INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 8 Batch, Run-to-Run, and Hybrid Control

CHAPTER 11

Switching a Three-Column Antibody Capture Train on Resin of Unknown Capacity

Pages 260–282

Edited by John D. Hedengren

How to cite this chapter. "Switching a Three-Column Antibody Capture Train on Resin of Unknown Capacity," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 260–282.

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

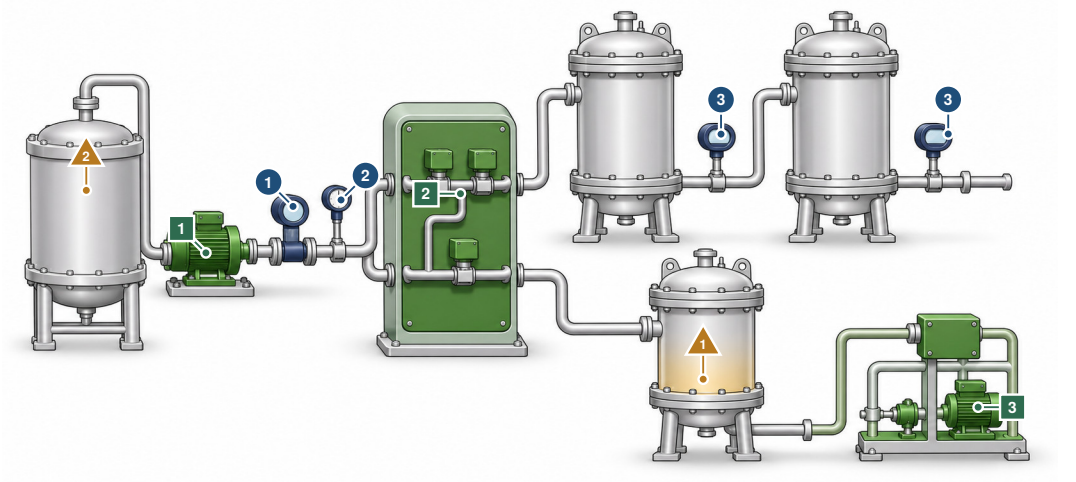
The app for this chapter
https://apmonitor.com/apps/mab-capture/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v08-batch-hybrid/#mab-capture

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

CHAPTER 11

Switching a Three-Column Antibody Capture Train on Resin of Unknown Capacity



The controller moves 1 the feed pump 2 the valve block that switches the columns 3 the pump that regenerates the offline column

The controller reads 1 the feed flow 2 the pressure at the train inlet 3 a detector at each column outlet

What disturbs it 1 the capacity the resin has left, which nothing measures 2 what the harvest brings

Notice that the two columns in the upper row run in series, so what passes the first is caught by the second, while the third is regenerated on its own pump. Illustration, not to scale.

IN PLAIN TERMS Catching a Medicine on Columns That Wear Out

The main idea. A protein medicine is caught from a broth by pumping it through columns of tiny beads that hold it. Two columns run in a row so the second catches what spills past the first, while a third is washed out.

The challenge. Nobody can measure how much a column can still hold, and it shrinks with every use. Switch too early and beads and wash liquid are wasted, too late and medicine leaves the second column.

Why it matters. The rule for when to switch sets how much medicine each column makes per hour and how much is lost.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the history in Section 11.2, the two try-its and the fact in Section 11.6, the fact in Section 11.7, and the key points
Study	40 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Sections 11.6 and 11.7), Technology in Practice (Section 11.8), Limitations (Section 11.9) and the open challenge

The control question A three-column Protein A train takes its lead column off while the next one is still being regenerated, and the resin capacity that decides when to switch is never measured. Which switching rule puts the most antibody through each liter of resin per hour without losing it, and what changes when the resin ages or the regeneration cannot keep up?

Learning objectives

- Explain why breakthrough the lag column catches is not loss, and read both on the scorecard.
- Check a three-column schedule for feasibility from the regeneration recipe and the time to load one column.
- Compute the capacity an aged column offers and the rate at which three columns can take antibody.
- Explain how a filtered UV estimate and a capacity learned from a two-minute average move the breakthrough a policy realizes.
- Choose a switching rule and a response to a late column for resin of uncertain capacity, and justify the choice with the measured trade-offs.

Prerequisites Sections 1.4.2 (guards that end a phase), 1.4.5 (EWMA), 1.4.7 (hybrid automata), 1.4.9 (cyclic steady state) and 1.6.2 (the PID of the two loops), and Chapter 8 for phases ended on a clock and on an integrated estimate.

The app https://apmonitor.com/apps/mab-capture/

See also Volume 19, Chapter 8 for the culture that secretes the antibody, fed on daily samples, and Volume 19, Chapter 9 for a recovery train whose bottleneck vessel limits the harvest it can take. Volume 5, Chapter 3 for control through a slow analyzer and Volume 5, Chapter 7 for a continuous line judged on a delayed at-line measurement.

11.1 Introduction

The lead column of a capture train can break through by 90 % without losing antibody. In the simulated train of this chapter, three Protein A columns of 0.503 L take a clarified harvest at 3.0 g/L. Two of them are in series, so the lag column catches what the lead lets through, while the third is washed, eluted, cleaned and re-equilibrated in 23.2 min. Switching when the lead's outlet reaches 40 % of the feed loses 0.06 % of the antibody, and switching at 90 % loses 0.09 % while the lag catches 78.3 instead of 21.8 g over the campaign. What a deep switch costs is time and capacity downstream, and the limit on it is the resin's capacity, the one number no instrument in the plant reads.

The simulator is built around a misconception that breakthrough measures yield loss. A batch column is loaded to a fraction of its dynamic capacity because whatever leaves it is lost, and the dynamic capacity of Protein A media depends strongly on residence time [1]. Periodic counter-current chromatography removes that penalty by loading the lead into its breakthrough [2], and the decision moves from how far to load to when to switch, under three uncertainties at once. The capacity falls with cycles of use, the titer varies from lot to lot, and the 280 nm signal that reports breakthrough also carries host-cell proteins and media components [3]. A schedule must also be feasible, because with three columns the regeneration has to finish within the loading time [4]. Fixed beds of ion exchanger follow the same flowing-exchange kinetics [5], and any train of regenerable beds run lead and lag shares the decision of when a loaded bed comes off while the next one is still being restored.

Continuous capture is now an established option for antibodies. A four-column system has run uninterrupted for 30 days coupled to a perfusion culture [6], optimal design has compared batch and multi-column processes [7], and a mechanistic model has been used for online control under changing upstream conditions and resin aging [8]. These studies were carried out on plants whose bound antibody, true breakthrough and remaining capacity are known only through samples, so a reader cannot load the same column twice with one decision changed and see where every gram went. This chapter ranks fixed-time, integrated-load, breakthrough-threshold and adaptive-capacity switching by recovery, resin utilization, buffer use, productivity and switching losses when each runs a simulated three-column Protein A train with the same resin, feed, regeneration recipe and information. It then

follows how the ranking moves as the resin's capacity departs from what its recipe assumes. The culture that secretes the antibody upstream of this train is the subject of Volume 19, Chapter 8.

The kinetic Langmuir law, the finite-volume transport, the switching rules and the learned capacity all have antecedents in [2, 3, 5, 8]. What the chapter adds are measurements on one synthetic train. Among them, a 25-minute clock that passes the challenge on new resin loses 8.86 % of the antibody when the resin holds 85 % of its capacity, and at 77.5 % no policy tested meets both the loss limit and the productivity target because the regeneration sets the pace. The simulator is a 380-state model whose antibody balance closes to 1.2×10^{-15} of the gross flux in each of 110 valve states. It is built on a column transport module shared with the peptide purification simulator of Volume 6, Chapter 9, and ten cases each credit one decision.

Section 11.2 follows the domain from the first immunosorbent to a rotating train. Section 11.3 sets out the train: transport and binding in each column, the capacity and its aging, the valves, the surge tank, the pressure and the instruments. Section 11.4 describes the two loops, the breakthrough estimate, the switching triggers and the supervisor that checks feasibility. Section 11.5 reports the code verification and the agreement with the literature, Section 11.6 the ten cases, and Section 11.7 the policies compared under equal resin and feed and over the resin's age. Section 11.8 describes the real detectors, valves and controllers, Section 11.9 states what the model leaves out, and the exercises of Section 11.11 work through the cases.

11.2 From an Immunosorbent to a Rotating Train

Affinity chromatography was introduced under that name in 1968, as a selective purification of enzymes on a ligand bound to the support [9]. Four years later Protein A from *Staphylococcus aureus* was itself isolated by affinity chromatography and used as an immunosorbent for the isolation of immunoglobulins [10]. The molecule became one of the most employed affinity purification systems, and a domain derived from it was improved to withstand the harsh environment of industrial purification [11]. By 2007 a platform process for monoclonal antibodies and Fc fusion proteins, developed on 14 molecules, rested on Protein A as its highly selective capture step [12].

Batch capture made dynamic capacity the design number. A comparison of 15 commercially available Protein A media measured equilibrium and dynamic capacities for human immunoglobulin G and found the dynamic capacity strongly dependent on residence time, with agarose media binding more at residence times above 3 min, and described the breakthrough profiles with the Thomas solution [1]. That solution goes back to a 1944 analysis of heterogeneous ion exchange in a flowing system [5], whose reversible second-order rate is the kinetic Langmuir law that this simulator integrates on a finite-volume grid.

Continuous operation reached capture next. A design method for periodic counter-current chromatography derived the residence time and number of columns from an empirical breakthrough model, verified them on two enzymes and an antibody, and attributed the economics to higher resin capacity utilization and lower buffer consumption [2]. The same group integrated a four-column system with a perfusion bioreactor and reported 30 days without a time-based decline in performance [6]. A decisional tool predicted direct-cost savings near 30 % for early clinical manufacture but near 5 % at commercial scale, and flagged the loss of binding capacity over long campaigns as a facility-fit risk [13]. A review of continuous downstream processing placed chromatography as the most advanced unit operation in continuous mode [14], and an optimal-design study found three-, four- and two-column processes with similar maximum capacity utilization, all well above batch [7].

Control of the load followed the measurements. A mechanistic model of affinity capture with transport into the beads was used both to design a two-column process and to change its operating conditions online, holding yield under changing upstream conditions and resin aging [8]. The measurements that next ended the load, and the guideline that followed, are told in the history box below.

Table 11.1 Variables by Role, With Values From the Reference Configuration

Variable	Role	Reference	Range
Q_f (FT-101)	CV of FIC-101	0.30 L/min	SP 0.10–0.80
u_P	MV, feed pump	31 %	0–100 %
P_{in} (PT-101)	CV of PIC-102	1.4 bar	limit 1.0–3.0 bar
Q_{lim}	MV of PIC-102	inactive	0.02–1.0 L/min
switch instant	discrete MV	policy	5 rules
Q_r	regeneration flow	0.40 L/min	0.1–0.8
$c_{P,h}$	DV, harvest titer	3.0 g/L	1–7
$c_{H,h}$	DV, UV background	1.2 g/L	0–4
r	DV, background drift	0	−50–150 %/h
a_0	DV, resin capacity left	1.0	0.4–1, hidden
$\widehat{BT}$	estimate, lead breakthrough	—	0–100 %
$\hat{L}, \hat{C}$	estimates, load, capacity	—	g/L of bed
c_i, q_i, φ	states, per cell	0 at $t = 0$	300
line tanks	states, 4 lines	feed line full	36
V_s, f_k, pumps, ledgers	scalar states	V_s 25 L	44
P	outcome, g/L/h	—	lead tile

HISTORY 2017 to 2022: The Load Ends on a Measurement

A Protein A load was long fixed before it began. Its volume came from a titer measured offline and from a conservative capacity found in resin-lifetime studies, and an automated chromatograph that ended a load at 1 % breakthrough still carried a sample's delay [3]. A study published in 2017 used a partial least squares model on UV and visible spectra of the effluent to end a Protein A load in real time at a chosen breakthrough, after subtracting an impurity background averaged once the signal had stopped changing [3]. A 2018 study wrote the three-column condition, a regeneration shorter than the loading time, and found the yield lower in the first cycle, whose preload is recovered only after the run [4]. The guideline on continuous manufacturing, adopted in November 2022, illustrates a therapeutic-protein process with continuous capture chromatography, and lists resin lifetime among its run-time considerations [15].

11.3 The Capture Train as Modeled

The simulated route starts at clarified harvest and ends at the eluate pool, as shown in Figure 11.1. Harvest from the clarification arrives in a surge tank, the feed pump P-101 under FIC-101 sends it through FT-101, PT-101 and the feed detector AT-101 to the lead column, the lead's outlet runs into the lag, and the lag's outlet is the flow-through, a loss. The offline column runs the regeneration recipe on its own pump, P-201, and each column's outlet line ends in a valve block that sends it to the next column, the flow-through, the wash waste, the eluate pool or the cleaning and equilibration waste. Low-pH virus inactivation, polishing and formulation are downstream context only. The route is a teaching route proposed from published evidence and does not reproduce any commercial process.

The variables appear by role in Table 11.1. Each column has 20 cells of five states (antibody and background in the liquid and on the resin, and the elution strength), each outlet line and the feed line three well-mixed tanks of three states, and 44 scalars carry the pumps, the surge tank, the destinations' inventories, the buffers and the per-column fouling, for 380 states in all. Engine time is in seconds and the display shows minutes and hours.

11.3.1 Transport in Each Column

Each column is the shared finite-volume module of the series, a packed bed with axial dispersion and a flux inlet, written for a peptide purification simulator and moved into the framework when this simulator became its second user, unchanged in arithmetic. The liquid balance for species i (antibody

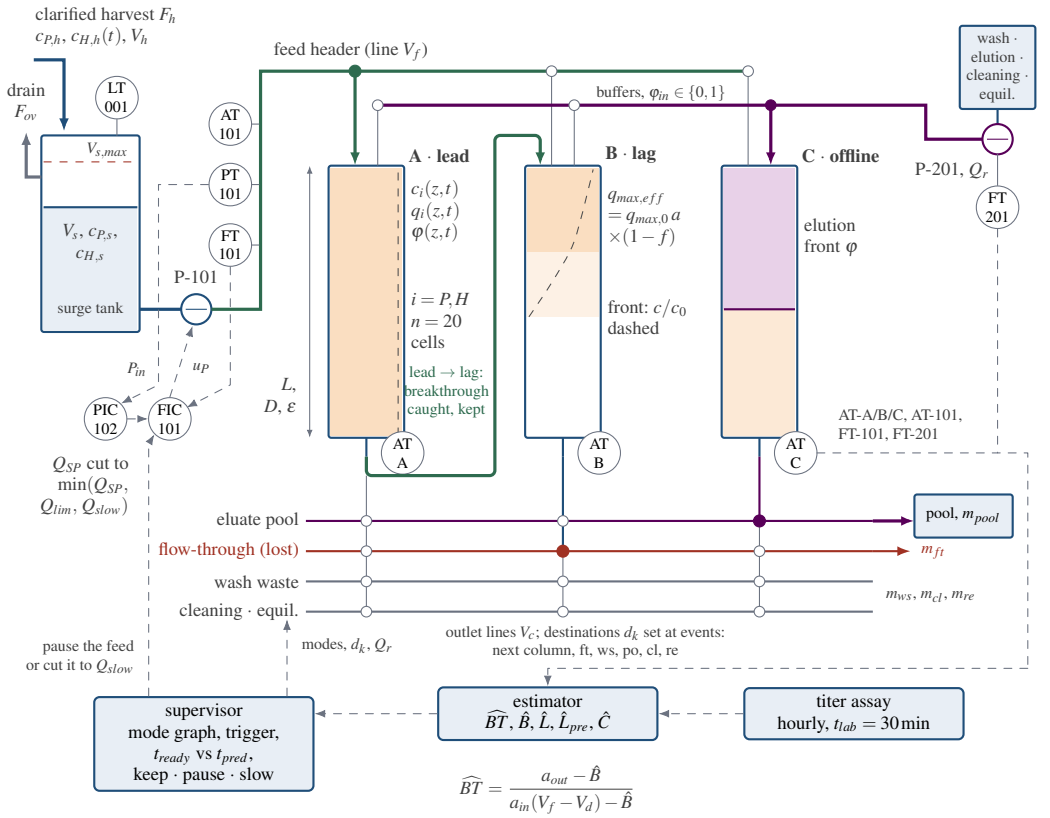


Figure 11.1 Three-column Protein A capture train as modeled, drawn at a switch. Harvest F_h fills the surge tank V_s, P-101 feeds the lead at Q_f under FIC-101 and the pressure override PIC-102, and each column of length L, diameter D and interstitial fraction ε carries $c_i(z,t)$, $q_i(z,t)$ and the elution strength $\varphi(z,t)$ on 20 cells, with a capacity $q_{max,eff}$ set by aging a and fouling f. The lead's outlet enters the lag, so its breakthrough is kept, while the lag's outlet joins the flow-through and is lost. P-201 drives the offline column through the buffers at Q_r, and the outlet manifold sends each line to the destination d_k set at events. The estimator and the supervisor read only AT-101, AT-A, AT-B, AT-C, FT-101, FT-201, PT-101, LT-001 and the titer assay. Process model only, not a piping diagram.

P or background H) and the elution strength φ are given in Equation (11.1).

$$\begin{aligned} \frac{\partial c_i}{\partial t} + \frac{1-\varepsilon}{\varepsilon}\frac{\partial q_i}{\partial t} + v\frac{\partial c_i}{\partial z} &= D_{ax}\frac{\partial^2 c_i}{\partial z^2}, \qquad v = \frac{Q}{\varepsilon A} \\ \frac{\partial \varphi}{\partial t} + v\frac{\partial \varphi}{\partial z} &= D_{ax}\frac{\partial^2 \varphi}{\partial z^2} \\ Qc_{i,in} = Qc_i(0) - \varepsilon A D_{ax}\frac{\partial c_i}{\partial z}\Big|_0, &\qquad \frac{\partial c_i}{\partial z}\Big|_L = 0 \end{aligned} \tag{11.1}$$

where q_i is mass per volume of stationary phase, $A = \pi D^2/4$, and $D_{ax} = D_{ax,ref}\, Q/Q_{ref}$ with $Q_{ref} = 0.30\,\mathrm{L/min}$. The inlet is the Danckwerts flux condition and the outlet a zero gradient [16]. The bed is

cut into $n = 20$ equal cells, and the flux through the face between cells j and $j+1$ is

$$\begin{aligned} F_{j+\frac{1}{2}} &= Q\Big(c_j + \tfrac{1}{2}\psi(c_j - c_{j-1},\, c_{j+1} - c_j)\Big) \\ &\quad - \varepsilon A D_{ax} \frac{c_{j+1} - c_j}{\Delta z} \\ \psi(a,b) &= \frac{2ab}{a+b} \text{ if } ab > 0, \text{ else } 0 \end{aligned} \tag{11.2}$$

the van Leer limited upwind value [17], with $F = Qc_{in}$ at the inlet face and no dispersive term at the outlet face. A cell's liquid changes by the difference of its face fluxes less what its resin takes up, $\varepsilon A \Delta z\, \dot{c}_j = F_{j-\frac{1}{2}} - F_{j+\frac{1}{2}} - (1-\varepsilon) A \Delta z\, \dot{q}_j$, so the column inventory closes to round-off.

11.3.2 Binding, Elution and Capacity

Adsorption follows a kinetic Langmuir law whose constants move with the elution strength, which is 0 in the load, wash and equilibration buffers and 1 in the elution and cleaning buffers, as shown in Equation (11.3).

$$\begin{aligned} \frac{\partial q_P}{\partial t} &= k_{on,0} \mathrm{e}^{-a\varphi} c_P (q_{max,eff} - q_P) - k_{off,0} \Big(\frac{k_{off,el}}{k_{off,0}}\Big)^{\varphi} q_P \\ \frac{\partial q_H}{\partial t} &= k_{on,H} \mathrm{e}^{-a\varphi} c_H (q_{max,H} - q_H) - k_{off,H} 10^{\varphi} q_H \end{aligned} \tag{11.3}$$

where φ is clipped to $[0,1]$ and a negative c is read as zero in the rate. Elution is therefore desorption where the elution front has arrived and transport of what desorbs, never $q = 0$ at a valve. In every elution of eleven audited runs the largest one-step fall of bound antibody is 1.17 % of the bound mass, and at most 0.020 % is left for the cleaning. The background H is one lumped species of host-cell protein and media components, weakly and reversibly bound.

> **✦ DID YOU KNOW? One Hundred Cycles on a Cleaner Feed**
>
> A four-column Protein A capture train was run for 100 cycles on two harvests. The harvest whose clarification also removed host-cell proteins and DNA gave 49 % more productivity than the depth-filtered one, and its eluate held 11 times less host-cell protein [4].

The capacity a column offers is its new capacity times two factors, as shown in Equation (11.4).

$$\begin{aligned} q_{max,eff} &= q_{max,0}\, a_k \left(1 - \min(f_k, 0.9)\right) \\ a_k &\leftarrow a_k (1-\delta) \quad \text{at the end of each cleaning} \\ \frac{\mathrm{d}f_k}{\mathrm{d}t} &= k_f \frac{\dot{m}_{H,k}}{V_{col}}\Big|_{\text{lead, lag}} - \frac{f_k}{\tau_{CIP}}\Big|_{\text{cleaning, hold}} \end{aligned} \tag{11.4}$$

where a_k is the resin's remaining life, starting at a_0, and f_k is fouling grown by the background mass rate $\dot{m}_{H,k}$ that enters column k while it loads. Cleaning removes fouling only while its solution is in contact and never restores a_k. The static capacity of new resin is $(1-\varepsilon) q_{max,0} = 65\,\mathrm{g/L}$ of packed bed.

11.3.3 Valves, Lines, Surge Tank and Pressure

A column's mode decides what feeds it, and its outlet destination d_k is set only at an event boundary. The flows follow the chain from the pumps, as shown in Equation (11.5).

$$\begin{aligned} Q_k &= Q_f[k \text{ lead}] + Q_r[k \text{ regenerating}] + \textstyle\sum_{j:\, d_j = k} Q_j \\ Q_k c_{i,in,k} &= Q_f c_{i,f}[k \text{ lead}] + \textstyle\sum_{j:\, d_j = k} Q_j c_{i,j}^{out} \\ \frac{\mathrm{d}c_{i,m}}{\mathrm{d}t} &= \frac{3Q_k}{V_c}\left(c_{i,m-1} - c_{i,m}\right), \quad m = 1,2,3 \end{aligned} \tag{11.5}$$

where the last line is the outlet line of volume V_c as three tanks, and $c_{i,f}$ leaves the feed line of volume V_f, built the same way. A switch moves valves and nothing else, and the liquid standing in a bed or a line goes wherever its valve now sends it.

The surge tank is well mixed, as shown in Equation (11.6).

$$\begin{aligned} \frac{\mathrm{d}V_s}{\mathrm{d}t} &= F_h - Q_f - F_{ov} \\ \frac{\mathrm{d}(V_s c_{i,s})}{\mathrm{d}t} &= F_h c_{i,h} - (Q_f + F_{ov}) c_{i,s} \\ c_{H,h}(t) &= c_{H,h,0}\,(1 + r t) \\ F_{ov} &= \max(0,\, F_h - Q_f) \end{aligned} \tag{11.6}$$

where F_h is the clarification's delivery until V_h has arrived and the overflow F_{ov} applies only while the tank overflows. A full tank either holds the clarification until it falls to 90 % or overflows to the drain, the learner's choice, and the feed pump stops at 1 L and restarts at 3 L. The pumps and the pressure at the train inlet are given in Equation (11.7).

$$\begin{aligned} \tau_f \frac{\mathrm{d}Q_f}{\mathrm{d}t} &= Q_{pump} \frac{u_P}{100}\left(1 - 0.04\, P_{in}\right) - Q_f \\ \Delta P_k &= \Delta P_{ref} \frac{Q_k}{Q_{ref}} (1 + k_p f_k) \\ P_{in} &= \Delta P_{line} \frac{Q_{lead}}{Q_{ref}} + \Delta P_{lead} + \Delta P_{lag} \end{aligned} \tag{11.7}$$

where P_{in} is in bar, $\tau_f = 4$ s, and P-201 is a stiff positive-displacement pump that delivers its command. The feed stops while P-201 pushes into the loading zone, so the train never carries both flows.

11.3.4 Instruments and the Outcome

Every UV detector reads the Beer-Lambert sum of both species, as shown in Equation (11.8).

$$a = \ell\,(\varepsilon_P c_P + \varepsilon_H c_H) + a_0^{UV} \tag{11.8}$$

through a 2 s lag, first-order colored noise of 0.0025 AU and saturation at 3 AU. FT-101 and FT-201 have 2 s lags and 0.5 % noise, PT-101 0.01 bar of noise and LT-001 0.3 %. The titer assay samples the feed every hour, reports 30 min later with 0.03 g/L of noise and holds its value until the next result. The manufacturing outcome is the antibody in the pool per liter of installed resin per hour of the whole campaign,

$$\begin{aligned} P &= \frac{m_{pool}}{3 V_{col}\, T} \\ m_{in} &= m_{held} + m_{pool} + \textstyle\sum_x m_x \end{aligned} \tag{11.9}$$

where T runs from the checkpoint to the last re-equilibration, m_{held} is the surge tank, lines and columns, and x runs over the flow-through, wash, cleaning, equilibration and drain exits. The ledger of the second line is checked at every step. El-Sabbahy et al. define productivity on the bound mass over the whole run time [4], and this simulator uses the pooled mass instead.

11.3.5 Parameters, Evidence and Initial State

Every parameter appears in Table 11.2 with the label of its source. Two values were read in the anchor papers, the antibody absorptivity at 280 nm [4] and the 2 mm UV path [3]. One, the nominal breakthrough shape of the adaptive policy, was fitted to this model's own fresh column. Everything else is a teaching assumption, and the kinetic constants were chosen only so that the breakthrough moves earlier and flattens as the residence time falls, the trend of El-Sabbahy et al.'s Figure 4, read

Table 11.2 Parameters. R Read in a Cited Source, F Fitted to This Model's Own Column, T Teaching Assumption, H Hidden From the Panel

Symbol	Meaning	Value	Unit	Src
L, D	bed length, diameter	0.10, 0.08	m	T
ε	interstitial fraction	0.35	—	T
$q_{max,0}$	new capacity (solid)	100	g/L	T
$k_{on,0}$	adsorption rate	1.5×10^{-3}	L/(g s)	T
$k_{off,0}$, $k_{off,el}$	desorption, load / elution	2×10^{-6}, 0.03	1/s	T
a	suppression by φ	12	—	T, H
$q_{max,H}$	background capacity	4	g/L	T, H
$k_{on,H}$	background adsorption	3×10^{-3}	L/(g s)	T, H
$k_{off,H}$	background desorption	3×10^{-3}	1/s	T, H
a_0	capacity left at start	1.0 (0.70, 0.85)	—	T
δ	capacity lost per cleaning	0.2 (0.4)	%	T
k_f	fouling gain	2×10^{-3}	L/g	T
τ_{CIP}	cleaning removal time	150	s	T
$D_{ax,ref}$	dispersion at 0.30 L/min	1.5×10^{-6}	m^2/s	T
V_c, V_f	outlet, feed line	0.02 (0.15), 0.05	L	T
ΔP_{ref}, ΔP_{line}	column, lines, 0.30 L/min	0.60, 0.20	bar	T
k_p	fouling on ΔP	2	—	T, H
Q_{pump}	pump capacity	1.0	L/min	T, H
ℓ	UV path length	0.2	cm	R
ε_P	antibody absorptivity	1.46	L/(g cm)	R
ε_H	background absorptivity	1.0	L/(g cm)	T, H
a_0^{UV}	detector zero	0.01	AU	T, H
$c_{P,h}$, $c_{H,h,0}$	titer, background	3.0, 1.2	g/L	T
F_h	harvest delivery	0.40	L/min	T
V_h	harvest still to come	90	L	T
$V_{s,0}$, $V_{s,max}$	surge start, working	25, 80	L	T
recipe	1.5, 3, 4, 2 CV, hold, 4 CV	5 min	0.40 L/min	T
m, s	nominal shape	1.02, 0.11	—	F
limits	loss, inlet pressure	1, 3	%, bar	T

qualitatively with no point digitized [4]. Each case starts from a documented checkpoint, three columns equilibrated in binding buffer with nothing bound, empty outlet lines, the feed line full of harvest, 25 L in the surge tank and the feed pump at its setpoint, with A as lead, B as lag and C ready.

The engine of Section 1.6.1 advances the state by fourth-order Runge-Kutta integration on a fixed step of 0.5 s. At 0.60 L/min the Courant number of a cell is 0.57, and the fastest binding relaxation, $k_{on,0} q_{max,0} (1 - \varepsilon)/\varepsilon$, is about 0.3 s^{-1}, both inside the method's stable range.

11.3.6 Assumptions

Film and pore diffusion in the beads are lumped into the kinetic rates, so the residence-time dependence of the breakthrough is a property of $k_{on,0}$ and not of a pore model, and it fails where intraparticle diffusion dominates. The buffers are one scalar φ, so pH, conductivity and the cleaning chemistry are not resolved, and the cleaning strips antibody exactly as elution does. Aging is a uniform loss of sites in every cell, although real aging also changes the kinetics and varies along the bed. The background is one species with no displacement by the antibody, which would bend the baseline near breakthrough. The outlet lines are three well-mixed tanks, whereas tubing is closer to plug flow. One regeneration pump serves one offline column at a time.

11.4 Loops, Switching Rules and the Supervisor

The controller executes every 2 s on the measurement packet of the instruments above, and the historian logs every 20 s (Section 1.6.1). No loop, estimator or live decision reads a model state.

11.4.1 FIC-101 and the PIC-102 Override

Both loops use the framework's velocity-form PID of Section 1.6.2, with back-calculation at a tracking time equal to τ_I. FIC-101 drives the feed pump from the applied flow setpoint, $\min(Q_{SP}, Q_{lim}, Q_{slow})$, with $K_c = 25\,\%$ per L/min, $\tau_I = 20\,\text{s}$, no derivative and an output of 0 to 100 %. PIC-102 is an override whose output is the flow limit Q_{lim}, with $K_c = 0.15$ L/min per bar, $\tau_I = 15\,\text{s}$ and a range of 0.02 to 1.0 L/min. While the inlet pressure is below its limit of 2.5 bar it tracks the request plus 0.02 L/min, so it cannot wind up. It takes over at 0.1 bar below the limit and lets go once its output exceeds the request by 0.02 L/min with the pressure 0.15 bar under the limit. Manual mode drives the pump by hand while the loop tracks it, so the return to automatic is bumpless (Section 1.6.2).

11.4.2 The Breakthrough Estimate

Total UV is not a product measurement, so the lead's breakthrough is estimated as an aligned ratio of product signals, as shown in Equation (11.10).

$$\begin{aligned} r_k &= \frac{a_{lead}(t) - \hat{B}}{\max\left(0.02,\; a_{in}(V_f - V_d) - \hat{B}\right)} \\ V_d &= \varepsilon V_{col} + V_c \\ \widehat{BT}_k &= \widehat{BT}_{k-1} + \left(1 - \mathrm{e}^{-\Delta t_c/\tau_{bt}}\right)\left(r_k - \widehat{BT}_{k-1}\right) \end{aligned} \tag{11.10}$$

where $a_{in}(V_f - V_d)$ is the feed UV that entered the lead one transit volume earlier ($V_d = 0.196\,\text{L}$), read from a buffer keyed by the totalized FT-101 volume, and $\tau_{bt} = 20\,\text{s}$. The background $\hat{B}$ is the lead outlet's UV captured once that outlet has risen by less than 0.004 AU/min for a minute (or after 10 min). It can be captured once for the campaign, the single-load practice of [3], at each new lead, or tracked through a 60 s filter on the lag's outlet. A plateau that leaves less than half the feed's product signal is breakthrough, not background, and the baseline the column showed as a fresh lag is kept instead. The lag's own outlet ratio against its baseline gives the end-of-train estimate.

11.4.3 Five Triggers and a Bound

The integrated load of the lead since it became lead is $\hat{L} = \sum Q_{f,m} c_{assay} \Delta t_c / V_{col}$, from the flow meter and the last reported titer. The triggers are a fixed time $t \geq t_{sw}$, an integrated load $\hat{L} \geq L^*$, a breakthrough threshold $\widehat{BT} \geq BT^*$ after 2 min of loading with the baseline captured, and the adaptive capacity of Equation (11.11). The fifth is a switch by hand, from a button on the FIC-101 card or the S key. The clock and the integrated load are the clock and estimate guards of Section 1.4.2, which Chapter 8 uses to end its coating layers.

$$\begin{aligned} x(b) &= m + s \ln \frac{b}{1-b} \\ \hat{C} &\leftarrow \hat{C} + \lambda\left(\frac{\hat{L}_{tot}}{x(\overline{BT}_{sw})} - \hat{C}\right) \\ \text{switch when } & \;\hat{L}_{pre} + \hat{L} \geq \hat{C}\, x(BT^*) \end{aligned} \tag{11.11}$$

where $\hat{L}_{pre}$ is the load the column caught as lag, integrated from the previous lead's $\widehat{BT}$ times the measured feed, $\hat{L}_{tot} = \hat{L}_{pre} + \hat{L}$ at the switch, $\overline{BT}_{sw}$ is the mean of $\widehat{BT}$ over the last 2 min before the switch, $\lambda = 0.5$ and $\hat{C}$ starts at 65 g/L. The capacity update is an exponentially weighted moving average of the observed capacities, updated once per switch, the estimator of the run-to-run controller of Section 1.4.5. An observation at either end of the shape's range, below 2 % or above 98 %, is used only as a bound. Because $\overline{BT}_{sw}$ is a two-minute average on a rising front, it reads below the breakthrough at the trigger, so the learned capacity is slightly high and the policy switches deeper than its target, as Section 11.7 measures. The true-breakthrough trigger reads the outlet line's concentration over the feed one transit volume earlier, which no instrument measures. It is a teaching bound, not a policy.

Table 11.3 Code Verification on the Shipped Engine

Check	Result
derivative identity, 11 runs, 110 valve states	worst 1.2×10^{-15} of gross flux
ledger of Eq. (11.9) at the end of each run	$\leq 7.9 \times 10^{-13}$ relative
antibody moved by a switch	0 in every switch
aging after n cycles against $a_0(1-\delta)^n$	within 2.2×10^{-16}
$t_{offline} \leq t_{available}$, hand against measured, 48 cells	agree in 48 of 48
fresh column, 10 % load at 10, 20, 40, 80 cells	49.63, 50.68, 50.94, 51.00 g/L
challenge, suggested, Δt against $\Delta t/2$	38.32 / 38.33 g/L/h, 0.706 / 0.705 %
challenge, suggested, 20 against 40 cells	38.32 / 38.34 g/L/h, 0.706 / 0.693 %
challenge ranking at Δt, $\Delta t/2$, 40 cells	identical
cycle to cycle, last three elutions against three before	0.01–0.69 % in 13 runs
campaign of 11.68 h, headless	1.7 s on one core

11.4.4 The Supervisor and Feasibility

Each column follows the mode graph ready, lag, lead, wash into the loading zone, wash, elution, cleaning, contact hold, re-equilibration and ready again, with a standby state used only at the end of the campaign. The graph is a hybrid automaton in the sense of Section 1.4.7. A step ends on the P-201 totalizer volume or on the hold clock, and a switch happens once, at an event boundary, when the trigger has fired and the third column is ready. The schedule is feasible only if $t_{offline} = \sum V_{step}/Q_r + t_{hold} \leq t_{available}$, 23.2 min for the shipped recipe and 23.6 min with the 20 to 22 s the wash into the train waits for the feed to stop. When the trigger fires early, the supervisor keeps loading the lead (its default), pauses the feed, or slows it ahead of time to $Q_{slow} = V_{rem}/(t_{ready} + 60\,\mathrm{s})$, where V_{rem} is the feed the policy predicts before its trigger and t_{ready} the regeneration still to run. The feed stops at the surge tank's low level, and when the harvest is through, the last two columns are eluted and cleaned in turn.

11.5 Code Verification and Agreement With the Literature

Code verification, agreement with the cited studies and educational evaluation are kept apart here, and only the first is quantitative. No learning study has been carried out, so no gain in learning is claimed. The checks in Table 11.3 run against the engine as shipped. An independent audit computed, at every step boundary of eleven disturbed runs, the rate of change of everything held plus every exit against the harvest inflow, from the raw state slots and the declared geometry. In each of the 110 valve states those runs exercised, the worst residual was 1.2×10^{-15} of the gross flux for the antibody. No antibody entered a column outside lead or lag, the pool received antibody only from an eluting column, and the count of elution steps equaled the cycles and the pool records in every run.

The feasibility condition was tested across titers of 2 to 5 g/L, flows of 0.2 to 0.5 L/min and resin at 1.0, 0.85 and 0.70 of new under integrated-load switching at 45 g/L. A hand prediction, the interval being the 1.88 min wash into the train plus the time the feed takes to bring 45 g/L of bed at the titer, agreed with the measurement on feasibility in all 48 cells, the measured margin lying 0.33 to 1.39 min above the hand value. At the reference flow of 0.30 L/min every column is ready with 4.7 min to spare. Cyclic operation, the cyclic steady state of Section 1.4.9, is shown rather than assumed. In each of 13 runs the mean of the last three elutions lies within 0.01 to 0.69 % of the three before, and successive elutions agree within 1 % from the second to the sixth cycle, after a first cycle that differs because the first lead starts fresh.

Agreement with the literature is qualitative, as shown in Figure 11.2. A fresh column at 0.125 L/min, 0.30 and 0.50, residence times of 4.0, 1.7 and 1.0 min, passes 10 % of its feed at 58.8, 50.7 and 41.1 g/L of bed, while the 50 % point stays at 66.1 g/L. The front moves earlier and flattens as the residence time falls, the trend that El-Sabbahy et al. report for residence times of 0.25 to 4 min [4] and that Hahn et al. report as a strong dependence of dynamic capacity on residence

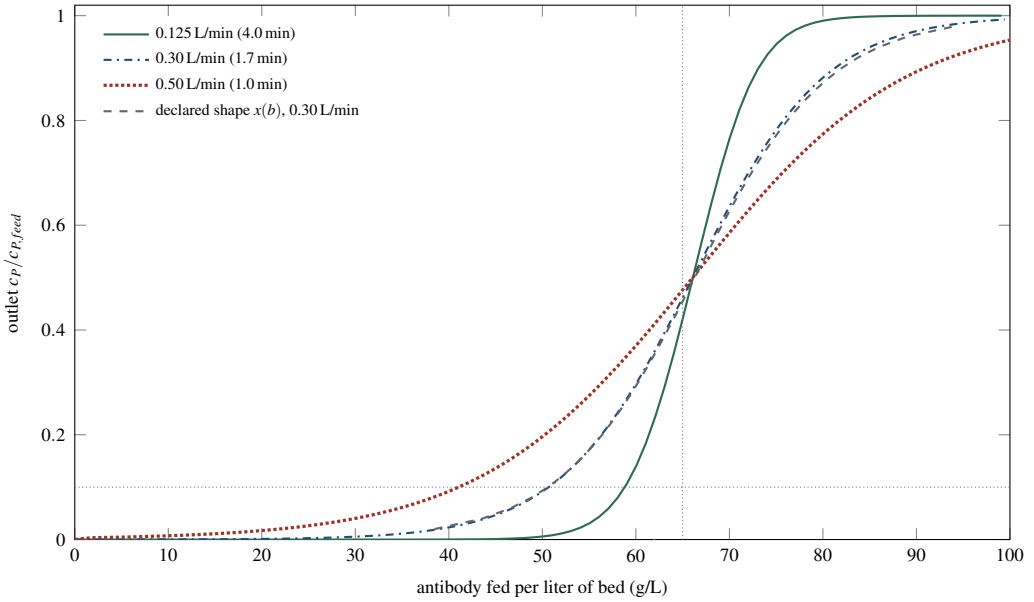


Figure 11.2 Breakthrough of one fresh column at three load flows, 0.125 L/min (solid), 0.30 L/min (dash-dotted) and 0.50 L/min (dotted). The front moves earlier and flattens as the residence time falls, while the 50 % point stays near the static capacity of 65 g/L. The dashed curve is the adaptive policy's declared shape, fitted to this model at 0.30 L/min.

time [1]. No breakthrough data were fitted, so the figure supports the direction of the effect and not its size for any resin.

11.6 Ten Cases, Measured

The cases appear in Table 11.4. All start from the reference configuration (integrated load at 45 g/L, 0.30 L/min, keep loading through a late column, the shipped recipe) and change only what their rows name. Every number in this section was measured headless at the default seed on the engine stepped exactly as in the browser, with any setting a row names entered before Run. Loaded at startup, selected after another case, or replayed after Reset, a case yields the same trajectory to the last logged digit. All 200 numbers that the app's own texts quote reproduce on the shipped engine. At the top playback speed of 900, the longest case, 11.68 h, plays in about 47 s. Each case is credited by a completion rule that requires both the learner's action and the resulting state of the train, so opening a case earns nothing, and a move made by the script never counts as the learner's. The challenge's verdict uses the same function in its scorecard tile, its completion rule and its tests.

11.6.1 Breakthrough Is Not Loss

Case 1 shows a conservative clock. At 0.20 L/min a 30-minute clock loads 31.9 g/L per switch, half the static capacity, and the campaign takes 11.68 h for 19.4 g/L/h, as listed in Table 11.5. The same clock at 0.30 L/min loads 48.1 g/L and makes 27.7 g/L/h, and at 0.40 L/min the lead takes 65.1 g/L and 1.52 % of the antibody leaves past the lag. Case 3 then separates breakthrough from loss. At a threshold of 90 % each lead comes off with 60.4 g/L instead of 54.6, the train switches 11 times instead of 12 and uses 0.278 instead of 0.299 L of buffer per gram, at a cost of 0.03 points of loss, as shown in Figure 11.3. The UV estimate runs about 2 points behind the truth at the trigger, because the 20 s filter sits on a front that climbs about 7 points a minute.

Table 11.4 The Ten Cases and the Decision Each Credits

#	Setup as shipped	Credited decision
1	fixed time 30 min, 0.20 L/min	raise the load flow without loss
2	fixed time 30 min, 0.35 L/min, titer +50 % at 1.5 h	switch on integrated load
3	breakthrough threshold 40 %	raise it to 85 to 95 %
4	load 45 g/L at 0.50 L/min	compare with 0.30 L/min
5	load 45 g/L on resin at $a_0 = 0.70$	switch on the column
6	threshold 20 %, background once, drift 60 %/h	keep the background current
7	0.40 L/min, 15 min cleaning hold, keep loading	pause or slow the feed
8	threshold 90 %, $V_c = 0.15$ L, no wash into the train	wash into the loading zone
9	150 L at 3.5 g/L arriving at 0.60 L/min, $a_0 = 0.85$	keep the antibody, hold upstream
C	120 L at 3.0 g/L, $a_0 = 0.85$, load 55 g/L at 0.45 L/min	meet the target

Table 11.5 Measured Outcomes, Default Seed, Whole Campaigns

#	Run	g/L/h	Lost (%)	Also
1	shipped, 0.20 L/min	19.4	0.03	31.9 g/L per switch, 11.68 h
	0.30 L/min	27.7	0.05	48.1 g/L, 8.18 h
	0.40 L/min	—	1.52	65.1 g/L per switch
2	shipped	35.6	8.57	up to 77.9 g/L per lead
	integrated load 45 g/L	39.1	0.10	waits 29.3 min
3	threshold 40 %	28.8	0.06	lag caught 21.8 g
	threshold 90 %	28.9	0.09	lag caught 78.3 g
4	shipped, 0.50 L/min	42.2	0.40	waits 49.1 min
	0.30 L/min	27.4	0.04	margin 4.7 min
5	shipped	26.8	2.10	lag caught 191.5 g
	threshold 60 %	26.8	0.10	17 switches, not 15
6	shipped (once)	23.1	0.04	20 switches, 0.47 L/g
	each lead	24.4	0.04	14 switches, 0.34 L/g
7	shipped	31.9	9.05	waits 90.4 min, 6.46 h
	pause / slow	23.3 / 22.6	0.05 / 0.04	9.75 / 10.01 h
8	shipped	30.0	3.01	2.95 in the wash
	wash 1.5 CV into the train	28.8	0.12	feed stopped 24.4 min
9	shipped	—	15.5	8.48 h
	pause / slow	—	0.17 / 0.09	held 91 / 121 min
C	shipped	38.4	4.62	fails on loss
	suggested	38.3	0.71	passes, 7.42 h

TRY IT IN THE APP 3 · The Breakthrough Threshold

Load the scenario and press Run. Watch the lead breakthrough climb to its 40 % target on the first trend strip, and the lead come off each time. At the end the scorecard shows 0.06 % of the antibody lost while the lag caught 21.8 g. Press Reset, set Breakthrough target to 90 % and run again: the lag catches 78.3 g and the loss is 0.09 %.

Pause. Try Exercise 11.1 before reading on.

11.6.2 Titer, Residence Time and Resin Age

A clock loads by time. In case 2 the richer lot puts up to 77.9 g/L on each lead against 56 before it and 8.57 % of the antibody is lost. Integrated load at 45 g/L follows the titer once the assay reports it, the first report coming at 2.5 h from the sample drawn at 2.0 h, and loses 0.10 %, but its trigger then comes in 21.6 min, faster than a column turns round, and the rotation waits 29.3 min. In case 4 a load flow of 0.50 L/min reaches the 45 g/L trigger in 15.1 min. With keep loading the leads run on to 58.2 g/L, the lag catches 96.9 g, and the campaign makes 42.2 g/L/h with 0.40 % lost. With the feed paused so that every lead carries only its 45 g/L, the lag still catches 21.6 g at 0.50 L/min against 5.3 g at 0.30, the per-pass cost of a shorter residence time. Case 5 runs the 45 g/L recipe on resin with 70 % of its capacity, 45.5 g/L of static capacity. The lag catches 191.5 g and still passes 1.96 % of the



Figure 11.3 Case 3 with the breakthrough target raised to 90 %, on the Columns view at $t = 2.50$ h, the step before the fourth switch, a composite of the app's header, its plant view and its first two trend strips. (1) Lead A is saturated (bound 58 of 58 g/L, true) and its last cell holds 92 % of the feed concentration, flowing to (2) lag B, whose front is a quarter of the way down (bound 14 of 61 g/L) while its outlet carries nothing. (5) The estimate reads 90 % against a true 91 %, and (3) C has been ready since the last regeneration. AT-101 reads 1.12 AU, AT-A 1.05 AU, AT-B 0.25 AU against a background estimate of 0.26 AU. The lag has caught 28.4 g, the pool holds 88.7 g in 6.0 L, and (4) 0.08 g has gone to the wastes and the flow-through. The trends show three earlier leads climbing to the target with the lag outlet near zero.

antibody, while a 60 % threshold loses 0.10 % at the same 26.8 g/L/h, because the feed flow sets the rate.

11.6.3 Background, Feasibility and Dead Volume

In case 6 the harvest background rises 60 % of its start per hour, reaching 2.2 times its start in the feed, and a baseline captured once reads 26 % for late leads whose true breakthrough is 1 %. The train switches 20 times instead of 14 at 33.8 g/L instead of 47.3, and uses 0.47 instead of 0.34 L/g of buffer, while losing only 0.04 %, because an estimate that reads high errs toward wasted resin. In case 7 a 15 min cleaning hold makes the recipe 33.2 min while the 45 g/L trigger comes after 18.9 min. Keep loading loses 9.05 %, pausing loses 0.05 % in 9.75 h and slowing 0.04 % in 10.01 h. In case 8 the liquid standing in a lead switched at 90 % is nearly feed, and with 0.15 L outlet lines and no wash into the train 3.01 % is lost, 2.95 % of it in the wash. Washing 1.5 column volumes into the new lead cuts that to 0.12 % at the cost of 24.4 min of paused feed.

🕮 **Pause.** Try Exercise 11.3 before reading on.

Table 11.6 Switching Policies on the Challenge Plant With the Feed Paused for a Late Column, Mean Over Six Paired Seeds

Policy	P g/L/h	Rec. %	FT %	Sw. %	Util. %	BT_{sw} %	Buffer L/g	Wait min
time 25 min	37.23	90.42	8.61	0.24	100.2	101.4	0.297	0
load 55 g/L	38.28	94.52	4.53	0.23	99.1	96.4	0.302	0
load 45 g/L	34.74	99.09	0.11	0.07	88.1	48.0	0.347	55.5
threshold 60 %	35.16	99.04	0.18	0.06	92.8	62.4	0.339	43.4
threshold 85 %	38.32	98.57	0.54	0.16	98.2	87.5	0.306	23.6
adaptive 60 %	36.59	98.94	0.26	0.08	94.8	70.9	0.322	38.2
adaptive 85 %	37.98	98.38	0.76	0.14	98.8	91.1	0.307	24.3
bound, true 85 %	36.54	98.79	0.42	0.07	97.4	83.0	0.323	28.1
bound, true 90 %	38.25	98.55	0.57	0.16	98.3	88.5	0.306	23.6

P captured per liter of resin per hour, Rec. recovery, FT flow-through loss, Sw. switching losses, Util. bound over static capacity at the switch (teaching truth), BT_{sw} true breakthrough at the switch, Wait the feed paused for a late column. Across seeds the standard deviation of P is at most 0.01 g/L/h for the clock and the UV-based rules and 0.25 and 0.89 g/L/h for the two load rules, which read the noisy assay.

> **TRY IT IN THE APP 6 · A UV Background That Drifts**
>
> Load the scenario and press Run. Late in the campaign the breakthrough estimate on the first trend strip reads far above the dashed truth, and the train ends with 20 switches and 0.47 L of buffer per gram while losing 0.04 %. Press Reset, set UV background (Advanced) to Each lead and run again: 14 switches and 0.34 L per gram.

11.6.4 A Harvest the Train Cannot Take

Case 9 delivers 150 L at 3.5 g/L at 0.60 L/min, 2.1 g of antibody a minute, onto resin at 85 %. Three columns that each turn round in 23.6 min with about 55 g/L of capacity can take about 1.2 g a minute. As shipped the train keeps loading and loses 15.5 %. Pausing keeps the antibody (0.17 % lost) but holds the clarification for 91 min, slowing holds it 121 min, and overflowing to the drain loses 15.47 %. A logged sweep of 60 settings (load flows from 0.30 to 0.54 L/min, integrated load, adaptive capacity, thresholds of 85 and 95 % and the true-breakthrough bound, with each of the three responses) finds none with the loss within 1 % and the surge tank never full. The closest, an 85 % threshold with the feed paused, loses 0.71 % and still holds the clarification 57 min. With this regeneration recipe the harvest rate is infeasible, and the case credits keeping the antibody and holding upstream, not a rate. The insulin train of Volume 19, Chapter 9 also counts an infeasible harvest among its cases, set by a refolding vessel that is the bottleneck of its train.

11.7 Policies Under Equal Resin and Feed

The policy comparison uses the challenge plant, a 120 L harvest at 3.0 g/L arriving at 0.33 L/min onto resin at $a_0 = 0.85$ that loses 0.4 % per cleaning, fed at 0.45 L/min. Every policy has the same instruments, recipe, feed and resin, and the same response to a late column, Pause the feed, so only the switching rule differs. Each row is the mean over six paired seeds (the default and 1 to 5), whose noise realizations every policy shares. Recovery is the pool over the antibody delivered. Utilization is the bound antibody in the outgoing lead at each switch over its static capacity at that moment, averaged over switches. Switching losses are the wash, cleaning and equilibration wastes, apart from the flow-through.

The measured policies appear in Table 11.6. The clock and the 55 g/L load run the lead past saturation, 100.2 and 99.1 % bound, and lose 8.61 and 4.53 % in the flow-through, so their productivity is high only because the antibody they lose is not slowing them down. The 45 g/L load is safe at 99.09 % recovery but leaves 12 % of the capacity unused, waits 55.5 min for regeneration, and uses 0.347 L of buffer per gram against 0.306 for the 85 % threshold. The 85 % threshold makes 38.32 g/L/h with 98.57 % recovery and 98.2 % utilization, and is the only measured rule that meets the challenge at all six seeds besides the adaptive 85 % rule (37.98 g/L/h, 98.38 %) and the 90 %

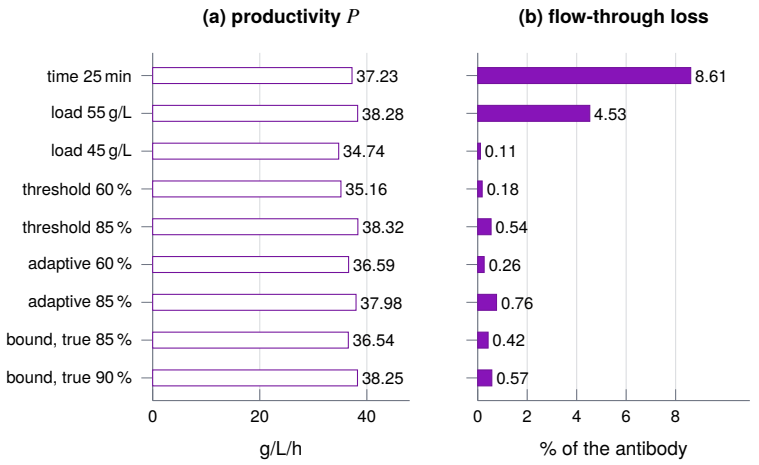


Figure 11.4 The switching policies of Table 11.6 on the challenge plant with the feed paused for a late column, mean over six paired seeds: (a) antibody captured per liter of resin per hour, (b) the share lost in the flow-through. Notice that the clock and the 55 g/L load are not more productive than the 85 % threshold, although they load the lead past saturation and lose far more.

bound. Switching losses stay between 0.06 and 0.24 % for every rule, so the flow-through, not the switch, carries the loss. Figure 11.4 draws the productivity and the flow-through loss of every policy side by side.

> ✦ DID YOU KNOW? **A Saturated Lead Does Not Pay**
>
> From the simulator: on the challenge plant a 25-minute clock takes each lead to 100.2 % of its static capacity and loses 8.61 % in the flow-through, yet makes 37.23 g/L/h, less than the 38.32 g/L/h of an 85 % breakthrough threshold that loses 0.54 % (Table 11.6).

Two departures from intuition follow from the information each rule has. The true-breakthrough bound at 85 % does worse than the UV estimate at 85 % (36.54 g/L/h against 38.32), because the estimate trails the front and so realizes 87.5 %. The bound is a bound only at a matched realized breakthrough, and at 90 % it realizes 88.5 % and matches the estimate within 0.07 g/L/h. The adaptive rule switches deeper than its target, at a true 91.1 % for a target of 85 %, because it learns from the two-minute average described in Section 11.4. Two ablations isolate the parts. With the learning gain at zero the adaptive rule keeps its starting guess of 65 g/L and loses 2.56 % at an 85 % target, against 0.90 % with learning. Dividing by the current rather than the aligned inlet changes nothing on this plant (0.70 % lost either way), because its titer is constant and the surge tank smooths the feed.

🕮 **Pause.** Try Exercise 11.2 before reading on.

The resin-age sweep shows what uncertain capacity changes, as shown in Figure 11.5. On new resin the 25-minute clock passes the challenge with 40.76 g/L/h and 0.28 % lost, and the 55 g/L load passes with 40.19 g/L/h and 0.13 %. At $a_0 = 0.925$ the clock already loses 2.43 %, and at 0.85 only the breakthrough, adaptive and bound rules pass. At $a_0 = 0.775$ none does. The 85 % threshold loses 1.05 % and makes 34.68 g/L/h, the 45 g/L load loses 0.65 % at 33.77 g/L/h, and the threshold's paused time grows from 0 on new resin to 50.7 min because the switches come faster than the regeneration as the capacity shrinks. The best rule therefore changes from a clock on new resin to a measured breakthrough on aged resin, and on still older resin no switching rule recovers the target with this recipe.

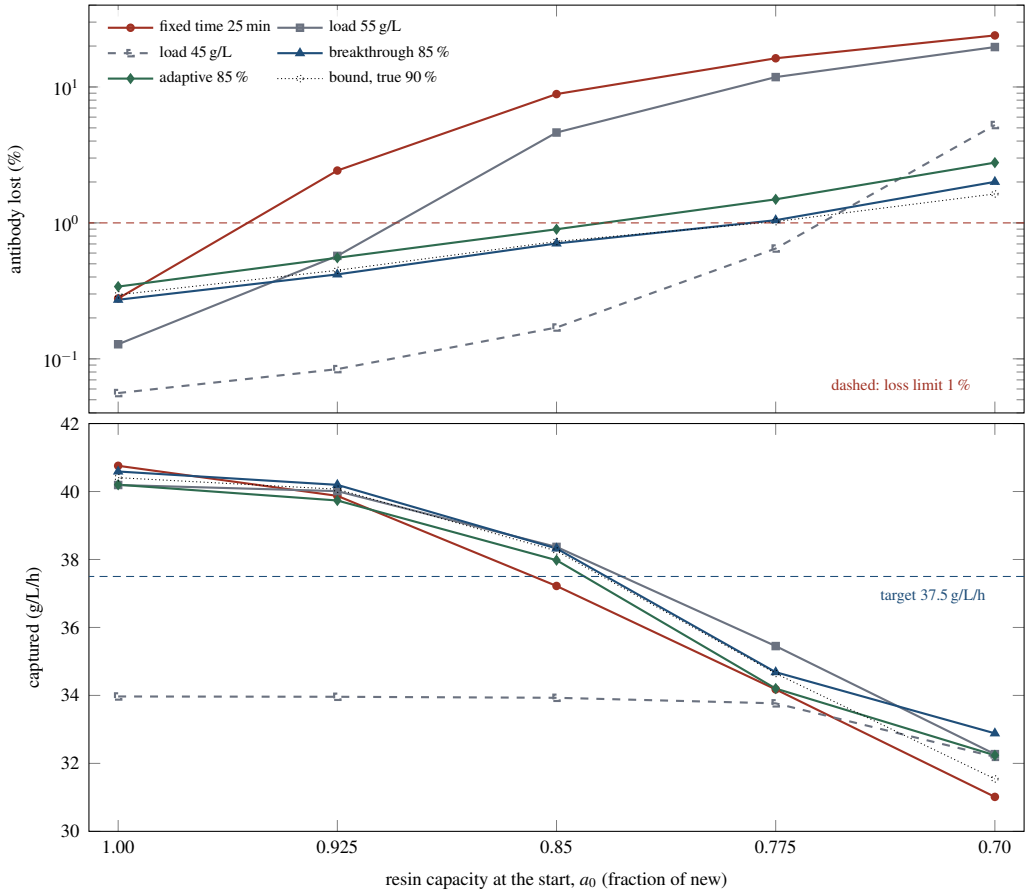


Figure 11.5 Loss and productivity against the resin's starting capacity, the challenge harvest with the feed paused for a late column, default seed. The fixed clock and the 55 g/L load pass only on resin near new. The breakthrough and adaptive rules keep the loss within 1 % down to $a_0 = 0.85$, and at 0.775 the regeneration waits, not the rule, decide the productivity.

11.8 Technology in Practice

What Measures It

The switching decision reads two kinds of measurement, in-line UV detectors around the columns and a titer assay. In-line UV flow cells that monitor therapeutic protein concentration are one of the examples the ICH guideline on continuous manufacturing gives of process analytical technology [15]. The titer was long measured offline for every lot to set the load volume. At-line monitoring carries the risk of human error, contamination, time delays and missing data, and an automated analytical chromatograph is expensive and error-prone, and its delay between sample and result risks a late reaction or forces a slower process [3]. The assay of this chapter, an hourly sample reported 30 min later, is that delay, and case 2 shows the integrated-load rule waiting for it.

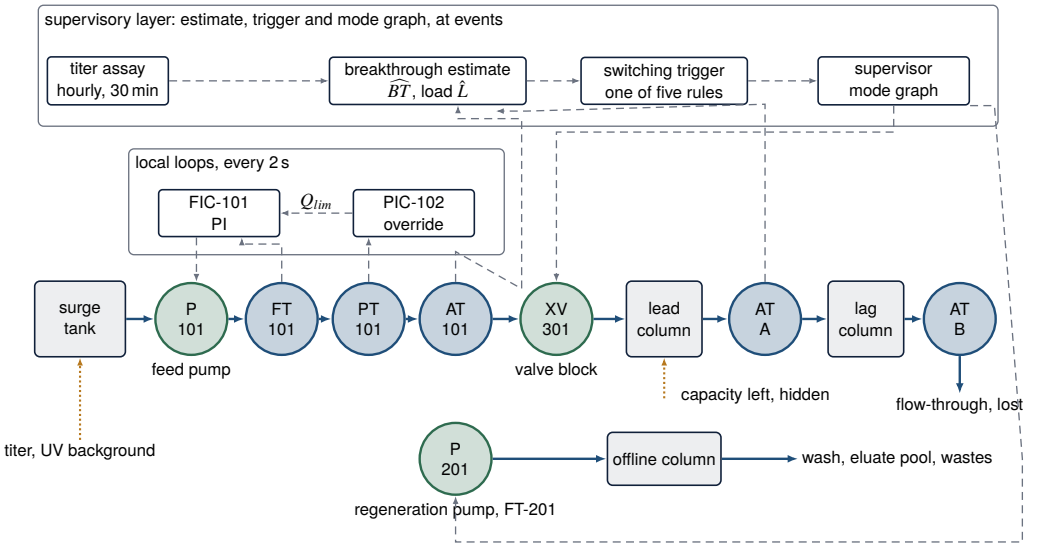


Figure 11.6 The chapter's loops on capture hardware. In the local layer FIC-101 drives the feed pump P-101 toward its flow setpoint, and the PIC-102 override lowers the flow limit when PT-101 nears its limit. In the supervisory layer the UV detectors before and after the lead and the titer assay feed the breakthrough and load estimates, a trigger decides the switch, and the supervisor moves the valve block and runs the regeneration pump P-201 on the offline column. Dotted arrows mark disturbances. XV-301 is the book's tag for the outlet manifold; the others are the chapter's.

TECHNOLOGY **Sensor: Two UV Detectors Around the Lead Column**

The antibody and the impurities of a harvest both absorb at 280 nm, so one wavelength cannot tell them apart. A UV method for continuous Protein A capture places one detector at the column inlet and one at its outlet. While the column loads, the outlet signal settles at an impurity baseline, and once the antibody breaks through it rises above that baseline by an amount that corresponds to the breakthrough level. Switching on that difference loads every column to the same percentage breakthrough whatever the titer or the capacity left. Two detectors can drift unequally, and impurities displaced by the antibody can keep the baseline from settling [3]. One laboratory setup carried a UV monitor with a 2 mm path and a diode array flow cell with a 0.4 mm path [3]. AT-101 and AT-A play the two detectors here.

What Moves It

Three final elements act on the train: the feed pump P-101, the regeneration pump P-201 and the valve block that sets where each column's inlet and outlet lead. Figure 11.6 draws them with the loops that move them. The model gives the feed pump a 4 s lag and a delivery that falls with the inlet pressure, and treats P-201 as a stiff positive-displacement pump that delivers its command.

> **TECHNOLOGY** **Final Element: The Valve Block of the Train**
>
> In a multi-column capture system the pumps feed the inlet ports of a valve block, and the block's software sets the inlets and outlets of each column, including two columns run in series. In one four-column laboratory system five pumps fed four inlet ports. Its steady-state cycle was split into four switches, each moving the feed on to the next column, and each switch into four steps that changed the wash, elution, strip or equilibration positions of the two columns not loading [4]. In this chapter the outlet manifold moves only at event boundaries, so a switch moves valves and no liquid, and case 8 measures what the liquid standing in the bed then costs.

Where the Controller Runs

In one laboratory study the chromatography system's own control software ran the method, and a script on a separate computer read the detector spectra and sent a signal back to end a phase once a set condition was met [3]. At plant scale a distributed control system holds a supervisory level that oversees localized controllers, PLCs serve as local controllers within such a scheme, and the loops of both levels run with cycle times from milliseconds to minutes [18]. The chapter's controller executes every 2 s and its historian logs every 20 s. Read as plant hardware, FIC-101 and PIC-102 belong in the local controller beside the pump, and the breakthrough estimate, the five triggers and the mode graph of the supervisor belong in the supervisory layer, which acts at events.

> **IN PRACTICE** **A State of Control That Is Not a Steady State**
>
> The ICH guideline on continuous manufacturing lets the size of a batch be defined by the quantity of output, the quantity of input or the run time at a defined mass flow rate. A continuous process can be in a state of control while its parameters change within specified ranges without being at steady state, and the plant should notice when they drift or trend inside those ranges and find the root cause. The guideline's own example is a change in the elution profile of a therapeutic-protein process that may be due to resin aging [15]. The cyclic steady state of Section 1.4.9 and the aging of Equation (11.4) are this chapter's versions of both ideas.

What Is Modernizing It

Three developments touch this train. The first is the selective soft sensor: a partial least squares model on UV and visible spectra predicted the antibody in the effluent with a root mean square error of 0.01 mg/mL and ended a load at 0.15 mg/mL, about 5 % breakthrough [3]. The second is new optics. A multi-pass Raman flow cell has monitored the breakthrough of a Protein A column, although advanced chemometrics and a large computational effort were needed to reach a usable model [19]. The third is model-based control: model predictive control has been developed for multi-column solvent-gradient purification, a continuous chromatography process [19], and process analytical technology that detects disturbances in real time makes continuous processes amenable to feedforward and feedback control [15]. Each would give the switching rule of this chapter a better reading than total UV, or a model to act on it.

11.9 Limitations

Every coefficient except the antibody absorptivity and the UV path is a teaching value, no breakthrough or elution data were fitted, and the agreement with the literature is the direction of the residence-time effect only. Productivity, utilization and buffer per gram are properties of this synthetic resin, recipe and column and do not predict any commercial medium. The pore and film resistances are lumped into the kinetic rates, the buffers are one scalar, and aging is uniform in the bed. The disturbance steps of the cases are larger than a plant would show (a 50 % titer step, 60 % more background per hour), but each isolates one cause. The background is one species, so a selective spectral estimate [3] would face a different and smaller error than the total-UV estimate here.

Capture alone is not a purification. Host-cell protein and DNA clearance, aggregates, leached ligand, low-pH virus inactivation and the polishing steps are outside the model, and viral safety of a continuous process follows the principles of ICH Q5A as Q13 restates [15]. The 1 % loss limit, the 3 bar pressure limit and the surge volume are teaching values. Real limits come from the product's specification, the column hardware and the site's validation, and passing in simulation does not release a batch. The adaptive rule's learning on a two-minute average makes it switch about 6 points deeper than its target, and the supervisor's surge-tank hold reads the true volume rather than LT-001. Each result holds for one plant and its seeds, and the six seeds move the UV-based rules by at most 0.01 g/L/h.

OPEN CHALLENGE **Reading the Capacity a Column Has Left**

A Protein A column loses capacity over its cycles, through leaching and degradation of the ligand and through pores and ligands blocked by impurities or product left behind, so batch plants load to a conservative capacity found in resin-lifetime studies [3]. Fluorescence through the bed has shown the fouling of Protein A resin, but a direct correlation with it is difficult [19]. In this chapter the capacity left is hidden from every instrument (Table 11.1), and the adaptive rule learns it from the breakthrough of earlier leads. Progress would be an in-line reading of the capacity a column has left, so that a switching rule could learn from a measurement of the resin itself.

11.10 Conclusion

This chapter has set out a three-column Protein A train as a 380-state simulator whose antibody balance closes to 1.2×10^{-15} of the gross flux in every valve state tested, and whose regeneration must finish within one loading interval. Its central lesson is that breakthrough the lag catches is not loss. Switching at 90 % instead of 40 % raises what the lag catches from 21.8 to 78.3 g for 0.03 points of loss. Under equal resin and feed on resin at 85 % of its capacity, an 85 % UV threshold makes 38.32 g/L/h with 98.57 % recovery, the learned capacity 37.98 g/L/h, a safe load 34.74 g/L/h with 12 % of the capacity unused, and a 25-minute clock loses 8.86 %. On resin at 77.5 % of its capacity no rule tested meets the target, because the regeneration sets the pace. The kinetics are qualitative and every limit is a teaching value. Further work would fit the binding law to measured breakthrough at several flows, add pore diffusion, a selective spectral estimate, a fourth column and the low-pH hold. The simulator shows that a capture schedule is as good as what it knows about the resin and the time the regeneration needs.

KEY POINTS

- Antibody that breaks through the lead column is caught by the lag, so breakthrough the lag catches is not loss.
- Switching at 90 % instead of 40 % raises what the lag catches from 21.8 to 78.3 g for 0.03 points of loss.
- On resin at 85 % of its capacity an 85 % UV threshold makes 38.32 g/L/h with 98.57 % recovery, and a 25-minute clock loses 8.86 %.
- On resin at 77.5 % of its capacity no rule tested meets the target, because the regeneration sets the pace.

Where else this applies. The clock and integrated-load triggers that end a lead here are the guards that end the coating layers of Chapter 8 (Section 11.4).

References

[1] R. Hahn, R. Schlegel, and A. Jungbauer, "Comparison of protein A affinity sorbents," *Journal of Chromatography B*, vol. 790, no. 1–2, pp. 35–51, 2003.

[2] R. Godawat, K. Brower, S. Jain, K. Konstantinov, F. Riske, and V. Warikoo, "Periodic counter-current chromatography – design and operational considerations for integrated and continuous purification of proteins," *Biotechnology Journal*, vol. 7, no. 12, pp. 1496–1508, 2012.

[3] M. Rüdt, N. Brestrich, L. Rolinger, and J. Hubbuch, "Real-time monitoring and control of the load phase of a protein A capture step," *Biotechnology and Bioengineering*, vol. 114, no. 2, pp. 368–373, 2017.

[4] H. El-Sabbahy, D. Ward, O. Ogonah, L. Deakin, G. M. Jellum, and D. G. Bracewell, "The effect of feed quality due to clarification strategy on the design and performance of protein A periodic counter-current chromatography," *Biotechnology Progress*, vol. 34, no. 6, pp. 1380–1392, 2018.

[5] H. C. Thomas, "Heterogeneous ion exchange in a flowing system," *Journal of the American Chemical Society*, vol. 66, no. 10, pp. 1664–1666, 1944.

[6] V. Warikoo, R. Godawat, K. Brower, S. Jain, D. Cummings, E. Simons, T. Johnson, J. Walther, M. Yu, B. Wright, J. McLarty, K. P. Karey, C. Hwang, W. Zhou, F. Riske, and K. Konstantinov, "Integrated continuous production of recombinant therapeutic proteins," *Biotechnology and Bioengineering*, vol. 109, no. 12, pp. 3018–3029, 2012.

[7] D. Baur, M. Angarita, T. Müller-Späth, F. Steinebach, and M. Morbidelli, "Comparison of batch and continuous multi-column protein A capture processes by optimal design," *Biotechnology Journal*, vol. 11, no. 7, pp. 920–931, 2016.

[8] F. Steinebach, M. Angarita, D. J. Karst, T. Müller-Späth, and M. Morbidelli, "Model based adaptive control of a continuous capture process for monoclonal antibodies production," *Journal of Chromatography A*, vol. 1444, pp. 50–56, 2016.

[9] P. Cuatrecasas, M. Wilchek, and C. B. Anfinsen, "Selective enzyme purification by affinity chromatography," *Proceedings of the National Academy of Sciences*, vol. 61, no. 2, pp. 636–643, 1968.

[10] H. Hjelm, K. Hjelm, and J. Sjöquist, "Protein A from *Staphylococcus aureus*. Its isolation by affinity chromatography and its use as an immunosorbent for isolation of immunoglobulins," *FEBS Letters*, vol. 28, no. 1, pp. 73–76, 1972.

[11] S. Hober, K. Nord, and M. Linhult, "Protein A chromatography for antibody purification," *Journal of Chromatography B*, vol. 848, no. 1, pp. 40–47, 2007.

[12] A. A. Shukla, B. Hubbard, T. Tressel, S. Guhan, and D. Low, "Downstream processing of monoclonal antibodies—Application of platform approaches," *Journal of Chromatography B*, vol. 848, no. 1, pp. 28–39, 2007.

[13] J. Pollock, G. Bolton, J. Coffman, S. V. Ho, D. G. Bracewell, and S. S. Farid, "Optimising the design and operation of semi-continuous affinity chromatography for clinical and commercial manufacture," *Journal of Chromatography A*, vol. 1284, pp. 17–27, 2013.

[14] A. Jungbauer, "Continuous downstream processing of biopharmaceuticals," *Trends in Biotechnology*, vol. 31, no. 8, pp. 479–492, 2013.

[15] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use, "ICH harmonised guideline Q13: Continuous manufacturing of drug substances and drug products," Adopted 16 November 2022, 2022. [Online]. Available: https://database.ich.org/sites/default/files/ICH_Q13_Step4_Guideline_2022_1116.pdf

[16] P. V. Danckwerts, "Continuous flow systems," *Chemical Engineering Science*, vol. 2, no. 1, pp. 1–13, 1953.

[17] B. van Leer, "Towards the ultimate conservative difference scheme. II. Monotonicity and conservation combined in a second-order scheme," *Journal of Computational Physics*, vol. 14, no. 4, pp. 361–370, 1974.

[18] 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.

[19] L. Rolinger, M. Rüdt, and J. Hubbuch, "A critical review of recent trends, and a future perspective of optical spectroscopy as PAT in biopharmaceutical downstream processing," *Analytical and Bioanalytical Chemistry*, vol. 412, no. 9, pp. 2047–2064, 2020.

11.11 Exercises

11.1 Concept check · 2 min

In case 3 the breakthrough target of the lead column is raised from 40 % to 90 % of the feed, with nothing else changed. What happens to the share of the antibody the train loses over the campaign?

(a) It rises about in proportion to the extra breakthrough, because breakthrough is antibody leaving the lead

(b) It falls, because a deeper switch means fewer switches and so less antibody sent to the wash

(c) It rises by several points, because the lag saturates once the lead passes most of its feed

(d) It barely changes, because the lag catches what the lead lets through on capacity still fresh

11.2 Concept check · 2 min

On the challenge plant, with the feed paused for a late column, the adaptive capacity rule at a breakthrough target of 85 % takes its leads off at a true breakthrough of 91.1 % on average. What makes it switch deeper than its target?

(a) It learns from a two-minute average of the estimate, which reads below the breakthrough at the trigger

(b) Its starting guess of 65 g/L is the capacity of new resin, and a learning gain of 0.5 never forgets it

(c) The titer assay reports 30 min late, so the integrated load it compares with the capacity runs high

(d) The UV background drifts upward through the campaign, so the estimate reads low on the late leads

11.3 Concept check · 1 min

In case 6 the UV background of the harvest rises through the campaign, and the breakthrough estimate subtracts a background captured once from the first lead. What does the train do?

(a) It switches later than it should and loses antibody in the flow-through

(b) It switches earlier than it should and wastes resin and buffer, with little loss

(c) It switches as before, because the aligned feed UV in the ratio carries the same background

(d) It switches as before, because the end-of-train estimate on the lag corrects the lead's estimate

11.4 Theory and calculation · 6 min

The shipped regeneration recipe pushes 1.5 column volumes into the loading zone, then 3 of wash, 4 of elution, 2 of cleaning, a 5 min contact hold and 4 of re-equilibration, at 0.40 L/min. A column holds 0.5027 L. The feed pauses during the first step, and the integrated-load trigger is set at 45 g/L with a titer of 3.0 g/L.

(a) Find the offline time $t_{offline}$ of one column. *Checkpoint:* 23.2 min.

(b) Find the time FT-101 takes to feed 45 g/L of bed at 0.30, 0.40 and 0.50 L/min. *Checkpoint:* 25.13, 18.85 and 15.08 min.

(c) The time a column has offline is the wash into the train plus the feed time of the next lead. Find the margin at each flow by this hand rule, and the highest load flow it calls feasible. *Checkpoint:* 3.8, −2.5 and −6.3 min, and about 0.35 L/min.

11.5 Theory and calculation · 8 min

New resin offers a static capacity $(1-\varepsilon)q_{max,0} = 65\,\mathrm{g/L}$ of bed, and Equation (11.4) scales it by the remaining life a. A column holds 0.503 L and turns round offline in 23.6 min.

(a) Find the static capacity of resin at $a_0 = 0.85$ and at $a_0 = 0.70$, and the antibody one column holds at $a_0 = 0.85$. *Checkpoint:* 55.25 and 45.5 g/L, and 27.8 g.

(b) Case 9 delivers 150 L of harvest at 3.5 g/L at 0.60 L/min onto resin at $a_0 = 0.85$, and the surge tank starts at 25 L with a working volume of 80 L. Bound the antibody per minute the train can take, and estimate when the surge tank fills if the train runs at that bound. *Checkpoint:* 1.18 g/min against 2.1 g/min, full after about 209 min of the 250 min the harvest takes to arrive.

(c) The adaptive rule starts at $\hat{C} = 65\,\mathrm{g/L}$ with $\lambda = 0.5$. If every observation were 55.25 g/L, after how many switches would $\hat{C}$ be within 1 g/L of it? *Checkpoint:* 4 switches.

11.6 Theory and calculation · 7 min

The breakthrough estimate of Equation (11.10) is filtered with $\tau_{bt} = 20\,\mathrm{s}$ at the 2 s scan, and the adaptive rule learns from its mean over the last two minutes. Assume the front passes the target

climbing steadily at 7 points a minute, as in case 3, and that the nominal shape $x(b) = m + s\ln[b/(1-b)]$ with $m = 1.02$ and $s = 0.11$ holds.

(a) Find the filter coefficient and how far the filtered estimate trails the true breakthrough on the ramp. *Checkpoint:* 0.0952, and about 2.3 points.

(b) The trigger fires when the filtered estimate reaches 85 %. Find the true breakthrough at that instant and the two-minute mean the rule learns from. *Checkpoint:* About 87.3 and 78 %.

(c) Find the factor by which the capacity learned from that switch is too high, taking x at the true breakthrough as right, and the true breakthrough at which a switch at $\hat{C}\,x(0.85)$ then comes. *Checkpoint:* 1.063, and about 92 %.

11.7 App challenge · 8 min

Load scenario 3, The breakthrough threshold, and change no setting. Set the speed to 900, at which one campaign of about 7.9 h plays in about 31 s, and use the default seed. Each run ends when the campaign finishes and the run pauses itself.

(a) Run the scenario as shipped, with the target at 40 %. Read the tiles Antibody lost, Breakthrough the lag caught, the switch count on Regeneration margin, worst switch, and Buffer per gram captured. *Checkpoint:* 0.06 %, 21.8 g, 12 switches, 0.30 L/g.

(b) Press Reset, set the Breakthrough target to 90 % and run again. Read the same four tiles. *Checkpoint:* 0.09 %, 78.3 g, 11 switches, 0.28 L/g.

Accept the tiles exactly as stated. Find how many grams more the lag caught and how many grams more were lost out of the 345 g the harvest delivered, and say what the deeper switch bought.

11.8 App challenge · 8 min

Load scenario 7, The next column is not ready, in which a 15 min cleaning hold makes the offline recipe 33.2 min while the 45 g/L trigger comes after 18.9 min. Set the speed to 900 and use the default seed. Run each case until the campaign finishes and the run pauses itself. For each run read Antibody lost and Captured per liter of resin per hour, and the campaign hours in the tooltip of the Captured tile.

(a) Run the scenario as shipped, with If the next column is not ready on Keep loading. *Checkpoint:* 9.05 %, 31.9 g/L/h, 6.46 h.

(b) Press Reset, set If the next column is not ready to Pause the feed and run again. *Checkpoint:* 0.05 %, 23.3 g/L/h, 9.75 h.

(c) Press Reset, set it to Slow the feed and run again. *Checkpoint:* 0.04 %, 22.6 g/L/h, 10.01 h.

Accept the tiles exactly as stated. Find the antibody that pausing saves against keeping on loading, out of the 345 g delivered, and the hours it costs.

11.9 App challenge · 9 min

Load the challenge, The most antibody per liter of resin per hour. Its target is at least 37.5 g/L/h with at most 1 % lost, the inlet under 3 bar and the surge tank never full. Set the speed to 900, use the default seed and do not change the model or the disturbances. Run each case until the campaign finishes and the run pauses itself, and read Captured per liter of resin per hour, Antibody lost and Challenge target.

(a) Run the challenge as shipped. *Checkpoint:* 38.4 g/L/h, 4.62 %, not met.

(b) Press Reset, set the Switching strategy to Breakthrough threshold, the Breakthrough target to 85 % and If the next column is not ready to Pause the feed, and run again. *Checkpoint:* 38.3 g/L/h, 0.71 %, met.

(c) Press Reset, set If the next column is not ready back to Keep loading, leaving the threshold at 85 %, and run again. *Checkpoint:* 38.7 g/L/h, 3.95 %, not met.

Accept the tiles exactly as stated. Say which decision fixed the loss.

11.10 Application area · 3 min

A plant manager prefers the integrated load at 45 g/L as the safe recipe for the challenge plant, against the 85 % breakthrough threshold. Use the policy comparison of Table 11.6.

(a) Find the extra buffer per kilogram of captured antibody that the 45 g/L load uses, as liters and as a share. *Checkpoint:* 41 L, 13 % more.

(b) Find the resin-hours each rule needs per kilogram captured, and the share of extra resin the 45 g/L load needs for the same output. *Checkpoint:* 28.8 against 26.1 liter-hours, about 10 % more.

11.11 Application area · 3 min

In case 8 the train switches 11 times at a 90 % breakthrough target, with outlet lines of 0.15 L and no wash into the loading zone, and 345 g of antibody enters the train. A column of 0.503 L has an interstitial fraction of 0.35, and the liquid in a lead at such a switch is close to the feed's 3.0 g/L.

(a) Estimate the antibody standing in the bed and in its outlet line at one switch. *Checkpoint:* About 0.53 and 0.45 g.

(b) Estimate the share of the campaign's antibody that goes to the wash over all switches, compare it with the measured loss, and say how plants recover it. *Checkpoint:* About 3.1 % against 2.95 % in the wash.

11.12 Review the agent · 5 min

An AI assistant comparing switching rules on the challenge plant writes: "The true-breakthrough trigger knows the lead's real breakthrough, which no instrument measures, so it is the best any rule can do. Set it at 85 % with the feed paused for a late column to find the ceiling. A UV threshold at the same 85 % only estimates the breakthrough, so it must make less antibody per liter of resin per hour." Find the flaw and say how you would check it.