INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 1

Introduction: Feedforward, Cascade, and Ratio Structures

Pages 1–21

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Feedforward, Cascade, and Ratio Structures," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 1, pp. 1–21.

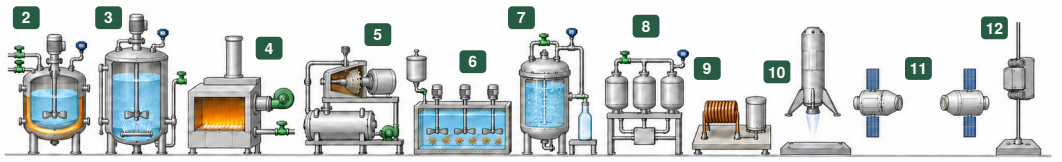
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 1

Introduction: Feedforward, Cascade, and Ratio Structures



The chapters 2 tank, two upsets 3 fermenter oxygen 4 fired heater 5 organic Rankine cycle 6 water treatment 7 soft drink carbonation 8 nanoparticle skid 9 boost converter 10 rocket landing 11 orbital rendezvous 12 space elevator climber

Notice that the first eight units read their feedforward from meters through a balance, while the three vehicles take theirs from a known force. Illustration, not to scale.

IN PLAIN TERMS **Act Before the Upset Arrives**

The main idea. A simple control loop waits until a change has pushed its reading off target, and only then does it act. The methods in this book give it an earlier look. It can watch what is coming in, hold two flows in a fixed proportion, or let a fast inner loop fix trouble at the valve.

The challenge. Each method helps only against an upset that travels its own path. A correction that comes too early or is the wrong size becomes an upset of its own.

Why it matters. The same few methods run water plants, drink lines, power supplies and spacecraft, so choosing them well pays almost everywhere.

THREE WAYS TO READ THIS CHAPTER

Explore	8 min	This page, Section 1.2 with its timeline and two History boxes, the Did you know? boxes in Sections 1.1 and 1.4.3, Figure 1.11, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, with each number of the heater example worked, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.7 with Table 1.3, Section 1.5 with Figures 1.6 and 1.9, and Section 1.8 with Figure 1.11.

1.1 The Question This Volume Answers

A single feedback loop learns about a disturbance only when the disturbance has moved the measurement. By then the plant has already paid for it. Tuning decides how quickly the loop answers after that moment. The structure of the loop decides whether the controller can answer earlier, and that choice is made before any tuning rule is applied. A loop is retuned many times in its life and restructured rarely, so a missing measurement is a lasting limit while a poor gain is a passing one [1]. This volume answers one question: where does each disturbance enter the plant, and which structure meets it there?

Three classical structures extend the single loop, and this volume measures each against the disturbance it was built for and against the ones it was not. Feedforward measures a disturbance and moves the actuator before the controlled variable responds. Ratio control holds one flow in proportion to another and lets a slow feedback loop correct the proportion. Cascade control closes a fast inner loop around the actuator, so that upsets entering there are removed before they reach the slow process. Cross-limiting adds a pair of selectors to a ratio, so that the error that is dangerous cannot occur while two flows of different speed move together.

DID YOU KNOW? Twice the Refinery for the Same Money

Cascade, selector and midranging structures appeared in the power and process industries, where much of PID control was also developed [2]. Between 1930 and 1940, automatic controls doubled the capacity that United States oil refineries could achieve for the same investment [3].

Eleven chapters test these structures on eleven plants, each with an app that runs in a browser. Chapter 2 opens with a tank that has one disturbance at its valve and one at its outlet, so that each structure can be seen to reject one and ignore the other. Chapters 3 to 8 are process plants: a fermenter, a fired heater, a waste-heat power cycle, a water treatment plant, a soft drink line and a nanoparticle formulation skid. Chapters 9 to 12 carry the same ideas to a power converter, a landing rocket, an orbital rendezvous and a space elevator climber, where the feedforward often comes from a known force and not from a meter.

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

- Trace each disturbance to the point where it enters, and say which measurement could see it first.
- Design a static feedforward from a balance and a lead–lag feedforward from two step tests, and predict what a wrong lag costs.
- Build a ratio with a feedback trim, and say why the trim multiplies the ratio and is not added to it.
- Decide whether a cascade pays, from the speed of its inner loop and the place where the disturbance enters.
- Arrange two flows of different speed so that the hazardous error cannot occur, and state what that arrangement costs.

This chapter develops each skill on one small example and names the chapter that uses it. PID tuning is not repeated here. Volume 2 develops it, and every loop in this chapter is a PI tuned from its model by the rules of that volume. Standard texts give fuller treatments of the structures [4, 5].

1.2 A Short History

The milestones of Figure 1.1 follow the three structures through the kinds of plant in this volume. Ratio control rests on proportions that physics sets. A paper of 1803 reported that water takes up twice or three times as much gas under two or three atmospheres as under the common pressure [6], and Chapter 7 doses carbon dioxide on that law in proportion to its product flow. A mechanical linkage that moves the fuel valve and the air damper from one actuator still sets the ratio of many simple burners, and adjusting it is a task for a skilled mechanic. A flow loop on one line and a flow transmitter on the other is the more automated approach [7].

First results on cascade control were published in the mid-1950s [17], among them the quantitative analysis of 1956 [8]. In a pneumatic boiler feedwater loop the output of the level controller is the setpoint of the flow controller, so a change of feedwater pressure is countered by the flow controller without the level controller having to react [7]. Computers arrived as supervisors over such loops.

HISTORY 1959: A Computer That Only Moved Setpoints

The first process control computers were supervisory. They adjusted the setpoints of pneumatic or electric controllers, so the plant kept working when the computer was off or busy with other tasks. The first computer built for process use went fully online in March 1959 on a refinery polymerization unit. Every hour it adjusted 19 parameters, and it raised the conversion of propylene to gasoline from 80 to 91 percent, paying for itself in less than three years. A process stream analyzer was the key instrument, because measuring product quality and closing the mass balance was the ultimate feedback if it could be applied quickly enough [3]. The slow trim of Section 1.4.4, which corrects a ratio from a late quality reading, is the same arrangement in a single loop.

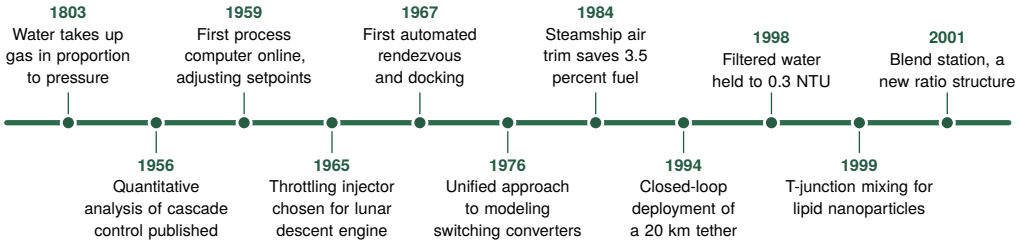


Figure 1.1 Milestones of feedforward, ratio and cascade control in the plants of this volume. The dates are from [6] (1803), [8] (1956), [3] (1959), [9] (1965), [10] (1967), [11] (1976), [12] (1984), [13] (1994), [14] (1998), [15] (1999) and [16] (2001). Notice that the vehicle and converter milestones sit among the process ones: the same structures served both. Events are evenly spaced, not to a time scale.

The vehicles of this volume take their feedforward from physics. A throttling concept for the lunar descent engine was selected in January 1965 [9], and two uncrewed spacecraft made the first automated rendezvous and docking in October 1967 [10]. A flight of 1994 deployed a tethered payload along the local vertical under a closed-loop law [13]. A conference paper of 1976 set out a unified approach to modeling the power stages of switching converters [11], and Chapter 9 uses its averaged form. In the process plants, a steamship fleet reported fuel savings of 3.5 percent from automatic combustion air trim in 1984 [12]. A rule of 1998 held the filtered water of large surface-water systems to 0.3 NTU in 95 percent of each month's readings [14]. T-junction mixing was introduced for lipid-based nanoparticles in 1999 [15], and a new ratio structure, the blend station, appeared in 2001 [16].

> ⧖ HISTORY **1963: Feedforward Beside Feedback**
>
> Feedback is reactive: an error must exist before it acts. When a disturbance can be measured before it enters, a controller can counteract it first, and this is feedforward. Because it tries to match two signals, feedforward needs a good process model, or its correction comes with the wrong size or at the wrong time. It runs open loop, acts on plans, is sensitive to model error and cannot destabilize the loop, while feedback is robust to model error and can. The value of systems with two degrees of freedom, which combine the two, was emphasized in 1963, and combining them is often the better choice [2]. The outer columns of Table 1.1 show a correction at the wrong time.

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. A tag such as TIC, FIC or FT names a temperature controller, a flow controller or a flow transmitter, and a subscript m marks a measured value. The controller settings are K_c and τ_I, and a process model has a gain K_p, a time constant τ_p and a dead time θ_p. A disturbance d reaches the PV through its own model, with gain K_d, time constant τ_d and dead time θ_d. In a cascade the outer or primary loop sets the setpoint of the inner or secondary loop, and τ_{in} is the closed-loop time constant of the inner loop. In a ratio, the wild flow is the one that is measured and not manipulated, and R is the ratio applied to it.

Each chapter keeps the time unit of its plant, from milliseconds in the converter to minutes and hours in the water plant. The running example of this chapter uses minutes, except Section 1.4.6, which uses seconds.

Four symbols need care across chapters. This chapter writes K_d for a disturbance gain, as Chapter 2 does. Chapter 3 uses K_{ff} for a confidence factor on its feedforward, where this chapter uses K_{ff} for

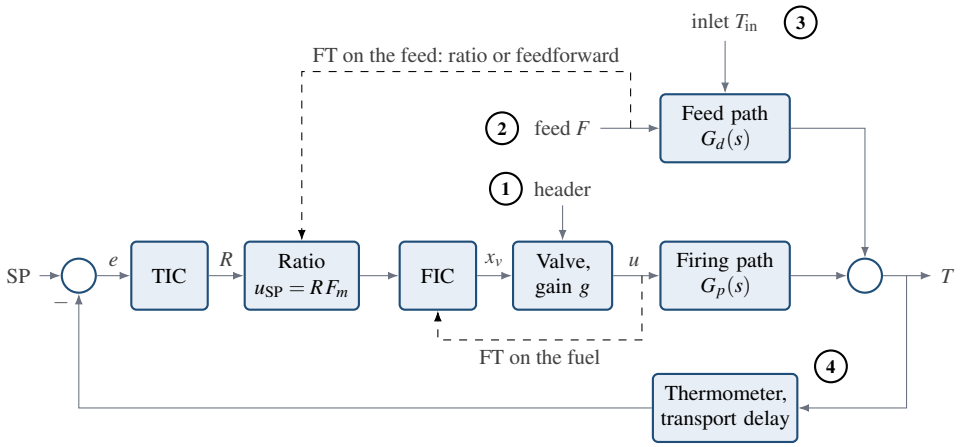


Figure 1.2 The Running Example With Every Structure of This Section. The TIC Trims the Ratio, the Ratio Station Sets the Firing Setpoint From the Measured Feed, and the FIC Moves the Fuel Valve. Dashed Lines Are Measurements Used Ahead of the PV. Circles Mark the Entry Points: (1) Header Pressure at the Valve, (2) Feed Flow, (3) Inlet Temperature, (4) the Measurement.

the feedforward gain itself. Chapter 4 writes λ for the burner air ratio and δ for the margin of its selectors, and Section 1.4.6 follows it. The shared framework uses a for the controller action and h for the integration step, so neither letter is used here for anything else.

1.4 Disturbance Paths and the Structures That Meet Them

1.4.1 The Running Example

The running example is a gas-fired heater that warms a stream of water. The MV is the firing demand u, in percent of the burner's maximum duty, and the PV is the outlet temperature T in °C. At design the feed is $F_0 = 10\,\mathrm{kg/s}$ at $T_{\mathrm{in}} = 20\,^{\circ}\mathrm{C}$, and a firing of 40 % gives 60 °C. The steady energy balance, scaled to these numbers, is Equation (1.1).

$$T = T_{\mathrm{in}} + u\,\frac{F_0}{F} \tag{1.1}$$

where F is the feed flow. At design one percent of firing is worth 1 °C, and one more kg/s of feed is worth −4 °C.

Two step tests give the dynamics, as shown in Equation (1.2).

$$G_p(s) = \frac{K_p\,\mathrm{e}^{-\theta_p s}}{\tau_p s + 1}, \qquad G_d(s) = \frac{K_d\,\mathrm{e}^{-\theta_d s}}{\tau_d s + 1} \tag{1.2}$$

with $K_p = 1\,^{\circ}\mathrm{C}/\%$, $\tau_p = 4\,\mathrm{min}$ for the firing path and $K_d = -4\,^{\circ}\mathrm{C}$ per kg/s, $\tau_d = 2\,\mathrm{min}$ for the feed path. Both paths share the dead time $\theta_p = \theta_d = 1\,\mathrm{min}$, the transport delay from the heater outlet to the thermometer. The firing path is slower because heat must pass the firebox and the tube wall, while a change of feed acts on the water at once. The dynamic runs below are linear in deviation variables about the design point. The temperature controller is a PI tuned by the SIMC rule with $\tau_c = \theta_p$, which gives $K_c = 4/(1 \times 2) = 2\,\%/^{\circ}\mathrm{C}$ and $\tau_I = 4\,\mathrm{min}$, and it executes every 0.1 min in the velocity form of Section 1.6.2.

Figure 1.2 shows the whole structure the section builds up. Each subsection adds one part of it. The numbered circles mark the four places where a disturbance can enter.

1.4.2 Where a Disturbance Enters

With a controller $C(s)$ in a single loop, a disturbance D that reaches the PV through G_d leaves the deviation of Equation (1.3).

$$Y(s) = \frac{G_d(s)}{1 + C(s)\,G_p(s)}\,D(s) \tag{1.3}$$

The numerator is the path from the entry point to the measurement. The denominator is the feedback, and it acts only on what has already arrived. Nothing in the controller can make G_d smaller. A structure helps only by measuring something that lies earlier on the disturbance's path than the PV does. Figure 1.2 gives four entry points, and each has a different remedy.

- *At the actuator* (1). A fall in header pressure lowers the fuel that a given valve lift passes. A flow meter downstream of the valve sees it within seconds, so a cascade with an inner flow loop removes it.
- *At the process, measured* (2). A change of feed flow enters the tubes. A flow meter on the feed sees it before the temperature does, so feedforward or a ratio can act on it.
- *At the process, unmeasured* (3). A change of inlet temperature has no transmitter here. Only feedback on the PV, or a slow trim, can correct it.
- *At the measurement* (4). A slow, biased or noisy thermometer corrupts the information every structure relies on. No rearrangement of controllers removes it. It needs a better instrument, a validation check or an estimator, which Volume 5 treats.

The same disturbance can change category with the instruments installed. A feed flow without a meter is entry point 3, and an inlet temperature with a transmitter becomes entry point 2.

Chapter 2 is built on this list. Its tank has a supply pressure at the valve and an outlet demand at the vessel, and the chapter shows that each structure rejects the disturbance on its own path and leaves the other nearly untouched. Chapter 3 sorts three lags of nearly equal size the same way: a cascade reaches the agitator drive, a feedforward reaches the feed, and nothing reaches the oxygen probe. The roadmap of Section 1.4.7 lists the paths of every chapter.

1.4.3 Static and Lead–Lag Feedforward

A feedforward adds $G_{ff}(s)$ times the measured disturbance to the MV. The PV is then $(G_d + G_p\,G_{ff})\,D$, which is zero for the ideal compensator of Equation (1.4).

$$G_{ff}(s) = -\frac{G_d(s)}{G_p(s)} = -\frac{K_d}{K_p}\,\frac{\tau_p s + 1}{\tau_d s + 1}\,\mathrm{e}^{-(\theta_d - \theta_p)s} \tag{1.4}$$

It has three parts. The gain $-K_d/K_p$ is the static feedforward, which holds the steady state. The ratio of lags is a lead–lag, which matches the timing. The delay term can be built only when the disturbance path is the slower one to reach the PV, $\theta_d \geq \theta_p$. Otherwise the ideal compensator would have to act before the disturbance is measured, and the design must settle for an approximation, for which tuning rules exist [18, 19].

For the heater the delays are equal and the compensator is $4\,(4s+1)/(2s+1)$ % per kg/s. The static part follows from Equation (1.1) without any test: holding T needs $u = (T - T_{\mathrm{in}})\,F/F_0$, whose slope at design is 4 % per kg/s. A static feedforward from a balance needs only the gains, which are known better than the dynamics.

Static feedforward alone leaves an error while the two paths respond at different speeds. For a step D and the exact static gain, the error integrates to Equation (1.5).

$$\int_0^\infty e\,\mathrm{d}t = K_d\,D\,\big[(\theta_p + \tau_p) - (\theta_d + \tau_d)\big] \tag{1.5}$$

The bracket is the difference between the mean response times of the two paths. When the firing path is the slower one, as here, the feedforward is late and the error keeps the sign of the uncompensated

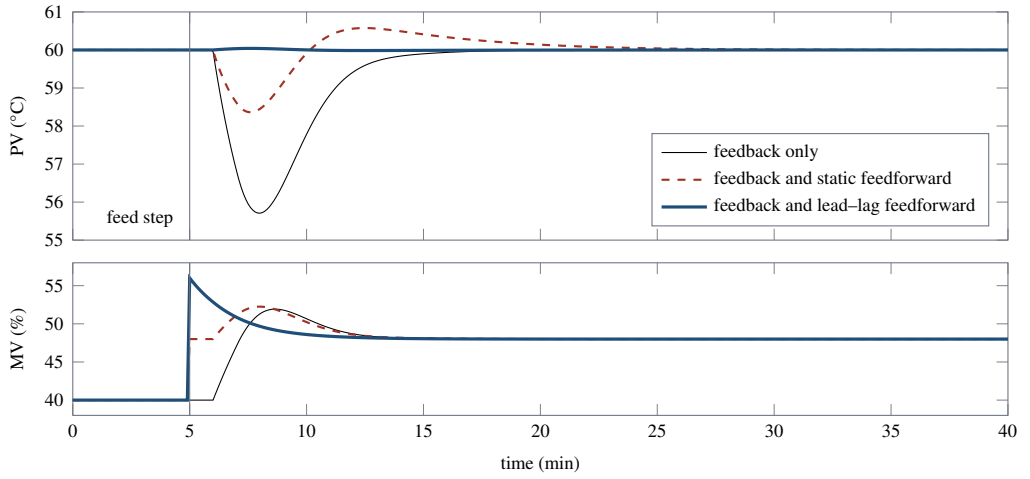


Figure 1.3 A Feed Step From 10 to 12 kg/s at 5 min on the Heater. Thin Solid: PI Feedback Only. Dashed: Feedback With the Static Feedforward of 4 % per kg/s. Thick Solid: Feedback With the Lead–Lag $4(4s+1)/(2s+1)$.

Table 1.1 Feedforward on the Feed Step of Figure 1.3. The Lead–Lag Is Designed for $\tau_d = 2$ min. The Two Outer Columns Run It on Plants Whose Feed Path Is Faster or Slower. Peak Error in °C, IAE in °C min Over 60 min.

	Plant $\tau_d = 1$ min		Plant $\tau_d = 2$ min		Plant $\tau_d = 3$ min	
Structure	Peak	IAE	Peak	IAE	Peak	IAE
Feedback only	5.89	16.8	4.29	16.0	3.39	16.0
Feedback and static feedforward	3.52	14.7	1.64	8.2	0.63	3.5
Feedback and lead–lag feedforward	1.92	7.0	0.04	0.2	1.06	5.0

disturbance. A step of 2 kg/s leaves −16 °C min, with a peak of 2.0 °C against 8 °C with no control. When the disturbance path is the slower one, the feedforward is early and the error takes the opposite sign. With $\tau_d = 8$ min the same static feedforward leaves 32 °C min of the opposite sign. An early feedforward is a disturbance of its own.

Figure 1.3 adds feedback to each feedforward, and Table 1.1 gives the measures. Feedback alone peaks at 4.29 °C. Its error integrates to 16 °C min, which is the firing change of 8 % times τ_I/K_c, the signature of any PI controller on a load. The static feedforward halves the IAE and cuts the peak to 1.64 °C. The lead–lag removes the error except for 0.04 °C left by the control period. It does so by moving the firing 16 % at once, twice its final change, because the lead of 4 min over a lag of 2 min doubles the initial response. That overshoot of the MV is what the slow firing path needs.

> **+ DID YOU KNOW? A Faster Upset Peaks Higher but Costs Nearly the Same**
>
> From the simulator: under feedback alone, the feed step of Table 1.1 peaks at 5.89 °C when its path lag is 1 min and at 3.39 °C when it is 3 min, yet the integrated error is 16.8 against 16.0. A PI pays for a load by its own settings, whatever the load's speed.

The outer columns of Table 1.1 show why a lead–lag is tuned with care. When the feed path is faster than the model, the lead–lag is still the best of the three. When it is slower, the lead–lag acts early and does worse than the static gain, which happens to match the slower path more closely. A lead–lag whose gain is 25 % low, three in place of four, leaves a quarter of the disturbance to

feedback, with a peak of 1.04 °C and an IAE of 4.0 °C min. Feedback is therefore kept beside every feedforward. It removes the residue of a wrong gain, and the feedforward removes the bulk of a measured disturbance that feedback would see late [20].

Every chapter of this volume uses a feedforward or a ratio, and most use a static one built from a balance or a known force. In Chapter 2 the outlet demand and the inlet flow enter the same integrator, so the ideal compensator is a pure gain. Chapter 4 inverts the heat balance of its heater, Chapter 9 computes a duty and a current from its power balance, and Chapter 10 and Chapter 12 hand the weight of the vehicle to the loop as a bias. Chapter 11 cancels the coupling terms of its equations of motion and measures how much a wrong orbital rate costs at two bandwidths. Timing decides two chapters. Chapter 3 must lag its feedforward to match the time the feed takes to reach the cells through a glucose pool, or the oxygen overshoots upward. Chapter 5 computes the lost heat from the source meters and lowers the pump's flow setpoint within seconds. These two chapters and Chapter 4 enter the feedforward as the output bias of the PID, so that switching it on or off moves the integral by the same amount and the output does not bump.

> ✦ DID YOU KNOW? **Feedforward in Markets, Firms and Muscles**
>
> A standard text compares feedforward with a planned economy and with running a company on strategic plans, and notes that it has been suggested as an essential, trained part of how humans control their motion [2]. The runner of Volume 19, Chapter 6 gets a measured feedforward of the incline, because the angle is known the instant it changes.

1.4.4 Ratio Control With a Feedback Trim

A ratio station sets the MV in proportion to a measured wild flow. For the heater the MV is the firing setpoint and the wild flow is the feed, as shown in Equation (1.6).

$$u_{\mathrm{SP}} = R\,F_m, \qquad R = R_0 + \Delta R \tag{1.6}$$

where R_0 is the design ratio and ΔR is the output of a slow feedback controller on the PV. Equation (1.1) gives $R_0 = (T - T_{\mathrm{in}})/F_0 = 4\,\%$ per kg/s. Linearized at the design point, the ratio is the static feedforward of Section 1.4.3. Ratio control is the oldest form of feedforward, long used on blending and combustion services [16, 21]. For the heater the balance makes the firing proportional to the feed at a fixed temperature rise. A feed of 15 kg/s needs 60 % of firing, and $R_0 F$ delivers it. The ratio and the static feedforward therefore agree at every feed, and they differ only in how a correction combines with them.

The trim exists to correct what the ratio does not know. Suppose the inlet water cools from 20 to 10 °C, with no transmitter to report it. The heater now needs $R = 5\,\%$ per kg/s. At the design feed the TIC finds the missing 10 % of firing either way. A trim that multiplies, as in Equation (1.6), settles at $\Delta R = 1$. A trim that adds a bias b to $R_0 F_m$ settles at $b = 10\,\%$. The two agree until the feed moves. At 15 kg/s the multiplied ratio gives the 75 % the heater needs, while the added bias gives 70 %, and the outlet runs 3.3 °C cold until the slow trim finds the difference. At 5 kg/s the added bias overfires by 5 % and the outlet runs 10 °C hot. An error in the ratio scales with the wild flow, so the correction should scale with it too.

The trim is slow on purpose. A trim as fast as the main loop would also react to every transient the ratio leaves while the wild flow changes, so it is tuned to follow composition, heating value or a fouled exchanger over many residence times. A ratio can also be applied to the setpoint of the wild flow in place of its measurement. The setpoint basis moves both flows at once, and the measurement basis follows the flow that actually arrived when a valve cannot keep up.

Five chapters build on this subsection. Chapter 6 computes the coagulant dose from the raw-water turbidity and lets the settled turbidity trim it through a multiplier. It argues the case made above, that a correction found at one raw turbidity remains right at another. Chapter 7 blends syrup into water at a ratio trimmed on a refractometer, doses carbon dioxide in proportion to the product flow, and

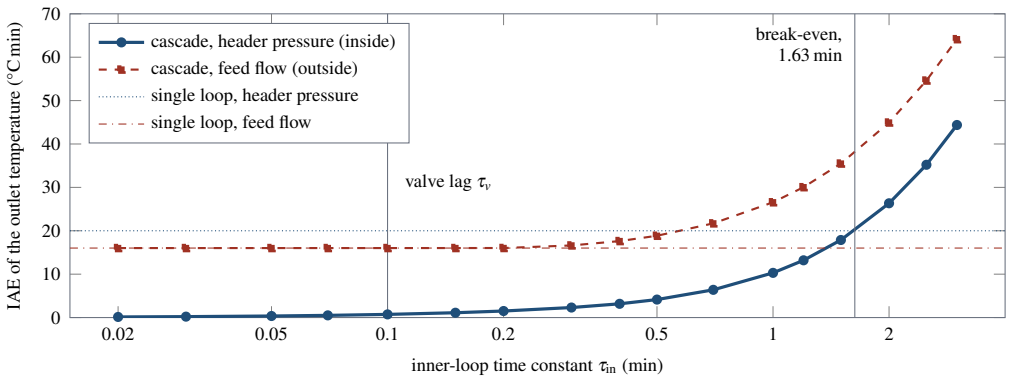


Figure 1.4 What the Speed of the Inner Loop Buys. The Heater's TIC Is Held at $K_c = 2$, $\tau_I = 4$ min. Thick Solid With Circles: Cascade Against a 36 % Fall in Header Pressure. Dashed With Squares: Cascade Against a Feed Step of 2 kg/s. Dotted and Dash-Dotted Lines: the Single Loop Against the Same Two Upsets.

compares the measured-water and setpoint bases when a water valve falls behind. Chapter 8 solves four feed flows that hold a total flow, two ratios and a buffer strength at once. Chapter 4 sets the air in ratio to the fuel and trims the ratio on stack oxygen, and Chapter 2 expresses its feedforward gain as a ratio of flow setpoint to demand when the cascade is selected.

1.4.5 Cascade Control and the Inner-Loop Speed Rule

The fuel valve of the heater passes a flow $u = g\,x_v$, where x_v is the lift and g grows with the square root of the pressure drop. When the header pressure falls by 36 %, g falls from 1 to 0.8 and a lift of 40 % passes 8 % less fuel. That upset enters at point (1) of Figure 1.2, upstream of the whole firing path. With the TIC on the valve, the temperature falls through G_p before the controller can respond. A cascade places a flow controller between them. The TIC asks for a fuel flow, and the FIC moves the valve until the meter reads it. The header pressure is then removed inside the inner loop, before it reaches the firebox.

The inner loop has dynamics of its own. Here the valve and the fuel line respond as a lag of $\tau_v = 0.1$ min. A PI on that lag with $\tau_I = \tau_v$, or a proportional controller alone, closes the inner loop with the time constant of Equation (1.7).

$$\tau_{in} = \frac{\tau_v}{K_{c,in}\,g} \quad \text{(PI)}, \qquad \tau_{in} = \frac{\tau_v}{1 + K_{c,in}\,g} \quad \text{(P only)} \tag{1.7}$$

The outer loop sees this closed inner loop in place of the valve. If it is faster than the valve was, the outer loop gains. If it is slower, the outer loop has a slower actuator than before.

Figure 1.4 measures the trade on the heater. The TIC keeps its tuning throughout, and only the inner gain changes. With a single loop the header pressure costs a peak of 3.09 °C and an IAE of 20.0 °C min. With an inner loop of 0.05 min the peak is 0.11 °C and the IAE 0.36 °C min, 56 times smaller. The feed step, which enters outside the inner loop, costs the same 16.0 °C min with or without the cascade, because no flow meter on the fuel sees it.

As the inner loop slows, the two curves separate the two parts of the speed rule. The rule, stated since the first quantitative analysis of cascade control, asks for an inner loop substantially faster than the outer one [8]. Against the header pressure the cascade keeps 93 % of its attainable benefit at $\tau_{in} = 0.2$ min, 80 % at 0.5 min and none at 1.63 min. Against the feed step it costs nothing until the inner loop is as slow as the valve lag it encloses, about 0.1 min. It then costs 18 % at 0.5 min and nearly three times the single-loop IAE at 2 min. The usual statement of the rule is three to five

Table 1.2 The Three Arrangements of Figure 1.5. The Fuel Time Is How Long the Fuel Takes to Cover 63 % of the Step.

	Increase, 40 to 60 %		Decrease, 60 to 40 %		Time short
Arrangement	Lowest λ	Fuel time (s)	Lowest λ	Fuel time (s)	of air (s)
Fuel leads	0.880	2.0	1.151	2.0	10.8
Air leads	1.150	12.2	1.095	12.2	0
Cross-limited, $\delta = 0$	1.150	12.2	1.152	2.0	0
Cross-limited, $\delta = 2\,\%$	1.133	10.9	1.151	2.0	0

times faster [4, 22]. The outer loop here has a closed-loop time constant of $\tau_c + \theta_p = 2\,\text{min}$, so that statement gives 0.4 to 0.67 min. It still collects 71 to 85 % of the benefit, and it already costs 10 to 32 % on the feed step outside the inner loop. A factor of ten costs nothing measurable here. A sharper statement for this heater is that the inner loop should be no slower than the actuator lag it replaces, since every disturbance that the inner loop cannot see then pays for its delay.

A cascade has two further uses that the curves do not show. The inner loop linearizes the actuator. The TIC asks for a flow, and the square root of the valve law, a sticking stem or a worn pump no longer reaches the outer loop. The inner loop is also where a limit can be enforced, since a flow or a current that is controlled can be capped. Tuning both loops from their models is treated in [23]. When the secondary measurement responds to the MV but does not lie on the path to the primary, the parallel arrangement applies [24].

Seven chapters measure this subsection. Chapter 2 finds a stability limit for an inner loop with no plant dynamics, set by one sample of delay, which the supply pressure carries in and out of range on every cycle at some gains. Chapter 3 sweeps the inner speed of an agitator loop, finds that a detuned cascade is three times worse than no cascade, and shows that the inner integral time moves the break-even point by a factor of 1.8. Chapter 5 wraps a flow loop around its feed pump. Chapter 9 closes an inner current loop, which enforces a current limit and lets the voltage loop run harder against a right-half-plane zero. Chapter 10 commands an attitude from a position loop and sweeps the separation of their bandwidths, and Chapter 4 and Chapter 7 put a flow loop under every ratio.

1.4.6 Cross-Limiting of Fuel and Air

A burner needs air in proportion to fuel, and the two errors are not equal. Excess air carries heat up the stack. Too little air leaves unburned fuel in a hot enclosure, which is a hazard. The heater burns gas at a design air ratio $\lambda_0 = 1.15$, where $\lambda = F_a/F_f$ is the air supplied over the air the fuel needs. The fuel flow loop closes in 2 s and the forced-draft fan in 10 s. Times in this subsection are in seconds.

Three arrangements are compared on a firing step from 40 to 60 %, the change the ratio station of Section 1.4.4 asks for when the feed rises from 10 to 15 kg/s. In *fuel leads* the fuel follows the demand and the air follows the measured fuel. In *air leads* the air follows the demand and the fuel follows the measured air. Cross-limiting combines the two with a low selector on the fuel and a high selector on the air, as shown in Equation (1.8).

$$F_f^{\text{SP}} = \min\left[u,\ \frac{(1+\delta)\,F_{a,m}}{\lambda_0}\right], \qquad F_a^{\text{SP}} = \max\left[\lambda_0\,u,\ (1-\delta)\,\lambda_0 F_{f,m}\right] \tag{1.8}$$

where δ is a small margin by which the fuel may run ahead of the air. The fuel can rise no faster than the measured air can burn it, and the air can fall no faster than the measured fuel allows. On an increase the air leads, and on a decrease the fuel leads [25].

Figure 1.5 and Table 1.2 give the result. With fuel leading, the fuel reaches its new flow in seconds and the fan follows five times more slowly. The burner runs short of air for 10.8 s, down to $\lambda = 0.880$. Air leading never runs short, but the fuel now waits for the fan in both directions, and on the decrease

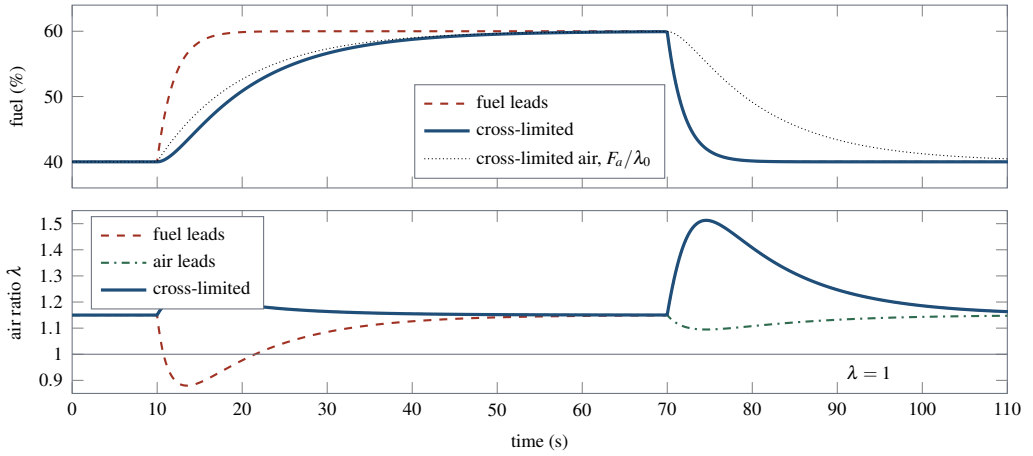


Figure 1.5 Firing Demand Stepped From 40 to 60 % at 10 s and Back at 70 s. Top: Fuel Flow, and the Cross-Limited Air in Units of the Fuel It Can Burn at λ_0. Bottom: Air Ratio. Dashed: Fuel Leads. Dash-Dotted: Air Leads. Thick Solid: Cross-Limited. Below the Thin Line at $\lambda = 1$ the Burner Is Short of Air.

the fuel stays high for 12 s after the demand has fallen. Cross-limiting takes the safe half of each. It is as fast as fuel leads on the way down and as safe as air leads on the way up. Its cost is the heat delay on an increase, the sum of the two loop lags, which the TIC sees as a slower actuator. A margin of $\delta = 2\,\%$ recovers 1.3 s of that delay and lets λ dip to 1.133. Every scheme runs with extra air for some seconds in at least one direction, which is the price of safety and is paid in fuel.

Chapter 4 builds this arrangement on a refinery heater with a fan six times slower than its fuel valve. It then trims the ratio on stack oxygen and adds a carbon monoxide override, a selector that Volume 4 develops. The asymmetry argument applies wherever one direction of a ratio error is a hazard, as with reactant feeds to an exothermic reactor.

1.4.7 Which Structure for Which Path

The section reduces to one rule. A structure helps against a disturbance only if it measures the disturbance, or a variable on its path, earlier than the PV does. Feedforward and ratio act on a measured disturbance at the process. A cascade acts on any upset between the controller output and the inner measurement. A trim corrects what was not measured, slowly. Nothing here corrects the measurement. A structure applied to the wrong path is not merely useless. A feedforward from a variable that the MV itself moves can oscillate, as in Chapter 2, and a slow inner loop delays every disturbance outside it.

Table 1.3 applies the rule to every chapter of the volume. It lists the main disturbances each chapter applies, where each enters, and the structure the chapter tests against it. The final chapter returns to this table with the measured result of each pairing.

Three patterns run through the table. The process plants of Chapters 3 to 8 combine a measured disturbance, met by feedforward or a ratio, with an unmeasured one, met by a slow trim. The vehicles of Chapters 10 to 12 take their feedforward from physics, so the model and not a meter is what can be wrong. Several chapters keep one disturbance that no structure reaches, such as the probe lag of Chapter 3 or the condenser capacity of Chapter 5, and those chapters state the limit flatly. A broader view of choosing controlled and secondary variables across a whole plant is given in [26], and Volume 6 develops it.

Table 1.3 Disturbance Paths in This Volume and the Structure Each Chapter Tests Against Them

Ch.	Disturbance	Where it enters	Structure tested
2	Outlet demand	Vessel, measured	Outlet feedforward (inlet-flow feedforward fails)
	Supply pressure	Valve	Level-to-flow cascade
3	Feed step	Culture, measured	Carbon-balance feedforward with a lag
	Gassing on the shaft	Actuator	Speed cascade
	Probe lag	Measurement	None
4	Process flow	Process, measured	Heat-balance feedforward
	Load steps	Fuel and air pair	Cross-limiting
	Heating value	Burner, unmeasured	Oxygen trim of the ratio
5	Source heat	Evaporator, measured	Measured-heat feedforward
	Pump wear, density	Actuator	Inner flow loop
	Warm cooling water	Condenser capacity	None
6	Raw-water storm	Intake, measured	Computed dose
	Organic carbon	Intake, laboratory only	Multiplying ratio trim
7	Filler rate	Throughput, measured	Throughput feedforward and ratio blend
	Syrup strength	Blend, unmeasured	Brix trim of the ratio
	Water header dip	Water valve	Ratio on the measured flow
8	Throughput change	Four feeds	Four-feed ratio block
	Size drift	Late laboratory result	Slow trim with a soft sensor
9	Load and supply voltage	Power stage, measured	Feedforward on both loops
	Overload	Inductor current	Inner current loop and its limit
10	Weight under tilt	Known force	Gravity feedforward
	Crosswind	Side force	Position-to-attitude cascade, no integral
11	Gravity gradient, Coriolis terms	Known coupling between axes	Model feedforward that cancels the coupling
12	Effective gravity along the route	Known force that changes sign	Gravity feedforward
	Ribbon mode	The climber's own moves	Input shaping

1.5 Instruments and Final Elements in This Volume

Every structure of this volume adds a measurement or a loop, so it can be no better than the instrument it adds. Figure 1.6 lists what the controllers of each chapter read and what they move. The flow meter and the control valve or flow loop lead, with six chapters each, because a ratio holds only what a meter reports and a flow loop is the usual inner loop. Analyzers and quality probes follow with five chapters, nearly always as the input of a slow trim.

A ratio station, or signal multiplying relay, multiplies the signal of the wild flow transmitter by a constant and sends the product to the captive flow controller as its setpoint [7]. The ratio is then only as good as the two meters. A Coriolis meter senses the inertial force of the fluid in its vibrating tubes, so it is a true mass flow instrument with no straight-pipe requirement. External vibration and gas bubbles trapped in bent tubes cause errors [7]. Figure 1.7 shows its principle, and the technology section of Chapter 2 places such a meter under its inner loop.

In a cascade the secondary loop must respond faster than the primary, and the scheme cannot work if that order is reversed [7]. The lowest loop is often in the valve itself. It is customary to close inner loops on valve positioners to reduce friction and nonlinearity at the lowest level of the hierarchy [2], so the fuel loop of Figure 1.2 on such a valve is three loops deep (Figure 1.8). A speed drive with a tachometer is likewise the secondary controller of a cascade, taking its setpoint from a process controller [7], as the agitator of Chapter 3 does.

A supervisory computer sends setpoints to regulatory controllers that run in automatic with a remote setpoint, which forms a two-layer control system [7]. Loops in such systems are often nested or cascaded, with the setpoint of one loop based on the process variable of another, and cycle times

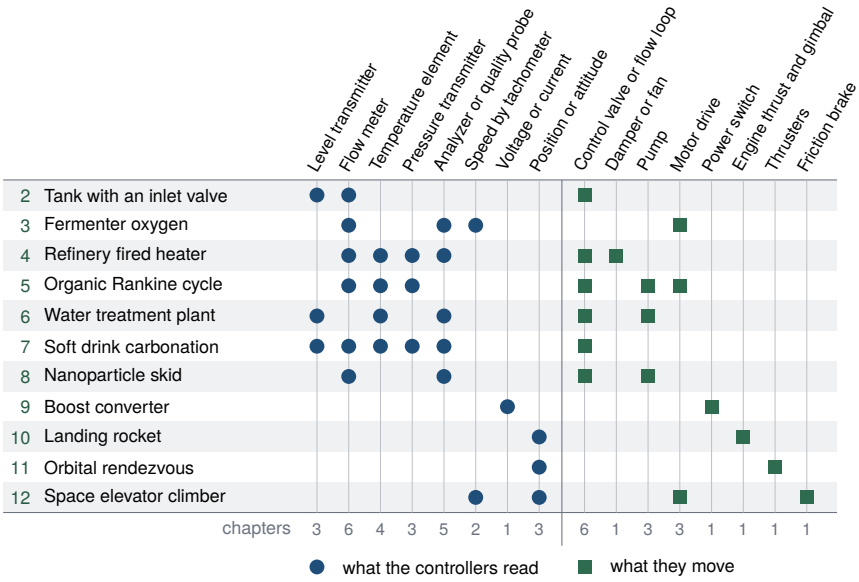


Figure 1.6 What the controllers of each chapter read and move: a round mark where a controller reads the instrument and a square mark where it moves the element, as the chapter's variables table and its technology section name them. A value computed from other readings, such as a superheat, is marked under the instruments it is computed from. Notice that the flow meter and the valve or flow loop lead the counts, with six chapters each.

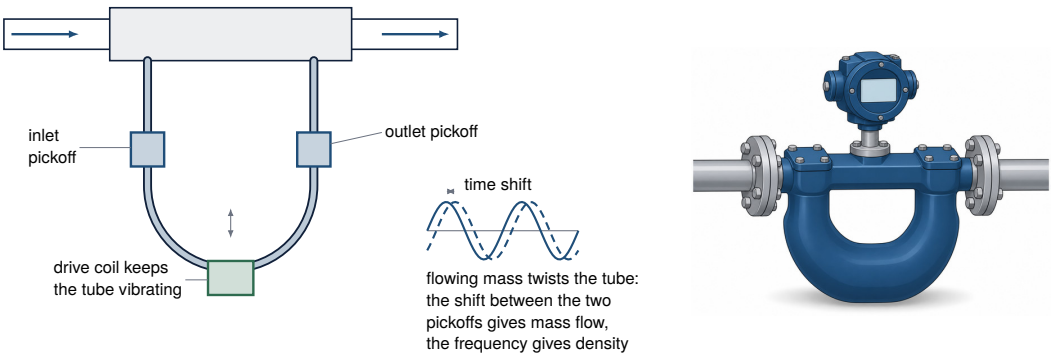


Figure 1.7 Principle of a Coriolis mass flow meter. Notice that it reads mass flow directly, which is what a ratio of two liquid streams needs, and that the reading depends on a steady vibration of the tubes.

run from milliseconds to minutes [27]. A converter splits its cascade across hardware. A gate array runs the fast controller, and an outer loop on a processor core supplies its reference at a lower rate [28]. Figure 1.9 places the loops of this volume on their platforms. The apps run from the 0.2 ms period of Chapter 9 to the 0.2 min period of Chapter 6, and the supervisor of Chapter 8 acts every 30 s.

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

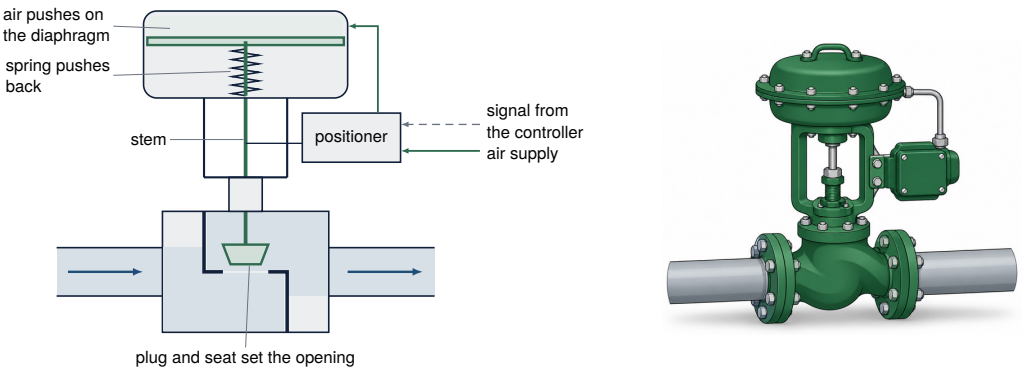


Figure 1.8 A globe control valve with its positioner. Notice that the positioner compares the controller's signal with the stem position and moves the stem itself, so it is the innermost loop of every valve cascade in this volume.

Layer	What it does in this volume	Where it runs	How often
Supervisor and slow trims	a trim on a late analyzer, a supervisor that picks a mixing loop, a predictive controller over a cascade	**Supervisory computer** that writes setpoints to the loops below	every 30 s in the formulation skid; trims over many residence times
Outer loop with feedforward, ratio	the primary PID with the feedforward as its bias, a ratio station, the cross-limiting selectors	**Distributed control system** or PLC function blocks; a **flight computer** on a vehicle	0.2 s to 0.2 min in the plants; 10 ms to 1 s in the vehicles
Inner loop	a flow loop under a ratio, a drive's own speed loop, a converter's current loop	**Flow controller**, transmitter or positioner; a **speed drive**; a **gate array**	0.25 s for a flow loop; 0.2 ms in the converter
Sensors and final elements	meters, analyzers and probes; valves, pumps, drives, a power switch, an engine, thrusters, a brake	wires or a fieldbus; a **positioner** closes the lowest loop on the valve stem	device lags from 0.2 ms to 20 s in the models

Figure 1.9 The control layers of this volume, the platform that hosts each and how often each acts. Platforms are from [7], [27], [2] and [28], and the app periods from the chapters' technology sections. Notice that each layer writes the setpoint of the one below it, so a cascade runs through the whole stack.

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 1.4 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.

Table 1.4 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.

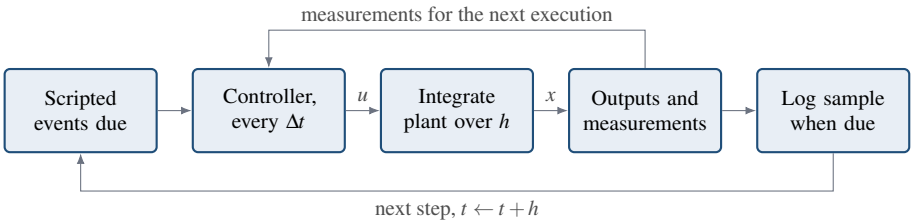


Figure 1.10 Work of the Engine in One Integration Step

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 (1.9) to (1.12) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.9}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.10}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.11}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.12}$$

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 (1.12) drives I toward the value that reproduces the output applied, which is back-calculation [29]. With no integral action, Equation (1.12) 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 (1.13).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.13}$$

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 (1.14) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.14}$$

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 [30]. 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 [31]. 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 [32].

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.

	Feedforward from a meter and a balance	Feedforward from a known force or model
Each structure on its own path	(2) feedforward from the wrong flow oscillates	
Feedforward gain and timing	(5) measured heat protects the expander	(11) coupling canceled exactly (12) a weight that fades and reverses
Ratio with a slow trim	(6) dose ahead of a slow basin (8) four ratios held (7) blend trimmed by an analyzer	
Inner-loop speed and limits	(3) a detuned cascade does worse (9) current loop and its limit	(10) position commands an attitude
Cross-limiting and override	(4) fuel and air of different speeds	

Figure 1.11 Method map of the volume. Each chapter is placed by the source of its feedforward, from Table 1.3, and by the structural lesson it centers on, shortened from Table 1.5. The placement is the editor's reading of the two tables. Notice that the three vehicles take their feedforward from physics, while the plants and the converter read it from meters.

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 13.

1.8 Chapter Roadmap

Table 1.5 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 where their feedforward comes from and by the structural lesson each centers on.

Table 1.5 Roadmap of the Volume

Ch.	Plant	The structural lesson	Level	Time
2	Tank on level with an inlet valve	Each structure rejects the disturbance on its own path. A feedforward from the wrong flow oscillates.	introductory	2.5 h
3	Fed-batch fermenter, dissolved oxygen	The inner-loop speed rule measured, and a feedforward that must wait for the culture.	intermediate	3 h
4	Refinery fired heater	Cross-limiting of fuel and air, oxygen trim of the ratio and a heat-balance feedforward.	intermediate	2 h
5	Waste-heat organic Rankine cycle	A flow cascade that needs a feedforward to protect the expander, and a limit no loop removes.	advanced	2.5 h
6	Municipal water treatment plant	A computed dose ahead of three hours of process, with a slow multiplying trim.	introductory	2 h
7	Soft drink carbonation line	Ratio blending with an analyzer trim, and a quality that is a pressure margin.	intermediate	2 h
8	Nanoparticle formulation skid	Four ratios held at once, and a slow trim on a result a minute old.	advanced	2.5 h
9	DC–DC boost converter	An inner current loop against a right-half-plane zero, and a current limit.	advanced	3 h
10	Rocket on a gimballed engine	A cascade whose outer loop commands an angle. Bandwidth separation and gravity feedforward.	advanced	3 h
11	Chaser approaching a target in orbit	A model feedforward that cancels coupling, and what a wrong orbital rate costs.	intermediate	3 h
12	Space elevator climber	Feedforward of a weight that fades and reverses, and shaping against a ribbon mode.	advanced	2.5 h

Chapter 2 is the first chapter to read, because it isolates each structure on one plant with two disturbances. A first course can continue with Chapters 6, 4 and 3, which add a dose computed ahead of a long process, cross-limiting and the inner-loop speed rule in that order. Chapters 5 to 8 combine several structures on one plant and can be read in any order once Sections 1.4.3 to 1.4.5 are familiar. Chapters 9 to 12 need the same sections and a little rigid-body or circuit dynamics. They form a second path for readers from electrical or aerospace engineering, best started with Chapter 11, whose feedforward is exact.

KEY POINTS

- A single loop learns of a disturbance only when the measurement has moved. A structure helps only if it measures the disturbance, or something on its path, earlier than the PV does.
- A static feedforward from a balance needs only the gains. A lead–lag also needs the timing, and an early feedforward is a disturbance of its own.
- A ratio trim should multiply the ratio, because an error in the ratio scales with the wild flow.
- A cascade pays when its inner loop is no slower than the actuator lag it replaces, and it cannot touch a disturbance outside that loop.

Where else this applies. Selectors and overrides in Volume 4, dead time in Volume 5, plant-wide structure in Volume 6.

References

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

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

[3] M. Pitt, "Taking a look back at control: Part 2," *The Chemical Engineer*, no. 989, Nov. 2023, thechemicalengineer.com/features/taking-a-look-back-at-control-part-2.

[4] 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.

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

[6] W. Henry, "Experiments on the quantity of gases absorbed by water, at different temperatures, and under different pressures," *Philosophical Transactions of the Royal Society of London*, vol. 93, pp. 29–274, 1803.

[7] 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.

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

[9] W. R. Hammock, Jr., E. C. Currie, and A. E. Fisher, "Apollo experience report: Descent propulsion system," NASA Manned Spacecraft Center, Houston, TX, NASA Technical Note TN D-7143, 1973, nASA Technical Reports Server 19730011150.

[10] J. L. Goodman, "History of space shuttle rendezvous," NASA Johnson Space Center, Mission Operations Directorate, Flight Dynamics Division, Houston, TX, Report JSC-63400, Revision 3, Oct. 2011, nASA Technical Reports Server 20110023479.

[11] R. D. Middlebrook and S. Cuk, "A general unified approach to modelling switching-converter power stages," in *1976 IEEE Power Electronics Specialists Conference*, 1976, pp. 18–34.

[12] J. J. Sweeney, "Boiler Flue Gas Analysis and Automatic Combustion Air Trim Control by Measurement of Carbon Monoxide, Unburned Hydrocarbons, and Opacity," *Marine Technology and SNAME News*, vol. 21, no. 01, pp. 74–79, 1984.

[13] "Tethers in space handbook," Smithsonian Astrophysical Observatory for NASA Marshall Space Flight Center, NASA Contractor Report NASA/CR-97-206807, 1997, nTRS 19980018321.

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

[15] J. Hong Hoang, M. Ott, E. Samaridou, M. Beck-Broichsitter, and J. Simon, "Industrial perspective on the manufacturing of lipid nanoparticles for nucleic acid delivery," *Pharmaceutics*, vol. 18, no. 4, p. 489, 2026.

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

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

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

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

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

[21] F. G. Shinskey and B. G. Lipták, "Ratio control," in *Process Control*. Elsevier, 1995, pp. 102–105.

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

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

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

[25] D. Lindsley, *Power Plant Control and Instrumentation: The Control of Boilers and HRSG Systems*. Institution of Engineering and Technology, 2000.

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

[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] P. Karamanakos, E. Liegmann, T. Geyer, and R. Kennel, "Model predictive control of power electronic systems: Methods, results, and challenges," *IEEE Open Journal of Industry Applications*, vol. 1, pp. 95–114, 2020.

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

[30] 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.

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

[32] 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 Concept check · 2 min

A steam-heated exchanger holds its outlet temperature with a steam valve. The steam header pressure swings by 20 % several times an hour, and a flow meter is installed on the steam line downstream of the valve. Which structure removes the header swing before it reaches the outlet temperature?

(a) Static feedforward from the measured process flow to the steam valve

(b) A cascade in which the temperature controller sets a steam flow controller

(c) A ratio of steam flow to process flow with no feedback trim

(d) A faster integral time on the temperature controller alone

1.2 Theory and calculation · 5 min

A process has the manipulated path $G_p(s) = 2\,\mathrm{e}^{-3s}/(6s+1)$ and a measured disturbance path $G_d(s) = 0.5\,\mathrm{e}^{-5s}/(3s+1)$, with time in minutes. Find the ideal feedforward compensator, say whether it can be built, and give the ratio of its initial response to its final response for a step disturbance.

1.3 Theory and calculation · 4 min

A static feedforward with the exact gain is applied without feedback. The manipulated path has $\theta_p + \tau_p = 5\,\mathrm{min}$, the disturbance path has $K_d = 1.5$ PV units per unit of D and $\theta_d + \tau_d = 3\,\mathrm{min}$. A disturbance step of $D = 2$ occurs. Find the integral of the error, and say whether the feedforward is early or late.

1.4 Concept check · 2 min

In a fermenter a feed step raises the oxygen demand of the culture only after a glucose pool has filled, about a minute later. A feedforward computes the agitator speed the new feed will need and applies it at once. What does it do to the dissolved oxygen?

(a) It holds the dissolved oxygen exactly, because the computed speed is right

(b) It lets the dissolved oxygen dip, as it would with no feedforward at all

(c) It drives the dissolved oxygen above the setpoint before the demand arrives

(d) It makes no difference, because the feedback trim removes the term

1.5 Theory and calculation · 6 min

(a) A heater fires $u = c\,(T_{\mathrm{SP}} - T_{\mathrm{in}})\,F$ with $c = 0.2\,\%$ per °C per kg/s, $T_{\mathrm{SP}} = 80\,^{\circ}\mathrm{C}$ and $T_{\mathrm{in}} = 20\,^{\circ}\mathrm{C}$. A ratio station sets $u = R\,F_m$. Find the design ratio R and the ratio needed after the inlet warms to 30 °C. *Checkpoint:* $R = 12$, then $R = 10$ % per kg/s

(b) With the inlet at 30 °C, a trim that adds a bias to $12\,F_m$ settles at the design feed of 5 kg/s. The feed then rises to 8 kg/s before the trim moves. Find the firing error and the outlet temperature error. *Checkpoint:* 6 % overfired, outlet 3.75 °C hot

1.6 Theory and calculation · 4 min

An agitator drive has a lag of 12 s. A proportional speed controller gives an inner loop gain $K_{c,\text{in}}K_p = 5$. Find the closed-loop time constant of the inner loop, and its ratio to an outer process whose fastest element has a time constant of 8 s.

1.7 Concept check · 2 min

A tank level is controlled by a cascade in which the level controller sets an inlet flow controller. The outlet is drawn by a metered pump whose demand steps by 50 %. What does the cascade do for this disturbance compared with a single level loop on the valve?

(a) Almost nothing, because the demand enters the tank and not the inlet line

(b) It removes most of the error, because the flow loop is fast

(c) It removes all of the error once the integral of the inner loop settles

(d) It doubles the error, because the inner loop always adds a lag

1.8 Theory and calculation · 5 min

A cross-limited burner uses the selectors of Equation (1.8) with $\lambda_0 = 1.2$ and $\delta = 0$.

(a) The demand has just risen to $u = 60$ while the measured fuel is 45 and the measured air is 60, in the units of the equation. Find both setpoints. *Checkpoint:* fuel 50, air 72

(b) Later the demand falls to $u = 40$ while the measured fuel is 55 and the measured air is 70. Find both setpoints. *Checkpoint:* fuel 40, air 66

1.9 Concept check · 2 min

A level loop adds a static feedforward from the inlet flow transmitter, which is downstream of its own control valve, with the gain that would be right for a disturbance at the outlet. What does this feedforward do?

(a) It rejects the outlet demand as a feedforward from the outlet would

(b) It rejects the supply pressure exactly, as a cascade does

(c) It has no effect, because the inlet flow is not a disturbance

(d) It feeds the valve's own effect back to the valve, which can oscillate

1.10 Theory and calculation · 4 min

The heater of Section 1.4.3 uses the lead–lag $4\,(4s+1)/(2s+1)$ % per kg/s, with time in minutes. The feed steps by 2 kg/s. Find the change of firing it asks for at once, after 2 min, and at steady state.