INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 6 Multivariable and Plantwide Control

CHAPTER 1

Introduction: Pairing, Interaction, and Plantwide Structure

Pages 1–21

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Pairing, Interaction, and Plantwide Structure," in *Multivariable and Plantwide Control*, John D. Hedengren, Ed., vol. 6 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 1, pp. 1–21.

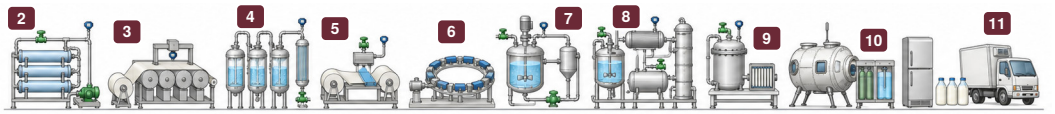
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v06-multivariable-plantwide/

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

CHAPTER 1

Introduction: Pairing, Interaction, and Plantwide Structure



The chapters 2 reverse osmosis train 3 paper machine 4 milk evaporator 5 coating line 6 storage ring 7 recycle plant 8 Tennessee Eastman plant 9 peptide purification 10 orbiting cabin 11 dairy chain

Notice that the recycle plant and the Tennessee Eastman plant each pipe material back to their reactor, while the coating line, the purification bench and the dairy chain pass every change on to the next unit downstream. Illustration, not to scale.

IN PLAIN TERMS **Which Handle Holds Which Target**

The main idea. A plant has several handles, such as valves, pumps and motor speeds, and several things worth holding steady. Before any loop is tuned, someone must decide which handle holds which target, and where the rate of production is set.

The challenge. Each handle moves several targets at once, so two loops can fight or undo each other's work. Streams that loop back, and tanks in a row, can also make a small change grow as it travels.

Why it matters. A wrong choice cannot be fixed by tuning, and the same few tests serve a seawater plant, a particle accelerator and a chain of milk depots.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	7 min	This page, Section 1.2 with its timeline and two History boxes, the Did you know? boxes in Sections 1.4.2, 1.4.6 and 1.4.8, Figure 1.11, and Key points in Section 1.8. No equations.
🕮	**Study**	40 min	The whole chapter in order, with each number of the recycle example worked, then Exercises 1 to 10.
⛰	**Practice**	5 min	This page, Section 1.4.9 with Table 1.2, Section 1.5 with Figures 1.7 and 1.9, and Section 1.8 with Figure 1.11.

1.1 The Question This Volume Answers

A single loop has one decision before tuning, the sign of its action. A plant with several actuators and several variables worth holding has a larger decision, and it is made before any loop is tuned. Which variables receive setpoints, which actuator holds each one, which flow sets the production rate, and which vessels absorb the upsets? A wrong answer cannot be repaired by tuning. A pairing can be unstable for every tuning, and a recycle can turn a small change of feed into a large change of load on the equipment inside it. This volume asks how such a structure is chosen, and which measures computed from a model predict whether it will work.

Ten chapters answer it, each on a plant with an app that runs in a browser. Chapter 2 and Chapter 3 pair two actuators with two product properties and add a static decoupler. Chapter 4 and Chapter 5 hold a quality along a chain of units that share one supply of heat or one line speed. Chapter 6 inverts a response of 96 monitors and 96 slow correctors whose weak directions amplify noise. Chapter 7 and Chapter 8 are recycle plants, where inventory control and the choice of production handle come first. Chapter 9 judges an upstream decision by its effect two units downstream, Chapter 10 holds two gas loops that share one total pressure, and Chapter 11 passes orders along a supply chain whose inventories amplify demand.

A reader who has worked through the volume can do the following on a plant that is new to the reader.

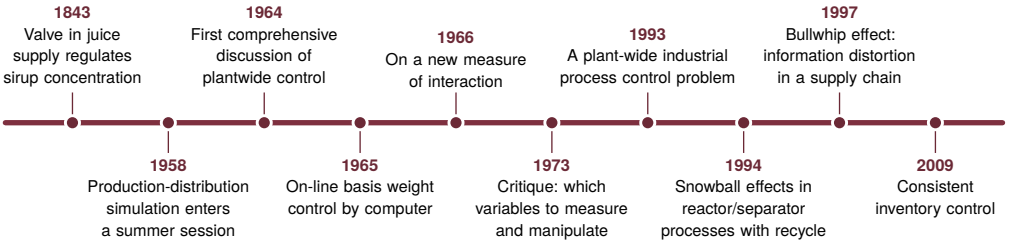


Figure 1.1 Milestones of pairing, interaction and plantwide structure in the systems of this volume. The dates are from [4] (1843), [5] (1958), [2] (1964 and 1973), [6] (1965), [7] (1966), [8] (1993), [9] (1994), [10] (1997) and [11] (2009). Notice that a sugar works, a paper mill and a classroom game sit beside the theory: each met a structural problem before a measure for it existed. Events are evenly spaced, not to a time scale.

- Count the degrees of freedom, and say which of them must stabilize inventories and which remain for production and quality.
- Place the production handle and arrange the inventory loops so that every vessel is held.
- Predict the snowball of a recycle and the handle that avoids it.
- Compute the relative gain array and the Niederlinski index, choose a pairing and reject one that no tuning can save.
- Design a static decoupler and say what it costs.
- Read the singular values of a gain matrix and decide how many directions a controller should try to correct.
- Say where a disturbance goes and whether a chain of inventories amplifies it.

This chapter develops each skill on one small plant and names the chapters that use it. PID tuning is not repeated here. Volume 2 develops it, and every loop below is a PI controller tuned by the rules of that volume. Standard texts treat the multivariable theory [1] and plantwide design [2, 3] in fuller detail.

1.2 A Short History

The milestones of Figure 1.1 follow the structural decisions of this volume through very different systems. A pairing was made long before it had a name. A patent of 1843 regulated the concentration of the sirup leaving the last of a train of sugar evaporators with a differential thermometer that turned a throttle valve in the juice supply, so the concentration was held with the feed [4], the pairing that Chapter 4 sets against the steam.

Interaction was met in practice before it was measured. On the first paper machine under computer control, whose on-line basis weight control first ran in 1965, the first identification experiment showed that the thick stock flow moved both the moisture and the dry basis weight, while the steam moved only the moisture [6]. A measure of interaction for multivariable process control, the relative gain, was proposed in 1966 [7], and the Niederlinski index followed in 1971 [12].

The plantwide questions grew with the plants. In the late 1970s and early 1980s one chemical company moved the production handle of numerous plants away from the process feeds, against a mindset that the feeds must stay constant and with pneumatic hardware that made the decision hard to reverse [13]. A plant-wide industrial process control problem was published in 1993 as a common test [8], the plant of Chapter 8, and in the same year the correction of a storage ring's orbit was written as a singular value decomposition of its response matrix [14]. A paper of 1994 studied snowball effects in reactor and separator processes with recycle [9].

> ⧖ HISTORY **1964: A Chapter on Overall Process Control**
>
> The first comprehensive discussion of plantwide control was a chapter on overall process control in a book of 1964 by Page Buckley. It introduced material balance control in the direction of flow and opposite to it, production rate control, and buffer tanks as low-pass filters. It discussed the recycle and the need to purge impurities, and it pointed out that inventory and flow cannot be controlled independently at one point of a plant, because the material balance relates them. A review of 2000 found the industrial approach still very much along those lines, with relatively little development in the area since [2]. Section 1.4.3 states the same rules on the example.

Supply chains have their own line. A production-distribution simulation with retail, distributor and factory stages was part of a university summer session from 1958 [5], an ancestor of the game behind Chapter 11, and a paper of 1997 described information distortion in a supply chain as the bullwhip effect [10]. Consistent inventory control, the rule of Section 1.4.3, received its own treatment in 2009 [11].

> ⧖ HISTORY **1973: The Structure Before the Theory**
>
> In the 1970s several articles criticized the gap between theory and practice in process control. The most famous, of 1973, observed that in many areas application was ahead of theory. It named the determination of the control system structure as the central issue: which variables should be measured, which inputs should be manipulated and which links should be made between the two sets. The gap was real, it said, but the theoretician must close it. A review of 2000 still found a systematic method for generating promising structures lacking, because the plantwide problem had not been well understood or even acknowledged as important [2]. Section 1.4 works through measures that can screen such structures.

1.3 Notation and Conventions

The chapters share the symbols of the notation page. The setpoint is SP, the process variable is PV and the manipulated variable is MV, with controller settings K_c and τ_I. A plant with inputs u_j and outputs y_i has the steady-state gain matrix K with elements $K_{ij} = \partial y_i / \partial u_j$ and, where dynamics matter, the transfer matrix $G(s)$. A gain carries the units of its output per unit of its input, so a matrix in engineering units mixes several units. A relative gain is written λ_{ij} and the array of them Λ. The singular values of a matrix are $\sigma_1 \geq \sigma_2 \geq \ldots$, and their ratio $\gamma = \sigma_1 / \sigma_{\min}$ is the condition number. The throughput manipulator (TPM) is the flow, speed or setpoint that sets the production rate.

Symbols differ between chapters. Chapter 2 writes the relative gain λ_{11} and Chapter 3 writes Λ_{11}. Chapter 10 uses κ for the cabin's moles per kilopascal, which is why this chapter writes the condition number as γ. Chapter 11 uses λ for the length of a supply line in days and Γ for the ratio of a stage's orders to the orders it receives. Chapter 6 uses λ for the pole of its per-mode loops and μ for the Tikhonov weight. Chapter 7 lists five quantities that each carry the name snowball gain and defines each by its formula. The running example uses hours, kilomoles and kilomoles per hour.

1.4 Structure Theory on One Plant

1.4.1 The Running Example

The running example is a small reaction plant with a recycle, drawn in Figure 1.2. Pure A enters at the fresh feed F_0 and meets the recycle D in a stirred reactor of holdup M, where it reacts to B at the rate kMz, with z the mole fraction of A in the reactor and $k = 1\,\mathrm{h}^{-1}$. The effluent F goes to a separator that is taken as ideal: it returns all of the A as the recycle and sends all of the B out as the product B, and it holds no material of its own. The reactor balances are Equation (1.1), with $n = Mz$ the kilomoles of A in the reactor.

$$\frac{\mathrm{d}M}{\mathrm{d}t} = F_0 + D - F, \qquad \frac{\mathrm{d}n}{\mathrm{d}t} = F_0 + D - Fz - kn, \qquad D = Fz \tag{1.1}$$

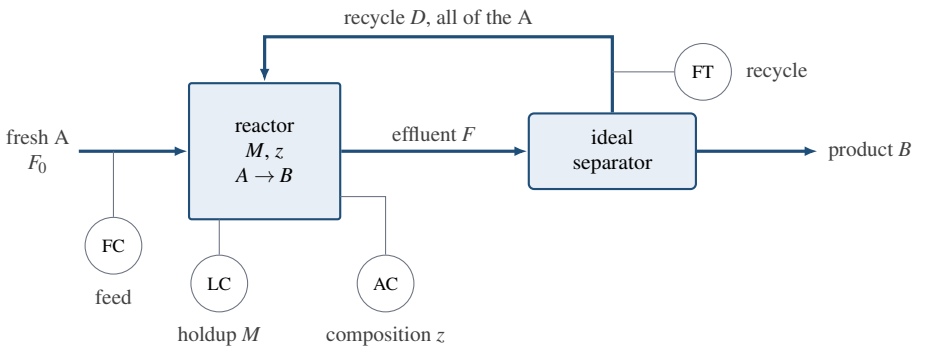


Figure 1.2 The Running Example. Fresh A Meets the Recycle in a Stirred Reactor. An Ideal Separator Returns All of the A and Sends the B Out as Product. The Instruments Measure the Fresh Feed, the Reactor Holdup, the Reactor Composition and the Recycle Flow.

At the design point $F_0 = 200\,\text{kmol/h}$ and $M = 300\,\text{kmol}$. Because every mole of A that leaves the reactor returns, A leaves the plant only by reaction, and the steady state needs $kn = F_0$. The reactor therefore holds $n = F_0/k = 200\,\text{kmol}$ of A whatever its size, which gives $z = 2/3$, a recycle $D = 400\,\text{kmol/h}$, an effluent $F = 600\,\text{kmol/h}$ and a product $B = F_0 = 200\,\text{kmol/h}$. The holdup only dilutes the A with B.

Two simplifications keep the arithmetic exact, and both are stated where they matter. The separator is ideal and holds nothing. The reactor holdup follows its level setpoint as a first-order lag of 1 h, the speed at which the level controller is allowed to move the effluent. Simulations use the velocity-form PI of Section 1.6.2, executed every 0.01 h, with the recycle flow and the reactor composition read through transmitter lags of 0.02 h. The plant is integrated by the fourth-order Runge–Kutta method on a step of 0.001 h. Chapter 7 runs the full version of this plant, with a twenty-tray column, a vapor limit and a recycle line.

1.4.2 Degrees of Freedom

A control structure starts with a count. Each valve, pump speed or heater duty that can be set independently is a dynamic degree of freedom. A real version of the example has seven: the fresh feed, the effluent, the recycle, the product, the separator's reflux and boilup, and the reactor coolant. Some of them must stabilize inventories, because a liquid level on a pumped outlet is an integrator and drifts unless a loop holds it. At steady state, a level whose value changes nothing else, such as the separator's drum and sump, uses up a degree of freedom without giving a choice. The steady-state count is therefore Equation (1.2) [3].

$$N_{ss} = N_{\text{valves}} - N_0 \tag{1.2}$$

where N_0 is the number of inventories with no steady-state effect. Here $N_{ss} = 7 - 2 = 5$, and the five choices are the production rate, the reactor holdup, the reactor temperature and the two purities of the separator. The reactor holdup is a choice because it changes the conversion.

> **+ DID YOU KNOW? Thousands of Loops, a Few Decisions**
>
> A chemical plant may have thousands of measurements and control loops. Plantwide control is not the tuning of each one but the structural decisions about the whole plant [2], and those decisions are in most cases made ad hoc, from experience and engineering insight [3].

The five are then spent. Some go to active constraints, such as a purity held at its sales limit or a temperature held at its upper limit, and the rest go to variables whose constant setpoints keep the

plant near its optimum as disturbances change, which is the idea of self-optimizing control [15]. In the example the reactor temperature and both purities are held, which leaves two handles: the fresh feed F_0 and the holdup M. The rest of this section is about those two.

The count is the first step of most chapters. Chapter 2 has two actuators and six variables worth holding, so two receive setpoints and four follow. Chapter 8 has twelve valves and several inventories that are integrating or unstable. Chapter 7 spends its handles on five level and purity loops and a flow. Chapter 9 shows that choosing the controlled variable is part of the count. Its chromatography cut holds the purity of the pool, while the filter two units later responds to one particular impurity that the purity does not single out.

1.4.3 Inventory Control and the Throughput Manipulator

Every inventory must be held, and the production rate must be set at one place only. These two requirements fix the direction of the level loops. A vessel downstream of the TPM is held by its outflow, in the direction of flow, and a vessel upstream of it by its inflow, against the flow. A structure that obeys the rule for every vessel and every group of vessels is consistent [11, 16]. Placing the TPM is the larger decision [17]. A TPM at the feed sends every change of production through each unit in turn, while a TPM inside the plant changes the units around it first and leaves the vessels upstream to adjust their inflows.

The example has two common structures, with the names Chapter 7 gives them. In CS1 the fresh feed is on flow control and is the TPM. The reactor level then sets the effluent, and the separator's levels set the recycle and the product, all in the direction of flow. In CS2 the effluent is on flow control. The reactor level sets the fresh feed, against the flow, and the production rate is moved by the reactor level setpoint. Both are consistent.

An inconsistent structure is easy to build. With the fresh feed and the product both on flow control, the plant's total holdup integrates $F_0 - B$, and any difference between the two settings fills or empties some vessel. Consistency must also hold for each component. In the example the reaction removes A at the rate kn, so the A inventory regulates itself at $n = F_0/k$. An inert that neither reacts nor leaves with a product has no such exit and accumulates until a purge removes it. That is the pressure of Chapter 8, which trips its interlock within hours with every valve fixed [8], and closes its inventories before anything else.

The chapters place the TPM in several ways. The production index of Chapter 8 and the line speed of Chapter 5 set the flows or the pump in ratio, so one setting moves the whole plant at once. Chapter 3 moves the machine speed to hold a tonnage, and the permeate tank of Chapter 2 cascades slowly onto the production setpoint. Chapter 4 can hold its product solids with the steam or with the feed, and the second choice makes the throughput the MV. Chapter 10 is a closed inventory whose machines can only add gas, so nothing removes an excess except the crew and the leak.

1.4.4 Recycle and the Snowball Effect

With the holdup fixed, the steady state of the example follows from the balances alone. The reactor holds $n = F_0/k$, so $z = F_0/(kM)$, and the recycle that the separator returns is Equation (1.3).

$$D = \frac{F_0\, z}{1 - z} = \frac{F_0^2}{kM - F_0}, \qquad \frac{\mathrm{d}\ln D}{\mathrm{d}\ln F_0} = 2 + \frac{F_0}{kM - F_0} \tag{1.3}$$

The denominator is the margin between the reactor's capacity to convert A, $kM = 300\,\mathrm{kmol/h}$, and the feed. At the design point the elasticity is $2 + 200/100 = 4$, so a percent of feed is four percent of recycle, and the slope is 8 kmol/h of recycle per kmol/h of feed. A 10 % step of fresh feed, from 200 to 220 kmol/h, takes the recycle from 400 to 605 kmol/h, 51.25 % more, and the separator's feed from 600 to 825 kmol/h. This is the snowball effect [9]. The recycle grows faster than the production because a reactor of fixed size can convert more A only by running richer in A, and the separator returns that richer effluent's A. At $F_0 = kM$ no steady state exists.

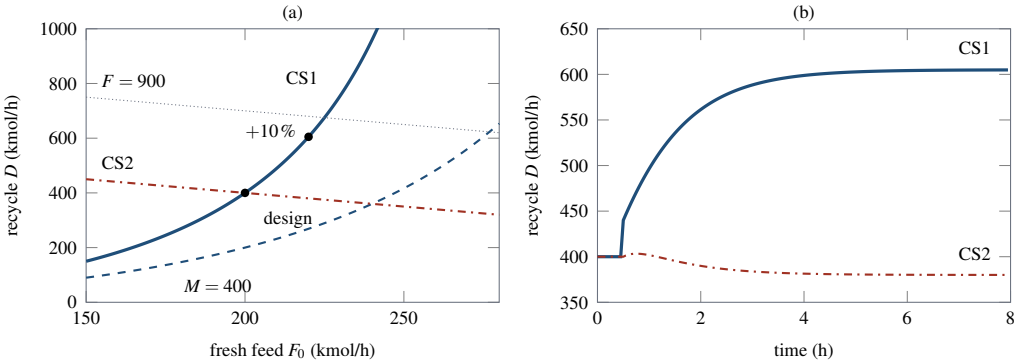


Figure 1.3 The Snowball on the Example. (a) Steady Recycle Against Fresh Feed: CS1 at the Design Holdup (Thick Solid) and at 400 kmol (Dashed), CS2 With the Effluent Fixed (Dash-Dotted) and the Separator Limit $F_0 + D = 900\,\text{kmol/h}$ (Dotted). (b) The Recycle After a 10 % Production Increase at 0.5 h, by the Fresh Feed in CS1 (Thick Solid) and by the Holdup Setpoint in CS2 (Dash-Dotted).

The snowball sets the capacity of the plant. If the separator can take 900 kmol/h, half again its design feed, CS1 reaches that limit at $F_0 = kMF/(kM+F) = 225\,\text{kmol/h}$, only 12.5 % above the design. The holdup is the other handle. In CS2 the effluent stays at 600 kmol/h, so $D = 600 - F_0$ and the recycle falls as production rises. The holdup must grow instead: 10 % more production needs $M = 347.37\,\text{kmol}$, 15.79 % more, and a vessel limited to 400 kmol allows 240 kmol/h. With both the vessel and the separator at their limits the plant makes 276.92 kmol/h, which is why the analysis of this plant in the literature places the reactor at its largest holdup [18]. Placing one flow inside the recycle loop on flow control, as CS2 does, is a general rule [19]. It does not remove the snowball. It moves it from the recycle into the reactor holdup. Figure 1.3(a) draws both structures.

A recycle also changes the dynamics, because it is a feedback path inside the plant [20]. Without the recycle the A in the reactor would settle with the time constant $M/(F+kM) = 0.333\,\text{h}$. With every mole of A returned, A can leave only by reaction, and the time constant becomes $1/k = 1\,\text{h}$, three times longer. A separator that holds some of the recycled material lengthens it further. With a lag of 0.1 h in the separator the slow time constant of the linearized plant is 1.73 h. Figure 1.3(b) shows the recycle after the production step. In CS1 it jumps from 400 to 440 kmol/h as the extra feed passes through, then climbs toward 605 kmol/h and covers 63 % of the change in 0.95 h. In CS2 it falls slowly toward 380 kmol/h, at the pace the holdup is allowed to grow.

Chapter 7 measures the snowball on its column, finds the vapor limit at 10.7 % more feed and switches among three structures during a run. The recycle of Chapter 8 carries gas and an inert, and its pressure is held by the purge.

1.4.5 The Relative Gain Array and Pairing

Suppose the operator of the example wants two things held. The recycle D is the load on the separator, and the reactor composition z sets the selectivity of a real reaction. The handles are the fresh feed F_0 and the holdup M. Differentiating the steady state at the design point gives the gain matrix of Equation (1.4), with z in mol %.

$$\begin{bmatrix}\Delta D\\ \Delta z\end{bmatrix} = \begin{bmatrix}8 & -4\\ 1/3 & -2/9\end{bmatrix}\begin{bmatrix}\Delta F_0\\ \Delta M\end{bmatrix}, \qquad G(s) = \frac{1}{s+1}\begin{bmatrix}2s+8 & -(2s+4)\\ 1/3 & -2/9\end{bmatrix} \tag{1.4}$$

The transfer matrix carries the dynamics. Extra feed reaches the recycle at once, and the rest follows the A inventory with its 1 h time constant. A larger holdup takes effluent while it fills, which lowers the recycle at once.

The relative gain of a pairing compares the gain of a loop with the other loop open against its gain with the other loop closed perfectly, as in Equation (1.5) [7].

$$\lambda_{11} = \frac{(\partial y_1/\partial u_1)_{u_2}}{(\partial y_1/\partial u_1)_{y_2}} = \frac{1}{1 - \dfrac{K_{12}K_{21}}{K_{11}K_{22}}}, \qquad \Lambda = K \circ \left(K^{-1}\right)^{\mathsf{T}} \tag{1.5}$$

where $\circ$ multiplies element by element. Each row and column of Λ sums to one, and Λ does not depend on the units. For the example $\lambda_{11} = 1/[1-(-4)(1/3)/(8\cdot(-2/9))] = 1/(1-3/4) = 4$, so

$$\Lambda = \begin{bmatrix} 4 & -3 \\ -3 & 4 \end{bmatrix} \tag{1.6}$$

Pairing 1 holds the recycle with the feed and the composition with the holdup, on the diagonal, with $\lambda = 4$. Each loop then has four times the gain alone that it has with the other loop closed, so closing the second loop makes the first one slower. Pairing 2 swaps them and has $\lambda = -3$. A negative relative gain means that the steady gain of a loop changes sign when the other loop is closed. With the recycle held by the holdup, more feed makes the holdup grow so much that the reactor becomes leaner, while with the holdup fixed more feed makes it richer.

The consequence is stated by the Niederlinski index [12], Equation (1.7), for a pairing whose loops have integral action and are each stable alone.

$$\mathrm{NI} = \frac{\det K}{\prod_i K_{ii}} \tag{1.7}$$

with the inputs ordered so that each output's paired input is on the diagonal. A negative index means the closed loop is unstable whatever the tuning. For two loops $\mathrm{NI} = 1/\lambda_{11}$, which is 0.25 for pairing 1 and $-1/3$ for pairing 2. The usual rules follow: pair on positive relative gains, prefer values near one, and treat values much larger than one as a sign of sensitivity to model error [21].

Table 1.1 and Figure 1.4 test the prediction. Every loop is tuned alone, with the other in Manual, by the IMC rule of Volume 2 on its own steady gain, a time constant of 1 h and $\tau_c = 0.5$ h, with $\tau_I = 1$ h. The recycle setpoint steps from 400 to 440 kmol/h at 0.5 h with the composition setpoint unchanged. The recycle loop of pairing 1 alone reaches 90 % of the step in 1.30 h, and the uncontrolled composition drifts up by 1.58 points. With both loops closed the recycle takes 5.69 h, 4.4 times as long, and its IAE rises from 18.0 to 72.1 kmol, the factor of about four that $\lambda = 4$ predicts. Pairing 2 runs away. The composition rises, the fresh feed falls to correct it, the recycle falls, and the holdup falls to restore the recycle, which raises the composition further. By 8 h the feed is down from 200 to 38 kmol/h, the holdup setpoint has reached its 60 kmol limit at 7.67 h, and the recycle is 216 kmol/h below its setpoint. Each loop of pairing 2 is stable alone. The recycle loop on the holdup alone settles with an IAE of 17.2 kmol.

The relative gain can also say nothing is wrong when something is. A triangular gain matrix, with one zero off the diagonal, has $\lambda_{11} = 1$ however large the other off-diagonal gain is. The coupling then runs one way. Chapter 2 reports $\lambda_{11} = 1.00005$ on a train whose pump moves the recovery strongly, and Chapter 10 has an oxygen loop and a total pressure loop coupled through one sum, where added oxygen raises the total and the nitrogen does not move the oxygen. Chapter 3 has two-way coupling with $\Lambda_{11} = 1.517$ and a swapped pairing at -0.517 that its run cannot hold. Chapter 8 runs a decentralized structure of seventeen PI loops, and Chapter 4 compares the steam and the feed as the handle on its product solids.

1.4.6 Decoupling

A decoupler adds to each actuator a correction computed from the other controller's output, so that each controller sees only its own loop. The simplified static form of Equation (1.8) needs only two

Table 1.1 Pairings and Decouplers on the Example. Left: the Recycle Setpoint Steps From 400 to 440 kmol/h. Right: the Catalyst Activity Falls 10 %. Both at 0.5 h, Runs of 8 h. t_{90} Is the Time to 90 % of the Step, IAE Is That of the Recycle and $|\Delta z|$ the Largest Composition Error.

	Recycle setpoint step			Activity loss	
Structure	t_{90} (h)	IAE (kmol)	$\|\Delta z\|$ (pts)	IAE (kmol)	$\|\Delta z\|$ (pts)
Pairing 1, recycle loop alone	1.30	18.0	1.58	61.1	2.73
Pairing 1, both loops	5.69	72.1	0.77	1.5	1.74
Both loops, decoupler	3.81	78.6	0.00	247.1	3.51
Both loops, decoupler, retuned	1.40	39.1	0.00	155.6	1.76
Both loops, half decoupler	1.62	25.4	1.11	69.7	2.39
Pairing 2, both loops	never	414.9	7.25	221.3	3.80

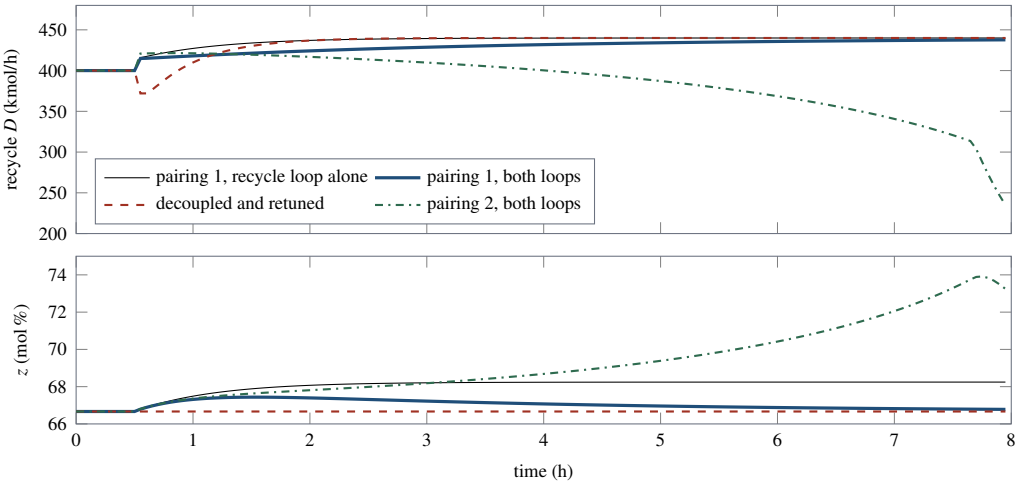


Figure 1.4 The Recycle Setpoint Step of Table 1.1. Thin Solid: Pairing 1, Composition Loop in Manual. Thick Solid: Pairing 1, Both Loops. Dashed: Pairing 1 With the Decoupler, Both Loops Retuned, Whose Recycle First Moves the Wrong Way. Dash-Dotted: Pairing 2, Which Drifts Away From Both Setpoints.

gain ratios [22].

$$u_1 = c_1 + d_{12}\,(c_2 - \bar{u}_2), \qquad u_2 = c_2 + d_{21}\,(c_1 - \bar{u}_1), \qquad d_{12} = -\frac{K_{12}}{K_{11}}, \quad d_{21} = -\frac{K_{21}}{K_{22}} \tag{1.8}$$

where c_i are the controller outputs and $\bar{u}_i$ the design values of the actuators. For the example $d_{12} = 0.5$ kmol/h of feed per kmol of holdup and $d_{21} = 1.5$ kmol of holdup per kmol/h of feed. At steady state each controller then sees the gain K_{ii}/λ_{ii}, which is 2 kmol/h per kmol/h for the recycle loop and $-1/18$ mol % per kmol for the composition loop. These are the gains each loop had with the other loop closed. A decoupler removes the interaction, not the loss of gain, so each loop must be tuned on the smaller gain.

The decoupler is static, and the two paths it balances do not have the same dynamics. In the example both paths to the composition share the 1 h time constant, so the composition is decoupled exactly at every instant. The paths to the recycle differ. The feed acts at once and the holdup's correction acts at once in the opposite direction, so the decoupled recycle loop sees $(2-s)/(s+1)$, a gain of 2 with an inverse response. Table 1.1 shows the result. With the old tunings the composition no longer moves at all, but the recycle reaches 90 % after 3.81 h. Retuned on the decoupled gains, with the inverse response counted as 0.5 h of dead time in the SIMC rule, the recycle loop takes

1.40 h and the composition still does not move. Half the decoupler, the term that moves the feed with the holdup, gives the recycle a t_{90} of 1.62 h and leaves the composition an error of 1.11 points.

The disturbance columns tell the other half. A 10 % loss of catalyst activity is cancelled exactly by raising the holdup to 333.3 kmol, because k and M enter the plant only as the product kM. The undecoupled pairing does exactly that, since its composition loop moves the holdup alone, and leaves the recycle an IAE of 1.5 kmol. Each decoupler tested here adds a feed move to that holdup move and so disturbs the recycle, from 69.7 kmol with half the decoupler to 247.1 kmol with the whole of it untuned. A decoupler designed for setpoint changes is not neutral to disturbances.

> ✦ DID YOU KNOW? **Retuned, Still Not Neutral**
>
> From the simulator: after the 10 % loss of catalyst activity, the decoupler retuned on its decoupled gains still leaves the recycle an IAE of 155.6 kmol, against 1.5 kmol for the plain pairing, while the largest composition errors are 1.76 and 1.74 points (Table 1.1).

Chapter 3 finds the same split. Its full decoupler cuts a moisture excursion from 1.80 to 0.23 points, one of its two terms carries nearly all of that, and the weight loop becomes slower. Chapter 2 needs only one term, because its plant is triangular, and its sweep of that term finds the least error near the design gain. Chapter 3 also notes that a static decoupler puts the right correction on at the wrong time.

1.4.7 Singular Values and Directionality

A gain matrix acts more strongly in some directions than in others. The singular value decomposition $K = U\Sigma V^{\mathsf{T}}$ finds them. Column j of V is a direction of the inputs, column j of U the direction of the outputs it produces, and σ_j the gain along it. The condition number $\gamma = \sigma_1/\sigma_{\min}$ is the ratio of the strongest to the weakest gain [1].

Singular values depend on units, so the matrix must first be scaled. Scaled by the design values, so that every entry is a percent per percent, the example's gain matrix and its inverse are Equation (1.9).

$$K_r = \begin{bmatrix} 4 & -3 \\ 1 & -1 \end{bmatrix}, \qquad K_r^{-1} = \begin{bmatrix} 1 & -3 \\ 1 & -4 \end{bmatrix}, \qquad \sigma_1 = 5.19, \quad \sigma_2 = 0.193 \tag{1.9}$$

The condition number is 27.0. The strong input direction is $(0.79, -0.61)$, more feed with less holdup, which closes the margin $kM - F_0$ and moves mainly the recycle. The weak direction is $(0.61, 0.79)$, feed and holdup together, which keeps that margin and moves mainly the composition, as Figure 1.5 shows. The inverse says the same in whole numbers. A percent more recycle at constant composition needs one percent more feed and one percent more holdup, a move of length 1.41. A percent more composition at constant recycle needs three percent less feed and four percent less holdup, a move of length 5.

Three consequences follow. A controller that inverts the plant asks for five times the move of the output change in the weak direction, so noise and model error in that direction are amplified five times and actuators saturate sooner. A disturbance that lies along the strong direction is cheap to reject and one along the weak direction is expensive. The catalyst loss lies exactly along a holdup move, which is why the plain pairing rejects it so well. Finally, the condition number depends on the scaling, while the relative gain does not. In engineering units the same matrix has $\gamma = 180$. No choice of units brings it below 13.9, and for two by two plants that floor is fixed by the relative gains, here $7 + \sqrt{48}$, where 7 is the sum of the magnitudes in one row of Λ [21]. A large relative gain means a large condition number in every scaling. The reverse does not hold. The triangular sum of Chapter 10, a matrix with rows $(1, 0)$ and $(1, 1)$, has $\gamma = 2.62$ and a floor of one, and the gains of Chapter 2 have $\gamma = 22.3$ in their units and a floor of 1.01.

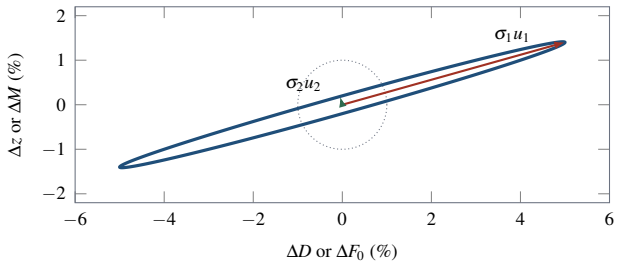


Figure 1.5 Directionality of the Example. Every Input Move of Length One Percent (Dotted Circle) Produces an Output Change on the Thick Ellipse, Whose Half-Axes Are $\sigma_1 = 5.19$ Along u_1 and $\sigma_2 = 0.193$ Along u_2.

Chapter 6 is this subsection at scale. Its 96 by 96 response has $\gamma = 405$, and the chapter chooses how many modes to correct by truncating the inverse or by Tikhonov regularization, trading residual orbit against corrector noise.

1.4.8 Disturbance Propagation

A control system does not remove a disturbance. It moves the variation from the controlled variables to the manipulated ones, and the structure decides which ones. After the 10 % catalyst loss with the holdup and feed fixed, CS1 sends the upset into the recycle, which settles at 571.43 kmol/h, 42.86 % above design. CS2 with the effluent and holdup fixed sends it into production, which falls to 186.21 kmol/h, 6.90 % below. Neither structure is better in general. The choice is which variable the plant can afford to move.

Inventories pass variation from unit to unit. A vessel whose level is held by its outflow passes its inflow to the next unit, filtered by its level loop. With proportional control the outflow follows the inflow as a first-order lag and never exceeds it. With PI control the level returns to its setpoint, and the outflow follows Equation (1.10). Tuning such loops to use the vessel as a buffer is averaging level control [23].

$$\frac{Q_{\text{out}}(s)}{Q_{\text{in}}(s)} = \frac{2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \qquad \left|\frac{Q_{\text{out}}}{Q_{\text{in}}}(j\omega_n)\right| = \sqrt{1 + \frac{1}{4\zeta^2}} \tag{1.10}$$

where $\omega_n = \sqrt{K_c'/\tau_I}$ and $\zeta = \frac{1}{2}\sqrt{K_c'\tau_I}$, with K_c' the controller gain in outflow per unit of held volume. Its amplitude ratio exceeds one at every frequency below $\sqrt{2}\,\omega_n$, whatever the damping. The peak is 1.47 at $\zeta = 0.5$ and 1.15 at $\zeta = 1$, and three such vessels in series pass 3.16 times a flow variation at $\zeta = 0.5$, as Figure 1.6 shows. The gain at zero frequency is exactly one, because a vessel must in the end pass on what it receives.

> **+ DID YOU KNOW? Less Buffer, More Structure**
>
> Setting production at the feeds and designing the loops unit by unit, in sequence, works for processes with much inventory between the units, including buffer tanks, and little recycle. As plants hold less inventory, the plantwide structure becomes a more important part of the design [13].

A chain whose every stage amplifies some frequency is not string stable, and a disturbance grows as it travels [24]. A supply chain is the same structure with orders in place of outflows, and the growth of order variation upstream is the bullwhip effect [10]. The recycle is a third path. In the example a slow change of feed reaches the recycle four times enlarged in relative terms, while at 1 rad/h the factor is 2.9 and at 10 rad/h it is 1.07, because the recycle's own dynamics take an hour.

Chapter 11 computes the amplitude ratio of each ordering rule, finds a peak of 1.89 for its human rule and no rule that is string stable, and measures the variance growing stage by stage. Chapter 4

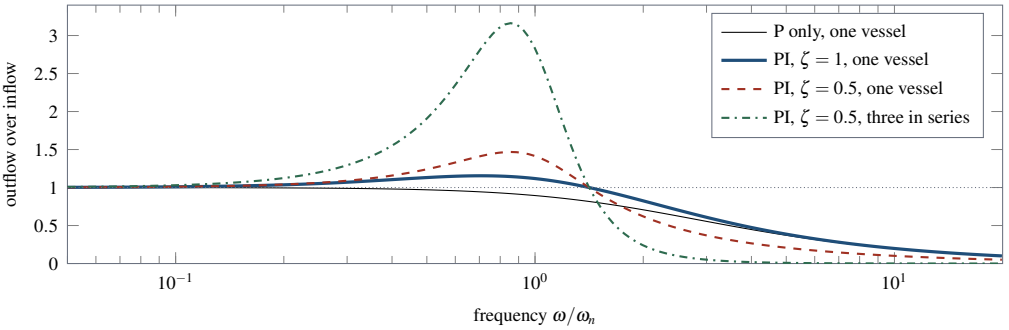


Figure 1.6 How Level Loops Pass Flow Variation Downstream. Thin Solid: Proportional Control, With the Same Gain as the PI at $\zeta = 1$. Thick Solid and Dashed: PI at $\zeta = 1$ and 0.5. Dash-Dotted: Three PI Vessels at $\zeta = 0.5$ in Series.

Table 1.2 The Structural Question of Each Chapter and the Measure It Reports

Ch.	Structural question	Measure reported
2	Which actuator holds production and which recovery	Relative gain of a triangular plant, one-way decoupler gain
3	Pairing of weight and moisture, and which half of a decoupler	Relative gain of two-way coupling, off-specification tonnage
4	Steam or feed on the product solids, tight or averaging levels	Steady cost of each handle, flow variation passed to the dryer
5	Line speed as the throughput handle, with feedforward to tension and coating	Scrap in meters after a speed change
6	How many directions of the inverse to correct	Singular values, condition number, noise gain per mode
7	Which flow in the recycle is fixed and where production is set	Snowball gain, vapor margin
8	Inventory loops before quality, one handle for production	Time to shutdown, operating cost
9	Which variable the upstream cut should hold	Late impurity in the pool against filter time
10	Two loops on one sum in a closed inventory	Valve opening as the signal of a leak
11	The ordering rule at each stage of a chain	Peak amplitude ratio, variance ratio by stage

compares tight and averaging level loops by the flow they pass to the dryer, and Chapter 8 finds that a stripper level held tight passes the disturbance on to production. Chapter 5 follows one speed change along two paths, and Chapter 9 follows one cut to the filter.

1.4.9 What Predicts a Structure

The measures of this section answer different questions. The degree-of-freedom count says how many choices exist. The consistency rule says whether every inventory is held. The snowball elasticity says how much a recycle magnifies a production change. The relative gain and the Niederlinski index say whether a pairing works and how much its loops slow each other. The decoupled gains K_{ii}/λ_{ii} say what a decoupler leaves. The singular values say which directions are expensive, and the amplitude ratio of an inventory says whether a chain grows a disturbance. On the example the measures computed from the gains predicted what the simulations then showed. Table 1.2 lists the structural question of each chapter and the measure it reports, as groundwork for the comparison in Chapter 12.

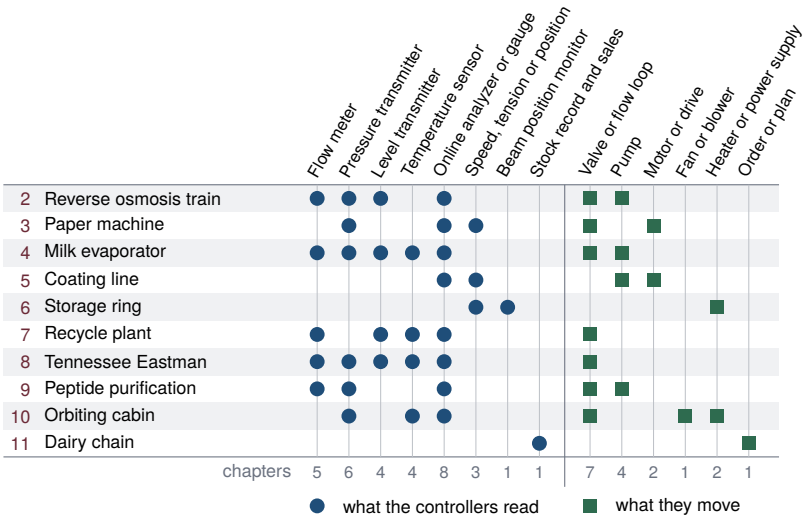


Figure 1.7 What the controllers of each chapter read and move: a round mark where a controller reads the instrument and a square mark where it moves the element, as the chapter's variables table and its technology section name them. A flow set through a flow station counts as a valve or flow loop, and a computed value, such as a recovery or a tonnage, is not marked. Notice that the online analyzers and gauges appear in more chapters than any other instrument.

1.5 Instruments and Final Elements in This Volume

A pairing joins an instrument to a final element, so the devices of a plant bound the structures it can have. Figure 1.7 lists what the controllers of each chapter read and what they move. An online analyzer or gauge appears in eight chapters and a pressure transmitter in six, and a valve or flow loop is the final element in seven. Two chapters stand apart. The controllers of Chapter 6 read beam position monitors and move the supplies of corrector magnets, and those of Chapter 11 read a stock record kept from deliveries and till sales and move an order.

The production handle is usually a flow or a speed. A flow handle sits on a flow loop. A magnetic flowmeter, top of Figure 1.8, reads the voltage that a moving liquid induces across a magnetic field, in proportion to its velocity, and asks only that the liquid's electrical conductivity meet a certain minimum [25]. Chapter 2 reads its recovery as the ratio of two flow measurements. A speed handle moves a machine. Varying the speed of the machine that pushes a fluid is more energy-efficient than throttling its discharge or recycling fluid to its suction, and it removes a control valve from the line [25]. The bottom of the figure shows such a drive on a pump, the arrangement of the high-pressure pump of Chapter 2. The machine speed of Chapter 3 and the line speed of Chapter 5 are handles of the same kind.

The quality measurements are the slow ones. A process chromatograph separates species in time, so cycle times of minutes are not uncommon even online, and some measurement dead time is inevitable [25]. The analyzers of Chapter 8 report every 0.1 h and 0.25 h, the scanning gauge of Chapter 3 posts one average per pass, and the cabin of Chapter 10 hears from its mass spectrometer once a minute. In the first two, the flow, pressure and level loops act many times between two results.

The structures sit in layers. The control system of a chemical plant is usually divided by time scale into scheduling in weeks, site-wide optimization in days, local optimization in hours, supervisory control in minutes and regulatory control in seconds [3], and a survey of predictive control draws basic control executing about every second under constraint control about every minute [26]. A distributed control system mediates a group of localized controllers through a supervisory level, programmable logic controllers can serve as those local controllers, and the loops run with cycle

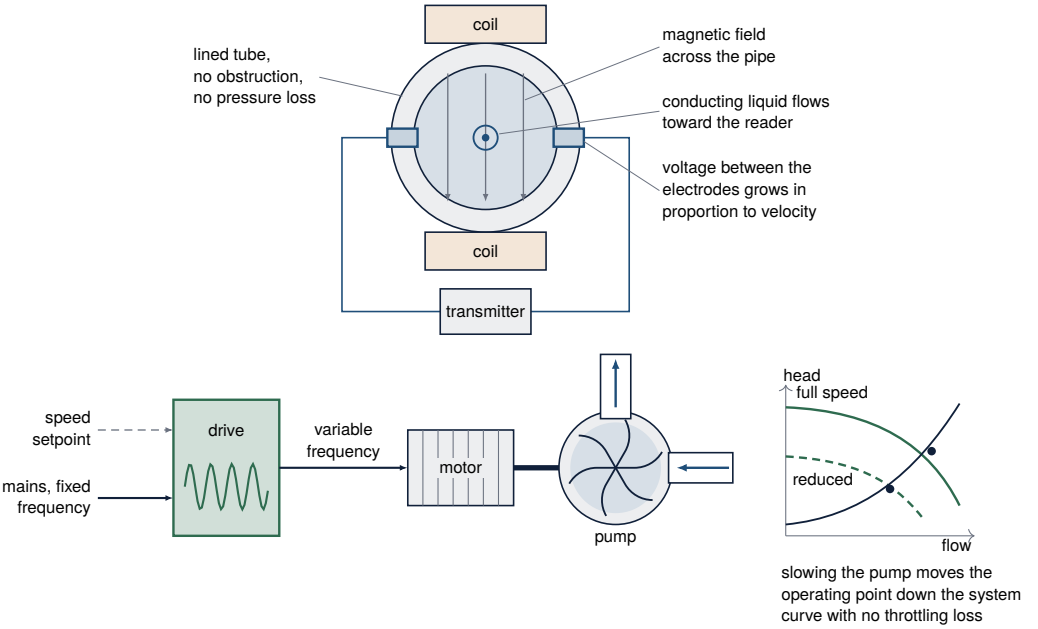


Figure 1.8 Top: principle of an electromagnetic flow meter. Bottom: principle of a variable-speed drive on a centrifugal pump. Notice that the meter reads the stream through electrodes at the wall of the tube, and that the drive moves the pump onto a new curve on the same system curve, so the flow changes with no valve in the line.

Layer	What it does in this volume	Where it runs	How often
Plan or ordering calendar	a power plan for the cabin's beds and oxygen, the orders of a dairy chain	**Local optimization** over the loops; a store's **stock management system**	at each swap or sunrise in the cabin; once a day in the dairy chain
Supervisory structure	a production index with ratio stations, selectors, a decoupler, a predictive controller	**Supervisory level** of a DCS over its localized controllers	0.05 s on the coating line to 0.05 h on the Tennessee Eastman plant
Base loops	PI loops on flows, levels and qualities, a controller per singular mode	**Localized controllers** and **PLCs**; **cell controllers** or arrays of signal processors at a ring	1 ms to 10 ms on the ring and the coating line; 0.5 s to 6 s in the plants; 1 min in the cabin
Sensors and final elements	flow meters, analyzers, gauges, monitors; valves, pumps, drives, magnet supplies	wires, sample lines, a scanning frame, a stock record and a truck	a gauge pass every 20 s; analyzer results every 0.1 h to 0.25 h

Figure 1.9 The control layers of this volume, the platform that hosts each and how often each acts. Platforms are from [27] and [3] and from the chapters' technology sections, and the periods from the chapters. Notice that the supervisory structure of a fast line can act more often than the base loops of a slow plant.

times from milliseconds to minutes [27]. Figure 1.9 places the loops of this volume on their platforms. They run from the millisecond of the fast orbit loop of Chapter 6 to the daily ordering decision of Chapter 11.

Table 1.3 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

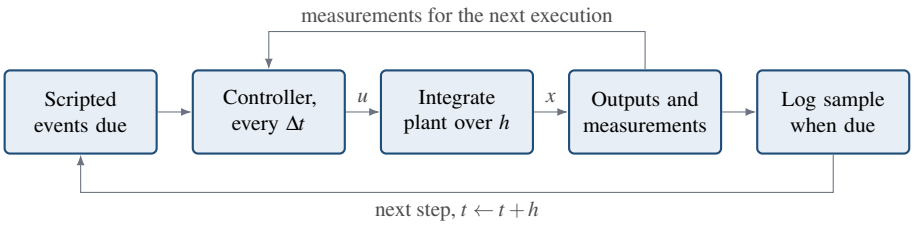


Figure 1.10 Work of the Engine in One Integration Step

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{dx}{dt} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.10. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings

are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.11) to (1.14) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.11}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.12}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.13}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.14}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.14) drives I toward the value that reproduces the output applied, which is back-calculation [28]. With no integral action, Equation (1.14) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.15).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.15}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.16) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^\top Q\,(y_k - r_k) + \sum_j \Delta u_j^\top R\,\Delta u_j + \sum_i \rho_i s_i^2 \tag{1.16}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [29]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [30]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [31].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according

to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed (the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 12.

Pairing: which handle, which target	Directions: how many to correct	Inventories: where production is set	Propagation: where an upset goes
(2) a relative gain of one on a plant that interacts	(6) an ill-conditioned inverse, modes to keep	(4) steam or throughput as the handle	(5) one speed change on two paths
(3) a static decoupler through a scanning gauge		(7) the snowball and the production handle	(9) a cut judged two units downstream
(10) two loops on one sum in a closed cabin		(8) inventories first, then one handle	(11) orders that grow along a chain

Figure 1.11 Method map of the volume. Each chapter is placed by the part of Section 1.4 that its structural question in Table 1.2 rests on, and labeled with its lesson, shortened from Table 1.4. The placement is the editor's reading of the two. Notice that three of the four questions draw three chapters each, and that only the storage ring turns on directions.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The structural lesson	Level	Time
2	Seawater reverse osmosis train	A pairing chosen against a gain that changes sign. A relative gain of one on a plant that interacts.	introductory	3 h
3	Paper machine with a scanning gauge	Pairing and a static decoupler through a sampled, averaged measurement.	intermediate	3.5 h
4	Three-effect milk evaporator	Steam or throughput as the handle, and averaging levels along a chain.	intermediate	2.5 h
5	Roll-to-roll coating line	One speed change on two paths, held by feedforward and compensation.	intermediate	2.5 h
6	Storage ring orbit	An ill-conditioned inverse and how many modes to correct.	advanced	2 h
7	Reactor, column and recycle	The snowball and the choice of the production handle.	introductory	2 h
8	Tennessee Eastman plant	Inventories first, then one handle for production, on a benchmark.	advanced	2.5 h
9	Peptide purification route	An upstream decision judged by a unit two steps downstream.	intermediate	2 h
10	Sealed orbiting cabin	Two loops on one sum, and a closed inventory that can only be added to.	intermediate	2 h
11	Perishable dairy chain	Inventories in series and the growth of variation upstream.	intermediate	2.5 h

1.8 Chapter Roadmap

Table 1.4 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case. Figure 1.11 places the same chapters by the part of Section 1.4 that each one's structural question rests on.

Chapters 2 and 3 form a first course in pairing and are best read in order, with Sections 1.4.5 and 1.4.6. Chapter 7 then carries Sections 1.4.3 and 1.4.4 to a full column, and Chapter 8 is best read after it. Chapter 6 needs Section 1.4.7 and stands alone otherwise. Chapters 4, 5, 9, 10 and 11 can be read in any order after Section 1.4.8. A reader who wants only the plantwide thread can take Chapters 7, 8, 4 and 11.

KEY POINTS

- The control structure of a plant is chosen before any loop is tuned, and a wrong pairing or a recycle that snowballs cannot be repaired by tuning.
- Every inventory must be held and the production rate set at one place only, which fixes the direction of every level loop.
- A negative relative gain means the closed loop is unstable whatever the tuning, and a decoupler removes the interaction but not the loss of gain.
- A control system does not remove a disturbance. It moves the variation to the manipulated variables, and level loops in a chain can amplify it.

Where else this applies. Recycle amplification and inventory pairing in an ammonia synthesis loop in Volume 13, Chapter 2, degrees of freedom and triangular interaction on a crude tower in Volume 11, Chapter 4, and offset-free predictive control of a boiler and turbine in Volume 9, Chapter 4.

References

[1] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd ed. Wiley, 2005.

[2] T. Larsson and S. Skogestad, "Plantwide control – A review and a new design procedure," *Modeling, Identification and Control: A Norwegian Research Bulletin*, vol. 21, no. 4, pp. 209–240, 2000.

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

[4] N. Rillieux, "Improvement in sugar-works," United States Letters Patent No. 3,237, Aug. 1843, patented August 26, 1843; image of the specification and drawings served by the United States Patent and Trademark Office.

[5] I. J. Martinez-Moyano, "History of the beer game," *System Dynamics Review*, vol. 40, no. 2, p. e1767, 2024.

[6] K. J. Åström, "Computer control of a paper machine—an application of linear stochastic control theory," *IBM Journal of Research and Development*, vol. 11, no. 4, pp. 389–405, 1967.

[7] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[8] J. J. Downs and E. F. Vogel, "A plant-wide industrial process control problem," *Computers & Chemical Engineering*, vol. 17, no. 3, pp. 245–255, 1993.

[9] W. L. Luyben, "Snowball effects in reactor/separator processes with recycle," *Industrial & Engineering Chemistry Research*, vol. 33, no. 2, pp. 299–305, 1994.

[10] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Management Science*, vol. 43, no. 4, pp. 546–558, 1997.

[11] E. M. B. Aske and S. Skogestad, "Consistent inventory control," *Industrial & Engineering Chemistry Research*, vol. 48, no. 24, pp. 10 892–10 902, 2009.

[12] A. Niederlinski, "A heuristic approach to the design of linear multivariable interacting control systems," *Automatica*, vol. 7, no. 6, pp. 691–701, 1971.

[13] 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.

[14] Y. Chung, G. Decker, and K. Evans, Jr., "Closed orbit correction using singular value decomposition of the response matrix," in *Proceedings of International Conference on Particle Accelerators*. IEEE, 1993, pp. 2263–2265.

[15] S. Skogestad, "Plantwide control: the search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[16] R. M. Price and C. Georgakis, "Plantwide regulatory control design procedure using a tiered framework," *Industrial & Engineering Chemistry Research*, vol. 32, no. 11, pp. 2693–2705, 1993.

[17] R. M. Price, P. R. Lyman, and C. Georgakis, "Throughput manipulation in plantwide control structures," *Industrial & Engineering Chemistry Research*, vol. 33, no. 5, pp. 1197–1207, 1994.

[18] T. Larsson, M. S. Govatsmark, S. Skogestad, and C. C. Yu, "Control structure selection for reactor, separator, and recycle processes," *Industrial & Engineering Chemistry Research*, vol. 42, no. 6, pp. 1225–1234, 2003.

[19] M. L. Luyben, B. D. Tyreus, and W. L. Luyben, "Plantwide control design procedure," *AIChE Journal*, vol. 43, no. 12, pp. 3161–3174, 1997.

[20] J. Morud and S. Skogestad, "Dynamic behaviour of integrated plants," *Journal of Process Control*, vol. 6, no. 2–3, pp. 145–156, 1996.

[21] P. Grosdidier, M. Morari, and B. R. Holt, "Closed-loop properties from steady-state gain information," *Industrial & Engineering Chemistry Fundamentals*, vol. 24, no. 2, pp. 221–235, 1985.

[22] E. Gagnon, A. Pomerleau, and A. Desbiens, "Simplified, ideal or inverted decoupling?" *ISA Transactions*, vol. 37, no. 4, pp. 265–276, 1998.

[23] K. A. McDonald, T. J. McAvoy, and A. Tits, "Optimal averaging level control," *AIChE Journal*, vol. 32, no. 1, pp. 75–86, 1986.

[24] D. Swaroop and J. K. Hedrick, "String stability of interconnected systems," *IEEE Transactions on Automatic Control*, vol. 41, no. 3, pp. 349–357, 1996.

[25] 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.

[26] 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.

[27] 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.

[28] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

[29] 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.

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

[31] 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.

1.9 Exercises

1.1 Theory and calculation · 6 min

A blending tank makes $10\,\mathrm{m^3/h}$ of product at $20\,\mathrm{g/L}$ from a concentrate at $100\,\mathrm{g/L}$, flow F_c, and water, flow F_w. The outputs are the product flow $F = F_c + F_w$ and its concentration c.

(a) Find F_c, F_w and the four steady-state gains of F and c to F_c and F_w. *Checkpoint:* $F_c = 2\,\mathrm{m^3/h}$, $F_w = 8\,\mathrm{m^3/h}$, gains 1 and 1 on F, $+8$ and -2 g/L per $\mathrm{m^3/h}$ on c.

(b) Find the relative gain of pairing F with F_c, and choose a pairing. *Checkpoint:* $\lambda = 0.2$, so pair F with F_w ($\lambda = 0.8$) and c with F_c.

1.2 Concept check · 2 min

A two-by-two pairing has a relative gain of -3. What does that predict when both loops have integral action and each is stable alone?

(a) The loops interact weakly and can be tuned one at a time

(b) The closed loop is unstable for every tuning of the two controllers

(c) The loops are stable but three times slower than alone

(d) The pairing works once a static decoupler is added

1.3 Theory and calculation · 4 min

A plant has the steady-state gain matrix

$$K = \begin{bmatrix} 2 & -1 \\ 1 & -1 \end{bmatrix}.$$

Find the relative gain of the diagonal pairing and the Niederlinski index of both pairings.

1.4 Theory and calculation · 5 min

For the plant of the previous exercise, paired on the diagonal, find the two gains of the simplified static decoupler $u_1 = c_1 + d_{12}(c_2 - \bar{u}_2)$, $u_2 = c_2 + d_{21}(c_1 - \bar{u}_1)$ and the steady gain each controller then sees.

1.5 Concept check · 2 min

What does a simplified static decoupler remove from a pair of interacting loops?

(a) The loss of gain each loop suffers when the other is closed
(b) The effect of every disturbance on both outputs
(c) The steady effect of each controller's moves on the other output
(d) The difference in speed between the two paths it balances

1.6 Theory and calculation · 6 min

In the running example the reactor's capacity is changed to $kM = 250\,\mathrm{kmol/h}$ with the fresh feed at 200 kmol/h.

(a) Find the reactor composition, the recycle and the elasticity of the recycle to the feed. *Checkpoint:* $z = 0.80$, $D = 800\,\mathrm{kmol/h}$, elasticity 6.

(b) Find the recycle after 5 % more feed. *Checkpoint:* 1102.5 kmol/h, 37.8 % more.

1.7 Theory and calculation · 3 min

An evaporator train has a feed valve, three effects whose sumps are each drained by a pump, and a steam valve. The sump levels affect nothing else at steady state. How many steady-state degrees of freedom does the train have, and what are they?

1.8 Concept check · 2 min

A plant has a fresh feed, three vessels in series and a product stream. Which structure cannot hold every inventory?

(a) Feed on flow control, each vessel's level on its outflow
(b) Product on flow control, each vessel's level on its inflow
(c) Second vessel's outflow on flow control, upstream levels on inflows, downstream levels on outflows
(d) Feed and product both on flow control, each level on its outflow

1.9 Theory and calculation · 4 min

Two gas loops share one total pressure, so that one unit of oxygen raises both the oxygen partial pressure and the total by one unit, and one unit of nitrogen raises only the total. Find the relative gain and the singular values of the gain matrix

$$K = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}.$$

1.10 Theory and calculation · 5 min

A vessel's level is held by its outflow with a PI controller of gain $K_c' = 2\,\mathrm{h}^{-1}$ (outflow per unit of held volume) and $\tau_I = 0.5\,\mathrm{h}$.

(a) Find ω_n and ζ of the outflow's response to the inflow. *Checkpoint:* $\omega_n = 2\,\mathrm{rad/h}$, $\zeta = 0.5$.

(b) Find the amplitude ratio of outflow to inflow at ω_n, for one vessel and for three identical vessels in series. *Checkpoint:* 1.41 and 2.83.