INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 3

Measuring the Inner Loop Speed a Fermenter Oxygen Cascade Requires

Pages 46–74

Edited by John D. Hedengren

How to cite this chapter. "Measuring the Inner Loop Speed a Fermenter Oxygen Cascade Requires," 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. 3, pp. 46–74.

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

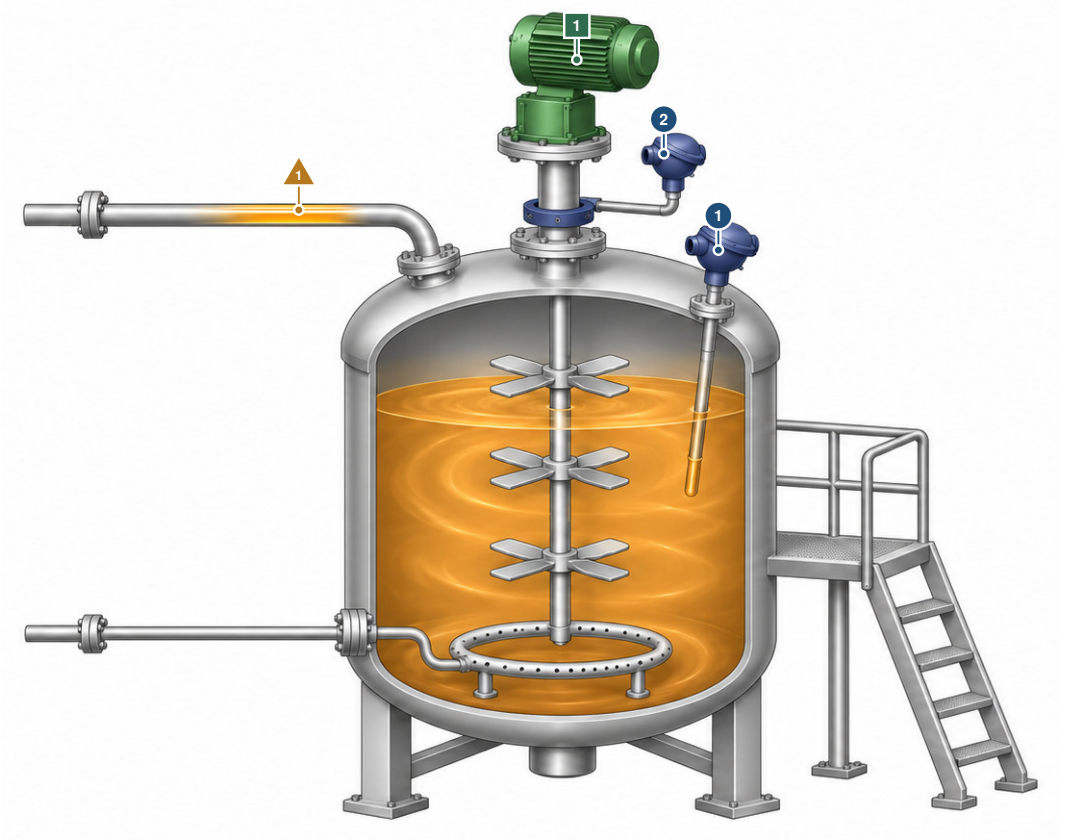
The app for this chapter
https://apmonitor.com/apps/bioreactor-do/

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

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

CHAPTER 3

Measuring the Inner Loop Speed a Fermenter Oxygen Cascade Requires



The controller moves 1 agitator motor and its drive, through the speed setpoint

The controller reads 1 dissolved oxygen probe in the broth 2 shaft speed sensor (tachometer)

What disturbs it 1 glucose feed rate

Notice the fast speed sensor on the shaft and the slow oxygen probe in the broth: the speed loop closes on the first, and no control structure can speed up the second. Illustration, not to scale.

IN PLAIN TERMS **An Inner Loop Must Be Fast to Help**

The main idea. Yeast in a big stirred tank uses oxygen all the time, and the liquid holds only a few seconds of it. One controller keeps the oxygen steady by changing how fast the stirrer turns, and an inner one makes the motor reach that speed.

The challenge. The motor, the liquid and the oxygen sensor all take several seconds to respond. A slow inner controller makes the pair worse than one controller alone, and nothing can speed up the sensor.

Why it matters. How fast the inner one must be is the main design question for any pair like this.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history of the oxygen probe (Section 3.2), the fact on oxygen-starved zones (Section 3.3.6), the fact on outer gain and the try-it segments of scenarios 6 and 8 (Section 3.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	30 min	the first page, the results (Section 3.6), Technology in Practice (Section 3.7), Limitations (Section 3.8) and the open challenges

The control question The agitator drive, the dissolved oxygen in the broth and the membrane probe of a fermenter respond within a factor of two of one another. How fast must the speed loop of an oxygen cascade be before it helps rather than hurts, and which of the three lags can any control structure reach?

Learning objectives

- Compute the oxygen inventory, the liquid-phase time constant and the drive lag of the fermenter from its model.
- Predict the closed-loop time constant of the speed loop from its gain, and state what that estimate leaves out.
- Measure the inner-loop speed at which the cascade stops paying, and compare it with the three-to-five rule.
- Explain why a feedforward from the glucose feed must wait for the culture.
- Tell a capacity limit or a biased probe from a tuning problem by reading the speed margin and the true dissolved oxygen.

Prerequisites Sections 1.4.5 (inner-loop speed), 1.4.3 (feedforward timing), 1.4.2 (entry points) and 1.6.2 (velocity form). Chapter 2 for the two structures on a simpler plant.

The app https://apmonitor.com/apps/bioreactor-do/

See also Volume 19, Chapter 8 for one oxygen loop that sequences agitation, enrichment and gas flow, and Volume 19, Chapter 12 for an oxygen cascade with a meal feedforward. Volume 2, Chapter 2 for the velocity form and back-calculation, Volume 2, Chapter 4 for the same three architectures on steam, and Volume 2, Chapter 3 for an actuator bound that no tuning removes.

3.1 Introduction

Every treatment of cascade control carries the same design rule. The inner loop must be substantially faster than the outer one, usually quoted as three to five times, or the arrangement is not worth building [1, 2]. The rule is rarely tested, because the standard teaching example is a flow loop inside a temperature loop. A flow loop settles in a second and a jacketed vessel settles in ten minutes, so the margin is six hundred to one, and detuning the inner controller by a factor of ten changes nothing visible on the outer trend. The rule is therefore delivered as an assertion rather than as a measurement. Section 1.4.5 states it and measures it once on a heater whose inner loop wraps a valve lag, and this chapter measures it on a plant where the margin is narrow.

Dissolved oxygen in a stirred fermenter is the counterexample. The liquid holds so little oxygen that its own time constant is 12.4 s at the design condition of the plant described here, the agitator drive changes speed with a time constant of 15.4 s, and the membrane probe reports what happened 20 s ago. Those three numbers are within a factor of two of each other, and the speed loop a cascade inserts closes in 2.9 s. The margin is eleven to one against the whole outer chain rather than six hundred to one, the inner tuning is an adjustable input, and the rule can be broken inside twenty simulated minutes. The consequence is a loss rather than a smaller benefit. A cascade whose inner loop settles in 12.3 s answers a feed step with an integrated error of 72.3, against 23.8 for the single loop it replaced.

This chapter targets the misconception that a cascade is structurally safe, so that adding an inner loop can improve performance but cannot degrade it. The measurement in Section 3.6 says otherwise, and it also identifies what the rule of thumb omits. Holding the inner proportional gain fixed and

shortening the inner integral time from 1.20 min to 0.10 min moves the break-even point from an inner time constant of 7.4 s out to 13.4 s, so the deciding quantity is the inner loop's whole closed-loop response and not the proportional gain alone.

The class of problem is any outer loop whose final control element has dynamics and a nonlinearity of its own, and whose own process is fast enough that the element cannot simply be ignored. Variable-speed pumps on flow or pressure, fans on temperature, compressors on suction pressure and smart positioners on valve stems are all the same structure, a fast measured inner loop wrapped around an actuator, so that the outer loop addresses an element that reaches its commanded value. The benefit of the arrangement was quantified against the ratio of primary to secondary dynamics decades ago, and both the size of the benefit and the region where it disappears depend on that ratio rather than on the tuning [2, 3]. Dissolved oxygen belongs to a second class as well, the aeration loops whose actuator power is the largest electricity consumer on the site, where the same measurement lag and the same capacity limits recur in wastewater treatment as in fermentation [4]. Volume 9, Chapter 10 treats such a plant, where a supervisory loop chooses the dissolved-oxygen setpoint itself.

The physical fundamentals underneath the experiments are compact. One oxygen mass balance carries an interphase transfer rate proportional to a concentration driving force against a metabolic uptake rate, with a stored quantity of a few seconds of demand between them, so the loop inherits a time-scale separation between a batch that runs for hours and an inventory that turns over in seconds [1, 5]. The experiments built on that balance cover cascade structure, measurement lag, actuator dynamics and nonlinearity, feedforward timing, anti-windup and capacity constraints, the principles a process dynamics and control course assembles around its cascade lecture [1, 6]. The behavior matters industrially because aeration power dominates the electricity bill of fermentation and wastewater sites alike, through loops that carry the same probe lags and capacity limits [4, 7].

Chapter 2 asks which structure answers which disturbance, using a tank level and two upsets that enter by different routes. This chapter takes the structure as given. The ordinary way to study what the inner loop must be is a pair of transfer functions with an assumed ratio between them, and that ratio is the quantity under study. The fermenter of this chapter has an oxygen holdup, a drive inertia, a probe lag and an inner loop tuning that are each adjustable on their own, so the speed an inner loop must reach becomes something measured on one plant instead of something assumed.

Section 3.2 traces how oxygen transfer in a stirred vessel became a measured quantity and how the loop that holds it was built. Section 3.3 sets out the nine-state fermenter with its transfer correlation, its drive torque balance and its instrument chain. Section 3.4 sorts the three lags of the outer loop by which control structure can reach each one. Section 3.5 gives the two controllers and the carbon-balance feedforward, three architectures that share one output unit. Section 3.6 carries the two central measurements, that a detuned cascade is three times worse than no cascade on the same disturbance, and that the crossover between harm and benefit moves by a factor of 1.8 when the inner integral time alone is changed. The same section runs the scenarios in which a capacity limit, an instrument bias and a derivative term each fail in a way the control error does not report. Section 3.7 sets the loop on real hardware. Sections 3.8 and 3.9 state what the model cannot support and what the error metrics leave out, and Section 3.10 collects the numbers. The correlation, the kinetics and the cascade structure are all standard.

3.2 Background

Aerobic fermentation at scale began without a dissolved-oxygen measurement of any kind. Submerged penicillin culture in the 1940s made the supply of oxygen the limiting design question of an entire industry, and the quantity of engineering interest was a transfer capacity rather than a concentration. The sulfite oxidation method measured how much oxygen an agitated vessel could deliver by letting a fast chemical reaction consume everything that crossed the interface, and correlating the result against impeller power and superficial gas velocity gave the first design equations for agitated gas-liquid contactors [8]. That work fixed the form the industry still uses, a volumetric coefficient k_La written as

a power law in gassed power per volume and in gas velocity, and it fixed the practice of sizing an agitator from a correlation rather than from a measurement on the vessel being built.

The correlation was refined for thirty years and then audited. A review of the measuring methods and the published results for nonviscous systems showed that the exponents cluster tightly while the leading coefficient does not, and separated coalescing from non-coalescing liquids as the largest single source of the spread [9]. The exponents used in this simulator, 0.70 on power per volume and 0.40 on gas velocity, come from that body of work, and so does the warning attached to them. The leading coefficient for a non-coalescing broth spans more than a factor of two across authors and vessels, so a scale-up that sizes a motor from a bench correlation is choosing the plant's oxygen capacity, its installed power and the loop's process gain in one calculation whose least trustworthy element is that coefficient [10].

Measurement of the dissolved oxygen itself came later, from medical instrumentation. The membrane-covered polarographic electrode reported in 1953 put a platinum cathode behind a permeable membrane so that the current was proportional to oxygen tension rather than to whatever else was in the solution [11]. Fermentation could not use it as built, because a fermenter is steam sterilized and a laboratory electrode is not, and the adaptation took another decade. Steam sterilizable probes with replaceable membranes and internal electrolyte reservoirs made continuous dissolved-oxygen measurement routine on production vessels [12]. The membrane that made the measurement possible is also what makes it slow, because the oxygen has to diffuse through it before the cathode sees any change, and a response time measured in tens of seconds has been a property of the instrument since its invention.

> ⧖ HISTORY **From Clark Cell to Optical Probe, Around 2005**
>
> Dissolved oxygen in bioreactors has traditionally been measured with Clark-type electrodes. Oxygen is reduced at a cathode with four electrons exchanged for each molecule, so the current is linear in the concentration [13]. A federal field manual for water-quality data used such amperometric sensors as its primary method until luminescence-based optical sensors became more common, around 2005 [14]. Its 2020 revision states that optical sensors have replaced the amperometric type and makes them the standard field method, while the older sensor remains acceptable [15]. Stirred fermenters changed more slowly. A review of 2013 found Clark-type electrodes still dominating in conventional larger-scale systems, the kind of probe this chapter models [13].

Closed-loop control followed once the measurement existed. Early practice fixed an airflow and an agitator speed from the batch recipe and accepted whatever the culture did to the reading, which is the manual operation scenarios 1 and 2 of this application reproduce. Automatic control through agitation speed was reported for a fermenter in the mid 1980s and is now the standard arrangement, because speed enters the transfer correlation through the power term with an exponent about twice that of the gas velocity term and is therefore the stronger actuator [16]. The two remaining levers are vessel back pressure, which raises the equilibrium concentration through Henry's law at negligible compressor duty, and airflow, which raises the gas velocity term while making foam. The industrial split that survives is agitation and pressure as the controlled actuators with airflow on a recipe schedule.

The fed-batch itself has a separate lineage that matters for the results in this chapter. Feeding a concentrated sugar to a substrate-limited culture rather than charging it at the start became standard once the respiratory capacity of yeast was understood as the quantity being protected, with overfeeding rather than oxygen starvation as the immediate failure mode [17]. In such a culture the feed pump sets the growth rate and dissolved oxygen limits nothing until it falls very low, which is why the biased probe of scenario 11 costs no product at all and still constitutes a serious fault. Product formation is written in the growth-associated and non-growth-associated form introduced for lactic acid fermentation, which closes the carbon balance rather than asserting a yield [18]. Fuel ethanol is the product because it is a commodity manufactured at very large scale with margins thin enough that the net energy balance of the process has been debated in the open literature [19], and because aeration

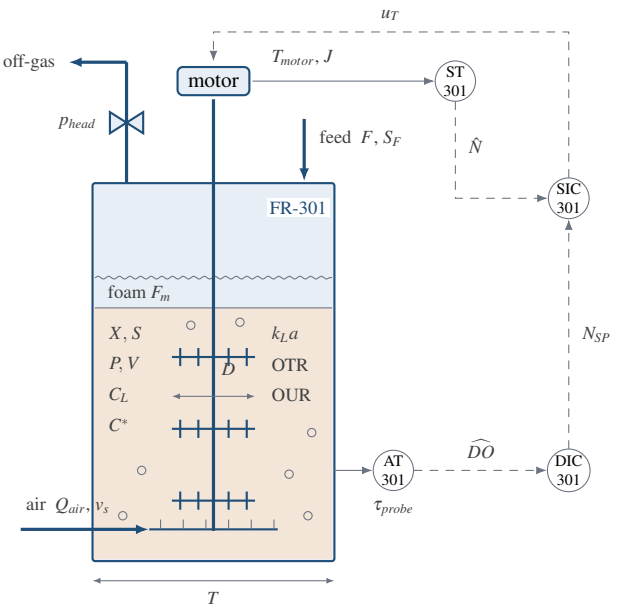


Figure 3.1 Fermenter FR-301 with its agitator drive, sparger, dissolved-oxygen probe and the two controllers of the cascade. Broth of volume V holds cells X, glucose S and ethanol P with dissolved oxygen C_L against an equilibrium concentration C^* set by the headspace pressure p_{head}. Three impellers of diameter D on a shaft of inertia J turn at speed N against motor torque T_{motor}, setting the transfer coefficient k_La and therefore the transfer rate OTR against the culture's uptake rate OUR. Glucose feed F at concentration S_F is the principal disturbance. AT-301 reports the dissolved oxygen $\widehat{DO}$ through a lag τ_{probe} to DIC-301, whose output is a speed demand N_{SP} for SIC-301, which reads tachometer ST-301 and writes the drive torque command u_T.

is by a wide margin the largest mechanical energy demand of a stirred production fermenter [7]. The agitator in this simulator draws 88 kW at the design condition and 281 kW at its speed limit on the same charge, and no control specification anywhere in the application mentions either number.

3.3 Process Model

Vessel FR-301 is a 4.0 m diameter stirred fermenter of 105 m^3 maximum working volume, fitted with three six-blade Rushton turbines of 1.40 m diameter on one shaft and a ring sparger below the bottom impeller. The design charge, from which nine of the thirteen scenarios start, is a fed-batch about half way through, 65 m^3 of broth carrying 15 g/L of cells and 40 g/L of ethanol at 32 °C, with a 550 g/L glucose syrup metered in at 24 L/min, air at 0.40 vvm and a headspace held at 0.5 barg. Dissolved oxygen is measured by a membrane probe and reported as a percentage of air saturation at one atmosphere, which is what a probe calibrated in air before the batch actually reads. Every symbol used below appears on Figure 3.1.

Table 3.1 separates what the loop holds from what it moves and from what moves underneath it. Time is in minutes throughout, because one time base must serve both a twelve-second oxygen transient and a fifteen-hour batch. Concentrations are g/L for the culture and mg/L for dissolved oxygen, volumes are m^3, agitator speed is integrated in rev/s and displayed in rev/min, torque is N m and pressure is Pa. Four annunciations bound normal operation, dissolved oxygen below 15 %, foam filling the headspace, the agitator at its speed limit, and the broth reaching the 105 m^3 working volume, after which the vessel overflows at the feed rate and the dilution terms carry cells and ethanol

Table 3.1 Variables of the Dissolved-Oxygen Loop. The Nominal Column Is the Design Charge That Opens Nine of the Thirteen Scenarios.

Symbol	Role	Nominal	Range
DO (AT-301)	CV, measured	25.0 %	5–90 % SP
N_{SP}	MV	64.1 rev/min	21–93 rev/min
$\hat{N}$ (ST-301)	inner CV	64.1 rev/min	0–93 rev/min
u_T	inner MV	43.5 %	0–100 %
F feed	DV, measured	24 L/min	0–90 L/min
Q_{air}	operator set	0.40 vvm	0.05–1.2 vvm
p_{head}	operator set	0.5 barg	0–1.5 barg
probe bias	DV, unmeasured	0 %	±20 %
probe noise	DV, unmeasured	0 %	0–4 % 1σ
X	state	15 g/L	≥ 0
S	state	0.047 g/L	≥ 0
P	state	40 g/L	≥ 0
V	state	65 m^3	≤ 105 m^3
C_L	state	1.65 mg/L	0–C^*
N	state	1.068 rev/s	0–1.55 rev/s
C_m	state	1.65 mg/L	≥ 0
F_m foam	state	0.22	0–1
A antifoam	state	0	≥ 0

out with the overflow. None of the four stops the run. Each latches and is recorded, so an excursion outside normal operation leaves a record. Nine states are integrated.

3.3.1 Oxygen Balance

The liquid oxygen balance is Equation (3.1), and every other equation in the model exists to supply one of its three terms.

$$\frac{\mathrm{d}C_L}{\mathrm{d}t} = k_L a\,(C^* - C_L) \;-\; \mathrm{OUR} \;-\; \frac{F}{V}\,C_L \tag{3.1}$$

where C_L is the dissolved oxygen concentration in mg/L, $k_L a$ is the volumetric transfer coefficient in min^{-1}, C^* is the equilibrium concentration at the working pressure in mg/L, OUR is the oxygen uptake rate in mg/(L·min), F is the glucose feed in m^3/min and V is the broth volume in m^3. Every term carries units of mg/(L·min). The dilution term is three orders of magnitude below the other two at every condition the application reaches and is kept only so that the balance closes exactly.

Equation (3.1) has the form of a first-order lag, although the size of its two large terms relative to the stored quantity produces different behavior. At the design charge the 25 % setpoint is 1.65 mg/L, which over 65 m^3 is 107 g of dissolved oxygen in the whole vessel. The culture consumes 39.6 mg/(L·min), which over the same volume is 2.57 kg every minute. The entire dissolved inventory is therefore consumed and replaced every 2.5 s. A tank holds minutes of throughput and a jacketed reactor holds minutes of thermal energy, while this vessel holds two and a half seconds of oxygen, so a one percent mismatch between the two rates moves the dissolved oxygen by one percent of its value every two and a half seconds. Dissolved oxygen therefore collapses rather than drifts when control is lost.

3.3.2 Oxygen Transfer and the Agitator

The transfer coefficient is written in the literature power-law form rather than directly in agitator speed, as shown in Equation (3.2).

$$k_L a = a\left(\frac{P_g}{V}\right)^{\alpha} v_s^{\beta}, \qquad \frac{P_g}{V} = \frac{2\pi f_g k_T N^3}{V} \tag{3.2}$$

where P_g is gassed shaft power in W, v_s is the superficial gas velocity in m/s, a, α and β are the correlation constants, f_g is the gassed-to-ungassed power ratio and $k_T = N_P \rho D^5 n_{imp}/2\pi$ is the ungassed load-torque coefficient in N m per $(\text{rev/s})^2$, built from the impeller power number N_P, the broth density ρ, the impeller diameter D and the number of impellers n_{imp}. Writing the correlation this way rather than as a fit in speed is deliberate, because airflow, working volume and gassing then act through the same equation and a change in any of them moves the loop gain with no control setting having been touched. Since gassed power per volume goes as N^3 at a fixed aeration condition, Equation (3.2) gives $k_L a \propto N^{3\alpha} = N^{2.1}$, and 2.1 is the process gain exponent of the outer loop. At the design point 64.1 rev/min delivers 1350 W/m^3 at a superficial gas velocity of 0.0345 m/s for a $k_L a$ of 4.84 min^{-1}, which by Equation (3.1) is a liquid-phase time constant $1/k_L a$ of 12.4 s.

The superficial gas velocity and the gassed power ratio are given by Equation (3.3).

$$v_s = \frac{Q_{air} V}{60 A_t}, \qquad f_g = \frac{1}{1 + k_g \text{Fl}}, \qquad \text{Fl} = \frac{Q_{air} V}{60 N D^3} \tag{3.3}$$

where Q_{air} is the airflow in vessel volumes per minute, $A_t = \pi T^2/4$ is the vessel cross-section in m^2, Fl is the aeration number and k_g is the gassing coefficient. Sparged gas forms cavities behind the blades and unloads the impellers, so the same speed costs less torque when the air is on. The reciprocal form is used rather than the customary linear fit in Fl because a linear fit returns a negative power ratio at the aeration numbers a 65 m^3 fermenter actually runs at. At the design condition $\text{Fl} = 0.148$ and $f_g = 0.849$.

Two multipliers ride on top of Equation (3.2). Antifoam divides the coefficient by $1 + k_{AF} A$, and a foam blanket multiplies it by $1 - 0.4\,\text{clamp}((F_m - 0.8)/0.4, 0, 1)$, which removes up to 40 % once the headspace is more than four fifths full and is ramped in over the last fifth so that the right-hand side stays smooth for the integrator. A fouled sparger multiplies it by a further 0.75.

The equilibrium concentration and the reading are Equation (3.4).

$$C^* = C_{ref}(T_b)\,\frac{p_{atm} + p_{head}}{p_{atm}}, \qquad DO\,\% = \frac{100\,C_L}{C_{ref}(T_b)} \tag{3.4}$$

where C_{ref} is the oxygen solubility of the broth under one atmosphere of air, taken as a cubic fit in broth temperature for clean water multiplied by a salting-out factor of 0.92. At 32 °C that gives $C_{ref} = 6.59$ mg/L. The reading of Equation (3.4) divides by C_{ref} and not by C^*, because the probe was calibrated in air at one atmosphere before the batch and nothing afterward changes the calibration. That convention is why back pressure raises the driving force at low compressor cost and why the reading can exceed 100 %. Taking the head from 0.5 barg to 1.5 barg moves C^* from 9.84 mg/L to 16.34 mg/L and multiplies the driving force in Equation (3.1) by about 1.8.

3.3.3 Culture Kinetics and Fed-Batch Balances

Growth is Monod on glucose with non-competitive inhibition by the ethanol the culture is making and a Monod term in dissolved oxygen, as shown in Equation (3.5).

$$\mu = \mu_{max}\,\frac{S}{K_S + S}\,\frac{K_P}{K_P + P}\,\frac{C_L}{K_O + C_L} \tag{3.5}$$

where μ is the specific growth rate in min^{-1}, K_S is the glucose half-saturation, K_P is the ethanol inhibition constant and K_O is the oxygen half-saturation. The oxygen term couples the dissolved-oxygen loop to the batch outcome, and its half-saturation of 0.35 mg/L is 5.3 % of air saturation, well below the 15 % the plant alarms at. The rates that follow from μ are Equations (3.6) and (3.7).

$$\text{OUR} = 1000 X \left(\frac{\mu}{Y_{X/O}} + m_O \right) \tag{3.6}$$

$$q_P = \alpha_{LP}\mu + \beta_{LP}, \qquad q_S = \frac{\mu}{Y_{X/S}} + m_S + \frac{q_P}{Y_{P/S}} \tag{3.7}$$

where X is the biomass in g/L, $Y_{X/O}$ is the cell yield on oxygen, m_O is the oxygen maintenance coefficient, q_P and q_S are the specific product and substrate rates in g/g/min, α_{LP} and β_{LP} are the growth-associated and non-growth-associated product coefficients [18], $Y_{X/S}$ and $Y_{P/S}$ are the cell and product yields on glucose and m_S is the glucose maintenance coefficient. The factor of 1000 in (3.6) carries g/L of cells into mg/L of oxygen. The maintenance term is the part of the demand that does not stop when growth does, which is why a stalled culture still needs the agitator. Writing q_S with three separate sinks in (3.7) closes the carbon balance, so substrate the culture cannot consume appears as rising residual glucose rather than disappearing from the balance.

The four culture states are Equation (3.8).

$$\begin{aligned} \frac{\mathrm{d}X}{\mathrm{d}t} &= \mu X - \frac{F}{V}X, & \frac{\mathrm{d}S}{\mathrm{d}t} &= \frac{F}{V}(S_F - S) - q_S X \\ \frac{\mathrm{d}P}{\mathrm{d}t} &= q_P X - \frac{F}{V}P, & \frac{\mathrm{d}V}{\mathrm{d}t} &= F \end{aligned} \tag{3.8}$$

where S_F is the glucose concentration in the feed. The disturbance this loop faces for the whole batch is in the first of those four. Oxygen demand is proportional to cell mass through Equation (3.6) and cell mass only ever grows, so the load on the loop is a ramp that runs from inoculum to harvest, 18.3 mg/(L·min) at 5 g/L of cells, 39.6 at 15 and 65.8 at 34.

3.3.4 Agitator Drive, Probe and Foam

The shaft speed follows from a torque balance rather than from a command that is directly obtained, as shown in Equation (3.9).

$$2\pi J \frac{\mathrm{d}N}{\mathrm{d}t} = T_{motor} - f_g k_T N^2, \qquad \tau_{drive} = \frac{2\pi J}{2 f_g k_T N} \tag{3.9}$$

where J is the effective drive inertia in kg m^2, lumping the rotating mass, the entrained broth and the drive's own current limit, and N is in rev/s. Linearizing the quadratic load about the running speed gives the drive time constant on the right, which is inversely proportional to speed. The drive is slowest when it is turning slowly, 23.5 s at the 43 rev/min of an early batch, 15.4 s at the design 64 rev/min and 11.6 s at the 86 rev/min of a late one. A fixed-gain outer controller cannot track that variation, while the speed loop removes it. Setting the derivative in (3.9) to zero whenever the speed rests on a stop with the torque imbalance pushing further outward keeps the fixed-step integrator from crossing a limit and coming back, and a clip applied after each completed step catches anything left over.

The probe is a first-order lag on the true concentration and is a state of the model rather than a filter in the controller, as shown in Equation (3.10).

$$\tau_{probe} \frac{\mathrm{d}C_m}{\mathrm{d}t} = C_L - C_m, \qquad \widehat{DO} = \frac{100\,C_m}{C_{ref}} + b + n(t) \tag{3.10}$$

where C_m is the probe state in mg/L, b is a calibration bias in percentage points and n is electrical noise added after the membrane. The noise is a Gaussian sequence taken through a first-order filter of 0.05 min (3 s) and scaled so that the standard deviation of the stationary sequence equals the set noise level. The tachometer is a separate first-order filter of time constant τ_N on the true shaft speed, with no bias and no noise.

Foam is one lumped height as a fraction of the available headspace, generated by gas velocity, power input and the protein a growing culture releases, as shown in Equation (3.11).

$$\begin{aligned} \frac{\mathrm{d}F_m}{\mathrm{d}t} &= G_F - \frac{F_m}{\tau_F} \\ G_F &= \frac{k_F}{1 + k_{AF,f} A} \left(\frac{v_s}{v_{s,ref}} \right)^{1.5} \frac{X}{K_{fx} + X} \sqrt{\frac{P_g/V}{(P/V)_{ref}}} \end{aligned} \tag{3.11}$$

where k_F is the generation rate constant, $v_{s,ref}$ and $(P/V)_{ref}$ are reference conditions, K_{fx} is the biomass half-saturation for foaming, $k_{AF,f}$ is the antifoam suppression strength and τ_F is the collapse time. Antifoam is a first-order washout, $\mathrm{d}A/\mathrm{d}t = -A/\tau_{AF}$, dosed as a unit step on the state, and each dose also knocks the foam down to at most 0.15 at once. The foam state is clamped to the interval $[0,1]$ after every step, and reaching 1 annunciates a foam-out. Equation (3.11) gives the airflow actuator a consequence that its own range does not show, which is the subject of scenario 10.

3.3.5 Constants, Initialization and Numerics

Table 3.2 gives every constant the model uses, and all of them are displayed and adjustable. The task here is not to recover a hidden constant but to compare four time constants the application reports, so full disclosure of the constants is part of the design.

Every scenario opens at a steady state of the oxygen subsystem. The initializer solves a coupled fixed point, because the feed sets the growth rate and the growth rate sets the oxygen demand while the oxygen in turn limits the growth. The residual glucose of a substrate-limited fed-batch is the root of $(F/V)(S_F - S) = q_S(\mu(S))X$, which is unique because the supply falls and the demand rises with S, and the dissolved oxygen is the root of Equation (3.1) set to zero, which is unique because transfer falls and uptake rises with C_L. Both are found by bisection to machine precision in sixty steps. Six passes alternate between them, after which the transfer correlation of Equation (3.2) is inverted for the agitator speed that balances the demand at the dissolved-oxygen setpoint, and that speed is clamped to $[N_{min}, N_{max}]$. The final dissolved oxygen is then re-solved at whichever speed survived the clamp, so a scenario whose agitator has run out opens at the resulting lower value rather than at the setpoint. The correlation inversion is itself a fixed point, since the gassed power ratio depends on the speed being solved for, and it converges in three or four passes because f_g is a mild function of speed. The initial torque is set to the load torque at the initial speed, so the shaft starts balanced. The first verification test the application runs on itself checks that the derivatives of C_L, N, C_m, F_m and A are all below 10^{-9} at the opening state.

States advance by classical fourth-order Runge-Kutta on a fixed step of 0.002 min (0.12 s), which places about a hundred steps across the 12.4 s oxygen lag. DIC-301 and SIC-301 share an execution period declared as 0.006 min. The engine runs a controller at the first step at or after each period, so both execute every third step, 0.006 min (0.36 s). The historian is declared at $1/0.018$ samples per simulated minute and for the same reason records every ninth step, 0.018 min (1.08 s). Noise is drawn from a seeded generator with seed 20260829, so scenario 12 repeats to the last digit. The engine and its three clocks are those of Section 1.6.1.

3.3.6 Assumptions

The model is detailed on the oxygen path and simplified elsewhere, and each simplification has a domain where it fails.

The broth is one perfectly mixed lump. A 100 m^3 vessel with three impellers carries a real dissolved-oxygen gradient from the sparger to the surface, and cells circulating through such a vessel experience substrate and oxygen fluctuations on a scale of seconds [20]. The probe sits at one height in that gradient. This is the largest single omission in the application, and it is one reason a plant runs a higher setpoint than the kinetics alone would justify.

> **DID YOU KNOW? Oxygen-Starved Zones at 22 Cubic Meters**
>
> Fed-batch cultures of a bacterium at 22 m^3 accumulated formate, a sign of zones short of oxygen. The study, published in 2001, suggests that moving zones of oxygen limitation explain part of the lower biomass yield at large scale [20]. One probe at one height sees none of them.

The gas phase carries no state. Oxygen leaves the sparger at its inlet composition and Equation (3.4) applies that composition everywhere, so the driving force is overstated near the top of a tall

Table 3.2 Constants of the Fermenter Model in the Units the Application Reports. Every Constant Is Displayed and Adjustable.

Symbol	Meaning	Value	Unit
T	vessel inside diameter	4.0	m
T_b	broth temperature	32	°C
—	broth solubility factor	0.92	—
V_{full}	maximum working volume	105	m^3
N_P	impeller power number	5.0	—
D	impeller diameter	1.40	m
n_{imp}	number of impellers	3	—
ρ	broth density	1050	kg/m^3
N_{min}	minimum agitator speed	21	rev/min
N_{max}	maximum agitator speed	93	rev/min
T_{rated}	rated motor torque	30	kN m
J	effective drive inertia	60000	$kg\,m^2$
$f_{g,nom}$	drive schedule calibration	0.78	—
k_g	gassed-power coefficient	1.2	—
a	transfer coefficient	2.0×10^{-3}	SI, s^{-1} basis
α	power exponent	0.70	—
β	gas velocity exponent	0.40	—
μ_{max}	maximum growth rate	0.55	h^{-1}
K_S	glucose half-saturation	0.05	g/L
K_P	ethanol inhibition	75	g/L
K_O	oxygen half-saturation	0.35	mg/L
$Y_{X/S}$	cell yield on glucose	0.45	g/g
$Y_{X/O}$	cell yield on oxygen	1.00	g/g
$Y_{P/S}$	ethanol yield on glucose	0.48	g/g
m_S	glucose maintenance	0.030	g/g/h
m_O	oxygen maintenance	0.015	g/g/h
α_{LP}	growth-associated ethanol	1.2	g/g
β_{LP}	non-growth ethanol	0.05	g/g/h
S_F	glucose in the feed	550	g/L
k_F	foam generation rate	0.024	min^{-1}
τ_F	foam collapse time	20	min
K_{fx}	foaming half-saturation	15	g/L
$v_{s,ref}$	reference gas velocity	0.035	m/s
$(P/V)_{ref}$	reference power input	1500	W/m^3
τ_{AF}	antifoam residence time	60	min
k_{AF}	antifoam k_La penalty	0.30	—
$k_{AF,f}$	antifoam suppression	4.0	—
τ_{probe}	probe time constant	20	s
τ_N	tachometer filter	0.5	s
—	critical dissolved oxygen	15	%
—	high foam alarm	0.75	—
X_0	biomass at start	15	g/L
P_0	ethanol at start	40	g/L
V_0	broth volume at start	65	m^3

vessel where the gas has already been stripped. A real plant closes the same balance on an off-gas analyzer, which is also the measurement a better feedforward than the one implemented here would use.

The transfer correlation has fixed exponents and no upper limit. Equation (3.2) says nothing about broth viscosity, which climbs through a fermentation and can halve the coefficient by the end, and nothing about impeller flooding, where the gas rate outruns the impeller's ability to disperse it and

k_La collapses instead of rising [9, 10]. Both set real limits that the model draws as a smooth curve.

The drive is a first-order torque balance. The 15 s response of Equation (3.9) is right for a machine of this size, although a real variable-frequency drive reaches speed under a ramp limit rather than an exponential, so a large speed change takes longer than the model says and a small one arrives sooner. The cascade argument survives that change while the exact crossover in Section 3.6 would move.

The kinetics are unstructured with constant yields. Real Crabtree-positive yeast switches metabolism with both glucose and oxygen [17], so neither the ethanol yield nor the oxygen demand per gram of sugar is truly constant. The values in Table 3.2 are a fair average over the middle of a batch and are not a prediction of a lag phase, a diauxic shift or a stationary phase. The oxygen demand modeled here is that of an aerated propagation-stage vessel rather than that of the anaerobic main fermenter of a fuel ethanol plant.

Broth temperature is a parameter. There is no temperature loop, although aerobic metabolism releases roughly 460 kJ of heat per mole of oxygen consumed, about 14 MJ per kilogram, so oxygen transfer and cooling duty are the same constraint seen twice and on many plants the cooling runs out first.

The probe is a pure lag with no dead time. A real membrane probe responds more like two lags in series with a small transport delay, and a loop tolerates a lag better than a delay of the same size, so the probe penalties measured in Section 3.6 are optimistic.

3.4 Dynamic Analysis

3.4.1 Lag Structure of the Outer Loop

Between the dissolved-oxygen controller's output and its input sit three first-order lags, and at the design condition they are within a factor of two of one another. The drive of Equation (3.9) contributes 15.4 s, computed from $2\pi J/(2f_g k_T N)$ with $k_T = 13\,481\,\mathrm{N\,m}$ per $(\mathrm{rev/s})^2$ and $f_g = 0.849$ at $N = 1.068\,\mathrm{rev/s}$. The liquid phase of Equation (3.1) contributes $1/k_La = 12.4\,\mathrm{s}$. The probe of Equation (3.10) contributes 20 s, and 60 s once the membrane has been in the vessel for six batches. The disturbance is not instant either, because a feed change reaches the pump immediately and reaches the cells through a glucose pool that takes about a minute to fill. On the cascade run of Section 3.6 the oxygen demand reaches 63 % of its change 1.3 min after the feed moves and 90 % of it at 4.3 min.

The three lags are not equally reachable by a control structure, and the comparison in Section 3.6 follows from that inequality. The drive lies between the controller output and the shaft, so a loop closed on a shaft measurement encloses it. The oxygen holdup lies between the shaft and the true dissolved oxygen, so no loop encloses it, because it is the process itself. The probe lies on the far side of everything, so no rearrangement of controllers touches it. A cascade reaches the actuator, a feedforward reaches the disturbance, and nothing reaches the measurement.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.4.2 Closed-Loop Response of the Inner Loop

The inner loop closes torque onto speed. Linearizing the load torque in (3.9) gives an inner process gain from percent of rated torque to speed, and proportional feedback divides the drive lag by the loop gain, as shown in Equation (3.12).

$$K_{p,N} = \frac{60\,T_{rated}/100}{2f_g k_T N}, \qquad \tau_{inner} = \frac{\tau_{drive}}{1 + K_{c,N}K_{p,N}} \tag{3.12}$$

where $K_{p,N}$ is in rev/min per percent of rated torque and $K_{c,N}$ is the inner proportional gain. The second expression is the proportional-only form of Equation (1.7), with the drive lag in place of the valve lag and $K_{p,N}$ in place of the valve gain. At the design condition $K_{p,N} = 0.736\,\mathrm{rev/min}$ per percent, so the shipped inner gain of 6 gives $\tau_{inner} = 2.85\,\mathrm{s}$ against a drive lag of 15.4 s and an

oxygen holdup of 12.4 s. Equation (3.12) is the number the application reports and it is deliberately incomplete, because it counts the inner proportional gain and not the inner integral time. Section 3.6 measures what that omission costs, and the answer is a factor of 1.8 on the very threshold the equation exists to locate.

The inner loop removes three separate effects. It removes the 15.4 s of drive lag. It removes the nonlinearity, because torque and speed are related by a square root in Equation (3.9), so the open-loop drive gain at 90 rev/min is only 39 % of what it is at 40 rev/min at the design volume and airflow (the square root alone gives 44 %, and the gassed power ratio rising with speed gives the rest), and a fixed outer gain has to survive both. And it removes every disturbance that enters at the shaft, which includes gassing, broth viscosity, a fouling impeller and a dragging bearing.

3.4.3 Steady-State Error of the Open-Loop Drive

The single-loop architecture is not a cascade with the inner controller removed but a drive running an open-loop torque schedule, as shown in Equation (3.13).

$$T_{motor} = f_{g,nom}\, k_T\, N_{dem}^2 \tag{3.13}$$

where $f_{g,nom}$ is the gassed power ratio the drive was commissioned at. Equation (3.13) represents a variable-frequency drive running without an encoder, commissioned at one gassing condition and retaining that calibration indefinitely. Setting (3.13) equal to the true load torque $f_g k_T N^2$ gives the steady-state error directly, since the shaft settles at $N = N_{dem}\sqrt{f_{g,nom}/f_g}$ with f_g evaluated at the settled speed. At the design airflow the commissioned 0.78 stands against a gassed power ratio of 0.855 at the settled speed, and a demand of 70.0 rev/min delivers 66.88 rev/min, a ratio of 0.955. Raising the airflow from 0.40 to 0.95 vvm at the same demand unloads the impellers further and the shaft runs up to 72.45 rev/min, an 8.3 % move in the manipulated variable caused by a disturbance the dissolved-oxygen controller was never told about. Both tests return 70.00 rev/min in cascade.

3.5 Controller Design

3.5.1 Velocity Form PID Algorithm

Both controllers are the velocity-form algorithm of Section 1.6.2, with the derivative on the measurement. The outer dissolved-oxygen controller DIC-301 is Equations (3.14) and (3.15).

$$\begin{aligned} e_k &= DO_{\mathrm{SP},k} - \widehat{DO}_k \\ \Delta N_k^P &= K_c\,(e_k - e_{k-1}) \\ \Delta N_k^I &= \frac{K_c}{\tau_I}\, e_k\, \Delta t_c \end{aligned} \tag{3.14}$$

$$\Delta N_k^D = -\,K_c \tau_D\, \frac{\widehat{DO}_k^f - 2\widehat{DO}_{k-1}^f + \widehat{DO}_{k-2}^f}{\Delta t_c} \tag{3.15}$$

where e_k is the error in percentage points of dissolved oxygen, K_c is the outer proportional gain in rev/min per percent, and τ_I and τ_D are the integral and derivative times in minutes. Both loops act direct, because more agitator speed and more motor torque each raise their own controlled variable. The derivative term reads a second difference of a measurement that has already been filtered, with $\widehat{DO}_k^f = \gamma\widehat{DO}_{k-1}^f + (1-\gamma)\widehat{DO}_k$ and $\gamma = \exp(-\Delta t_c/\tau_f)$, the filter constant τ_f defaulting to 0.05 min. The three increments are combined and the demand is clipped to the speed limits, as shown in

Equation (3.16).

$$\begin{aligned} s_k &= s_{k-1} + \Delta N_k^P + \Delta N_k^D \\ c_k &= N_{ff,k} + s_k + I_{k-1} \\ N_{SP,k} &= \min(N_{max},\ \max(N_{min},\ c_k)) \\ I_k &= I_{k-1} + \Delta N_k^I + \frac{\Delta t_c}{\tau_t}\left(N_{SP,k} - c_k\right) \end{aligned} \tag{3.16}$$

where Δt_c is the realized 0.006 min controller execution period, $N_{ff,k}$ is the feedforward speed of Section 3.5.3, $[N_{min}, N_{max}] = [21, 93]$ rev/min and $\tau_t = \tau_I$ is the tracking time of the back-calculation. In the notation of Equations (1.10) to (1.12), s_k is the running sum S_k, c_k is the output before limiting v_k, $N_{ff,k}$ is the bias u_b and τ_t is T_t. The back-calculation corrects the integral alone and never truncates the running sum [6]. At the shipped $\tau_I = 0.6$ min each execution removes $1/100$ of the clipped excess, so an integral wound against a speed limit unwinds over about τ_I rather than in one execution. Nothing rate-limits the output, and its only bound is that clamp. Manual mode and bumpless transfer work as Section 1.6.2 describes. The manual tracking also runs on whichever controller is idle, which is what makes a change of architecture mid-run bumpless as well.

The inner speed controller SIC-301 is proportional-integral with no derivative term, as shown in Equation (3.17).

$$\begin{aligned} e_{N,k} &= N_{SP,k} - \hat{N}_k, \qquad \Delta u_{N,k}^I = \frac{K_{c,N}}{\tau_{I,N}}\, e_{N,k}\, \Delta t_c \\ s_{N,k} &= s_{N,k-1} + K_{c,N}\left(e_{N,k} - e_{N,k-1}\right) \\ c_{N,k} &= s_{N,k} + I_{N,k-1} \\ u_{T,k} &= \min\bigl(100,\ \max\bigl(0,\ c_{N,k}\bigr)\bigr) \\ I_{N,k} &= I_{N,k-1} + \Delta u_{N,k}^I + \frac{\Delta t_c}{\tau_{I,N}}\left(u_{T,k} - c_{N,k}\right) \end{aligned} \tag{3.17}$$

with $T_{motor} = (u_T/100)\, T_{rated}$, where $\hat{N}$ is the filtered tachometer reading and u_T is clamped to 0–100 % of rated torque with the same back-calculation of (3.16), here with the tracking time $\tau_{I,N}$. Both loops execute on the same 0.006 min period. Tuning is bounded to $0 \le K_c \le 8$ rev/min per percent, $0.1 \text{ min} \le \tau_I \le 20$ min, $0 \le \tau_D \le 1$ min, $0 \le K_{c,N} \le 25$ percent of torque per rev/min and $0.01 \text{ min} \le \tau_{I,N} \le 3$ min.

3.5.2 A Common Output Unit Across Architectures

The outer controller's output means an agitator speed demand in rev/min in all three architectures, which is what makes the comparison in Section 3.6 fair. In the single loop that demand goes to Equation (3.13), in the cascade it is the setpoint of Equation (3.17), and in the feedforward architecture the same cascade is retained while N_{ff} in Equation (3.16) becomes nonzero. One set of tuning constants therefore carries across the whole ladder. The choice is not free, because in many cascades the outer output changes units when the structure changes and the gains have to change with it. The steam-heated exchanger of Volume 2, Chapter 4 offers the same ladder, a single loop, a cascade to steam flow and a heat-balance feedforward with a feedback trim.

🕮 **Pause.** Try Exercise 3.1 before reading on.

3.5.3 Feedforward From a Carbon Balance

The feedforward is calculated rather than fitted. At quasi steady state every gram of glucose the pump delivers is consumed, the share of it that becomes new cells follows from Equation (3.7) with the maintenance terms dropped, and that new cell mass has to be paid for in oxygen at $Y_{X/O}$, as shown in

Equation (3.18).

$$\mathrm{OUR}_{ff} = \frac{1000\,K_{ff}}{Y_{X/O}} \cdot \frac{F S_F / V}{1/Y_{X/S} + \alpha_{LP}/Y_{P/S}}$$
$$k_L a_{req} = \frac{\mathrm{OUR}_{ff}}{C^* - C_{L,\mathrm{SP}}} \tag{3.18}$$

where K_{ff} is an operator confidence factor defaulting to 1 and $C_{L,\mathrm{SP}}$ is the setpoint expressed as a concentration. Equation (3.2) is then inverted for the speed that delivers $k_L a_{req}$, and that speed is passed through a first-order lag of time constant τ_{ff} and enters (3.16) as N_{ff}. Entering as the controller bias rather than as a summed output means saturation, anti-windup and bumpless transfer all account for the term with no extra arithmetic. A change of K_{ff} is itself treated as a transfer. The lag state is shifted by the step the new gain makes in the calculation and I in (3.16) takes back the same step, so a feed change already in transit keeps its timing and N_{SP} does not jump.

Dropping the maintenance terms from both halves of (3.18) is deliberate, because a feedforward that requires the biomass concentration requires a measurement that is rarely available online. The result is a little high, asking for 66.5 rev/min where the plant needs 64.1 rev/min at the design feed and 83.7 rev/min where it needs 82.2 rev/min at 40 L/min, errors of 3.8 % and 1.8 %. Errors of that size are typical of a calculated feedforward, and the feedback trim removes them. The lag τ_{ff} is a different matter entirely, and Section 3.6 measures what happens without it.

Pause. Try Exercise 3.3 before reading on.

3.5.4 Implementation

The app runs on the framework of Section 1.6, with manual and automatic modes and bumpless transfer in every architecture. Integrated error is scored against the true dissolved oxygen rather than the probe reading, so a biased probe scores badly although its trend is flat. Runs export one row every 1.08 s of simulated time for offline analysis, carrying the true dissolved oxygen, the probe state and the biased noisy signal the controller acted on as three separate series, which is enough to repeat every comparison in Section 3.6. Twelve built-in verification checks cover the opening steady state, the oxygen and carbon balances, the exponents of Equation (3.2), the sign of every control action and the four time constants of Section 3.4. Figure 3.2 shows scenario 5 after the feed step is rejected.

3.6 Results

Table 3.3 lists thirteen scenarios in instructional order. Each one selects a starting charge, a mode, an architecture, both sets of tuning constants, the disturbance settings and whatever timed events it needs. Every integrated absolute error quoted below is in percent-DO minutes accumulated over a twenty-minute run with the feed step at $t = 4$ min, except where another window is stated, and every travel figure is the total distance the speed demand moved in rev/min over the same window.

3.6.1 Open-Loop Response

Scenario 1 opens four hours into a fed-batch, with 5 g/L of cells in 48 m^3 on an 8 L/min feed and the agitator speed demand parked at 42.7 rev/min, where transfer balanced uptake when the run started. No input changes for twenty minutes. The single-loop drive runs the open-loop torque schedule of Equation (3.13), so the shaft sags to 41.4 rev/min within a minute and $k_L a$ falls by 6 %. Cell mass grows from 5.00 to 5.32 g/L, but the uptake term of Equation (3.1) stays near 18.3 mg/(L·min), because the feed pump rather than the cell mass sets the growth rate of a substrate-limited fed-batch. The dissolved oxygen falls from 25.0 to a minimum of 16.0 % at 15 min and reads 16.1 % at twenty minutes. The fall is arrested by the plant rather than by any control action. As the dissolved oxygen drops, the driving force $C^* - C_L$ widens until transfer meets uptake again, and by sixty minutes the reading has recovered to 18.1 % as the glucose pool runs down and the uptake eases. The plant is



Figure 3.2 Scenario 5 twenty simulated minutes after the start, with the cascade in automatic and the feed step at $t = 4$ min rejected: the plant view of the capture with three scorecard tiles of the same capture pasted into its empty middle (a composite). (1) The agitator speed, moved from the design 64.1 rev/min of Table 3.1 to 82.2 rev/min. (2) Transfer and uptake, level at 68.2 mg/(L·min). (3) The agitator bar, with 10.8 rev/min of margin and 191 kW of shaft power listed below it, against the 88 kW of the design condition (Section 3.2). (4) The setpoint marker, where the probe is back at the 25 % setpoint. The true dissolved oxygen dipped 9.34 % below it on the way (Table 3.4). (5) The pasted tiles: an integrated error of 18.17, 18.19 rev/min of demand travel and no time below the critical value.

Table 3.3 The Thirteen Scenarios and the Mechanism Each One Isolates

#	Scenario	What it isolates
1	Early batch, open loop	a variable with no inventory, self-regulation
2	Late batch, open loop	identical CV, transformed MV
3	Single loop on a feed step	three lags in series, gain limit
4	Single loop at excessive gain	actuator travel as a cost
5	Cascade at the same tuning	what an inner loop removes
6	Cascade with detuned inner loop	the crossover from benefit to harm
7	Probe at 60 s	a lag no structure reaches
8	Feedforward from the feed	wrong in size against wrong in time
9	Feed raised at the speed stop	capacity, not tuning
10	Airflow against back pressure	a constraint that is not a range
11	Probe biased 14 % high	a perfect trend at the wrong value
12	Probe noise with derivative on	differentiating a slow measurement
13	Graded challenge, three faults	the whole ladder under time

stable at whatever value it arrives at rather than at a chosen one, so control supplies the selected operating point rather than stability. With the cascade's speed loop holding the shaft at 42.7 rev/min, the same twenty minutes stay between 24.7 and 25.3 %.

Scenario 2 repeats the experiment ten hours later at 34 g/L of cells in 92 m^3 on a 58 L/min feed, where the oxygen demand is 3.6 times larger. The dissolved oxygen reaches a minimum of 15.2 % at 4.4 min and reads 16.5 % at twenty minutes, against 16.0 and 16.1 % in scenario 1, which is nearly

Table 3.4 Integrated Absolute Error, Peak Deviation and Demand Travel on the Same 24 L/min to 40 L/min Feed Step, Read at $t = 20$ min With $\tau_I = 0.6$ min Throughout.

K_c	Single loop IAE / peak / travel	Cascade IAE / peak / travel	+ Feedforward IAE / peak / travel
0.6	23.8 / 11.83 / 25	18.2 / 9.34 / 18	5.9 / 2.90 / 18
1.8	10.6 / 8.82 / 44	6.0 / 6.19 / 18	2.4 / 2.43 / 18
3.2	12.0 / 5.18 / 95	**3.5** / 4.82 / 23	**2.1** / 2.13 / 22
4.5	57.1 / 21.02 / 683	11.7 / 3.93 / 152	10.1 / 2.69 / 137

the same excursion. The demand does not set the excursion. The open-loop drive loses about the same fraction of its speed at both ends of the batch, from 86.0 to 83.0 rev/min here against 42.7 to 41.4 rev/min early, and the same fractional loss of $k_L a$ costs the same few points of dissolved oxygen. The changes underneath are not visible on that trace. The agitator is at 83.0 rev/min instead of 41.4, drawing 188 kW instead of 23, with 10.0 rev/min of margin left instead of 51.6. The same batch run the same way is eight times as expensive to aerate and has spent more than four fifths of its actuator range.

3.6.2 Three Architectures on One Feed Step

Table 3.4 carries the central comparison. The feed steps from 24 to 40 L/min at $t = 4$ min from the design charge, the integral time is held at 0.6 min throughout, the inner loop is the shipped fast one, and every figure is read at $t = 20$ min.

Three results come out of Table 3.4. At the shipped tuning the cascade improves the peak from 11.83 to 9.34 % and the integrated error from 23.8 to 18.2, a modest gain, because it removed 15.4 s from a chain that still contains 12.4 s of liquid and 20 s of probe. The larger difference is in what the outer loop can then be tuned to. The single loop reaches its best integrated error of 10.6 at $K_c = 1.8$ and spends 44 rev/min of travel getting there, while the cascade reaches 3.5 at $K_c = 3.2$ for 23 rev/min, a factor of three on the error and a factor of two on the wear. And the single loop's best peak sits at a different gain from its best integrated error, which is the usual sign that a loop is being pushed against a lag rather than tuned against a process. Figure 3.3 draws the integrated error and the travel of Table 3.4 as bars.

Scenario 4 is what happens past that point, with K_c at 4.5 and the integral time also shortened to 0.5 min. At the design operating point this tuning is unstable. The oscillation grows from the start, the demand swings between its speed limits, and the dissolved oxygen reaches 23.77 % above the setpoint at 3.9 min and 21.47 % below it at 4.3 min, just after the feed step. The demand is close to a square wave and the actual speed is a rounded version of it running behind, so every correction arrives while the previous one is still in transit. Once the feed step has moved the operating point to 82 rev/min, where the drive lag and the oxygen holdup are both shorter, the same tuning is stable and the oscillation decays to ± 1.1 % over the last four minutes. The run ends with an integrated error of 72.5, 1.06 min below the critical value and 859 rev/min of demand travel against 25 at the shipped tuning. The significant number is the travel rather than the error, because 859 rev/min of demand movement in twenty minutes is gearbox and mechanical-seal wear spent to obtain a worse result than a smaller gain delivers. A tuning that is unstable at one point of the batch and stable at another is the operating-point dependence of the three lags seen from the outer loop.

> ✦ **DID YOU KNOW?** **One Outer Gain, Two Very Different Bills**
>
> From the simulator: at an outer gain of 4.5 on the same feed step, the single loop moves its speed demand 683 rev/min for a peak deviation of 21.02 %, while the cascade moves it 152 rev/min for 3.93 % (Table 3.4). The same gain costs the drive far less once a speed loop sits under it.

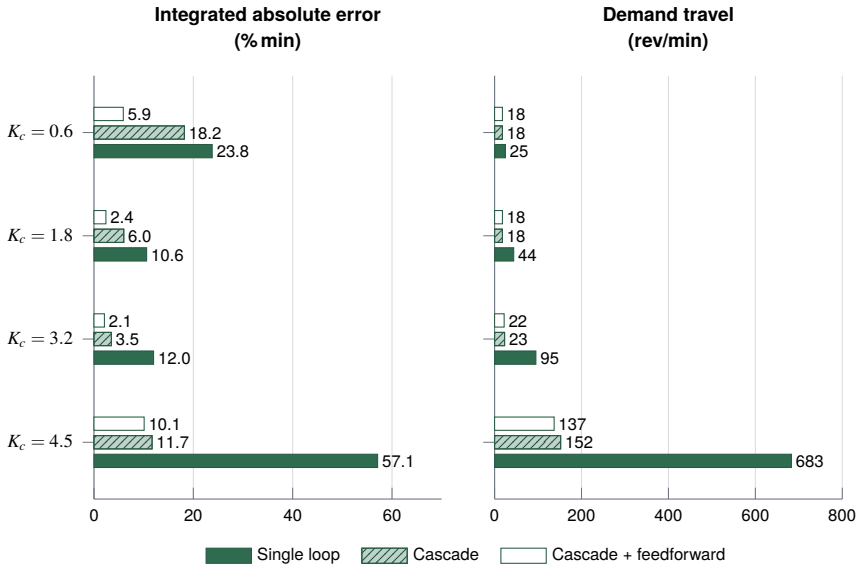


Figure 3.3 Table 3.4 redrawn: integrated absolute error and demand travel of the three architectures at four outer gains on the same feed step, the single loop in solid bars, the cascade in hatched bars and the cascade with feedforward in open bars, each panel on its own scale. Notice that the single loop's error rises again above $K_c = 1.8$ while its travel climbs at every step, and that the cascade keeps its travel between 18 and 23 rev/min up to $K_c = 3.2$.

3.6.3 Inner Loop Speed Requirement

Scenario 6 changes one thing about scenario 5. The inner gain is cut from 6 to 0.35 and the inner integral time is stretched from 0.10 to 1.20 min, so that by Equation (3.12) the speed loop settles in 12.3 s instead of 2.9 s. The detuned inner loop still takes the shaft to its setpoint. The result is worse than no cascade at all. Peak deviation is 14.84 % against the single loop's 11.83 %, the integrated error is 72.3 against 23.8, and the batch spends 2.76 min under the critical dissolved oxygen where the single loop spent 0.53. Detuning the inner loop turned a 23 % improvement into a threefold degradation, because an inner loop slower than the process it sits inside is not a cascade but an extra lag with an extra integrator in it.

> **TRY IT IN THE APP** **6 · The Same Cascade With a Slow Inner Loop**
>
> Load the scenario, press Run in Auto and let twenty simulated minutes pass. Watch the IAE and Below critical tiles. The chapter reports an integrated error of 72.3 and 2.76 min under the critical value, against 23.8 and 0.53 for the single loop of scenario 3, from a speed loop that still reaches its setpoint.

Sweeping the inner gain locates the crossover, and repeating the sweep at a different inner integral time shows what the rule of thumb leaves out. Figure 3.4 carries both sweeps with the outer controller held at $K_c = 0.6$ throughout.

The upper sweep holds the inner integral time at the 1.20 min scenario 6 ships and walks $K_{c,N}$ through 0.15, 0.25, 0.35, 0.50, 0.75, 1.0, 1.5, 2.0, 4.0 and 6.0, giving integrated errors of 172, 102, 72.3, 52.6, 38.3, 30.5, 23.4, 20.3, 18.2 and 18.2. It crosses the single loop's 23.8 at an inner time constant of about 7.4 s, which is 0.6 times the 12.4 s oxygen holdup the loop is nested inside. Below that speed the cascade is a liability. Above it the improvement continues until roughly 4 s and then stops. Going from 6.2 to 3.9 s takes the integrated error from 20.3 to 18.2, and the last step to 2.8 s

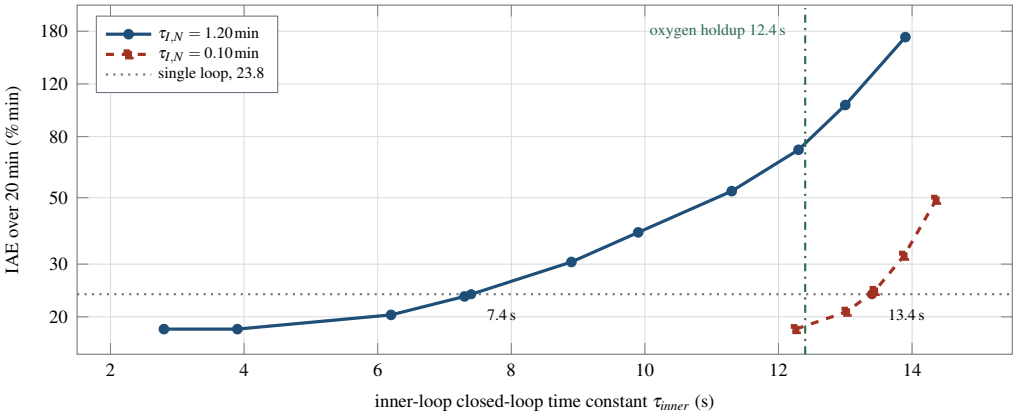


Figure 3.4 Where a cascade stops paying, measured on the scenario 5 feed step with the dissolved-oxygen controller held at $K_c = 0.6$ and $\tau_I = 0.6$ min. The inner gain $K_{c,N}$ is walked from 0.15 to 6.0 and the abscissa is the closed-loop time constant Equation (3.12) returns for it. Solid line with circles: $\tau_{I,N} = 1.20$ min. Dashed line with squares: $\tau_{I,N} = 0.10$ min. Below the dotted horizontal line the cascade beats the single loop and above it the cascade is worse than no cascade. The dash-dotted vertical line marks the oxygen holdup. The crossover moves from 7.4 s to 13.4 s when only the inner integral time is changed.

changes it by less than 0.01. The inner loop has become invisible to the outer one, which is what a cascade is trying to achieve.

The second sweep identifies what the conventional rule omits. Restoring the inner integral time to 0.10 min and repeating with $K_{c,N}$ at 0.10, 0.15, 0.20, 0.25 and 0.35 gives 49.2, 32.0, 24.4, 20.8 and 18.2. Equation (3.12) assigns those same gains time constants of 14.36, 13.88, 13.44, 13.02 and 12.26 s, so the crossover has moved out to about 13.4 s, nearly twice as slow as before, and everything available is collected by 12.3 s. A fast integral term is doing work the proportional term was doing in the first sweep. The number that decides whether a cascade pays is the inner loop's whole closed-loop response and not the proportional gain alone, a quantity that a rule expressed in gains does not contain.

Stated as a rule the measurement supports, a cascade stops paying at about the moment its inner loop becomes as slow as the fastest dynamics in the outer process. Breaking even needs an inner loop about 1.7 times faster than that fastest element, 7.4 s against the 12.4 s holdup, while collecting the whole benefit needs about three times faster, near 4 s. The conventional three to five is under-specified rather than wrong, because the whole outer chain is 12.4 s of liquid plus 20 s of probe and three to five times faster than that 32 s total is 6.5 to 11 s, which brackets the measured crossover only because the probe dominates the sum. A comparison against the wrong time constant errs by a factor of three. Two practical consequences follow. At the crossover the cascade is actively harmful rather than neutral, so the rule has no safe side to err toward at the slow end, and inner speed beyond the point where it stops mattering costs nothing, which is the argument for specifying an inner loop generously rather than exactly.

The inner measurement imposes its own limit. Running scenario 5 with the tachometer filter at 0.5, 1, 2, 4 and 6 s gives an integrated error of 18.2 every time with 18 rev/min of demand travel. At 8 s the loop goes unstable, at an integrated error of 73.9 and 120 rev/min of demand travel with a shaft that moves 570 rev/min in response to its own filtered measurement. An inner loop with a gain of 6 and a measurement filtered at 8 s is not a slow cascade but an unstable one, and the failure appears in the outer loop as an oxygen excursion that resembles a tuning problem. The tachometer is specified at 0.5 s for that reason, and half a second against twenty seconds of dissolved-oxygen probe states

the cascade argument in two numbers.

3.6.4 Probe Lag

Scenario 7 triples the probe time constant from 20 to 60 s and changes nothing else. The cascade's peak deviation goes from 9.34 to 12.01 % and the batch now spends 0.77 min under the critical value where before it spent none. The single loop goes from 11.83 to 14.12 %. The penalty is 2.7 points for the cascade and 2.3 for the single loop, which within the resolution of the experiment is the same penalty paid twice, and the cascade absorbed none of it. Set beside the 15.4 s of drive lag the same cascade removed outright, the asymmetry is the central result of the comparison.

A second-order effect deserves attention because it appears beneficial. A slow probe is also a filter. Pushing the cascade to $K_c = 3.2$ with the probe at 10 s makes the loop oscillate, at an integrated error of 38.7 for 696 rev/min of demand travel, where the shipped 20 s probe gives 3.5 for 23 and the aged 60 s probe gives 3.4 for 18. The slow probe delivers stability by discarding the information that would have made the loop fast, so a loop tuned against an aged probe becomes unstable when a new one is fitted. No control-system diagnostic reports that the membrane has aged, because a slower probe makes the trend appear calmer, and the record that would catch it is a calibration sheet carrying a response time.

3.6.5 Feedforward Magnitude and Timing

Scenario 8 adds the carbon-balance term of Equation (3.18) to the same cascade. Peak deviation falls to 2.90 % and the integrated error to 5.9, against 9.34 and 18.2 for the cascade alone and 11.83 and 23.8 for the single loop, and demand travel is unchanged at 18 rev/min, so the improvement came free of actuator wear. Setting K_{ff} to zero before the run, or carrying that setting through a restart, reproduces scenario 5 to within 1.6×10^{-13} percentage points of dissolved oxygen, which is the cleanest way to separate what the calculation contributes from what the structure contributes. Made during a run, the change is bumpless. The bias falls to the 21 rev/min speed floor in one execution and the integral takes back the same step, so the speed demand does not move and the loop continues as the plain cascade.

Setting the feedforward lag τ_{ff} to zero is the more instructive mistake. The agitator reaches the new speed the instant the feed changes while the oxygen demand is still a minute or more behind, and the dissolved oxygen overshoots 22.86 % upward with an integrated error of 15.1 and twice the travel, so a correct but early feedforward acts as a disturbance in its own right, the early case of Section 1.4.3. Matching τ_{ff} to the glucose pool at 0.8 min brings the error back to 5.9 and makes it nearly symmetric, at 2.90 % low and 1.75 % high. A feedforward admits two errors, a static calculation from incorrect data and a dynamic mismatch between when the term acts and when the disturbance arrives. Feedback corrects the first and only the lag corrects the second.

> **TRY IT IN THE APP** **8 · Feedforward From the Feed the Operator Is Already Setting**
>
> Run the scenario as loaded, then press Reset, switch to Advanced, set Feedforward lag to 0 and run again. Watch which way the dissolved oxygen leaves its setpoint after the feed step. The chapter reports a dip of 2.90 % and an integrated error of 5.9 with the lag at 0.8 min, and 22.86 % upward with 15.1 without it.

3.6.6 Transfer Capacity Limits

Scenario 9 raises the feed from 58 to 88 L/min at 5 min late in the batch, with the agitator already at 86.0 rev/min and 7.0 rev/min of margin. The agitator reaches 93.0 rev/min at 5.6 min and stays there with zero margin, the dissolved oxygen falls to 6.0 % by twenty minutes and keeps sinking slowly, to 3.2 % at three hours, and the batch spends 14.8 of the first 20 minutes below the critical value. No tuning changes any of it. A higher gain, a shorter integral time or the feedforward moves the time below critical by less than 0.4 min, because oxygen uptake is asking for 89.7 mg/(L·min) and the

Table 3.5 Vessel Back Pressure at the Design Charge, Holding 25 % Dissolved Oxygen After the 24 L/min to 40 L/min Feed Step.

p_{head}	C^*	Agitator	Shaft power
0 barg	6.59 mg/L	93.0 rev/min, pinned	281 kW
0.5 barg	9.84 mg/L	82.2 rev/min	191 kW
1.0 barg	13.09 mg/L	70.6 rev/min	118 kW
1.5 barg	16.34 mg/L	63.0 rev/min	83 kW

drive cannot transfer that much at this volume and airflow. A demand pinned against its limit while the error stays large is the signature of a capacity problem rather than a tuning one.

The cost of the fault is not primarily lost product. Growth in a substrate-limited fed-batch is set by the feed pump rather than by the oxygen, so the titer changes little for a long while and the unconsumed sugar accumulates instead. Residual glucose goes from 0.04 to 1.00 g/L at twenty minutes, 2.85 at an hour and 7.53 g/L at three hours. That is the first clear symptom of the fault and it appears on the culture trend rather than on the oxygen trace.

Airflow appears to be the remedy and fails. Scenario 10 applies the same 88 L/min feed at 4 min, and at 10 min, with the agitator pinned at its speed stop, more than doubles the air from 0.40 to 0.95 vvm. The instant the flow changes, the gas velocity term of Equation (3.2) lifts k_La from 9.53 to 11.83 min^{-1}, a 24 % gain, and the gassed power falls from 269 to 223 kW because the extra gas unloads the impellers through Equation (3.3). Then the foam of Equation (3.11) arrives. The headspace is full 3.5 min after the change, the sparged gas then passes through foam rather than through broth, and twenty minutes after the change k_La is back to 9.37 min^{-1}, below where it started. The dissolved oxygen ends at 5.6 % with 25.8 of the 30 minutes below the critical value. The airflow range extends to 1.2 vvm, while the effective limit arrived at 0.95 vvm on a different variable.

Back pressure succeeds where airflow does not, because it raises the equilibrium concentration in Equation (3.4) rather than the transfer coefficient. Table 3.5 sweeps it at the design charge on the same feed step.

At atmospheric pressure the plant cannot hold its own setpoint at the design charge once the feed has stepped, and it burns 281 kW failing to, while one and a half bar of head holds it for 83 kW, and the compressor duty that buys the difference is a rounding error beside a 200 kW agitator. On the capacity disturbance of scenario 9, with 1.5 barg set before the run starts, the same lever holds 25.0 % dissolved oxygen at 80.7 rev/min with 12.3 rev/min of margin at twenty minutes. Run for three hours, the pinned batch makes 83.38 g/L of ethanol for 792 kWh of agitator energy while the pressurized one makes 85.31 g/L for 466 kWh, with residual glucose at 0.05 g/L instead of 7.53 and half a minute below the critical value instead of 175. Raised at the 5 min feed step instead, the pressure gives the same titer for 476 kWh and no time below the critical value. Both batches reach the 105 m^3 working volume at 149 min, so their last half hour is made under overflow rather than fed-batch. The pressurized batch returns 2.3 % more product for 41 % less electricity, and the lever is a pressure setpoint rather than a controller tuning.

3.6.7 Instrument Faults

Scenario 11 puts 14 percentage points of positive bias on the probe of Equation (3.10) and changes nothing else. The controller holds the reading at 25.0 % and the true dissolved oxygen at 11.0 %, below the 15 % critical value for the whole run, and the trend remains flat. The batch loses nothing. After four hours it has made 49.24 g/L of ethanol against 49.25 with an honest probe, from the same 21.3 g/L of cells, and it used 279 kWh of agitator energy against 325 because it was holding a lower real dissolved oxygen.

The reason is structural. The feed pump sets the growth rate in a substrate-limited fed-batch, and oxygen becomes the limit only once it can no longer support the growth the feed supplies [17].

At this charge that crossover is at a true dissolved oxygen of 3.5 %, not 15, so the bias consumed margin rather than product, and it stays invisible until a second fault such as a fouled sparger removes that margin. This is the one fault in the collection that a better controller makes strictly worse. The defenses are outside the loop, a second probe, an off-gas balance that would fail to close, or calibration discipline, and the same conclusion has been reached for aeration control in wastewater treatment, where a biased sensor is invisible to every diagnostic built on control error [21].

Scenario 12 fails the other way, costing no product but a large amount of actuator wear. The probe carries 1.5 % of electrical noise, the feed is steady, the vessel is otherwise at steady state, and derivative action is switched on at 0.15 min with no filter. Twenty idle minutes cost 13 943 rev/min of demand travel. Filtering the derivative at 0.05 min drops that to 2350 and at 0.20 min to 1076, although the filtered runs pay for it in integrated error, 34.3 and 30.0 against 25.4. Removing the derivative altogether gives 613 rev/min with the best integrated error of the four at 21.4. Dropping the noise from 1.5 to 0.5 % takes the 613 to 204, and removing it entirely leaves 0.25 rev/min of travel. Derivative action supplies anticipation, and on a 20 s probe it has nothing to anticipate, because a real change is no longer predictive by the time it reaches the derivative term and everything the term responds to promptly is noise. The general test is whether the measurement is faster than the quantity the derivative is asked to predict, and a dissolved-oxygen probe fails it on both counts.

3.6.8 Combined Fault Scenario

Scenario 13 applies three faults over 45 min to a cautious single loop ($K_c = 0.25$, $\tau_I = 2.0$ min), a feed increase from 24 to 44 L/min at $t = 6$, a probe membrane fouling that triples the probe lag at $t = 18$ and a quarter of the sparger plugging at $t = 28$. The case is passed when the true dissolved oxygen stays within ± 4 % of the setpoint over the last third of the run, from 29.7 min, and the agitator finishes off its speed limit. With the shipped settings the batch spends 26.0 of the 45 minutes below the critical value, ends at 9.1 %, burns 137 kWh and leaves the band by up to 15.9 %, nearly four times the band, although the agitator ends 5.5 rev/min below its limit. The case therefore fails on the band alone. Tuning alone removes most of the error and then meets a capacity limit. The feedforward architecture with a fast inner loop ($K_{c,N} = 8$, $\tau_{I,N} = 0.08$ min) and $K_c = 2.6$ at $\tau_I = 0.5$ min cuts the integrated error from 471 to 158 and the time below critical to 4.32 min, but the shaft reaches its speed limit at 28.3 min, just after the sparger fault, and finishes with no margin at all. Adding 1.0 barg of back pressure to that same tuning, which starts the plant at the 1.0 barg steady state with the agitator at 55.1 rev/min, gives a worst deviation of 0.19 % over the judged window, inside the band with 3.8 points to spare, an integrated error of 7.1, 0.17 min below critical, 9.0 rev/min of margin still in hand and 110 kWh against the 162 kWh the tuning-only configuration spent. The rule reads only the judged window and the final margin, so the back pressure with $K_c = 2.6$ and $\tau_I = 0.5$ min on the single loop also passes, at 2.4 %, after a 28.7 % swing and 696 rev/min of demand travel early in the run, while neither the pressure (12.5 %) nor the tuning (10.6 %) passes alone. The ranking of the contributions is the result to retain. The inner loop and the architecture buy the response, the feedforward buys the feed disturbance, the pressure buys the capacity, and no measure recovers the probe lag.

3.7 Technology in Practice

Figure 3.5 draws the cascade and its feedforward on real hardware, where the speed loop lives inside the drive.

What Measures It

The outer loop reads a membrane probe, AT-301, which the model of Equation (3.10) reduces to one first-order lag of 20 s with a bias and a noise term. The inner loop reads the shaft. A tachometer is typically a small generator whose voltage is proportional to speed, often on a 0 to 10 V scale [22], and the model's ST-301 is that reading through a 0.5 s filter with no bias and no noise.

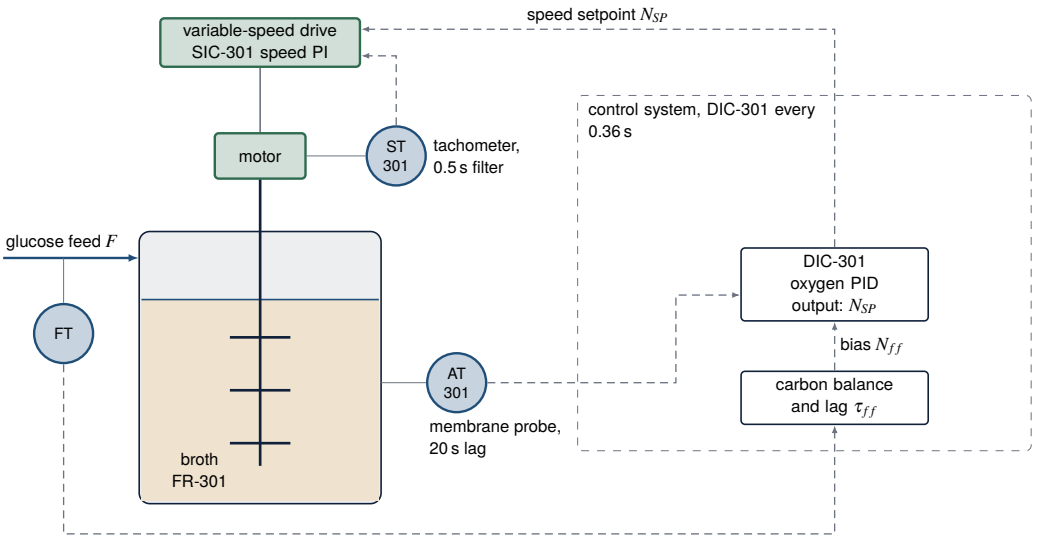


Figure 3.5 The oxygen cascade and its feedforward on real hardware. The membrane probe AT-301 feeds the oxygen controller DIC-301, whose output is a speed setpoint for the speed loop SIC-301, which runs inside the variable-speed drive, reads the tachometer ST-301 and sets the motor torque. The measured glucose feed enters DIC-301 as its bias through the carbon balance and its lag. Solid arrows are process streams and dashed gray arrows are signals. Notice how the slow probe and the fast tachometer close their loops in two different devices.

> **TECHNOLOGY** **Sensor: A Membrane Electrode That Consumes What It Reads**
>
> A Clark-type electrode reduces oxygen at a cathode behind a membrane, so it consumes the oxygen it measures. Its signal then depends on the flow past the membrane, where a depleted boundary layer forms, and its response is close to first order with a time constant that is generally not small against $1/k_La$ [13]. One probe on a 15 L cell-culture vessel measured 49.6 s, and microbial fermentations run transfer coefficients more than ten times higher [23]. A field manual asks for about one foot per second of water across such a membrane, and lists a fouled membrane as the cause of an instrument too slow to react [14]. On a trend that fault looks like a calmer loop, the change from 20 s to 60 s of scenario 7.

What Moves It

The final element is the agitator motor on its variable-speed drive. The model's drive is a torque balance with a 15.4 s lag at the design speed, clamped to 0 to 100 % of rated torque and to 21 to 93 rev/min. Real drives carry the same limits as settings: speed limiting, torque limiting, current limiting, and limits on acceleration and deceleration [22]. The last of these is the ramp that the Assumptions of Section 3.3.6 say a real drive follows, so a large speed change takes longer than the model shows.

> **TECHNOLOGY** **Final Element: A Drive That Closes Its Own Speed Loop**
>
> A variable-speed drive that reads a tachometer on the motor shaft throttles power until the shaft turns at the commanded speed. That built-in loop makes the drive the secondary controller of a cascade, taking its speed setpoint from the process controller. In the direct-current drives that rely on this feedback, a lost tachometer signal reads as zero speed and the motor runs away to full speed [22]. An alternating-current drive sets speed through its output frequency and knows the speed only approximately without feedback, and precise control, above all at low speed where slip is a larger share, may need a speed sensor [22]. Equation (3.13) is such a drive without its sensor.

Where the Controller Runs

DIC-301 would run in a distributed control system or a programmable logic controller, and SIC-301 inside the drive itself, which takes the speed demand as its setpoint. A host system can often update a networked drive's limits while the motor runs [22]. Loops in such systems are sometimes nested or cascaded, with cycle times from milliseconds to minutes [24], and the app's 0.36 s execution period for both controllers sits inside that range. The feedforward of Section 3.5.3 enters DIC-301 as its bias, so it belongs in the same controller as the oxygen loop.

What Is Modernizing It

Three developments touch this loop. First, optical probes measure the lifetime of a quenched luminescence, consume no oxygen, need no flow past the sensor and hold their calibration longer than amperometric ones [14]. Second, soft sensors compute the oxygen uptake from the probe reading, a characterized transfer coefficient and the solubility, which removes the need for an off-gas analyzer for that purpose [23]. Third, such soft sensors commonly run in the supervisory systems above the control level, because scripts written into the control system itself are the exception, and in an industrial plant they must first pass a functional and risk assessment [25].

> **IN PRACTICE** **Inspect the Membrane, Check the Zero, Then Calibrate**
>
> A federal field manual sets a routine for membrane probes. Inspect the membrane for wrinkles, tears and bubbles under it before use. Read the probe in a fresh zero-oxygen sodium sulfite solution, and if more than 0.2 mg/L persists, replace the membrane and the electrolyte. Let a new membrane condition for two to six hours, overnight for greater stability, before calibrating. Calibrate at one point, full saturation in air or in air-saturated water, with the sensor and its thermistor at the same temperature, and expect an amperometric sensor to need recalibration every week or two [14]. The zero check is the one that exposes a constant offset like the 14 points of scenario 11.

3.8 Limitations

The application is not a design tool. With one perfectly mixed lump, no gas-phase state, fixed correlation exponents, unstructured kinetics and no cooling model, the application cannot support a claim about a real vessel's oxygen capacity, its agitator sizing or the yield a real batch would make, and Section 3.3.6 says where each of those simplifications breaks. The upsets are idealized as well. A feed step from 24 to 40 L/min arriving in one instant would not occur under a fed-batch recipe, and a sparger that loses a quarter of its holes at once isolates a loss of transfer capacity from everything else. The twenty-minute scenarios compress a batch that runs fifteen to sixty hours, so no number in the application is a batch yield or a real cycle time.

The control scope is one dissolved-oxygen loop with one cascade and one feedforward. There is no respiratory quotient control, no substrate feedforward from an off-gas balance, no model predictive control and no supervisory optimizer choosing the feed profile, and a modern plant would have at least the off-gas balance. Airflow and back pressure are operator settings rather than controlled variables, and on a real vessel both would be loops of their own. The cell-culture vessel of Volume 19, Chapter 8 sequences agitation, oxygen enrichment and gas flow on one dissolved-oxygen loop. The feed profile is an input rather than a decision the application optimizes, although the real question on a fed-batch is what feed trajectory maximizes titer against the oxygen the vessel can transfer. There is no electricity price and no product price, so the reported metrics state what a decision cost in kilowatt hours and in grams per liter without saying which way the decision should go. Nothing fails outright, because the critical dissolved oxygen, the foam-out, the speed limit and the working volume annunciate and the run continues past the point where a real batch would have been shut down. Perfect repeatability is itself an idealization, because a fixed integration step and a fixed execution period give one trajectory on every machine while a distributed control system scanning about once a second gives a slightly different one every batch.

The model of this chapter also supports control and estimation methods beyond the implemented PID, cascade and feedforward loops. The nine states with the speed limit, the torque limit and the foam constraint form a natural plant for Model Predictive Control, which would coordinate agitation, airflow and back pressure as one constrained multivariable problem and could weight the agitator electricity against the dissolved-oxygen target as an economic objective. The same structure supports state and parameter estimation, because a moving horizon estimator or a Kalman filter driven by the probe signal and the measured shaft power can reconstruct the oxygen uptake rate and the biomass concentration that no instrument measures directly, a soft sensor that addresses the biased probe of scenario 11 from outside the loop. The APMonitor framework implements MPC and moving horizon estimation for differential algebraic models of exactly this form [26], and the Gekko Optimization Suite provides the same methods in Python [27], so the constants of Table 3.2 transfer directly to a graduate study of advanced control.

> ⚠ **OPEN CHALLENGE** **Gradients That One Probe Cannot See**
>
> A large stirred tank develops gradients when its power input is limited. Computational fluid dynamics can predict them, but one turbulent simulation may take hours, too slow for routine optimization or control. A 2026 study trained surrogate models on 760 such simulations of a 170 m^3 vessel with four impellers and predicted the three-dimensional substrate field in milliseconds, at a test R^2 near 0.91. Its simulations were single-phase, without experimental validation, for one geometry at one scale [28]. Progress would be a validated model that carries the gas phase, so that a dissolved-oxygen probe could be read against a known gradient.

3.9 Discussion

The dissolved-oxygen setpoint is an economic decision as much as a control decision, and no control-error metric reports the cost. Agitator power goes as the cube of speed through Equation (3.2), while transfer goes as the 2.1 power, so the last few points of dissolved oxygen are the expensive ones. At the design charge, holding 10 % costs 72.9 kWh over an hour, 25 % costs 86.0 and 50 % costs 118.8 kWh, and all three make the same ethanol to within 0.01 g/L, 42.41 g/L. That is 63 % more electricity for no additional product, committed by the choice of setpoint. Aeration is the dominant mechanical energy demand of a stirred production fermenter [7], and on a commodity fuel whose net energy balance has been debated [19] the difference between 52.0 kWh and 90.4 kWh per tonne that scenario 9 measures is not a rounding error.

Two quantities that no control-error metric contains are therefore reported beside it. Actuator travel is a maintenance budget, and the 859 rev/min of scenario 4 and the 13 943 rev/min of scenario 12 represent gearbox, mechanical seal and drive wear rather than abstract numbers. Speed margin is capacity, and no tuning creates more of it. The application reports six competing numbers rather than one composite score for that reason, because a run that scores well on error by pinning the agitator has exhausted its capacity margin rather than solved the problem.

> ⚠ **OPEN CHALLENGE** **A Soft Sensor Inherits Its Probe's Faults**
>
> A soft sensor joins hardware signals to a model to predict a quantity no instrument reads, such as the oxygen uptake of a culture. A critical review notes that a faulty input very probably gives a faulty output, yet studies that detect sensor faults in bioprocesses, or tolerate them, are rare [25]. An uptake estimate built on the biased probe of scenario 11 would start from a reading 14 points high. Progress would be estimators that test their own inputs, for example against a redundant probe, before the plant acts on them.

3.10 Conclusion

This chapter used an aerated fed-batch fermenter to measure how fast the inner loop of a cascade has to be. The simulator integrates nine states, the four culture balances, the dissolved oxygen, the agitator speed, the probe, foam and antifoam, with oxygen transfer from a Van 't Riet power law of

exponents 0.70 and 0.40. The vessel holds 107 g of dissolved oxygen against a 2.57 kg per minute uptake, so four lags decide the loop, a drive at 15.4 s, a liquid phase at 12.4 s, a probe at 20 s and a speed loop at 2.9 s. On one feed step the single loop reaches an integrated error of 10.6 for 44 rev/min of demand travel, the cascade 3.5 for 23 and the calculated feedforward 2.1 for 22. Detuning only the inner loop takes the cascade to 72.3 against the single loop's 23.8, and sweeping the inner gain places the crossover between benefit and harm at 7.4 s, which is 0.6 times the oxygen holdup, while the same sweep at a shorter inner integral time places it at 13.4 s. Tripling the probe lag costs 2.7 points of peak deviation in cascade and 2.3 in the single loop, which is the boundary of what any structure buys.

A fuller model would add a gas-phase state with an off-gas balance, a compartment model for the mixing gradient, a ramp limited drive and a cooling loop so that the two capacity constraints can be studied together. As it stands, the fermenter measures a cascade design rule instead of quoting it.

KEY POINTS

- A cascade helps only when its inner loop is fast enough, and a slow inner loop makes the oxygen loop worse than no cascade at all.
- Detuning only the inner loop takes the integrated error to 72.3 against the single loop's 23.8.
- The crossover between benefit and harm sits at an inner time constant of 7.4 s, 0.6 times the oxygen holdup, and at 13.4 s with a shorter inner integral time.
- A cascade reaches the actuator, a feedforward reaches the disturbance, and nothing reaches the measurement.

Where else this applies. Oxygen cascades return with a meal feedforward in Volume 19, Chapter 12 and under a supervisory setpoint in Volume 9, Chapter 10.

References

[1] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*. John Wiley & Sons, 2016.

[2] 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.

[3] G. K. McMillan, "Effect of cascade control on loop performance," in *1982 American Control Conference*, 1982, pp. 363–368.

[4] L. Åmand, G. Olsson, and B. Carlsson, "Aeration control – a review," *Water Science and Technology*, vol. 67, no. 11, pp. 2374–2398, 2013.

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

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

[7] D. Humbird, R. Davis, and J. D. McMillan, "Aeration costs in stirred-tank and bubble column bioreactors," *Biochemical Engineering Journal*, vol. 127, pp. 161–166, 2017.

[8] C. M. Cooper, G. A. Fernstrom, and S. A. Miller, "Performance of agitated gas-liquid contactors," *Industrial & Engineering Chemistry*, vol. 36, no. 6, pp. 504–509, 1944.

[9] K. Van't Riet, "Review of measuring methods and results in nonviscous gas-liquid mass transfer in stirred vessels," *Industrial & Engineering Chemistry Process Design and Development*, vol. 18, no. 3, pp. 357–364, 1979.

[10] F. Garcia-Ochoa and E. Gomez, "Bioreactor scale-up and oxygen transfer rate in microbial processes: An overview," *Biotechnology Advances*, vol. 27, no. 2, pp. 153–176, 2009.

[11] L. C. Clark, Jr., R. Wolf, D. Granger, and Z. Taylor, "Continuous recording of blood oxygen tensions by polarography," *Journal of Applied Physiology*, vol. 6, no. 3, pp. 189–193, 1953.

[12] M. J. Johnson, J. Borkowski, and C. Engblom, "Steam sterilizable probes for dissolved oxygen measurement," *Biotechnology and Bioengineering*, vol. 6, no. 4, pp. 457–468, 1964.

[13] T. V. Kirk and N. Szita, "Oxygen transfer characteristics of miniaturized bioreactor systems," *Biotechnology and Bioengineering*, vol. 110, no. 4, pp. 1005–1019, 2013.

[14] S. A. Rounds, F. D. Wilde, and G. F. Ritz, "Dissolved oxygen (ver. 3.0)," U.S. Geological Survey, Techniques of Water-Resources Investigations, book 9, chap. A6, sec. 6.2, 2013. [Online]. Available: https://pubs.usgs.gov/twri/twri9a6/twri9a62/twri9a6_6.2_ver3.pdf

[15] U.S. Geological Survey, "Dissolved oxygen," U.S. Geological Survey, Techniques and Methods, book 9, chap. A6.2, 2020.

[16] T. A. Clark, T. Hesketh, and T. Seddon, "Automatic control of dissolved oxygen tension via fermenter agitation speed," *Biotechnology and Bioengineering*, vol. 27, no. 10, pp. 1507–1511, 1985.

[17] B. Sonnleitner and O. Käppeli, "Growth of Saccharomyces cerevisiae is controlled by its limited respiratory capacity: Formulation and verification of a hypothesis," *Biotechnology and Bioengineering*, vol. 28, no. 6, pp. 927–937, 1986.

[18] R. Luedeking and E. L. Piret, "A kinetic study of the lactic acid fermentation. Batch process at controlled pH," *Journal of Biochemical and Microbiological Technology and Engineering*, vol. 1, no. 4, pp. 393–412, 1959.

[19] A. E. Farrell, R. J. Plevin, B. T. Turner, A. D. Jones, M. O'Hare, and D. M. Kammen, "Ethanol can contribute to energy and environmental goals," *Science*, vol. 311, no. 5760, pp. 506–508, 2006.

[20] S.-O. Enfors, M. Jahic, A. Rozkov, B. Xu, M. Hecker, B. Jürgen, E. Krüger, T. Schweder, G. Hamer, D. O'Beirne, N. Noisommit-Rizzi, M. Reuss, L. Boone, C. Hewitt, C. McFarlane, A. Nienow, T. Kovacs, C. Trägårdh, L. Fuchs, J. Revstedt, P. C. Friberg, B. Hjertager, G. Blomsten, H. Skogman, S. Hjort, F. Hoeks, H.-Y. Lin, P. Neubauer, R. van der Lans, K. Luyben, P. Vrabel, and Å. Manelius, "Physiological responses to mixing in large scale bioreactors," *Journal of Biotechnology*, vol. 85, no. 2, pp. 175–185, 2001.

[21] O. Samuelsson, G. Olsson, E. Lindblom, A. Björk, and B. Carlsson, "Sensor bias impact on efficient aeration control during diurnal load variations," *Water Science and Technology*, vol. 83, no. 6, pp. 1335–1346, 2021.

[22] 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.

[23] M. Pappenreiter, B. Sissolak, W. Sommeregger, and G. Striedner, "Oxygen uptake rate soft-sensing via dynamic kLa computation: Cell volume and metabolic transition prediction in mammalian bioprocesses," *Frontiers in Bioengineering and Biotechnology*, vol. 7, p. 195, 2019.

[24] 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.

[25] V. Brunner, M. Siegl, D. Geier, and T. Becker, "Challenges in the development of soft sensors for bioprocesses: A critical review," *Frontiers in Bioengineering and Biotechnology*, vol. 9, p. 722202, 2021.

[26] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[27] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, "GEKKO optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.

[28] V. Puig-I-Laborda, J. Carrasco Muriel, A. Castañeda Tena, T. Groves, J. le Nepvou de Carfort, U. Krühne, K. V. Gernaey, C. Haringa, I. Puig-de Dou, and L. K. Nielsen, "Machine-learning surrogates enable rapid prediction of 3-D substrate gradients in industrial bioreactors," *Nature Communications*, vol. 17, p. 11636, 2026.

3.11 Exercises

3.1 ○ Concept check · 1 min

In the cascade of this chapter, what does the dissolved-oxygen controller DIC-301 manipulate?

(a) The motor torque command u_T, in percent of rated torque

(b) The glucose feed rate F, in liters per minute

(c) The agitator speed setpoint N_{SP}, in revolutions per minute

(d) The vessel back pressure p_{head}, in bar gauge

3.2 Concept check · 2 min

Three lags lie between the output of DIC-301 and its measurement: the agitator drive (15.4 s), the oxygen holdup of the broth (12.4 s) and the membrane probe (20 s). Which of them does a speed loop closed on the tachometer remove?

(a) The probe lag, because the tachometer reads far faster than the probe

(b) The oxygen holdup, because the agitator speed sets the transfer coefficient

(c) All three, as long as the speed loop is three to five times faster

(d) The drive lag only, since only the drive lies inside the speed loop

3.3 Concept check · 2 min

The carbon-balance feedforward of Equation (3.18) is computed from the glucose feed and drops the maintenance terms. Why is it not computed from the cell concentration X, which sets the uptake through Equation (3.6)?

(a) The feed rate is measured, and X would need a measurement rarely available online

(b) The maintenance terms are larger than the growth terms, so X adds nothing

(c) An off-gas analyzer already supplies the uptake, so X is not needed

(d) Dropping X makes the calculation exact, so no feedback trim is needed

3.4 Theory and calculation · 5 min

At the design charge the broth holds $65\,\mathrm{m^3}$ at 25 % dissolved oxygen, with $C_{ref} = 6.59\,\mathrm{mg/L}$, $C^* = 9.84\,\mathrm{mg/L}$ at 0.5 barg, an uptake of 39.6 mg/(L·min) and $k_La = 4.84\,\mathrm{min^{-1}}$.

(a) Find the dissolved oxygen concentration C_L and the mass of dissolved oxygen in the vessel. *Checkpoint:* 1.65 mg/L and 107 g.

(b) Find the oxygen the culture consumes per minute and the time in which it would consume the whole dissolved inventory. *Checkpoint:* 2.57 kg/min and 2.5 s.

(c) Find the liquid-phase time constant and check that transfer balances uptake at this state. *Checkpoint:* 12.4 s, and a transfer rate of 39.6 mg/(L·min).

3.5 Theory and calculation · 8 min

The agitator drive has $J = 60\,000\,\mathrm{kg\,m^2}$, $k_T = 13\,481\,\mathrm{N\,m}$ per $(\mathrm{rev/s})^2$ and a rated torque of 30 kN m. At the design point $N = 1.068\,\mathrm{rev/s}$ and $f_g = 0.849$.

(a) Find the open-loop drive time constant from Equation (3.9). *Checkpoint:* 15.4 s.

(b) Find the inner process gain $K_{p,N}$ and the closed-loop time constant of Equation (3.12) at the shipped $K_{c,N} = 6$ and at the detuned $K_{c,N} = 0.35$ of scenario 6. *Checkpoint:* 0.736 rev/min per percent, 2.85 s and 12.3 s.

(c) Find the inner gain at which Equation (3.12) gives the measured break-even time constant of 7.4 s, and say what the equation omits. *Checkpoint:* $K_{c,N} \approx 1.47$.

3.6 Theory and calculation · 6 min

Back pressure raises C^* through Equation (3.4). Take $C_{ref} = 6.59\,\mathrm{mg/L}$, $p_{atm} = 1.013\,25\,\mathrm{bar}$, a setpoint of 25 % and the uptake of 68.2 mg/(L·min) that the culture reaches after the scenario 5 feed step.

(a) Find C^* at 0.5 and 1.5 barg, and the ratio of the two driving forces $C^* - C_L$. *Checkpoint:* 9.84 mg/L, 16.35 mg/L and 1.79.

(b) Find the transfer coefficient that holds the setpoint at each pressure. *Checkpoint:* $8.33\,\mathrm{min^{-1}}$ and $4.64\,\mathrm{min^{-1}}$.

(c) At 0.5 barg the agitator turns at 82.2 rev/min for 191 kW. Use $k_La \propto N^{2.1}$ and a shaft power proportional to N^3 to estimate the speed and power at 1.5 barg, and compare with Table 3.5. *Checkpoint:* About 62.2 rev/min and 83 kW.

3.7 App challenge · 8 min

Load scenario 3, Single-loop DO control on a feed step. The glucose feed steps from 24 to 40 L/min at 4 min. Use the default seed and run each case for 20 min. For each run read the IAE and MV travel tiles of the scorecard and the last row of the exported column `do_max_error_pct`, the largest deviation of the true dissolved oxygen.

(a) Run the scenario as shipped, then set DO proportional gain to 3.2, press Reset and run again. *Checkpoint:* IAE 23.8, peak 11.83 %, travel 25 rev/min. At 3.2: 12.0, 5.18 % and 95 rev/min.

(b) Load scenario 5, Cascade: a speed loop under the DO loop, and repeat both runs. *Checkpoint:* IAE 18.2, peak 9.34 %, travel 18 rev/min. At 3.2: 3.5, 4.82 % and 23 rev/min.

(c) Explain why the cascade improves the IAE by a quarter at the shipped gain and by a factor of more than three at the higher gain.

Accept an IAE within 0.1 and a peak within 0.02 %.

3.8 App challenge · 8 min

Load scenario 6, The same cascade with a slow inner loop. Leave the Speed loop integral time at 1.20 min and the outer tuning as shipped. Use the default seed and run each case for 20 min, pressing Reset between runs.

(a) Set Speed loop gain to 1.0, then 1.5, then 2.0, and read the IAE tile after each run. *Checkpoint:* 30.5, 23.4 and 20.3, within 0.1.

(b) Compute the inner time constant for each gain from Equation (3.12) with $\tau_{drive} = 15.4\,\text{s}$ and $K_{p,N} = 0.736\,\text{rev/min}$ per percent. *Checkpoint:* 8.9 s, 7.3 s and 6.2 s.

(c) The single loop gives an IAE of 23.8 on the same step. Interpolate linearly for the inner time constant at which the cascade breaks even, and compare it with the 12.4 s oxygen holdup. *Checkpoint:* About 7.4 s, 0.6 times the holdup.

3.9 App challenge · 8 min

Load scenario 8, Feedforward from the feed the operator is already setting. The feed steps from 24 to 40 L/min at 4 min. Use the default seed and run each case for 20 min. For each run read the IAE and MV travel tiles and the extremes of the exported column `do_error_pct`, the true dissolved oxygen less the setpoint.

(a) Run the scenario as shipped ($K_{ff} = 1$, $\tau_{ff} = 0.8\,\text{min}$). *Checkpoint:* IAE 5.9, −2.90 % and +1.75 %, travel 18 rev/min.

(b) Set Feedforward lag to 0, press Reset and run again. *Checkpoint:* IAE 15.1, +22.86 % at 4.25 min, travel 38 rev/min.

(c) Press the reload button, set Feedforward gain to 0 and run again. Compare with scenario 5. *Checkpoint:* IAE 18.2, the scenario 5 value.

Accept an IAE within 0.1 and an error within 0.05 %.

3.10 Application area · 4 min

At the design charge, an hour at a dissolved-oxygen setpoint of 10 % costs 72.9 kWh of agitator energy, 25 % costs 86.0 kWh and 50 % costs 118.8 kWh, and all three make the same ethanol to within 0.01 g/L.

(a) Find the extra energy of the 25 % and the 50 % setpoints over the 10 % one. *Checkpoint:* 18 % and 63 %.

(b) The oxygen half-saturation is 5.3 % of air saturation. Give two reasons from the chapter for not running at 10 %.

3.11 Application area · 4 min

In scenario 11 the probe reads 14 percentage points high. The controller holds the reading at 25.0 % while the true dissolved oxygen is 11.0 %. After four hours the batch has made 49.24 g/L of ethanol against 49.25 with an honest probe, for 279 kWh of agitator energy against 325.

(a) Find the energy saved, and explain why the product barely changed. *Checkpoint:* 14 %.

(b) Name two checks outside the loop that would expose the bias, and explain why a better controller makes this fault worse.

3.12 Review the agent · 5 min

An AI assistant is asked to cut wear on the agitator drive and replies as follows. “Detune the speed loop to $K_{c,N} = 0.35$ with $\tau_{I,N} = 1.20$ min. Equation (3.12) gives 12.3 s. The outer process is 12.4 s of liquid plus 20 s of probe, 32 s in all, so the inner loop is still 2.6 times faster than the outer chain, close to the three-to-five rule, and the cascade keeps most of its benefit with gentler torque.” Find the flaw and say how to check it.