INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 19 Control of Biological Systems and Human Performance

CHAPTER 1

Introduction: Controlling Living Systems

Pages 1–22

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Controlling Living Systems," in *Control of Biological Systems and Human Performance*, John D. Hedengren, Ed., vol. 19 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 1, pp. 1–22.

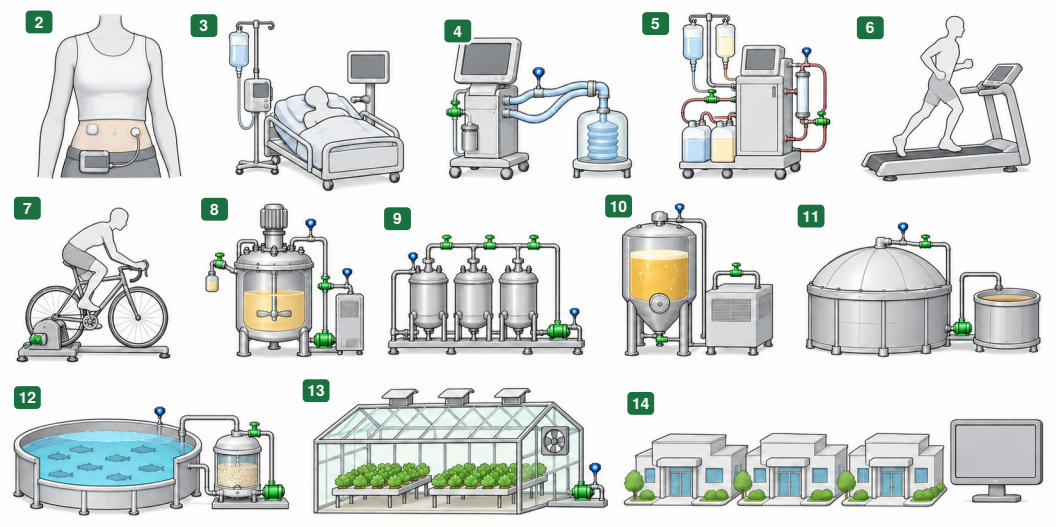
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v19-biological-human/

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

CHAPTER 1

Introduction: Controlling Living Systems



The chapters 2 artificial pancreas 3 closed-loop anesthesia 4 mechanical ventilator 5 kidney replacement 6 treadmill pacing 7 cycling exertion 8 antibody culture 9 insulin refolding 10 lager fermenter 11 anaerobic digester 12 recirculating fish farm 13 greenhouse climate 14 epidemic response

Notice that the top row is the bedside and a runner, the middle row a rider and the vessels in which cells make medicines, beer and gas, and the bottom row a fish farm, a glasshouse and a city. Illustration, not to scale.

IN PLAIN TERMS **Keeping Living Things Within Safe Limits**

The main idea. A pump gives a patient a drug, a feed keeps cells growing, and a cold valve holds a tank of beer at its temperature. In each case the thing being controlled is alive, and no two of them are quite the same.

The challenge. The value that matters most is often measured once a day or not at all. A pump can add a drug but cannot take it back, and some limits must never be crossed, such as a blood sugar that falls too low.

Why it matters. Good control keeps patients safe, makes medicines and food reliably, and helps a city keep its hospital from filling up.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.5, 1.4.6 and 1.4.7, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the examples P-1 and C-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.7 with Table 1.1, Section 1.5 with Figures 1.7 and 1.9, and Section 1.8.

1.1 The Question This Volume Answers

A patient, a runner, a culture of cells, a tank of fish and a city have properties that a heat exchanger does not. No two of them are alike, and the model the controller carries describes an average one. The variable that matters is often measured once a day, a week late or not at all. The plant grows, so the demand it puts on the actuator doubles while the controller waits for its next sample. Many of the

actuators push one way only, and the limits that the plant must not cross are hard on one side and soft on the other. This volume answers one question: how does a loop hold a living plant inside limits it may not exceed when the individual differs from the model, the measurement is sparse and the actuator cannot take back what it has given?

Seven control problems define the field, and every chapter meets several of them.

- *The variable that matters is not measured.* No probe reads the drug at its site of action in the brain, the breath that reaches the lung behind a leak, the nitrifiers on a biofilter or the infections in a city.
- *Each individual differs from the model.* The propofol an adult needs varies by a factor of three, the plant gain of a runner on a belt differs by a factor of 2.1 between four runners, and a yeast strain's diacetyl runs to 1.96 times what its estimator predicts.
- *Measurements are sparse and late.* A culture is counted once a day, a digester is titrated each morning and reported six hours later, and an epidemic is reported a week after the onsets it counts.
- *The plant grows.* Cells double, their appetite doubles with them, a biofilter must be grown before it can carry its load, and an epidemic doubles in days.
- *Actuators push one way.* A pump can add insulin or propofol and cannot withdraw it, a glycol valve can only cool, and a feed pump can only feed.
- *Limits are hard and asymmetric.* A deep low glucose is an immediate danger and a high one a slow one, and a plateau pressure, a nitrite level or a count of hospital beds is a line the loop must not cross.
- *Some actuators are stores.* The power a rider buys with effort above a sustainable level is drawn from a work capacity that runs out, and a digester's buffer spends itself on the acids it neutralizes.

Thirteen chapters answer the question on thirteen plants, each with an app that runs in a browser. Chapters 2 to 5 are at the bedside: an insulin pump, an anesthetic infusion, a ventilator and a kidney replacement circuit. Chapters 6 and 7 pace a runner and a rider. Chapters 8 to 10 make an antibody, refold insulin and ferment a lager. Chapters 11 to 13 feed a digester, a fish farm and a glasshouse crop, and Chapter 14 steers an epidemic.

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

- Write a compartment model, find its time constants and the infusion that holds a concentration, and say why that infusion keeps falling.
- Write a growth law, predict how fast demand grows, and find the fold beyond which an inhibited culture washes out.
- State what a population model gets wrong on an individual and which structure removes the error.
- Correct a model from a sample at the time the sample was drawn, and say what interval a growing plant allows between samples.
- Size the margin a one-way actuator needs in front of a hard limit, and place an override where a limit binds.

This chapter develops each of these on two small examples and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, feedforward and cascade from Volume 3, selectors and overrides from Volume 4, delays and estimation from Volume 5, multiplicity from Volume 7, batch operation from Volume 8 and predictive control from Volume 9.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from a pandemic and a purified extract to the first automatic loops at the bedside, in the glasshouse and in the digester.

In 1918 seventeen cities of the United States brought in measures against pandemic influenza, among them the closure of schools, churches and theaters. Cities that put several measures in place early in their epidemic had peak death rates about half those of cities that did not [1], the steering problem of Chapter 14. An apparatus for the prolonged administration of artificial respiration,

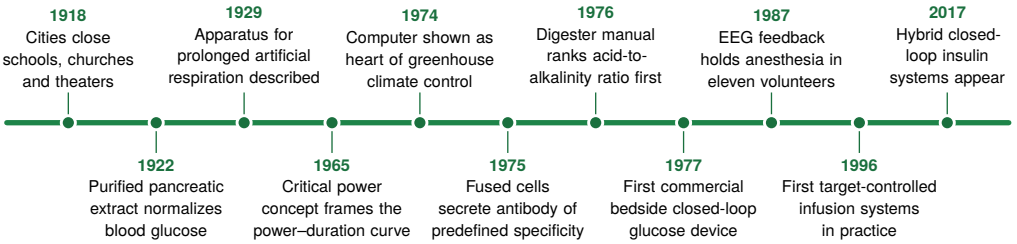


Figure 1.1 Milestones of biomedical and biological control. The dates are from [1] (1918), [2] (1922 and 2017), [3] (1929), [4] (1965), [5] (1974), [6] (1975), [7] (1976), [8] (1977), [9] (1987) and [10] (1996). Notice that the first loops closed by machine on living plants all fall within the 1970s and 1980s. Events are evenly spaced, not to a time scale.

described in 1929, enclosed the body below the neck in a tank and breathed for a child paralyzed by poliomyelitis almost continuously for 122 hours [3], a forerunner of the ventilator of Chapter 4. In 1965 the concept of critical power gave the power a person can hold over a given time a simple frame of two parameters, a fixed energy reserve and the critical power itself, and in 1981 the concept was applied to cycling [4]. That reserve is the store of Chapter 6 and Chapter 7.

> ⧖ HISTORY **1974 to 1987: First Automatic Loops on Living Plants**
>
> At the Dutch horticultural fair of 1974 a computer was shown for the first time as the heart of a climate control system for greenhouses, and within two years systems from five producers were on sale [5]. In 1977 the first continuous arteriovenous hemofiltration was given in Göttingen. Its removal was first set by hand, by hanging the filtrate bag higher or lower, and later systems balanced the fluid automatically with scales and electronic clamps [11]. A report of 1978 found such balancing accurate enough to replace the bed scales that bedridden patients had needed [12]. In 1987 a loop held anesthesia in 13 volunteers by feedback on the median frequency of the electroencephalogram, with a setpoint of 2 to 3 Hz, and succeeded in 11 of them [9]. The climate loop of Chapter 13, the balance of Chapter 5 and the depth loop of Chapter 3 are the modern forms of these three problems.

A report of 1975 described continuous cultures of fused cells that secreted antibody of a predefined specificity [6], the first monoclonal antibodies, a class of product that Chapter 8 makes in cultured cells. In February 1976 a federal operations manual for sludge digesters ranked the ratio of volatile acids to alkalinity first among the indicators that warn of an upset [7], the measurement on which Chapter 11 feeds. The first target-controlled infusion systems came into practice in 1996, with a model for propofol [10], the model inversion of Section 1.4.2.

> ⧖ HISTORY **1922 to 2017: Insulin, From an Extract to a Closed Loop**
>
> On 23 January 1922 a purified pancreatic extract was given under the skin to a patient with diabetes. The sugar and ketones in his urine disappeared, his blood glucose became normal, and he lived 13 more years on daily injections [2]. Half a century later the dose was put under feedback at the bedside. Five teams reported closed loops between 1974 and 1978 that measured glucose in a vein and computed the insulin from the glucose and its rate of change, a proportional-derivative law, and in 1977 one of them became the first commercial device [8]. The first commercially used insulin pump followed in 1978, glucose sensors under the skin were joined to pumps from the early 2000s, and in 2017 hybrid closed-loop systems combined a continuous monitor, a control algorithm and a pump [2]. Chapter 9 makes the insulin, and Chapter 2 closes the loop on it.

1.3 Notation and Conventions

Concentrations are in the units of each field: glucose in mg/dL at the bedside and in g/L in a culture, drugs in mg/L, which equals µg/mL, and nitrogen in mg N/L. A volume V is in liters unless a chapter gives cubic meters, a clearance Cl is a volume cleared per unit time, and a growth rate μ is per hour or per day as the chapter's clock runs. The clocks differ. Chapter 2 runs a model clock in hours on rate constants per minute. Chapter 3 counts minutes, Chapter 4 seconds and milliseconds, and Chapter 6 and Chapter 7 seconds. Chapter 5 closes its balance in hours over drive loops in seconds. Chapter 8, Chapter 9 and Chapter 10 count hours and days, Chapter 12 days, hours and minutes, Chapter 13 minutes and hours, and Chapter 11 and Chapter 14 days.

The framework's h is its integration step and its a the controller action. Loop action follows Section 1.6.2: $a = +1$ is reverse acting, for a plant whose measurement rises with the manipulated variable, and $a = -1$ is direct acting. An insulin pump on glucose is direct acting, because insulin lowers glucose, and so is the depth loop of an anesthetic, because more propofol lowers the depth index. So are a ventilator's expiratory valve on airway pressure, a glycol valve on a fermenting zone, a glasshouse vent on air temperature, a blower on dissolved carbon dioxide and a dial of distancing measures on a hospital census. Base on pH, a jacket on a culture's temperature, gas on dissolved oxygen and a belt on a runner's heart rate are reverse acting. Several chapters write the action with their own symbol: $\mathscr{A}$ in Chapter 3, Chapter 4 and Chapter 5, γ in Chapter 10, γ_a in Chapter 13 and α in Chapter 12.

The chapters keep their own symbols, and many collide. ρ is the slack price of the framework's predictive controller, a patient's sensitivity in Chapter 3, the share of infections reported in Chapter 14 and a ration fraction in Chapter 12. γ is also the Hill slope in Chapter 3, the inverse infectious period in Chapter 14 and the psychrometric constant in Chapter 13. λ is a move suppression in Chapter 2, a bias weight in Chapter 10 and Chapter 8 (λ_X) and the latent heat in Chapter 13. σ is the monitor noise in Chapter 2, the lactate switch in Chapter 8 and the inverse latent period in Chapter 14. D is a dilution rate in Chapter 8 and Chapter 11, the diacetyl concentration in Chapter 10 and the vaccine doses a day in Chapter 14. φ is a central-volume multiplier in Chapter 3, a throttle parameter in Chapter 5 (which writes the filtration fraction FF), the relative humidity in Chapter 6 and Chapter 7, and the air exchange in Chapter 13, while ϕ is the accumulated fatigue in Chapter 6 and Chapter 7. C is a compliance in Chapter 4 and a concentration nearly everywhere else. k_e is the plasma insulin clearance in Chapter 2 and an expiratory valve coefficient in Chapter 4, and E is the exertion in Chapter 7, the exposed class in Chapter 14 and the ethanol in Chapter 10. a itself is a pressure coefficient in Chapter 3, the airflow in Chapter 6 and Chapter 7, the vent opening in Chapter 13 and a correlation coefficient in Chapter 14, none of them the action. μ is a specific growth rate except as the peak-power factor of Chapter 6 and the mean daily report of Chapter 14. This chapter writes the clearance ratio of an individual ν, a letter that no chapter uses.

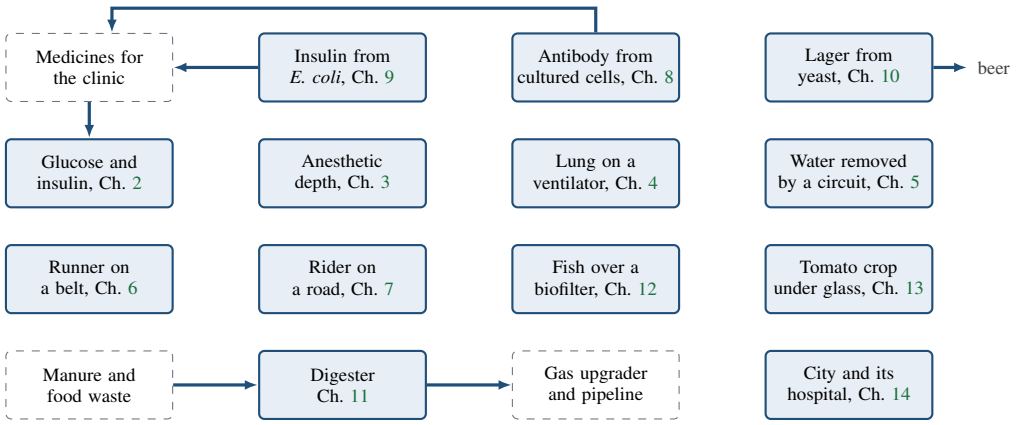


Figure 1.2 The Plants of the Volume. Filled Blocks Are the Plants of Chapters 2 to 14, Dashed Blocks Lie Outside the Volume. Top Row: Cells Making Medicines and Food. Second Row: the Patient at the Bedside, Who Receives the Insulin. Third Row: People Who Exercise and a Farm's Living Stock. Bottom Row: Microbes That Turn Waste Into Gas, and a City in an Epidemic.

1.4 Models, Variability and Limits of Living Plants

1.4.1 The Flow Scheme

Figure 1.2 places each chapter's plant. The top row is the manufacture of medicines and food by cells: an antibody secreted by cultured mammalian cells, insulin refolded from a bacterial precursor, and beer fermented by yeast. The insulin goes to the second row, the patient at the bedside, where a pump delivers it under a glucose monitor and other pumps deliver an anesthetic, a ventilator breathes and a circuit removes water. The third row holds people who exercise and a farm's living stock, fish over a biofilter and a tomato crop under glass. The bottom row is microbes that turn waste into gas, and a city whose epidemic fills its hospital.

Two running examples carry the rest of this section, because no single plant has both a bloodstream and a growing culture. *P-1* is a patient receiving a drug by intravenous infusion. Its central compartment has $V_1 = 10\,\mathrm{L}$ and an elimination clearance $\mathrm{Cl}_1 = 0.5\,\mathrm{L/min}$, its peripheral compartment has $V_2 = 40\,\mathrm{L}$ and exchanges with the central one at $\mathrm{Cl}_2 = 0.4\,\mathrm{L/min}$, and its effect site follows the plasma with $k_{e0} = 0.25\,\mathrm{min}^{-1}$. The target is 2 mg/L, inside a teaching window of 1.5 to 4 mg/L. *C-1* is a fed-batch culture that starts at 10^6 cells per milliliter and 4 g/L of glucose, doubles in 24 h and takes 0.5 g/L of glucose for each 10^6 cells per milliliter it grows. It is counted at 08:00 each day, and the count is reported 4 h later. Every number of both is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Compartments and Pharmacokinetics

A compartment is a volume that is treated as well mixed. Its balance in concentration is Equation (1.1) for an infusion u in mg/min into one compartment that is cleared at Cl_1.

$$V_1 \frac{\mathrm{d}C_1}{\mathrm{d}t} = u - \mathrm{Cl}_1 C_1, \qquad \tau = \frac{V_1}{\mathrm{Cl}_1}, \qquad C_{1,ss} = \frac{u}{\mathrm{Cl}_1} \tag{1.1}$$

On P-1 alone, the central compartment would have $\tau = 20\,\mathrm{min}$, a maintenance rate of 1 mg/min for the target and a loading dose of $V_1 C^* = 20\,\mathrm{mg}$ to reach it at once. Without the loading dose the plasma would take three time constants, 60 min, to reach 95 % of the target.

Drug also leaves the blood for tissue and returns from it, and the effect is produced at a site that the blood reaches with a lag. The model then adds a peripheral compartment, and an effect site of

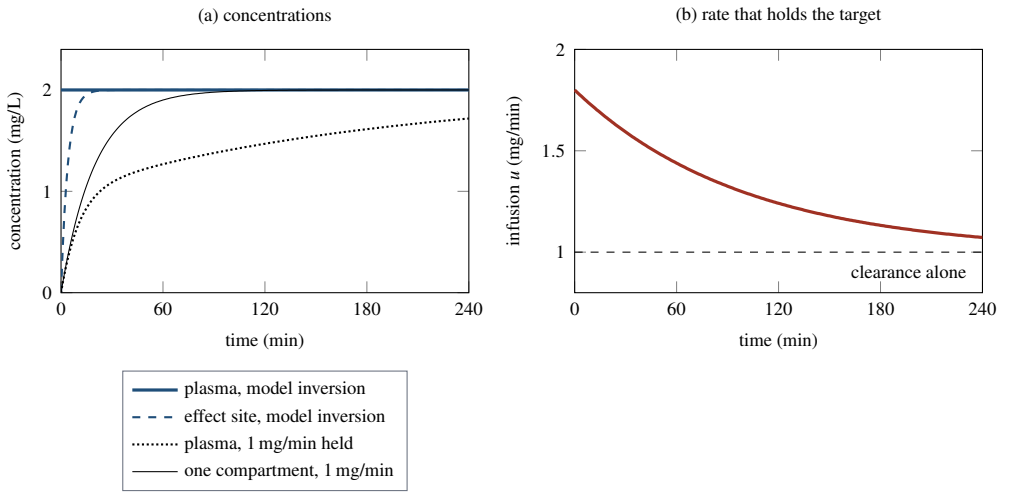


Figure 1.3 P-1 Started by Model Inversion and by a Constant Infusion. (a) Under Model Inversion the Plasma Is at 2 mg/L From the Start (Thick Solid) and the Effect Site Follows It (Dashed). A Constant 1 mg/min Without a Bolus Reaches the Target in About an Hour in One Compartment (Thin Solid) but Is Still Short After Four Hours in Two (Dotted). (b) The Rate That Holds the Target Falls for Hours Toward the 1 mg/min That Clearance Alone Needs.

negligible volume linked to the plasma [13], as in Equation (1.2).

$$\begin{aligned} V_1 \frac{\mathrm{d}C_1}{\mathrm{d}t} &= u - \mathrm{Cl}_1 C_1 - \mathrm{Cl}_2 (C_1 - C_2) \\ V_2 \frac{\mathrm{d}C_2}{\mathrm{d}t} &= \mathrm{Cl}_2 (C_1 - C_2), \qquad \frac{\mathrm{d}C_e}{\mathrm{d}t} = k_{e0} (C_1 - C_e) \end{aligned} \tag{1.2}$$

The two compartments of P-1 have time constants of 10.6 min and 189 min. A constant 1 mg/min without a loading dose gives 1.27 mg/L after an hour and 1.72 mg/L after four, because the peripheral compartment is still filling.

A target-controlled pump inverts the model instead [14]. It gives the bolus $V_1 C^*$ and then the rate of Equation (1.3), which holds $C_1 = C^*$ exactly while C_2 fills with the time constant V_2/Cl_2.

$$u(t) = \mathrm{Cl}_1 C^* + \mathrm{Cl}_2 C^* \, \mathrm{e}^{-\mathrm{Cl}_2 t / V_2} \tag{1.3}$$

On P-1 the rate starts at 1.80 mg/min, falls to 1.44 mg/min after an hour and to 1.07 mg/min after four, and only then nears the 1 mg/min that clearance alone needs, as Figure 1.3 shows. The plasma is at the target from the first second, while the effect site reaches 95 % of it at $\ln 20 / k_{e0} = 12.0$ min.

Chapter 3 is this example at full scale, with three compartments and an effect site for each of two drugs. Its slow compartment fills with a time constant of 285 min, so the rate that holds a maintenance concentration falls from 3.43 to 2.84 mg/kg/h over four hours toward 2.39 mg/kg/h, and the pump inverts a copy of the model that never sees the patient. In Chapter 2 the remote insulin action follows plasma insulin exactly as an effect site follows plasma, with a time constant $1/p_2 = 81.3$ min behind a plasma clearance of $1/k_e = 11.1$ min, on the glucose and insulin compartments of the minimal model [15]. In Chapter 4 the lung is one compartment of compliance C behind a resistance R, with $RC = 0.5$ s, and Chapter 5 closes its balance on a weighed running total that integrates the net removal. The physiological states of Chapter 6 and Chapter 7 follow the belt or the effort through lags in series, reaching 63 % of a belt step in 40, 96, 257 and 475 s in the first. The compartments of Chapter 14 hold people rather than drug, from susceptible to removed [16], and Chapter 10 splits one tank into three exchanging zones.

1.4.3 Growth Kinetics

Cells grow at a specific rate μ that depends on their substrate S, in the saturating form of Equation (1.4) [17], and they consume substrate in proportion to their growth through a yield Y.

$$\mu = \mu_{max}\frac{S}{K_S + S}, \qquad \frac{\mathrm{d}X}{\mathrm{d}t} = (\mu - k_d)X, \qquad q = \frac{\mu}{Y} \tag{1.4}$$

where X is the cell density, k_d a death rate and q the uptake per cell. A specific rate is the property that makes a living plant unlike a reactor. The demand for substrate is proportional to the cells, which multiply, so a constant feed that suits today falls short tomorrow. C-1 grows at $\mu = \ln 2/24 = 0.0289\,\mathrm{h}^{-1}$ while glucose is plentiful, and with $K_S = 0.05\,\mathrm{g/L}$ it still grows at 95 % of that rate at 1 g/L. The glucose it takes in day k is $X_0(2^{k+1} - 2^k)/Y$, so the daily uptake runs 0.5, 1, 2, 4 and 8 g/L over its first five days. Chapter 8 measures the same appetite on a real growth law, 0.53 to 4.27 g/L a day over five days with each day 44 to 94 % above the last, and Chapter 10 shows the heat of fermentation growing with the yeast: with the valves shut its beer warms 0.56 °C in the first day and runs away later.

Some substrates inhibit the organism that consumes them. The Haldane form of Equation (1.5) rises, peaks at $S_M = \sqrt{K_S K_I}$ and falls [18]. In a continuous vessel with dilution rate D, a steady population needs $\mu(S) = D + k_d$.

$$\mu = \frac{\mu_{max}\,S}{K_S + S + S^2/K_I}, \qquad \mu(S_M) = \frac{\mu_{max}}{1 + 2\sqrt{K_S/K_I}}, \qquad D_c = \mu(S_M) - k_d \tag{1.5}$$

Below D_c there are two roots. The lower one is a stable steady state and the upper one bounds the basin from which the population washes out. Above D_c there are none, and the steady states on either side of D_c form a fold. With the methanogen constants of Chapter 11, $\mu_{max} = 0.164\,\mathrm{d}^{-1}$, $K_S = 20$, $K_I = 130\,\mathrm{mmol/L}$ and $k_d = 0.02\,\mathrm{d}^{-1}$, the equation gives $S_M = 51.0\,\mathrm{mmol/L}$, $\mu(S_M) = 0.0919\,\mathrm{d}^{-1}$ and $D_c = 0.0719\,\mathrm{d}^{-1}$, a feed of 287.6 m^3/d to its 4000 m^3 tank. That is the fold the chapter reports, 1.80 times its design feed. At the design feed of 160 m^3/d the two roots are 12.2 and 213 mmol/L of acids, and they meet at S_M as the feed approaches the fold.

Two other chapters use rates of a different shape. In Chapter 12 the biofilter is two nitrifier populations that grow to the load they carry, so a larger ration waits for the slower population to grow and the nitrite returns days after the feed changes. In Chapter 14 the infections grow at a rate r that satisfies $(1 + r/\sigma)(1 + r/\gamma) = R_0$ for exponential latent and infectious stages [19], which doubles them every 4.69 d at $R_0 = 3$. Chapter 9 sets a first-order fold against a second-order aggregation of the same protein, so a charge of unfolded protein U_0 added at once folds the share $\ln(1 + \alpha)/\alpha$ with $\alpha = k_a U_0/k_f$, 75.3 % at $\alpha = 0.72$, and a charge fed slowly folds more.

1.4.4 Variability Between Individuals

A population model describes the typical individual and the spread around it [20], and a loop that inverts it is exact only on the typical individual. Let an individual's clearance be ν times the population's, with the volumes the same. A pump that holds the model at C_T then delivers, in the steady state, a concentration C_T/ν, as in Equation (1.6). The window $[C_{lo}, C_{hi}]$ holds every individual with ν between ν_{lo} and ν_{hi} only if one target satisfies both ends, which needs the population's spread to be no wider than the window.

$$C_{1,\infty} = \frac{C_T}{\nu}, \qquad \nu_{hi}C_{lo} \le C_T \le \nu_{lo}C_{hi} \;\Rightarrow\; \frac{\nu_{hi}}{\nu_{lo}} \le \frac{C_{hi}}{C_{lo}} \tag{1.6}$$

Take three patients of P-1 with $\nu = 0.6$, 1.0 and 1.6. Held at the 2 mg/L target by the population model, they head for 3.33, 2.00 and 1.25 mg/L, and the third leaves the window within the first hour. The population's spread, $1.6/0.6 = 2.67$, equals the window's, $4/1.5 = 2.67$, so open-loop dosing could hold all three only at a target of exactly 2.4 mg/L, with no margin at either end.

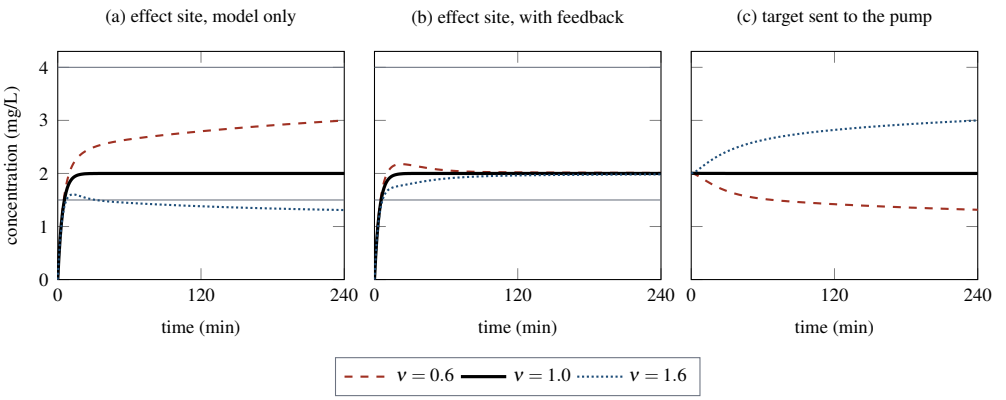


Figure 1.4 Three Patients of P-1 Whose Clearance Is 0.6 (Dashed), 1.0 (Thick Solid) and 1.6 (Dotted) Times the Population's. (a) The Pump Inverts the Population Model and Holds It at 2 mg/L: Only the Population Patient Reaches the Target, and the Fast Clearer Leaves the Window (Thin Lines at 1.5 and 4 mg/L). (b) An Outer PI Loop on the Measured Effect Site Brings All Three to the Target. (c) The Target That Loop Sends to the Pump Moves Toward ν Times 2 mg/L.

Feedback from a measurement of the effect removes the error. In Figure 1.4 an effect monitor reads the effect site once a minute, and an outer PI loop ($K_c = 0.5$, $\tau_I = 8$ min) moves the target sent to the model inversion. It compares the reading with the effect site the model predicts for the population patient, so it acts only on what the model did not predict and leaves the population patient alone. All three patients are within 5 % of the target for good from 45 min, 12 min and 58 min, and after four hours the targets the loop sends are 1.32, 2.00 and 3.00 mg/L, on their way to the νC^* of 1.2, 2.0 and 3.2 mg/L. The commanded target is a parameter estimate, made by an integral that nobody told which parameter was wrong. On the slow clearer the loop's corrections stop the pump for 4.8 min in all, a reminder that a downward correction is made by the patient's own elimination.

Chapter 3 measures exactly this. Under model inversion alone its settled depth index spans 52.1 to 70.4 as the patient's sensitivity moves from 0.70 to 2.0 times the population's, and with the depth loop closed it spans 50.6 to 51.7 while the commanded concentration moves from 1.90 to 5.03 µg/mL. Identifying a patient takes about half an hour, because the outer loop must be slower than the effect-site lag, a pace published closed loops also show [21]. Other chapters use other remedies. Chapter 6 carries its gains as a fraction of the runner's own threshold pace, a schedule on the one number known for each runner, because its plant gain moves by a factor of 2.1 across four runners. Chapter 7 scales the power loop by the rider's critical power so that its loop gain below the sustainable exertion is 0.667 for every rider. Chapter 4 identifies the lung on every breath, and Chapter 10, Chapter 8 and Chapter 9 correct a declared model from laboratory samples, the subject of the next subsection. Variability that no measurement reveals is the hardest kind. Exercise raises insulin sensitivity in Chapter 2 with no change in any measured quantity, and under a hidden inhibition in Chapter 11 a predictive feed controller loses to the sampled PI loop.

1.4.5 Sparse and Late Measurements

A culture that doubles between samples takes more in each interval than in all the earlier ones together. A bolus sized from a sample restores the substrate that the sample had, not what the broth has when the report arrives, so the broth ends each interval below the target by what the cells took since the draw. The interval's uptake must fit inside the band, as in Equation (1.7).

$$\frac{1}{Y}\int_{t_k}^{t_{k+1}} \mu X \, \mathrm{d}t \;\le\; G_{hi} - G_{lo} \tag{1.7}$$

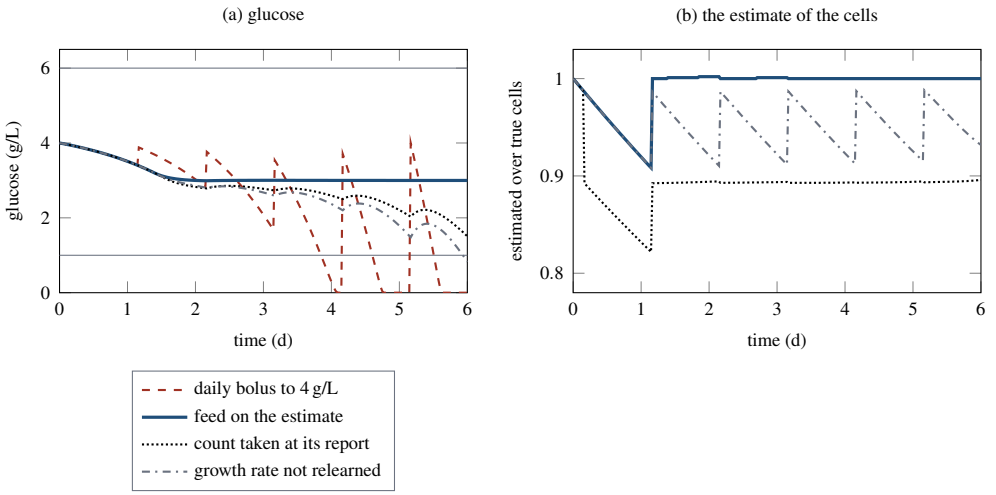


Figure 1.5 C-1 Over Six Days. (a) A Daily Bolus Sized From the Morning Sample (Dashed) Leaves the Band of 1 to 6 g/L (Thin Lines) on Day 3.85 and Runs Out on Day 4.06. A Feed on the Estimated Uptake, Corrected at the Time Each Count Was Drawn, Holds 3 to 3.5 g/L (Solid). The Same Feed With Each Count Applied as if Taken at Its Report Runs Low (Dotted), and Without Relearning the Growth Rate It Leaves the Band on Day 5.93 (Dash-Dotted). (b) The Estimated Cells Over the True Cells for the Three Feeds: Exact After Day 1, About 11 % Low Throughout, and a Sawtooth Corrected at Each Report.

C-1 fed a daily bolus to 4 g/L fails that test on its fifth day, when the day's uptake of 8 g/L exceeds the 5 g/L band of 1 to 6 g/L. The test is necessary and not sufficient. With its target below the top of the band and a four-hour report, the bolus fails a day earlier. Its morning samples read 4.00, 3.51, 3.03, 2.08 and 0.29 g/L, and in Figure 1.5 it leaves the band on day 3.85 and runs out of glucose on day 4.06. Starved for 21.9 h in six days, it ends with 31.2 million cells per milliliter against 59.9 million for the feed described next.

The remedy is to feed what the cells are estimated to take, from a model that runs between samples and is corrected by each one at the instant the sample was drawn [22, 23]. Equation (1.8) is the correction in its logarithmic form, with the sample drawn at t_s and reported later.

$$b \leftarrow b + \lambda \left[\ln X_{lab} - \ln \hat{X}(t_s) - b \right], \qquad \hat{X}_c(t) = \hat{X}(t)\, \mathrm{e}^{b} \tag{1.8}$$

where b is the bias and $0 < \lambda \leq 1$ its weight. With a constant model error the bias closes the fraction $1-(1-\lambda)^k$ of it after k samples. C-1's feed uses $\lambda = 1$, starts from a model growth rate of $0.025\,\mathrm{h}^{-1}$, which is 13 % slow, and relearns the rate from the last two counts. It holds glucose between 2.99 and 3.51 g/L from day 1 to day 6. The time stamp matters. Applying each count as if it had been drawn at the report leaves the estimate low by $\mathrm{e}^{-\mu\tau}$, here $\mathrm{e}^{-0.0289\times 4} = 0.891$, and the glucose then falls to 1.51 g/L. Keeping the slow model growth rate instead of relearning it leaves the estimate up to 9 % low before each report, a share of an uptake that doubles daily, and the glucose leaves the band on day 5.93.

Several chapters read a key variable this way. Chapter 8 counts its cells once a day with a four-hour report, and a bolus to 4 g/L leaves its culture outside the band for 34.8 h, while a culture held warm until day 8 eats up to 5.93 g/L in a day, more than its band is wide. Chapter 10 corrects a diacetyl model by a result that arrives a day after its sample, and Chapter 9 corrects its unfolded protein by each assay at the hour of the draw. In Chapter 11 the titration drawn at 08:00 is reported six hours later, the gas of an overloaded tank peaks on day 2.63 and the titration warns on day 3.58, and a loading gain of 30 swings the feed between its limits where a gain of 8 settles. The reports of

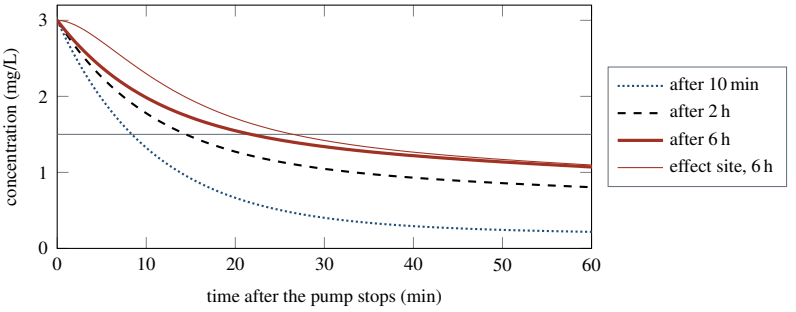


Figure 1.6 P-1 Held at 3 mg/L for Ten Minutes (Dotted), Two Hours (Dashed) and Six Hours (Thick Solid), Then Left With the Pump Stopped. The Plasma Reaches 1.5 mg/L (Thin Gray Line) After 8.4, 14.5 and 21.8 min, and After Six Hours the Effect Site (Thin Solid) Takes 26.5 min.

Chapter 14 count infections about eleven days old, scaled by an unmeasured share, and the hospital census, counted precisely, is what lets all three of its estimators tell a testing cut from a turn in the epidemic [24]. Chapter 5 shows the limit of any one measurement: a balance loop that holds a drifting scale moves the patient to keep the scale right, and only the pumps' own counts, a second measurement of the same flows, reveal the drift.

> **DID YOU KNOW? Five to Ten Minutes**
>
> A federal manual of 1976 told the operators of sludge digesters to run the pH and temperature tests within 5 to 10 min of drawing a sample. A sample left standing too long releases carbon dioxide, and its pH rises [7]. A late result can also be a wrong one.

1.4.6 One-Sided Actuators, Stores and Safety Limits

When a pump can only add, the way down is the plant's own removal. In one compartment a stopped pump lets the concentration fall as in Equation (1.9), and with a peripheral compartment the fall slows the longer the infusion has run, because the tissue returns what it took [25].

$$t_{lim} = \tau \ln \frac{C_0}{C_{lim}} \tag{1.9}$$

For P-1 held at 3 mg/L, one compartment would need $20\,\text{min} \times \ln 2 = 13.9\,\text{min}$ to fall to 1.5 mg/L. With the peripheral compartment, Figure 1.6 gives 8.4 min after a ten-minute infusion, 14.5 min after two hours and 21.8 min after six, and the effect site takes 4.7 min longer again. Raising the target from 1.5 to 3 mg/L takes a 15 mg bolus, under two seconds at a pump capacity of 500 mg/min. The actuator moves the plant up in seconds and leaves it to come down in tens of minutes.

The chapters meet the same asymmetry. In Chapter 2 exercise drops glucose from 74.9 to 47.8 mg/dL within 0.04 h, before the controller's next execution, and the pump then sits at zero for 0.84 h while the insulin already given goes on acting. Deployed systems therefore suspend delivery on a prediction [26] or bound the insulin on board [27]. In Chapter 3 the depth index is lowest a quarter of a minute after the stimulation stops, sooner than any infusion change can reach the effect site. The glycol valves of Chapter 10 can only cool a tank whose yeast make the heat, and the inspiratory valve of Chapter 4 can only add gas, so no controller holds the airway pressure against a leak larger than the open valve passes.

> ✦ DID YOU KNOW? **A Twentieth of a Unit an Hour**
>
> Insulin pumps let the programmed delivery be changed in increments as small as 0.05 to 0.25 units of insulin an hour [2]. Like the pump of P-1, it still cannot take back a unit it has delivered.

Because the two sides of a limit carry different costs, a symmetric error integral misjudges such a loop. Glucose is reported as time in range with separate thresholds below [28], and an objective of the form of Equation (1.10) weights the low side by w.

$$J = \int \max(0,\, y - y_{hi})\,\mathrm{d}t + w \int \max(0,\, y_{lo} - y)\,\mathrm{d}t \tag{1.10}$$

Chapter 2 finds that a run that falls well below its low flag beats a run that never does on the symmetric score, and that the ranking reverses at $w = 7.9$, the ratio of the two runs' excursion areas. The margin a loop keeps from a hard limit has to cover both the spread of individuals of Section 1.4.4 and the slow way back of Equation (1.9). Where a limit is approached through another variable, an override takes the actuator, as Volume 4 develops. Chapter 6 and Chapter 7 share one actuator among limits on a low selector, Chapter 5 caps its blood pump on access pressure and its ultrafiltration on transmembrane pressure, taking the replacement fluid before the removal, Chapter 12 cuts the ration on a low select of four limits, and Chapter 13 opens a screen gap, a vent and then heat as the humidity nears its limit.

Some actuators are stores. A rider's effort above a sustainable level draws on a work capacity W' above critical power [29], and when the power available falls in proportion to the capacity left [30], a held effort E above the sustainable E_c fades onto critical power with the time constant of Equation (1.11).

$$\tau_W = \frac{W'_0\,(1 - E_c)}{(P_{pk} - CP)\,(E - E_c)} \tag{1.11}$$

For the recreational rider of Chapter 7, with $W'_0 = 16\,\mathrm{kJ}$, $P_{pk} = 950\,\mathrm{W}$, $CP = 220\,\mathrm{W}$ and $E_c = 0.75$, it gives 21.9 s at full effort and 54.8 s at $E = 0.85$, the chapter's values. The actuator loses its upper range without moving. Chapter 6 drains the same reserve at exactly the excess power, and Chapter 11 spends its alkalinity on the acids it neutralizes.

1.4.7 How Each Loop Handles Variability and Limits

Each chapter's plant can be read in three parts: what varies between individuals or during the run, the limit that may not be exceeded, and the structure that keeps the loop inside it. Table 1.1 lists the three for each chapter. Chapter 15 adds the measured results and compares how each loop handled variability and limits it could not exceed.

> ✦ DID YOU KNOW? **Half a Liter an Hour**
>
> A clinical report of 1978 on automatic hemofiltration found that low blood pressure, nausea and muscle cramps during treatment could be reduced, provided that the effective fluid withdrawal did not exceed 0.5 L/h [12]. The reference prescription of Chapter 5 removes 100 mL/h.

1.5 Instruments and Final Elements in This Volume

Figure 1.7 lists what the controllers of each chapter read and what they move. Eight chapters close a loop on a model's estimate rather than on a reading, and six wait for a laboratory sample, an assay or a report, more than read any single kind of sensor. A pump or feeder is the commonest handle, in seven chapters, and two chapters move a person: a rider who follows a cue and a city that adopts a measure.

A monitor on the body reads a signal that stands for the variable that matters. A continuous glucose monitor reads the glucose of the fluid under the skin every 1 to 5 min, and its reading lags blood

Table 1.1 What Varies, the Limit and the Structure That Holds It, by Chapter

Ch.	What varies	Limit	Structure
2	Insulin sensitivity with exercise, unannounced meals, basal drift	Low glucose flag, pump floor at zero	Direct-acting PI with back-calculation, MPC told of meals, MHE with a meal state
3	Drug requirement by a factor of three, surgical stimulation	Surgical depth band, arterial pressure	Depth loop correcting a model-inverting pump, pressure loop, MPC with learned potency
4	Compliance, resistance, effort, leak	Plateau pressure, volume per kilogram	Per-breath least squares, protective cap from the identified compliance
5	Pump tubing, wear, scale drift, filter fouling	Transmembrane and access pressure, filtration fraction	Weighed balance loop, overrides, each scale checked against its pump
6	Runner, heat, fuel	Lactate, heart rate, core temperature	Gains scheduled on threshold pace, low selector with external reset, MPC with terminal value
7	Rider, a work capacity that drains	Heart rate, lactate, exertion range	Power loop, two overrides with margin-scaled reset, MPC
8	Appetite that doubles, cooled cells, probe drift	Glucose band never at zero, osmolality	Feed on estimated uptake, daily correction, hourly Raman trim, MPC
9	Folding constants, harvest load	Unfolded ceiling, hold times, lot limits	Scheduler on an assay-corrected estimate, transfer and endpoint rules
10	Yeast strain, mixing that stops	Recipe temperatures, diacetyl, shared glycol	Time-proportioned PI per zone, corrected estimator, glycol coordinator
11	Feed, temperature, hidden inhibition	Titration ratio at the fold	Sampled PI on a daily titration, economic supervisor, MPC
12	Biofilter capacity, fish biomass	Nitrite, ammonia, oxygen	Oxygen cascade with meal feedforward, feed governor, soft sensor, film observer
13	Weather and its forecast	Hourly humidity limit	Pipe cascade with external reset, three-stage humidity override, planner over the forecast
14	Reported share, testing, a variant	Hospital beds	Estimators anchored by the census, PI on the census, MPC with a margin, scenario tree

glucose by 5 to 10 min, a delay that matters most when the reading is used for decisions in real time [31].

Weight is the measurement a kidney replacement machine trusts most. Weighing with one or more scales is the commonest way to balance its fluids, because it is the most reliable over long treatments, and the weight stands in for the flow [32]. A scale rests on a load cell, strain gauges bonded to a metal element that stretches and compresses with the applied force [33] (Figure 1.8), the gauge that also sits in the power meter of Chapter 7.

In United States regulation an infusion pump moves fluid into a patient in a controlled manner by a piston, roller or peristaltic mechanism, and it may detect air in the line or a blocked line and sound an alarm [34]. Each of those mechanisms only pushes, which is why the pumps of Chapters 2, 3 and 5 are the one-way actuators of Section 1.4.6. The glycol valves of Chapter 10 are on/off solenoid valves, opened and closed by the current through a coil [33], so the loop on each zone works through the share of time its valve is open.

Figure 1.9 places the volume's loops on their platforms. Supervisory and lower loops in industry run with cycle times from milliseconds to minutes [35]. The apps span a wider range, from the 2 ms servo of Chapter 4 to the daily policy of Chapter 14, and several loops have a person above them. While a target-controlled pump runs, the clinical response and the depth of anesthesia are monitored, and closed loops on such an endpoint are still experimental [10].



Figure 1.7 Sensors and final elements by chapter: a round mark where the chapter's controllers read the device and a square mark where they move it, as its variables table and control structure name them. A flow set through a flow loop counts as a valve, and the last column of reads marks a controlled variable that is a model's estimate rather than a reading. Notice that the estimate and the sample outnumber every sensor.

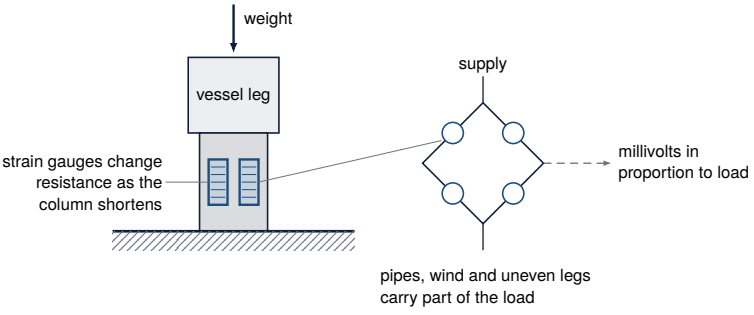


Figure 1.8 Principle of a strain-gauge load cell. Notice that a load that reaches the support by another path is not weighed, one way a scale can drift from the fluid it holds, the fault that Chapter 5 finds by comparing each scale with its pump.

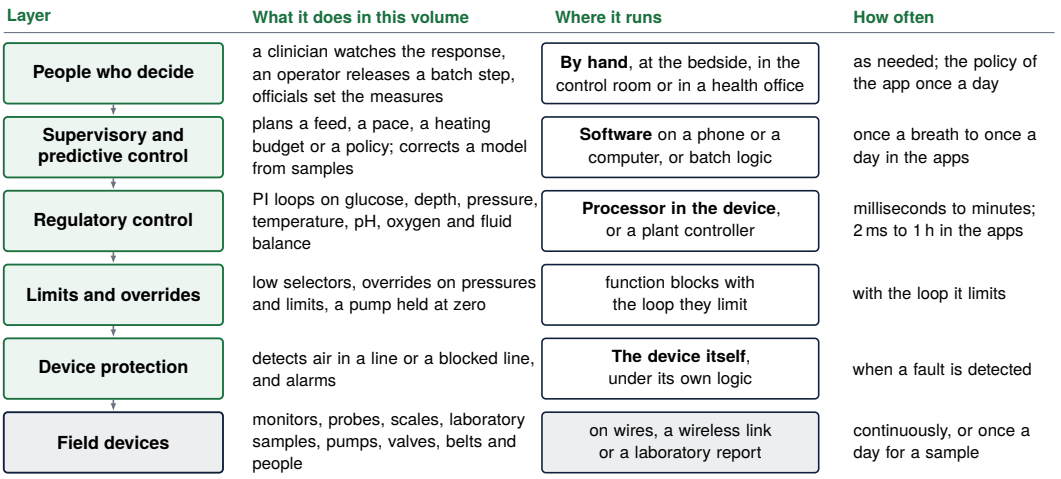


Figure 1.9 The control layers of this volume, the platform that hosts each and how often each acts. Loop periods are from [35] and the app periods from the chapters' technology sections. Notice that the top layer is a person, not a program.

Table 1.2 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.

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 1.2 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.

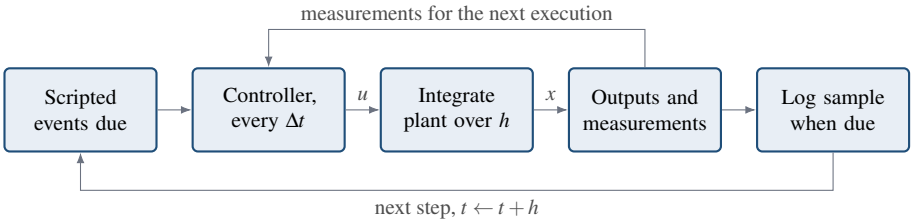


Figure 1.10 Work of the Engine in One Integration Step

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 (1.12) to (1.15) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.12}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.13}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.14}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.15}$$

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 (1.15) drives I toward the value that reproduces the output applied, which is back-calculation [36]. With no integral action, Equation (1.15) 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 $I_k = u_{\text{man}} - u_b - S_k$, 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 (1.16).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.16}$$

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 (1.17) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.17}$$

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 [37]. 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 [38]. 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 [39].

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.

Table 1.3 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Artificial pancreas	A pump that only adds, and a cost that is not symmetric.	introductory	3 h
3	Intravenous anesthesia	An exact loop on a model is not a correct loop on a patient.	intermediate	3 h
4	Mechanical ventilator	The measured breath and the delivered breath part company.	advanced	2.5 h
5	Kidney replacement	A loop that holds its measurement need not hold the patient.	intermediate	2 h
6	Runner on a treadmill	One actuator, four limits and a deadline.	intermediate	3.5 h
7	Road cyclist	An actuator whose authority is a store.	advanced	4 h
8	Antibody culture	A doubling appetite fed on a sample a day old.	advanced	3 h
9	Insulin refolding	A better fold can cost the campaign.	intermediate	2 h
10	Lager fermenter	Triggers and probes decide the batch more than tuning.	introductory	2 h
11	Anaerobic digester	More gas can mean a tank near its fold.	introductory	2.5 h
12	Recirculating fish farm	A biofilter is an actuator that has to be grown.	advanced	2.5 h
13	Glasshouse crop	A daily temperature budget under an hourly humidity limit.	advanced	2.5 h
14	City in an epidemic	Estimate before steering on counts a week old.	advanced	2 h

- *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 15.

1.8 Chapter Roadmap

Table 1.3 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.

Chapters 2, 10 and 11 form a first course: a one-way actuator, an on-off valve with a day-old result, and a fold watched through a daily titration. They rest on Sections 1.4.3, 1.4.5 and 1.4.6. Chapters 3, 5, 6 and 9 add a model inside the loop, redundancy, a selector and a schedule. The six advanced chapters combine estimation or prediction with constraints and are best left until Section 1.4.7 is

familiar. A reader at the bedside can read Chapters 2 to 5 in order, a reader of sport Chapters 6 and 7, and a reader of bioprocesses Chapters 8 to 12, ending with the glasshouse and the city.

KEY POINTS

- A living plant differs from its model, is measured sparsely and late, and is often moved by an actuator that cannot take back what it has given.
- A loop that inverts a population model is exact only on the typical individual, and feedback from a measurement of the effect removes the error.
- A bolus sized from a sample restores what the sample had, so a model between samples is corrected at the time each sample was drawn.
- When a pump can only add, the way down is the plant's own removal, and the margin from a hard limit must cover the spread of individuals and that slow way back.

Where else this applies. Selectors and overrides are the subject of Volume 4, estimation from late samples of Volume 5 and batch operation of Volume 8.

References

[1] R. J. Hatchett, C. E. Mecher, and M. Lipsitch, "Public health interventions and epidemic intensity during the 1918 influenza pandemic," *Proceedings of the National Academy of Sciences*, vol. 104, no. 18, pp. 7582–7587, 2007.

[2] C. Lambert and E. Delgado, "100 years since the discovery of insulin, from its discovery to the insulins of the future," *Biomedicines*, vol. 12, no. 3, p. 533, 2024.

[3] P. Drinker and L. A. Shaw, "An apparatus for the prolonged administration of artificial respiration: I. A design for adults and children," *Journal of Clinical Investigation*, vol. 7, no. 2, pp. 229–247, 1929.

[4] A. Chorley and K. L. Lamb, "The application of critical power, the work capacity above critical power (W'), and its reconstitution: A narrative review of current evidence and implications for cycling training prescription," *Sports*, vol. 8, no. 9, p. 123, 2020.

[5] I. Th. H. Gieling, "Commercial greenhouse computersystems: Inventarisation of implemented research, carried out by commercial firms," *Acta Horticulturae*, no. 106, pp. 59–66, 1980.

[6] G. Köhler and C. Milstein, "Continuous cultures of fused cells secreting antibody of predefined specificity," *Nature*, vol. 256, no. 5517, pp. 495–497, 1975.

[7] C. Zickefoose and R. B. J. Hayes, "Operations manual: Anaerobic sludge digestion," United States Environmental Protection Agency, Office of Water Program Operations, Washington, DC, Tech. Rep. EPA 430/9-76-001, 1976. [Online]. Available: https://permanent.fdlp.gov/gpo81586/00000IKT.pdf

[8] C. Cobelli, E. Renard, and B. Kovatchev, "Artificial pancreas: Past, present, future," *Diabetes*, vol. 60, no. 11, pp. 2672–2682, 2011.

[9] H. Schwilden, J. Schüttler, and H. Stoeckel, "Closed-loop feedback control of methohexital anesthesia by quantitative EEG analysis in humans," *Anesthesiology*, vol. 67, no. 3, pp. 341–347, 1987.

[10] P. U. Bidkar, A. Dey, P. Chatterjee, R. Ramadurai, and J. J. Joy, "Target-controlled infusion – past, present, and future," *Journal of Anaesthesiology Clinical Pharmacology*, vol. 40, no. 3, pp. 371–380, 2024.

[11] C. Ronco, "Continuous renal replacement therapy: Forty-year anniversary," *The International Journal of Artificial Organs*, vol. 40, no. 6, pp. 257–264, 2017.

[12] P. Kramer, W. Wigger, D. Matthaei, C. Langescheid, J. Rieger, C. Fuchs, K.-W. Rumpf, and F. Scheler, "Clinical experience with continuously monitored fluid balance in automatic hemofiltration," *Artificial Organs*, vol. 2, no. 2, pp. 147–149, 1978.

[13] L. B. Sheiner, D. R. Stanski, S. Vozeh, R. D. Miller, and J. Ham, "Simultaneous modeling of pharmacokinetics and pharmacodynamics: Application to d-tubocurarine," *Clinical Pharmacology & Therapeutics*, vol. 25, no. 3, pp. 358–371, 1979.

[14] S. L. Shafer and K. M. Gregg, "Algorithms to rapidly achieve and maintain stable drug concentrations at the site of drug effect with a computer-controlled infusion pump," *Journal of Pharmacokinetics and Biopharmaceutics*, vol. 20, no. 2, pp. 147–169, 1992.

[15] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity," *American Journal of Physiology-Endocrinology and Metabolism*, vol. 236, no. 6, p. E667, 1979.

[16] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A*, vol. 115, no. 772, pp. 700–721, 1927.

[17] J. Monod, "The growth of bacterial cultures," *Annual Review of Microbiology*, vol. 3, no. 1, pp. 371–394, 1949.

[18] J. F. Andrews, "A mathematical model for the continuous culture of microorganisms utilizing inhibitory substrates," *Biotechnology and Bioengineering*, vol. 10, no. 6, pp. 707–723, 1968.

[19] J. Wallinga and M. Lipsitch, "How generation intervals shape the relationship between growth rates and reproductive numbers," *Proceedings of the Royal Society B: Biological Sciences*, vol. 274, no. 1609, pp. 599–604, 2007.

[20] L. B. Sheiner, B. Rosenberg, and V. V. Marathe, "Estimation of population characteristics of pharmacokinetic parameters from routine clinical data," *Journal of Pharmacokinetics and Biopharmaceutics*, vol. 5, no. 5, pp. 445–479, 1977.

[21] A. R. Absalom, N. Sutcliffe, and G. N. Kenny, "Closed-loop control of anesthesia using bispectral index," *Anesthesiology*, vol. 96, no. 1, pp. 67–73, 2002.

[22] R. Luttmann, D. G. Bracewell, G. Cornelissen, K. V. Gernaey, J. Glassey, V. C. Hass, C. Kaiser, C. Preusse, G. Striedner, and C.-F. Mandenius, "Soft sensors in bioprocessing: A status report and recommendations," *Biotechnology Journal*, vol. 7, no. 8, pp. 1040–1048, 2012.

[23] D. Dochain, "State and parameter estimation in chemical and biochemical processes: A tutorial," *Journal of Process Control*, vol. 13, no. 8, pp. 801–818, 2003.

[24] E. L. Haseltine and J. B. Rawlings, "Critical evaluation of extended Kalman filtering and moving-horizon estimation," *Industrial & Engineering Chemistry Research*, vol. 44, no. 8, pp. 2451–2460, 2005.

[25] M. A. Hughes, P. S. A. Glass, and J. R. Jacobs, "Context-sensitive half-time in multicompartment pharmacokinetic models for intravenous anesthetic drugs," *Anesthesiology*, vol. 76, no. 3, pp. 334–341, 1992.

[26] B. W. Bequette, "A critical assessment of algorithms and challenges in the development of a closed-loop artificial pancreas," *Diabetes Technology & Therapeutics*, vol. 7, no. 1, pp. 28–47, 2005.

[27] C. Ellingsen, E. Dassau, H. Zisser, B. Grosman, M. W. Percival, L. Jovanovič, and F. J. Doyle, "Safety constraints in an artificial pancreatic β cell: An implementation of model predictive control with insulin on board," *Journal of Diabetes Science and Technology*, vol. 3, no. 3, pp. 536–544, 2009.

[28] T. Battelino, T. Danne, R. M. Bergenstal, S. A. Amiel, R. Beck *et al.*, "Clinical targets for continuous glucose monitoring data interpretation: Recommendations from the international consensus on time in range," *Diabetes Care*, vol. 42, no. 8, pp. 1593–1603, 2019.

[29] A. M. Jones, A. Vanhatalo, M. Burnley, R. H. Morton, and D. C. Poole, "Critical power: Implications for determination of $\dot{V}O_2$max and exercise tolerance," *Medicine & Science in Sports & Exercise*, vol. 42, no. 10, pp. 1876–1890, 2010.

[30] R. H. Morton, "A 3-parameter critical power model," *Ergonomics*, vol. 39, no. 4, pp. 611–619, 1996.

[31] G. Cappon, M. Vettoretti, G. Sparacino, and A. Facchinetti, "Continuous glucose monitoring sensors for diabetes management: A review of technologies and applications," *Diabetes & Metabolism Journal*, vol. 43, no. 4, pp. 383–397, 2019.

[32] G. Villa, M. Neri, R. Bellomo, J. Cerda, A. R. De Gaudio, S. De Rosa, F. Garzotto, P. M. Honore, J. Kellum, A. Lorenzin, D. Payen, Z. Ricci, S. Samoni, J.-L. Vincent, J. Wendon, M. Zaccaria, and C. Ronco, "Nomenclature for renal replacement therapy and blood purification techniques in critically ill patients: practical applications," *Critical Care*, vol. 20, p. 283, 2016.

[33] 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.

[34] U.S. Food and Drug Administration, "Infusion pump," Code of Federal Regulations, Title 21, Section 880.5725, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[35] 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, Tech. Rep. NIST Special Publication 800-82r3, 2023.

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

[37] 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.

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

[39] 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 · 4 min

A drug is infused into one well-mixed compartment of 20 L that is cleared at 1 L/min. The target concentration is 3 mg/L.

(a) Find the time constant, the maintenance rate and the loading dose. *Checkpoint:* 20 min, 3 mg/min, 60 mg.

(b) Without the loading dose, how long does the concentration take to reach 95 % of the target? *Checkpoint:* about 60 min.

1.2 Theory and calculation · 4 min

Patient P-1 of this chapter has $V_1 = 10\,\mathrm{L}$, $V_2 = 40\,\mathrm{L}$, $\mathrm{Cl}_1 = 0.5\,\mathrm{L/min}$ and $\mathrm{Cl}_2 = 0.4\,\mathrm{L/min}$. A pump gives the bolus that brings the plasma to 3 mg/L and then holds it there exactly.

(a) Find the infusion rate just after the bolus and 100 min later. *Checkpoint:* 2.70 mg/min and 1.94 mg/min.

(b) Find the rate the pump approaches after many hours, and say where the difference goes in the meantime. *Checkpoint:* 1.50 mg/min, the rest fills the peripheral compartment.

1.3 Concept check · 2 min

A target-controlled pump inverts a population model and holds the model’s plasma at 2 mg/L exactly. The patient’s volumes match the model, but the patient clears the drug 1.6 times as fast as the population. After four hours, what is true of the patient’s effect-site concentration?

(a) It equals 2 mg/L, because the pump’s loop is exact.

(b) It lies above 2 mg/L, because the pump gives extra drug.

(c) It lies below 2 mg/L and heads for about 1.25 mg/L.

(d) It oscillates about 2 mg/L at the effect-site rate.

1.4 Theory and calculation · 4 min

A drug’s window is 2 to 4 mg/L. A population model is used open loop, and patients’ clearances range from 0.8 to 1.25 times the population’s.

(a) A pump holds the model at 2.5 mg/L. Find the steady concentrations of the two extreme patients. *Checkpoint:* 3.125 mg/L and 2.0 mg/L.

(b) Find the range of model targets that keeps every patient in the window, and say whether open-loop dosing is possible. *Checkpoint:* 2.5 to 3.2 mg/L, yes.

1.5 Theory and calculation · 5 min

A culture doubles every 20 h and takes glucose in proportion to its growth. Over its first 20 h it takes 0.4 g/L. It is sampled once a day and fed a daily bolus, and its glucose must stay inside a band 5 g/L wide.

(a) By what factor does the culture grow in one day, and how much glucose does it take on day 0 and day 1 (the first and second days)? *Checkpoint:* 2.30, 0.52 g/L and 1.19 g/L.

(b) On which day does the daily uptake first exceed the band, so that no daily bolus can hold it? *Checkpoint:* day 3, with 6.29 g/L.

1.6 Theory and calculation · 5 min

Methanogens grow by a Haldane law with $\mu_{max} = 0.5\,\mathrm{d}^{-1}$, $K_S = 4\,\mathrm{mmol/L}$, $K_I = 100\,\mathrm{mmol/L}$ and a death rate of $0.05\,\mathrm{d}^{-1}$, in a stirred tank that keeps no biomass back.

(a) Find the acid level of the peak, the peak growth rate and the critical dilution rate. *Checkpoint:* 20 mmol/L, $0.357\,\mathrm{d}^{-1}$, $0.307\,\mathrm{d}^{-1}$.

(b) At a dilution rate of $0.2\,\mathrm{d}^{-1}$, find the two steady acid levels and say which one the tank should run at. *Checkpoint:* 4.17 mmol/L and 95.8 mmol/L, the lower.

1.7 Theory and calculation · 4 min

A cell count is drawn at 08:00 and reported at 14:00. The culture grows at $0.03\,\mathrm{h}^{-1}$.

(a) An estimator resets its cells to the count when the report arrives, as if the count had been drawn then. How far low is its estimate? *Checkpoint:* 16.5 %.

(b) An estimator instead updates a bias with weight $\lambda = 0.4$, compared at the hour of the draw. Against a constant model error, what fraction of the error has it removed after three counts? *Checkpoint:* 0.784.

1.8 Concept check · 2 min

A loop holds a weighed running total of fluid removed on its target. One scale begins to drift slowly. The loop's own error stays small. What reveals that the patient's true removal is falling behind?

(a) A second, independent measurement of the same flows, such as the pumps' own counts.

(b) A faster integral action on the balance loop.

(c) A derivative term on the weighed running total.

(d) A larger proportional gain so that the error stays smaller still.

1.9 Theory and calculation · 5 min

Two actuators push one way.

(a) A drug in one compartment with a time constant of 30 min stands at 5 mg/L, above a limit of 4 mg/L. The pump stops. How long until the concentration is back at the limit, and how long until it reaches 2 mg/L? *Checkpoint:* 6.7 min and 27.5 min.

(b) A rider with $W_0' = 20\,\mathrm{kJ}$, $P_{pk} = 1000\,\mathrm{W}$, $CP = 250\,\mathrm{W}$ and $E_c = 0.75$ holds an exertion of 0.9 from a full reserve. Find the fade time constant, the power at the start and the power a minute later. *Checkpoint:* 44.4 s, 700 W, 367 W.

1.10 Concept check · 2 min

Run A spends an area of 40 mg/dL h above its glucose band and none below it. Run B spends 5 mg/dL h below the band and none above. An objective weights the area below the band by w and the area above by one. For which weights does the objective rank A as the better run?

(a) For every w, because A never goes low.

(b) For w below 8 only.

(c) For no w, because B has the smaller total area.

(d) For w above 8.