE APP 9 · MPC When the Splitter Reshapes the Plant

Load the case and run it to 300 s. The splitter opens to 50 % at 10 s and the setpoint steps to 0.60 m at 60 s. Watch the pump: the predictive controller, whose model knows only the shut splitter, throws it from limit to limit. Then run case 10, whose model follows the splitter, and the pump stays off its limits.

The loss belongs to the model, not to the tuning. With λ anywhere on its range of 0 to 3 and q_d set by script to 1, 10 or 100 times r/K_m^2, the true-level IAE of case 9 stays between 6.20 and 9.43 m s, and horizons of 40 to 600 s move it by under 1 %. Told the splitter setting, the same controller takes its 50 % entry when the splitter opens and scores 1.53 m s on the events of case 9 (case 10), against 8.68 fixed and 3.14 for the PID, as shown in Figure 3.4. After the step its pump stays between 29.6 and 91.7 %, never on a limit, and from 180 s the true level holds within 0.596 to 0.606 m. Over seeds 1001 to 1010 the scorecard IAE spans 1.48 to 1.65 m s scheduled, 8.48 to 9.13 fixed and 3.09 to 3.13 for the PID. With the splitter shut the two models give identical runs in cases 3, 4, 8 and 11.

Frozen by script on its shut entry, the scheduled controller repeats the fixed one pump for pump, so switching, carry-over and the stricter plan check add nothing. Frozen on its 50 % entry it wins case 9 (1.54 m s) but loses case 3 at 14.96 m s with TK-1 at its rim and fails the challenge (worst error 0.066 m). The gain is the model identified at the setting in force, kept in place by the schedule, as Table 3.5 found for the two IMC tunings. Volume 9, Chapter 3 makes the same point with two predictive controllers that differ only in the step test that filled their model. The fixed model at $\lambda = 3$ still scores 8.20 m s, and the PID scheduled by the IMC rule on first-order fits of the same step tests (K_c from 67 to 354 %/m) scores 4.20 m s and fails the challenge, the shipped 300 %/m being already near the open-splitter gain.

On the challenge the MPC, which does not read K_c, repeats the first 120 s of its case 4 run. Over the scored final third, from 79.2 to 120 s, its largest measured error is 0.012 m, against 0.072 m for the shipped K_c, which fails the ±0.05 m band, and 0.023 m for the suggested one. With the shipped gain the true level falls to 0.398 m at 47 s and is still 0.015 m short at 120 s. The suggested tuning, which raises K_c alone, stays inside the band with 0.027 m to spare. None of the 660 plans of the five cases was rejected. Only the three plans that answered the event in each of cases 3, 4 and 11 had a bound active, against 145 in case 8 and 112 in case 9, with up to 12 active-set iterations. A plan took 0.08 to 0.18 ms on average under Node.js and 211 and 217 µs in two Chrome runs of case 9, at most 2.7 ms, or 0.011 % of one processor core in real time, wall-clock values from one desktop computer. The MPC wins cases 3, 4 and 11, trails case 8 by 1.0 % and loses case 9 by a factor of 3.3 unless its model is scheduled.

The schedule omits the splitter's own transient, which still lifts the level to 0.532 m in case 10. With the splitter at 25 to 100 % before the step of case 3 it scores 1.30 to 1.44 m s, the PID 1.38 to 1.96 and the fixed model 3.54 to 16.24, though the PID leads on the true level from 60 % up (0.68 to 0.71 against 0.80 to 0.82 m s). Between table entries the schedule loses little. At splitter settings of 37 and 60 %, where it picks the 25 and 50 % entries, case 3 scores 1.43 and 1.31 m s scheduled, 1.68 and 1.46 under the PID, and 5.41 and 12.13 under the fixed model, whose pump sits at a limit for 36 and 92 % of the time after the step. An entry ten or twelve percentage points away therefore costs far less than the shut-splitter entry. Over 40 runs (every case plus extremes of splitter motion, load, noise and geometry) under the PID and both models with either estimator, the scheduled controller never overflows and is never longer outside ±0.03 m than the safest shipped structure, while the fixed model overflows TK-2 for 60.9 s with a 0.65 m rim. Its TK-1 peak exceeds the fixed model's in three runs, 0.907 against 0.825 m (PID 0.978 m) for a splitter taken to 50, 100 and 0 %, 0.891 against 0.874 m when it shuts at 150 s, and 0.470 against 0.443 m at 25 % throughout. A scheduled plan took a median of 85 µs under Node.js on a loaded host.

Table 3.9 Shipped PID, Relay-Written PI, Two Counterfactual PI Loops and the MPC on Cases 3 and 4 Over 900 s With One Noise Seed

Case	Controller	IAE (m s)	IAE_h (m s)	e_{pk} (m)	$h_{1,\max}$ (m)	Travel (%)	Travel^0 (%)
3, setpoint	PID, shipped	6.17	4.31	0.049	0.807	2424	114.5
	PI, relay	5.92	4.22	0.012	0.674	1084	30.9
	PI[a]	5.43	3.45	0.034	0.777	2428	95.1
	PI[b]	6.19	4.22	0.031	0.704	1082	35.9
	MPC	4.30	2.60	0.018	0.898	1973	216.1
4, draw-off	PID, shipped	5.17	3.10	0.060	0.849	2396	64.3
	PI, relay	9.40	8.17	0.082	0.752	1072	21.5
	PI[a]	5.48	3.81	0.061	0.811	2400	49.1
	PI[b]	6.67	4.67	0.080	0.800	1070	28.1
	MPC	3.82	2.12	0.039	0.983	1914	223.5

PID, shipped: $K_c = 300$, $\tau_I = 60$ s. PI, relay: $K_c = 134$, $\tau_I = 149$ s. [a]$K_c = 300$, $\tau_I = 149$ s. [b]$K_c = 134$, $\tau_I = 60$ s. Neither counterfactual is shipped or written. e_{pk} is the true overshoot of the new setpoint in case 3 and the largest true deviation after the draw-off in case 4. Travel^0 repeats Travel with $\sigma_n = 0$. Other columns as in Table 3.8. Either estimator gives the same MPC rows.

3.8.1 Estimator Choice and the Relay Test

Each run here lasts 900 s (three trend windows), the challenge 120 s, uses the shipped noise seed and has all of its settings entered before it starts. Over the 1799 executions of case 9 the two estimates of the bottom level differ by at most $2.35 \times 10^{-7}\sigma_n$, a residue of the filter's Riccati iteration stopped 3.3×10^{-9} short of convergence, and in every case, case 10 included, the two estimators give pump records within 1.15×10^{-5} % of one another with no bound active, no window rejected and no failed plan, the filter's own estimate coming no closer than 0.053 m to a bound, so no trend compares them.

Against the true bottom level, read by the measuring script alone, the two are equally accurate, with a root-mean-square error of 1.55 to 2.36 mm (under $0.6\sigma_n$) and a largest error of 5.6 to 8.8 mm in the nine cases with the splitter shut. The open bypass of case 9 raises them to 24.4 and 44.8 mm. The model level $\bar{h}_m + \hat{z}_1$ misses the true top level by up to 0.354 m in case 4, where $\hat{d}$ carries the draw-off, and by 0.753 m in case 9. On case 9 an estimate took a median of 1410 µs under Node.js and a plan 96 µs, about 0.3 % of each 2 s decision for four estimates and a plan, and in Chrome estimates and plans took up to 0.11 % of a core at 1× and 6.3 % at 10×, within a factor of two of the budget of a tenth of a core.

The relay test was started at 30 s of case 3 with its setpoint step removed, the bottom level at 0.50 m and u_0 at 35.34 %, against the 35.36 % that holds the level, so δ_r was the full 5 %. It converged after 8 cycles and 280 s with $\alpha = 0.0191$ m, $D = 0.503$, $P_u = 67.5$ s and $K_{cu} = 428$ %/m, and it wrote $K_c = 134$ %/m and $\tau_I = 149$ s in place of the shipped 300 and 60 s, labeled because $\psi = 38.9°$. All ten further seeds (1001 to 1010) converged as well, giving K_{cu} from 271 to 383 %/m, P_u from 63.9 to 73.0 s and ψ from 27.1° to 35.8°.

The measured point belongs to the experiment more than to the tanks. Two lags in series approach −180° without reaching it, and sampled with a zero-order hold at Δt_c the model of Equation (3.19) first crosses −180° at a period of 13.1 s, where the ultimate gain is 10 100 %/m. A relay with hysteresis settles where the imaginary part of the frequency response equals $-\pi\varepsilon_r/(4\delta_r)$, which on the sampled model predicts 410 %/m at 77.3 s, and the same relay on the noise-free model returns 407 %/m at 77.1 s. The measured cycle is 5.2 times slower than the ultimate one and its K_{cu} a twenty-fourth of the ultimate gain, although ψ is only 39°. With $\sigma_n = 0$ the simulator's test starts from 35.36 %, runs at $D = 0.498$ and returns 416 %/m at 77.0 s, 3 percent below the noisy gain at a period 14 percent longer.

Started at 300 s of case 3, settled at 0.60 m, the relay centers on 38.46 % against the 38.73 % that holds the level, spends $D = 0.506$ of each cycle high and converges after five cycles on $K_{cu} = 293$ %/m and $P_u = 74.4$ s, writing $K_c = 91.6$ %/m and $\tau_I = 164$ s. With $\sigma_n = 0$ it starts at 38.67 %, cycles at $D = 0.502$ and returns 389 %/m at 82.1 s. In both runs $\sin(\pi D)$ of Equation (3.24) departs from one

by less than 0.001, so the ratio 0.75 of the two gains comes from the wider noisy swing, 0.0248 against 0.0203 m.

The written constants win the setpoint step narrowly and lose the draw-off in Table 3.9. On the setpoint step the true-level IAE falls from 4.31 to 4.22 m s, and the loop reaches 0.60 m with an overshoot of 12 mm against 49 mm. On the draw-off it rises from 3.10 to 8.17 m s and the largest deviation from 0.060 to 0.082 m, because the integral gain K_c/τ_I falls from 5.00 to 0.899 %/(m s), 18 percent of the shipped value. The counterfactual rows show that neither constant alone accounts for that loss. The shipped K_c with the written τ_I scores 3.81 m s and the written K_c with the shipped τ_I scores 4.67 m s. The relay-tuned loop saves pump travel and top-tank inventory. Its travel is 1084 against 2424 % on case 3, a ratio of 0.447 equal to the gain ratio $134/300 = 0.447$, because transmitter noise passed through K_c makes up 95 to 98 percent of the travel of both loops, as Travel^0 shows, and TK-1 peaks at 0.674 m against 0.807 m.

3.9 Technology in Practice

The loop of Section 3.3 reads LT2 through white noise and moves a pump whose flow is strictly proportional to its signal. Figure 3.5 draws it with the devices a plant pair of tanks would carry in the place of the schematic's symbols.

What Measures It

A level in an open tank can be read from the pressure at its floor, from an echo timed from its roof, or from a float [24]. The magnetostrictive transmitter is a float instrument that times a pulse in a rod. It can carry two floats, one on the surface and one on an interface between two liquids, and with digital communication it reports both levels over one pair of wires [24]. No range, accuracy or response time is given here, because none was found in an open source that names no maker. The noise σ_n of Equation (3.7) stands for the moving surface of TK-2, and the model's transmitters have no lag, drift or fouling. LT1 gives an indication only, so no loop watches the vessel that comes closest to its rim.

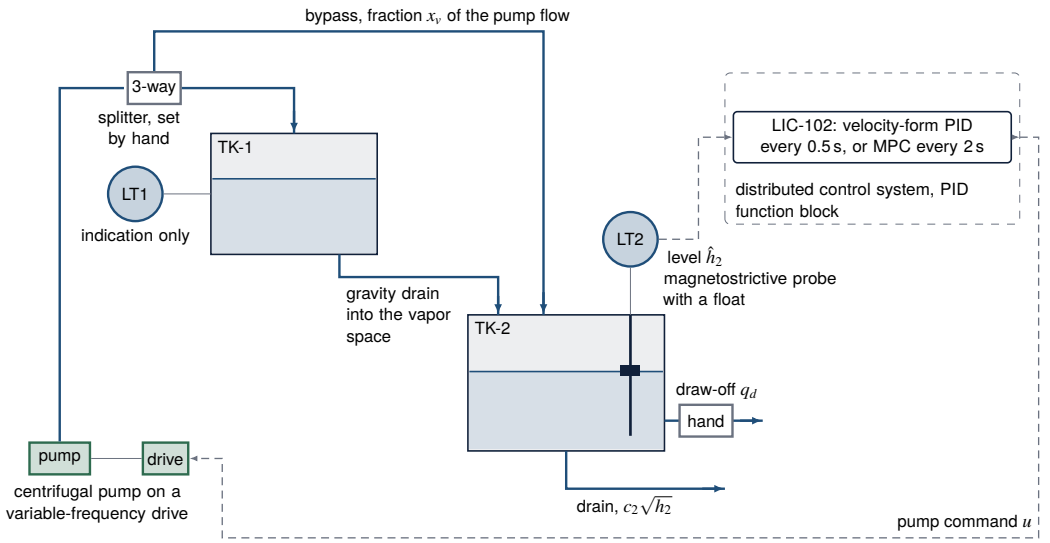


Figure 3.5 The level loop of the chapter with plant devices in place of the schematic's symbols. LT2 reads the bottom level with a float on a magnetostrictive probe, and LIC-102 on the control platform sets the speed of the pump through its drive. Notice that the splitter and the draw-off are hand valves the loop never moves, and that LT1 on the top tank only indicates, so no loop guards the vessel that the bypass can empty and a load can fill.

TECHNOLOGY Sensor: A Magnetostrictive Level Transmitter

A magnetostrictive transmitter hangs a metal rod, the waveguide, into the vessel. A light ring-shaped float with a magnet rides on the rod and follows the liquid surface. A current pulse sent along the rod meets the float's magnetic field and starts a twisting pulse that travels at the speed of sound to a sensing head at the top, and its travel time gives the distance to the float. Sound keeps a steady speed in a metal rod over wide ranges of temperature and pressure, which gives the method a stable calibration and fine resolution. Its weak point is the close fit of float on rod: solids or sludge can bind the float, and the reading then stops following the level [24].

What Moves It

The pump of Equations (3.1) and (3.2) has no head dependence, no minimum flow and no slew limit. A plant pump on a variable-frequency drive has a flow that falls as the head rises and a speed that changes at the rate its drive allows [24]. As a final element a pump also differs from a valve. A stopped pump does not necessarily stop the flow through it, and a drive offers fewer choices of a fail-safe state [24]. The splitter is a three-way valve, a body with three ports that divides one stream between two directions [24]. In the model it is set by hand and moves instantly, and its two paths have no resistance of their own.

TECHNOLOGY Final Element: A Centrifugal Pump on a Variable-Frequency Drive

The centrifugal pump is the most common pump in the process industries. On a variable-frequency drive, the line power is rectified and then rebuilt at the frequency the controller asks for, and the motor's speed follows that frequency. Running a pump only as fast as needed saves power and wear. Speed alone does not set the flow, because a centrifugal pump's discharge pressure falls as its flow rises. A pump also changes flow more slowly than a valve, since drives limit acceleration and deceleration, and a drive shuts down on overcurrent or overvoltage whatever the controller asks [24].

Where the Controller Runs

A distributed control system runs production at one site in refining, chemical manufacturing, water treatment and power generation. Its loops may be nested or cascaded, and they run with cycle times from milliseconds to minutes [25]. LIC-102 executes every 0.5 s and the predictive controller decides every 2 s, both inside that range. The velocity-form PID of Equation (3.17) is the PID function block of such a platform, with the bumpless transfer of Section 3.5. The predictive controller and its table of models would run in the supervisory level that oversees the loop controllers [25].

IN PRACTICE Ask How Far the Level May Move

An instrument handbook treats a change of tuning as seriously as the reconfiguration of a field instrument. Before any change, it advises asking the operators how far and how fast the process variable may move, and keeping clear of alarm and shutdown limits, because few loops in a plant stand alone. Its own example is three vessels in series, drained by gravity, with one transmitter on the last. That loop is hard to control, not least because the vessels in between may overflow while nothing measures them [24].

What Is Modernizing It

Three developments touch this loop. The first is variable-speed pumping, which saves the energy that a throttling valve dissipates, although a pump changes flow more slowly than a valve and cannot positively stop it [24]. The second is the digital level instrument that reports more than one level from one probe, such as the two-float transmitter above [24]. The third is the correction of a predictive controller's model from routine operating data, without a new step test, the subject of the open challenge in Section 3.11 [26].

3.10 Limitations

The simulator is a teaching instrument rather than a design tool. Two of its elements are stylized on purpose. A drain that empties a 1 m vessel in 50 s compresses the two-lag response into minutes, and no structural result depends on it. A draw-off reported as a level velocity keeps every term of Equation (3.2) in one unit and makes the ceiling of Equation (3.25) a one-line calculation. The non-interacting discharge law is imposed rather than checked, so the model does not detect the flooded-nozzle condition under which the same two vessels would become a genuinely coupled pair, and it must not be used to study that case. Both tanks share one area, one drain coefficient and one rim height, so mismatched vessels cannot be represented. The pump is ideal and its flow does not fall with discharge head, which removes the one physical mechanism that would couple the bottom level back to the pump. The splitter is a pure fraction with no dynamics and no pressure interaction between its branches, so the plant order changes exactly when the fraction moves and no real valve does that. Mass is not conserved during an overflow and the spill is not routed anywhere. Transmitter error is white noise and nothing more, which leaves out the slow drift and surface effects that dominate a real level reading. No flow controller sits under the pump, no second actuator exists and the PID's tuning is never scheduled, so the remedy that Section 3.6 names runs only by script. The noise generator is seeded, so even a noisy run repeats exactly at every speed setting. The square-root drains, the four hard limits, the triangular structure and the 0.5 s sample period are retained because each produces behavior the tuning has to accommodate, and Section 3.3.5 says where each one stops being true.

With the fixed model the predictive controller plans on two lags fitted at 0.50 m with the splitter shut. Without the bypass path its predicted response over one decision period is 18.6 times too small at the half-open splitter, and without the square-root drains the plant gain at 0.60 m is 9.5 % above its own. It has no top-tank level or rim, so nothing protects TK-1, which reached 0.983 m in case 4, and the splitter, the draw-off and the leak reach it only through $\hat{d}$. Its Kalman gains assume the default noise and are never recomputed, its q_d sits at the top of the tuning grid, and its comparison rests on one noise sequence. The scheduled model knows the splitter only as the setting it is told, so a valve moved without telling it leaves the plan on the wrong entry, and every entry is fitted at 0.50 m and none holds the splitter's own transient. The bounds of the moving horizon estimator act on the states of that fitted model, whose first state misses the true top level by up to 0.753 m in case 9, so they could not protect TK-1 even when active. The relay keeps its center fixed during a test and, on delay-free lags, measures its own band as much as the plant.

3.11 Discussion

The situation this application reproduces is ordinary. A surge drum is bypassed for maintenance, a spare vessel is lined up in parallel, or a plant runs a summer and a winter lineup, and in each case the piping between two units changes while the material balance does not, so every steady-state indication looks normal while every dynamic indication has moved. Inventory between units exists so that upstream variation need not pass downstream at once, and the structure deciding which flow regulates which holdup is settled before any loop is tuned [6]. Taking a vessel out of that chain removes a lag and a buffer together. The last two rows of Table 3.6 show both, because opening the splitter turned a load the loop could not reject into one it rejects to within 0.2 mm in 700 s. A gain and a time constant identified on one lineup belong to that lineup, as the scheduling surveys [8, 9] and the longer history of feedback design [27] both record, and the splitter separates what a single step test cannot. The gain is a property of the material balance and survives the change, while the phase is a property of the path and does not.

⚠ OPEN CHALLENGE Keeping a Controller's Model True to the Plant

A predictive controller is only as good as its model, and a plant drifts from the model identified at commissioning as feedstocks, equipment and operating conditions change. A new step test disturbs production and keeps operators watching, while routine data carry little excitation because setpoints may stay fixed for months. Progress would be a model corrected from routine data with no injected signal, as one simulation study has shown for a pilot distillation column [26].

Methods beyond the two implemented controllers can be built on the model of Section 3.3. A predictive controller built on Equations (3.1) and (3.2) rather than on fitted lags would hold the two rims and the two dry bottoms as output constraints and would recognize the ceiling of Equation (3.25) as unreachable before the pump saturates. Moving horizon estimation or Kalman filtering on both level readings would recover the splitter fraction and the leak as parameters entering through known structure, which the estimator of Section 3.5.6 does not attempt and which would free the schedule of Section 3.5.5 from what the operator reports. A multivariable formulation could move the pump and the splitter together, trading top-tank inventory against bottom-level regulation. The APMonitor framework implements this combination of nonlinear modeling, estimation and predictive control [28], and the Gekko Optimization Suite brings it to Python so that the exported runs feed the estimator directly [29].

3.12 Conclusion

Two gravity-drained tanks behind one pump and a three-way splitter valve form an interactive laboratory in which the operator changes the order of the plant rather than its operating point. Linearization gives a lower-triangular pair of non-interacting lags with an upper time constant of $(1-x_v)\tau_2$, a pole and a zero that cancel when the splitter is wide open, and a process gain $c_1\tau_2/100$ free of the splitter, and step tests in the application return 0.0323 m/% at every setting while the apparent dead time falls from 18.5 to −0.6 s. Internal model control tunings from the two ends differ by a factor of 4.6 in proportional gain, and each on the other structure gives 58.7 % overshoot in one direction and 159.7 s of settling against 12.3 s in the other. A linear predictive controller fitted with the splitter shut cuts the true-level IAE of the shipped PID from 4.00 to 1.65 m s on the setpoint step but cycles against the pump limits half open, 8.68 against 2.65 m s, while the same controller scheduled on the told splitter setting scores 1.00 m s there, a gain that freezing the schedule traces to the model. A bounded moving horizon estimate leaves every pump record within 1.2×10^{-5} % of the filter's, and a relay test writes $K_c = 134\,\%/\mathrm{m}$ and $\tau_I = 149\,\mathrm{s}$, which raise the true-level IAE of the draw-off from 3.10 to 8.17 m s. A 0.020 m/s draw-off that the shut-splitter plant cannot reject above 0.2500 m is rejected completely with the splitter open, so the valve bounds the load as well as the bandwidth. A fuller model would add a PID gain schedule on the splitter, a flooded-nozzle option that makes the pair interacting, a head-dependent pump curve and unequal tank areas. The application is an accessible browser instrument for level control experiments in which the flow path, not the operating point, sets the plant order.

KEY POINTS

- The splitter changes only the path the water takes, not where the bottom level settles.
- Step tests return 0.0323 m/% at every setting, while the apparent dead time falls from 18.5 to −0.6 s.
- Each end's tuning on the other structure overshoots 58.7 %, or settles in 159.7 s against 12.3 s.
- A predictive controller fitted with the splitter shut cycles against the pump limits half open unless its model follows the splitter.

Where else this applies. Volume 15, Chapter 7 strings five such lags in series.

References

[1] 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.

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

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

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

[5] K. H. Johansson, "The quadruple-tank process: a multivariable laboratory process with an adjustable zero," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 3, pp. 456–465, 2000.

[6] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[7] K. J. Åström, "Limitations on control system performance," *European Journal of Control*, vol. 6, no. 1, pp. 2–20, 2000.

[8] D. J. Leith and W. E. Leithead, "Survey of gain-scheduling analysis and design," *International Journal of Control*, vol. 73, no. 11, pp. 1001–1025, 2000.

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

[10] E. P. Gatzke, E. S. Meadows, C. Wang, and F. J. Doyle, III, "Model based control of a four-tank system," *Computers & Chemical Engineering*, vol. 24, no. 2-7, pp. 1503–1509, 2000.

[11] H. Pan, H. Wong, V. Kapila, and M. S. de Queiroz, "Experimental validation of a nonlinear backstepping liquid level controller for a state coupled two tank system," *Control Engineering Practice*, vol. 13, no. 1, pp. 27–40, 2005.

[12] C. C. Ko, B. M. Chen, J. Chen, Y. Zhuang, and K. C. Tan, "Development of a web-based laboratory for control experiments on a coupled tank apparatus," *IEEE Transactions on Education*, vol. 44, no. 1, pp. 76–86, 2001.

[13] E. J. Ionascu, "Drainage time and shape: inequalities from Torricelli's law," arXiv preprint 2511.00023, 2025.

[14] O. Tramis, E. Merlin-Anglade, G. Paternoster, M. Rabaud, and N. M. Ribe, "Torricelli's curtain: morphology of horizontal laminar jets under gravity," *Physics of Fluids*, vol. 33, no. 6, p. 062116, 2021.

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

[16] 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.

[17] 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.

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

[19] 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.

[20] B. Aufderheide and B. W. Bequette, "Extension of dynamic matrix control to multiple models," *Computers & Chemical Engineering*, vol. 27, no. 8-9, pp. 1079–1096, 2003.

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

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

[23] 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.

[24] 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.

[25] 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.

[26] Y. Kono and M. Koizumi, "Model life extension for continuous process: non-invasive correction of model-plant mismatch with regularization," in *2023 European Control Conference (ECC)*, 2023, pp. 1–8.

[27] K. J. Åström and P. R. Kumar, "Control: A perspective," *Automatica*, vol. 50, no. 1, pp. 3–43, 2014.

[28] 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.

[29] 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.

3.13 Exercises

3.1 Concept check · 2 min

Case 1 holds the bottom level at 0.50 m in manual with the pump at 35.36 % and the splitter shut. The splitter is opened to 50 % and the pump is left where it is. What does the bottom level do?

(a) It settles about 0.18 m higher, because half of the pump flow now reaches TK-2 without delay

(b) It falls at first, because the top tank drains away the flow it used to pass on to TK-2

(c) It stays at 0.50 m throughout, because the process gain does not depend on the splitter

(d) It rises for about half a minute and returns to 0.50 m, because the splitter acts through a zero at the origin

3.2 Concept check · 1 min

The same 10 % pump step is made from rest at 0.5013 m with the splitter at 0, 50 and 100 %, as in Table 3.3. Which result is the same in all three tests?

(a) The fitted process gain, because every unit of pump flow leaves through the bottom drain

(b) The apparent dead time, because the two tanks and their drains are the same in every test

(c) The fitted time constant, because the bottom tank's time constant does not move with the splitter

(d) The top-tank level, because the pump that holds the bottom level does not move with the splitter

3.3 Concept check · 2 min

The shut-splitter step test of Table 3.3 returns a fitted gain of 0.0323 m/%, while Equation (3.11) gives 0.0283 m/% at the starting point. Which explanation is right?

(a) The two-point fit reads part of the apparent dead time as extra gain

(b) The step lifts the level to 0.8245 m, and the local gain grows with level along the way

(c) Transmitter noise biases the final level that the fit reads upward

(d) The 0.05 s Euler substep accumulates error over the long response

3.4 Theory and calculation · 5 min

Two tanks have the default coefficients $c_1 = 0.08\,\mathrm{m/s}$ and $c_2 = 0.04\,\mathrm{m}^{1/2}/\mathrm{s}$, no leak and no draw-off, and hold the bottom level at 0.50 m.

(a) Find the pump position that holds this level and the resting top-tank level at splitter settings of 0 and 25 %. *Checkpoint:* 35.36 %, then 0.500 m and 0.281 m.

(b) At $x_v = 0.25$, find τ_1, τ_2, τ_z and K_p of Equation (3.11). *Checkpoint:* 26.52 s, 35.36 s, 6.63 s and 0.0283 m/%.

(c) Write the denominator of Equation (3.11) in the standard form of Equation (1.11) and find the natural period and the damping factor at $x_v = 0$ and at $x_v = 0.25$. *Checkpoint:* 35.36 s and 1.000, then 30.62 s and 1.010.

3.5 Theory and calculation · 8 min

The plant of the previous exercise holds 0.50 m with the splitter shut and $h_{\max} = 1$ m. A draw-off of $q_d = 0.010\,\mathrm{m/s}$ is applied to the bottom tank.

(a) The pump is held at 35.36 %. Find the level at which the bottom tank comes to rest. *Checkpoint:* 0.2089 m.

(b) A PI loop restores 0.50 m. Find the pump it settles at and the top-tank level that this pump requires. *Checkpoint:* 47.86 % and 0.916 m.

(c) Find the largest draw-off that the shut-splitter plant can reject at a bottom level of 0.50 m, and the highest level it can hold against 0.020 m/s. *Checkpoint:* 0.0117 m/s, then 0.250 m.

3.6 Theory and calculation · 5 min

The PI tunings of Section 3.6 come from the internal model control rule $\tau_c = \max(0.1\hat{\tau}_p, 0.8\hat{\theta}_p)$, $K_c = \hat{\tau}_p/[\hat{K}_p(\tau_c + \hat{\theta}_p)]$ and $\tau_I = \hat{\tau}_p$, applied to the rows of Table 3.3.

(a) Compute K_c and τ_I from the shut-splitter row and from the half-open row. *Checkpoint:* 66.6 %/m and 71.6 s, then 174.0 %/m and 51.3 s.

(b) The open-splitter row gives $K_c = 309.4\,\%/\mathrm{m}$ and $\tau_I = 42.95$ s. Find the ratios of the proportional gains and of the integral gains K_c/τ_I between the open and the shut tunings. *Checkpoint:* 4.6 and 7.7.

(c) The splitter position is not reported to the control room. Using Table 3.5, say which of the two end tunings to install. *Checkpoint:* The shut-splitter tuning.

3.7 App challenge · 9 min

Load case 2, Open-loop actuator step. The pump steps from 35.36 to 45.36 % at 10 s. Run to 160 s at 5× speed with the default seed and export the CSV file. Press the reload button, set the Splitter bypass to 100 % before Run, and repeat.

(a) From Equation (3.3), find the level at which the new pump settles and the gain of the step. *Checkpoint:* 0.823 m and 0.0323 m/%.

(b) In each file find the times after the step at which the true bottom level, column `true_pv`, has covered 28.3 % and 63.2 % of the change toward that level. Interpolate between rows. *Checkpoint:* Shut 42.3 and 90.0 s, open 13.7 and 42.3 s, each within 0.5 s.

(c) Fit $\hat{\tau}_p$ and $\hat{\theta}_p$ to both tests by the two-point method of Section 3.4. *Checkpoint:* Shut 71.6 s and 18.5 s, open 42.9 s and −0.6 s, each within 1 s.

3.8 App challenge · 8 min

Load case 3, Closed-loop setpoint step. Before Run set $K_c = 309$ and $\tau_I = 42.95$ s, the open-splitter tuning of Section 3.6, set the Lower level sensor noise to zero in Show Model and the Splitter bypass to 100 %. Run to 410 s at 10× speed with the default seed. Then set the splitter to 0 %, press Reset and run to 410 s again, exporting the CSV file of the second run.

(a) Read the scorecard IAE and overshoot of the first run. *Checkpoint:* 0.419 m s within 0.005 m s, no overshoot.

(b) Read the same tiles for the second run, and from column `true_pv` the period between the first two peaks. *Checkpoint:* 4.950 m s within 0.01 m s, 58.7 %, then 79.5 s within 1 s.

(c) Find the highest true top-tank level of the second run in column `level1`, and compare it with the rim. *Checkpoint:* 0.835 m, 84 % of the rim.

3.9 App challenge · 8 min

Load case 4, Disturbance rejection, which applies a draw-off of 0.006 m/s at 10 s under $K_c = 300\,\%/\mathrm{m}$. Before Run set the integral time to 500 s, which switches integral action off, and the Lower level sensor noise to zero. Run to 300 s at 5× speed with the default seed and export the CSV file.

(a) The loop now acts about the fixed bias of 35.36 %. Solve $c_1(\bar{u} + K_c e)/100 = c_2\sqrt{0.5 - e} + q_d$ for the steady offset e and the pump. *Checkpoint:* $e = 0.0223\,\mathrm{m}$ and 42.06 %.

(b) Read `true_pv` and `mv` at 300 s. *Checkpoint:* 0.4777 m within 0.001 m and 42.06 % within 0.05 %.

(c) Find the level at which the draw-off would leave the bottom tank with the pump held at 35.36 %, and the share of that fall the proportional loop recovers. *Checkpoint:* 0.3104 m and 88 %.

3.10 Application area · 3 min

On a pair of plant tanks the outlet nozzle of TK-1 enters TK-2 at 0.40 m above its floor, and the loop holds TK-2 at 0.50 m.

(a) Does the discharge law of the model hold at this operating point? Name the entry of the state matrix of Equation (3.8) that it sets to zero and say what replaces it. *Checkpoint:* No. The upper-right entry is no longer zero.

(b) Does the steady bottom level still depend on the pump alone, whatever the splitter setting? Answer from the balances. *Checkpoint:* Yes.

3.11 Application area · 4 min

A surge drum between two units is bypassed for two weeks of maintenance, the way the splitter bypasses TK-1.

(a) Which results of a step test made before the bypass still hold, and which must be measured again? *Checkpoint:* The gain holds. The time constant and dead time do not.

(b) Using the 0.020 m/s draw-off of Table 3.6, name one gain and one cost of the bypass. *Checkpoint:* A higher load ceiling and the lost buffer.

3.12 Review the agent · 5 min

An AI assistant reviews a run of case 9 and writes: “The process gain is 0.0283 m/% at every splitter setting, so the MPC’s model is still right after the splitter opens. The cycling between the pump limits after the setpoint step is a tuning problem: a move suppression of 0.3 is too small. Raise it to 3, the top of its range, and the MPC will beat the PID’s 3.14 m s.” Find the flaw and say how you would check the claim in the app.