INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 6 Multivariable and Plantwide Control

CHAPTER 9

GLP-1 Peptide Purification Where a Chromatography Cut Is Judged at the Filter

Pages 207–228

Edited by John D. Hedengren

How to cite this chapter. "GLP-1 Peptide Purification Where a Chromatography Cut Is Judged at the Filter," in *Multivariable and Plantwide Control*, John D. Hedengren, Ed., vol. 6 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 207–228.

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

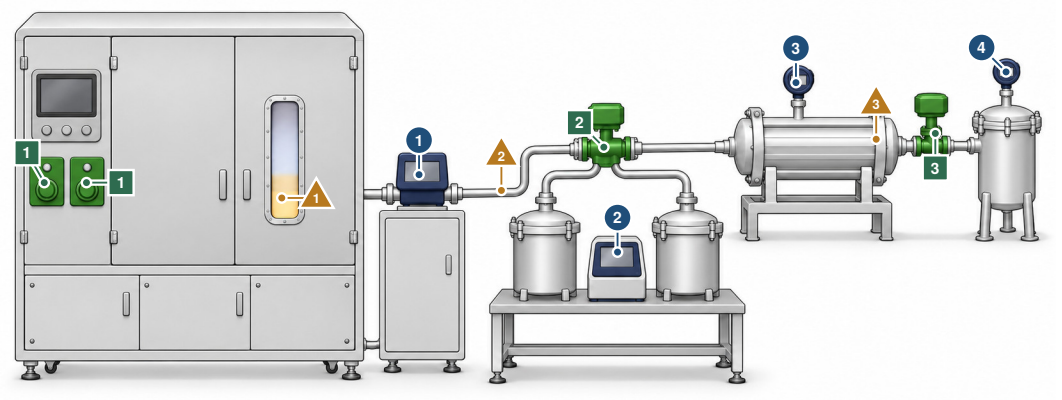
The app for this chapter
https://apmonitor.com/apps/glp1-manufacturing/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v06-multivariable-plantwide/#glp1-manufacturing

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

CHAPTER 9

GLP-1 Peptide Purification Where a Chromatography Cut Is Judged at the Filter



The controller moves 1 column pump and solvent blend 2 cut valve to pool, side cut or waste 3 membrane retentate valve

The controller reads 1 UV detector at the column outlet 2 laboratory assay of the pool 3 membrane pressure 4 filter pressure

What disturbs it 1 crude purity and its late variant 2 extra line before the cut valve 3 membrane passes or fouls more

The cut valve sits a line's length after the detector, and what it lets into the pool decides, hours later, how fast the cake on the filter drains. Illustration, not to scale.

IN PLAIN TERMS **One Cut Decides What the Filter Gets**

The main idea. This chapter follows a small protein drug through the steps that clean it. A tube of fine grains lets each part come out at its own time, and a valve keeps the middle part.

The challenge. A bit more of the flow can still pass the purity test, yet an impurity that comes out last makes the solid grains small. Small grains pack into a cake that drains slowly, and the batch runs out of filter time.

Why it matters. A choice made at the first step has to be judged by what the last step needs.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 9.2, the try-its and the two facts in Section 9.6, and the key points at the end of Section 9.9
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 (Section 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the open challenge

The control question A chromatography cut is usually judged by the purity of its pool, but the filter two unit operations later answers to one late-eluting variant that the purity does not single out. Which variable should the cut hold, and what does a collection rule driven by a soft sensor add over a fixed window?

Learning objectives

- Explain why a pool that passes its purity assay can lose the batch on the filter clock, and name the variable that separates the cuts that pass from those that fail.
- Compute retention factors, selectivity, diafiltration losses and the split of a cake resistance from the chapter's model equations.
- Correct a cut set in minutes for an unmeasured holdup between the detector and the valve.
- Predict what a higher filter pressure buys on a compressible cake.
- Judge a soft-sensor collection rule against a fixed window on the same crude and on a crude outside its calibration.

Prerequisites Sections 1.4.2 (choice of controlled variable), 1.4.8 (disturbance propagation), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios, scorecard and export).

The app https://apmonitor.com/apps/glp1-manufacturing/

See also Volume 8, Chapter 11 for a finite-volume column switched on a capacity nobody measures, and Volume 8, Chapter 4 for a crystal population carried in its leading moments. Volume 5, Chapter 7 for material tracked in volume to a diversion valve, and Volume 5, Chapter 10 for a soft sensor whose bias a delayed laboratory result corrects.

9.1 Introduction

The purity assay of a chromatography pool answers one question about the material, how much of it is product. In the simulated campaign of this chapter, a cut from 23 to 31 min of a 60 min gradient pools 61.7 g of peptide at 98.92 % purity, and 56.2 g of it leaves the filter as acceptable product, 9.9 g/h of campaign time. Shutting the cut 0.75 min later adds 4.4 g to the pool, and the laboratory assays it at 98.35 %, still inside the 98.5 % limit of the isolated solid. The batch is nevertheless lost. The extra minute carried the pool's late-eluting variant from 0.358 to 0.939 %, that variant slows the growth of the precipitating solid, the mean particle size falls from 11.6 to 9.0 µm, and the cake on the filter needs 79.3 min against the 60 min its schedule allows.

The simulator is built around a belief that this case contradicts, that a chromatography cut is right when its pool meets the purity limit. The pool assay is a sufficient statistic only when the next unit operation responds to the total impurity and not to which impurity it is. A cooling crystallizer whose impurity rejection depends on the impurity, a filter whose cake depends on the particle habit, and a dryer whose endpoint depends on the solvent the cake occluded have the same shape. The filtration of an active pharmaceutical ingredient (API) can become the bottleneck of its whole process train, and its specific cake resistance is estimated from laboratory filtrations before scale-up [1]. Gradient purification of biomolecules needs a central cut between weakly and strongly adsorbing impurities, and a single batch step often cannot purify satisfactorily [2]. Kinetic models describe how an impurity in solution slows crystal growth [3], which is the link from the first problem to the second.

However, these effects are usually studied one unit at a time. Preparative chromatography models are fitted and optimized against the band and the pool [4, 5], and filtration studies estimate the cake a given slurry makes [1]. This chapter sets side by side, in a simulated campaign whose column, membrane and filter process the same material, local purity-maximizing cuts, recovery-maximizing cuts and a downstream-aware collection rule, and reports what each costs in membrane loss, solvent use, filtration time and isolated mass. In the terms of Section 1.4.2, the question is which variable the cut should hold. The finite-volume transport, the competitive isotherm, the population balance, the compressible-cake law and the velocity-form proportional-integral (PI) loops are established methods, and every comparison below is measured on a single declared teaching model. The model has 677 states, and its product ledger closes to 2×10^{-11} of the charge in every case measured. Controllers and collection rules read instruments only.

Section 9.2 follows the domain from peptide synthesis to isolation. Section 9.3 gives the route and the model, with the equations and defaults from which the simulator can be rebuilt, and Section 9.4 the recipe, the loops and the three collection rules. Section 9.5 separates code verification from agreement with sources. Section 9.6 reports the nine cases, the challenge and the comparison of cut policies, Section 9.7 describes the real detectors, pumps and controllers, Section 9.8 states the limitations, and Section 9.9 draws the conclusions. The cases and the challenge carry completion rules, and the exercises of Section 9.10 are built on them.

9.2 Peptides on Resin, Peptides on Silica

Peptide manufacture by synthesis begins with the solid-phase method of 1963, in which the chain is assembled step by step on an insoluble resin and the excess reagents of each coupling are washed away [6]. Every coupling that falls short leaves a deletion sequence and every side reaction a related variant, so a crude peptide of thirty residues is a family of molecules rather than one. GLP-1, the hormone behind this class, has its own history, told in the box below. Drug design then turned the native hormone into acylated analogues that bind serum albumin, and semaglutide was selected as a once-weekly candidate with two amino acid substitutions and a derivatized lysine at position 26 [7].

> ⧖ **HISTORY** **Early 1980s to 2012: From a Gene Sequence to a Weekly Medicine**
>
> In the early 1980s the proglucagon sequences of the anglerfish and then of the rat, hamster, cow and human showed that glucagon is cut from a larger precursor. The mammalian precursors hold glucagon and two related peptides, named GLP-1 and GLP-2 [8, 9]. The work of the following three decades established GLP-1 as an insulin-stimulating glucoregulatory hormone. Injected, it proved disappointingly ineffective: an enzyme cleaves it, and in people with diabetes its half-life in the circulation is only 1.5 to 2 minutes. Analogues that resist the enzyme last longer. The first receptor agonist approved for clinical use, a synthetic form of a peptide isolated from lizard venom, reached approval for twice-daily injection in 2005, and a once-weekly form followed in 2012 [9].

Purification grew up with high-performance liquid chromatography, whose gradient elution was given a theoretical treatment in 1979 [10]. Preparative work rests on a theory of nonlinear, nonideal chromatography whose models are used to optimize operating conditions, and by 2002 preparative chromatography was becoming an established purification method in the pharmaceutical industry [4]. For peptides the dependence on the modifier is steep. A heptadecapeptide's retention factor changed by more than 400 % between 16 and 19 % acetonitrile, and its overloaded gradient band was modeled with reasonable accuracy by a Langmuir isotherm whose parameters follow the linear solvent strength model [5]. Because the product elutes between weakly and strongly adsorbing impurities, a batch gradient run needs a central cut, and multicolumn countercurrent solvent gradient purification was introduced to recycle the impure side streams internally [2].

The public record of GLP-1 manufacture is mostly in patents, which disclose process options rather than validate a train. One describes multidimensional reversed-phase purification of a GLP-1 analogue, desalting by tangential-flow filtration (TFF) on a membrane of 1 kDa cutoff, and the propensity of glucagon-like peptides to aggregate at acidic pH [11]. A second purifies semaglutide in three reversed-phase passes from a crude of about 60 to 70 % purity and names lyophilization, spray drying and antisolvent addition among the isolation options [12]. A third, for a dual GLP-1R/GIPR agonist, exchanges acetonitrile for 2-propanol/water by diafiltration over 4 to 8 diavolumes and isolates the peptide by antisolvent precipitation followed by filtration and drying, by spray drying or by freeze drying [13]. The volumes involved are large. A perspective written for an industry roundtable team describes peptide manufacture as dependent on large amounts of hazardous reagents and solvents and names greener processes a critical unmet need [14].

The isolation end has older foundations. General equations for cake filtration were derived in 1935 [15], membrane processes became a reviewed technology of bioprocessing in their own

Table 9.1 Variables by Role, With Defaults and Ranges

Symbol	Role	Default	Range
Q	CV, column flow (FIC-201)	1.5 L/min	0.3–2.5
φ_{in}	CV, modifier at the inlet (AIC-202)	program	10–70 %
pH	CV, conditioning (AIC-401)	7.3	3.5–8.5
TMP	CV, membrane (PIC-501)	2 bar	0.5–6
ΔP	CV, filter (PIC-601)	1.5 bar	0.2–4
u_P, u_B	MV, feed pump, B valve	—	0–100 %
u_{pH}	MV, acid (−) or base (+)	—	±100 %
u_V, u_G	MV, retentate valve, gas regulator	—	0–100 %
t_1, t_2	decision, cut opens, shuts	23, 31 min	10–75
rule	decision, times, UV or downstream aware	times	—
$\hat{x}_{L,max}$	decision, shut limit of the soft sensor	0.80 %	0.1–3
m_{load}	decision, crude charged	120 g	20–260
N_D	decision, diavolumes	4	0–12
F_s	decision, antisolvent rate	0.06 L/min	0.01–0.5
x_c, x_ℓ	DV, crude purity, late share	70, 23 %	45–90, 5–60
ΔV_c	DV, extra collector holdup	0 L	0–3
$s_\times$, $f_\times$	DV, sieving, fouling multiples	1, 1	0.5–12, 0–6
c_i, q_i, φ	states, 60 cells	—	—
V, Z, C_T, M_k	states, vessel	—	—
M_c, V_{cake}	states, filter	—	—

right [16], and the population balance gave particle technology a way to carry a size distribution through a dynamic model [17]. What none of these supply is a peptide's cake resistance, and no source consulted for this chapter reports one. The route below is therefore a proposed teaching route assembled from these disclosed options, and every filtration coefficient in it is a teaching value.

9.3 The Route as Modeled

One campaign charges 120 g of crude at 70 % peptide, dissolved in 3 L, onto a column 0.20 m across and 0.25 m long (7.85 L of bed), as shown in Figure 9.1. A gradient from 22 to 45 % organic over 60 min at 1.5 L/min elutes the bands past a UV detector to the cut valve XV-203, which routes the stream to the pool, a side cut or waste. The pool is diluted with 4 L of buffer in the conditioning vessel T-401, set to pH 7.3, concentrated on a 1 m^2 membrane to 3 L, diafiltered with four diavolumes of 60 % 2-propanol in water, and concentrated again to 100 g/L. Antisolvent dosed at 0.06 L/min to a 60 % fraction precipitates the peptide in R-701, and an agitated Nutsche filter of 0.05 m^2 forms, washes and dries the cake at 1.5 bar. The train is a teaching assembly of the options in Section 9.2, not a reproduction of any manufacturer's process, and its quality scores describe simulated attributes, not batch release.

The variables are classified in Table 9.1. The model carries 677 states, 60 column cells of four mobile concentrations, four stationary loads and the modifier, two holdup lines of five and four tanks, the pump and mixer, three collection vessels, the vessel's dissolved species, solvents, buffer, solids, four moments, membrane adsorption and fouling, the permeate, the filter with its free liquor, suspended solids, cake and cake liquor, and seven ledgers. Physics is in SI units, and the display shows minutes, liters, grams, g/L and bar.

9.3.1 The Column and Its Binding Law

The column is one-dimensional transport with axial dispersion and a lumped kinetic exchange for each species i (P, E, L and the unretained smalls X), written on the mobile-volume basis, as shown in

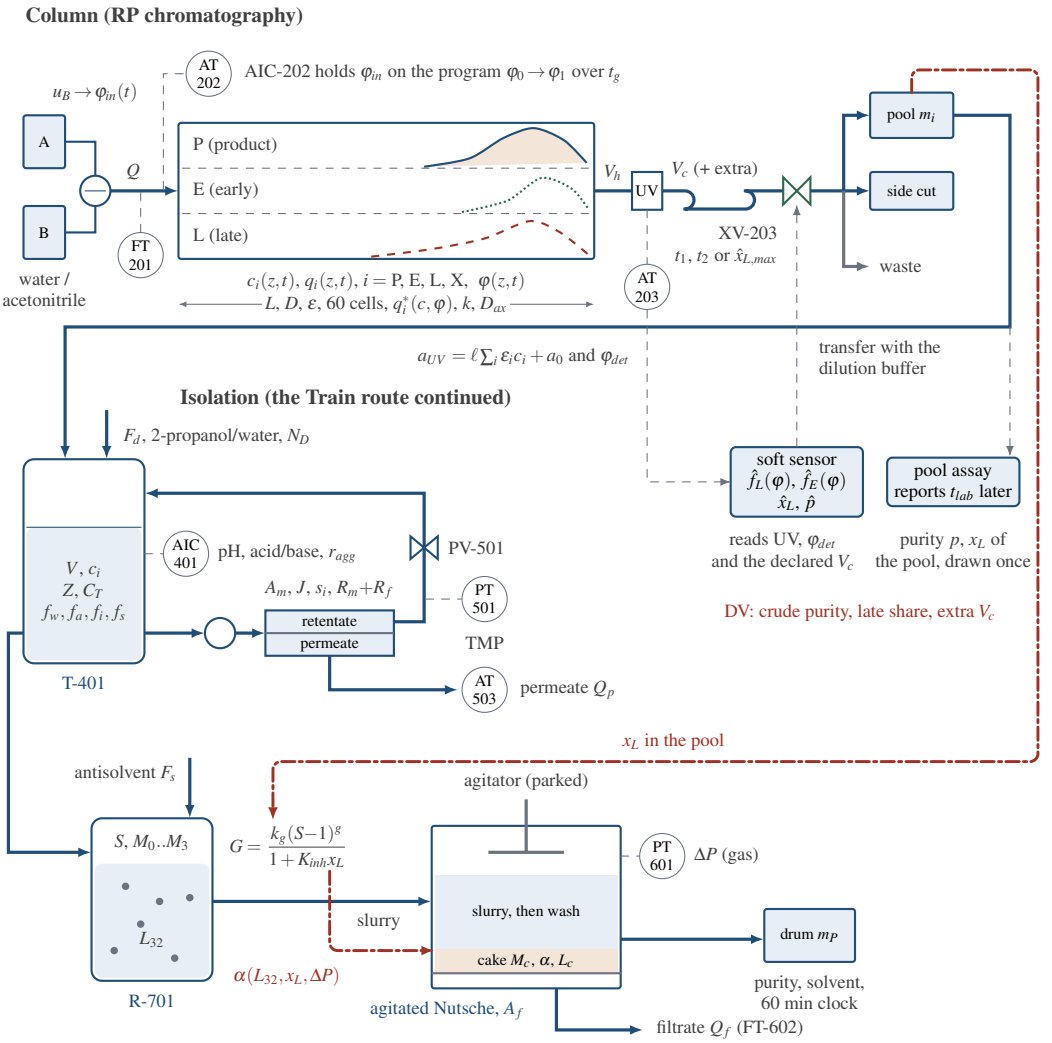


Figure 9.1 GLP-1 purification and isolation as modeled, the Column view's rig continued along the Train route. The gradient pump blends A and B to $\varphi_{in}(t)$ at flow Q (FT-201, AT-202), the bed of length L and diameter D carries the bands $c_i(z,t)$ and loads $q_i(z,t)$ of product P, early variant E and late variant L, and the stream passes the holdup V_h, the UV cell AT-203 and the collector holdup V_c to XV-203, which cuts on t_1 and t_2 or on the soft sensor's $\hat{x}_L$. The pool's assay reports after t_{lab}. T-401 holds V, c_i, the strong ion Z and buffer C_T, and recirculates through the membrane (A_m, J, s_i) under PV-501 and PT-501 with permeate UV AT-503 and diafiltration buffer F_d. R-701 precipitates with antisolvent F_s, and the Nutsche at ΔP (PT-601) holds the cake M_c of resistance α. The thick dash-dotted line (red) is the coupling this chapter studies, the late variant x_L slowing growth G, which sets the Sauter size L_{32} and through it α. In the bed the product band is solid, the early variant dotted and the late variant dashed. Process model only, not to scale.

Equation (9.1).

$$\begin{aligned}
&\frac{\partial c_i}{\partial t}+\frac{1-\varepsilon}{\varepsilon}\frac{\partial q_i}{\partial t}+v\frac{\partial c_i}{\partial z}=D_{ax}\frac{\partial^2 c_i}{\partial z^2}, \qquad v=\frac{Q}{\varepsilon A_c}\\
&\frac{\partial q_i}{\partial t}=k_i\left(q_i^{*}(c,\varphi)-q_i\right), \qquad \frac{\partial \varphi}{\partial t}+v\frac{\partial \varphi}{\partial z}=D_{ax}\frac{\partial^2 \varphi}{\partial z^2}\\
&v c_{i,in}=v c_i(0)-D_{ax}\frac{\partial c_i}{\partial z}(0), \qquad \frac{\partial c_i}{\partial z}(L)=0
\end{aligned} \tag{9.1}$$

where c_i is mass per mobile volume, q_i mass per stationary volume, $\varepsilon = 0.60$ the effective mobile fraction with pore liquid included, A_c the bed area and $k_i = k$ except $k_X = \max(k, 0.2\,\mathrm{s}^{-1})$. The modifier φ travels with the same transport and does not adsorb, so each cell sees the gradient when the liquid carrying it arrives. The isotherm is a competitive Langmuir with one capacity whose affinities fall exponentially with the modifier, as shown in Equation (9.2).

$$q_i^{*}=\frac{q_{max} b_i c_i}{1+\sum_j b_j c_j}, \qquad b_i(\varphi)=b_{i,ref}\,\mathrm{e}^{-S_i(\varphi-\varphi_{ref})} \tag{9.2}$$

At vanishing load the retention factor is $k_i' = q_{max} b_i (1-\varepsilon)/\varepsilon$, which is 10.0 for the product at the reference 30 % organic and 1215 at the 22 % the charge is loaded at. The bed saturates near 157 g of peptide. The ratio $b_L/b_P = 1.73$ at φ_{ref} is the selectivity that separates the product from its late variant. The charge enters at $c_{i,in} = m_{load} f_i / V_{feed}$ during the load and zero after it, where the crude fractions are $f_P = x_c$, $f_E = 0.27(1-x_c)$, $f_L = x_\ell(1-x_c)$ and f_X the rest. The column starts with $c = q = 0$ at φ_0.

The equations are discretized as 60 equal finite volumes with a van Leer limited advective flux, the Danckwerts flux at the inlet face and a zero gradient at the outlet. At the loading modifier the exchange relaxes in milliseconds while the band moves over minutes, so the exchange is solved in closed form per cell, with the competition frozen on a midpoint secant, and split Strang-wise around a fourth-order Runge-Kutta step of the transport (half an exchange, the transport, half an exchange). The step conserves each cell's inventory to round-off for any step length. The column outlet passes a holdup $V_h = 0.4\,\mathrm{L}$ of five well-mixed tanks to the detector and a holdup $V_c = 1.2\,\mathrm{L}$ of four tanks, plus any unmeasured extra ΔV_c, to the valve. Each destination integrates $\dot m_{i,d} = r_d\, Q\, c_{i,out}$ with $r_d \in \{0,1\}$ set by the rule, so every collected gram is a gram that left the column. The antibody capture train of Volume 8, Chapter 11 uses the same finite-volume column with a binding law of its own.

9.3.2 Conditioning, Aggregation and the Membrane

The transfer moves the pool and its phosphate into T-401 once, with a dilution buffer. The vessel integrates the strong-ion amount Z and the buffer total C_T, and the pH is the bracketed root of the charge balance in Equation (9.3), never an integrated state.

$$\begin{aligned}
&H+\frac{Z}{V_\ell}-\frac{K_w}{H}-\frac{C_T}{V_\ell}\frac{K_a}{K_a+H}=0\\
&\mathrm{pH}=-\log_{10}H+0.9\,f_{org}\\
&r_{agg}=k_{agg}\,\frac{V}{1+10^{(\mathrm{pH}-\mathrm{pH}_{agg})/0.5}}\left(\frac{c_P}{c_{ref}}\right)^2
\end{aligned} \tag{9.3}$$

where V_ℓ is the volume in liters, f_{org} the organic fraction and $0.9 f_{org}$ a declared apparent-pH offset. Acid and base enter as 1 M reagents at up to 0.05 L/min, and r_{agg} converts product into an insoluble aggregate G, half activated at $\mathrm{pH}_{agg} = 5.6$.

During concentration and diafiltration the vessel recirculates through the membrane, as shown in Equation (9.4).

$$\begin{aligned} J &= \left(\frac{1}{J_{\Delta P}} + \frac{1}{J_{lim}}\right)^{-1}, \qquad J_{\Delta P} = \frac{\mathrm{TMP}}{\mu (R_m + R_f)} \\ J_{lim} &= k_m \ln \frac{c_g}{c_{pep}}, \qquad \frac{\mathrm{d}R_f}{\mathrm{d}t} = k_f J c_{pep} \\ \frac{\mathrm{d}(V c_i)}{\mathrm{d}t} &= -A_m J s_i c_i - (\dot{a}_P + r_{agg})\, \delta_{iP} \\ \frac{\mathrm{d}V_k}{\mathrm{d}t} &= F_{d,k} - A_m J f_k \\ \frac{\mathrm{d}a_P}{\mathrm{d}t} &= \frac{1}{\tau_{ads}} \left(\frac{a_{max} A_m c_P}{K_{ads} + c_P} - a_P\right) \end{aligned} \tag{9.4}$$

where s_i is the observed sieving ($s_P = 0.010$, $s_{E,L} = 0.025$, $s_X = 0.95$, times the disturbance multiple), V_k the volume of solvent k (water, acetonitrile, 2-propanol, antisolvent) with fraction f_k, μ a logarithmic volume-fraction blend of the solvent viscosities, and a_P the peptide adsorbed on the membrane, kept as an inventory. The diafiltration buffer follows the permeate, $F_d = A_m J$, so the exchange runs at constant volume and N_D counts buffer volumes over the vessel volume. The soft minimum caps the flux at $J_{lim} = 8.25\,\mathrm{L/(m^2\,h)}$ at 100 g/L whatever the pressure.

9.3.3 Precipitation

Antisolvent lowers the solubility exponentially, residual acetonitrile raises it, and the solid is born and grown as described by extensive moments, as shown in Equation (9.5). The batch crystallizer of Volume 8, Chapter 4 carries its crystal population in the same leading moments.

$$\begin{aligned} c^* &= c_0^* \exp(-\kappa f_s + \kappa_a f_{a,rel}), \qquad S = c_P / c^* \\ B &= k_b (S-1)^b, \qquad G = \frac{k_g (S-1)^g}{1 + K_{inh} x_{L,liq}} \\ \frac{\mathrm{d}M_0}{\mathrm{d}t} &= VB, \qquad \frac{\mathrm{d}M_k}{\mathrm{d}t} = kGM_{k-1} + VBL_0^k, \quad k = 1,2,3 \\ \dot{m}_{solid} &= \rho_s k_v \frac{\mathrm{d}M_3}{\mathrm{d}t}, \qquad L_{32} = M_3 / M_2 \end{aligned} \tag{9.5}$$

where $f_{a,rel}$ is acetonitrile as a fraction of the non-antisolvent liquor, $x_{L,liq} = c_L/(c_P + c_E + c_L)$ is the late variant's share of the dissolved peptide and L_{32} is the Sauter mean size. The new solid is exactly $\dot{m}_{solid}$, split among the species in proportion to $c_P - c^*$, $1.35 c_L$, $0.60 c_E$ and $50 c_G$, so the variants enter by incorporation and every aggregate reports to the solid. The growth inhibition $K_{inh} = 700$ is the coupling of this chapter, a teaching value chosen so that the lesson appears at a realistic cut and not fitted to any data. Its form is a simple saturating inhibition in the spirit of impurity-adsorption models of growth [3].

9.3.4 The Nutsche Filter and the Discharge

The cake obeys Ruth's constant-pressure relation [15] with a specific resistance that follows the Kozeny scaling on particle size, a compressibility law, and a gel contribution of the late variant and the aggregate, as shown in Equation (9.6).

$$\begin{aligned} \alpha &= \alpha_{ref} \left(\frac{L_{ref}}{L_{32}}\right)^2 \left(\frac{\Delta P}{\Delta P_{ref}}\right)^s (1 + \gamma_L x_L + \gamma_G x_G) \\ Q_f &= \frac{A_f \Delta P}{\mu\,(R_{cl} + \alpha M_c / A_f)}, \qquad L_c = \frac{M_c}{\rho_s (1 - \varepsilon_c) A_f} \end{aligned} \tag{9.6}$$

Table 9.2 Parameters and Their Evidence. S Anchored in a Read Source, T Teaching Value, F Fitted Inside the App, N Numerics

Symbol	Meaning	Value	Unit	Ev.
L, D, ε	bed length, diameter, mobile fraction	0.25, 0.20, 0.60	m, —	T
q_{max}	saturation capacity	50	kg/m^3	T
b_P, b_E, b_L	affinities at $\varphi_{ref} = 0.30$	0.30, 0.16, 0.52	m^3/kg	T
S_P, S_E, S_L	modifier sensitivities	60, 50, 72	—	T
k, D_{ax}	exchange rate, dispersion	0.05, 3×10^{-7}	s^{-1}, m^2/s	T
V_h, V_c	holdups to detector, collector	0.4, 1.2	L	T
x_c	crude purity	70	%	S
C_T, pH	phosphate, mobile-phase pH	20, 7.5	mM, —	S
φ_0, φ_1, t_g	gradient	22, 45, 60	%, min	T
k_{agg}, pH$_{agg}$	aggregation at 100 g/L, half point	2.6×10^{-3}, 5.6	kg/(m^3 s)	T
A_m, R_m	membrane area, resistance	1.0, 1.5×10^{13}	m^2, m^{-1}	T
k_m, c_g	polarization coefficient, gel	2.5×10^{-6}, 250	m/s, g/L	T
s_P, $s_{E,L}$, s_X	observed sieving	0.010, 0.025, 0.95	—	T
a_{max}, τ_{ads}, K_{ads}	membrane adsorption	0.35, 300, 5	g/m^2, s, g/L	T
k_f	fouling per peptide carried	2×10^{14}	m/kg	T
TMP, N_D	pressure, diavolumes	2, 4	bar, —	S
c_0^*, κ, κ_a	solubility, antisolvent, acetonitrile	150, 9.5, 5.0	g/L, —	T
k_b, b	nucleation	2×10^{11}, 2.4	m^{-3} s^{-1}	T
k_g, g	growth	1.6×10^{-7}, 1.3	m/s	T
K_{inh}	growth inhibition by L	700	—	T
ρ_s, k_v, L_0	solid density, shape, birth size	1150, 0.52, 1	kg/m^3, —, µm	T
α_{ref}, s	cake resistance (1 bar, 20 µm), compressibility	2×10^{12}, 0.60	m/kg	T
γ_L, γ_G	gel weights of L and G	200, 55	—	T
A_f, ε_c, R_{cl}	filter area, porosity, cloth	0.05, 0.60, 10^{11}	m^2, —, m^{-1}	T
f_{occ}, τ_{dew}, k_{dis}	occlusion, drying, redissolution	0.015, 240, 7×10^{-5}	—, s, m/s	T
$\hat{f}_L$ constants	φ_{L50}, w_L, $f_{L,max}$	0.3303, 0.00329, 0.785	—	F
Δt, cells	integration step, finite volumes	1, 60	s, —	N
limits	purity, acetonitrile, solvent, clock	98.5, 0.010, 2.5, 60	%, min	T

where x_L and x_G are mass fractions in the cake, R_{cl} the cloth and L_c the cake depth, 2.46 mm for the 56.6 g cake of the reference case. In cake-dominated flow $Q_f \propto \Delta P^{1-s}$, so with $s = 0.60$ more pressure gives diminishing returns and never less flow. The filtrate carries the suspended solids into the cake, which takes up liquor at $\varepsilon_c/(1-\varepsilon_c)$ of its solid volume. The wash is one well-mixed stage that displaces the pore liquor and redissolves cake at $k_{dis} a_s (c^* - c)$ over the solid surface $a_s = 6M_c/(\rho_s L_{32})$. The drying relaxes the pore liquor toward 1 % of the pore volume at the rate $\lambda = \max(0.1, \Delta P/1\,\text{bar})/\tau_{dew}$. A share $f_{occ}(1 + 3(x_L + 2x_G))$ of the solid volume is occluded mother liquor that no wash reaches, so the residual acetonitrile of a washed cake is set by the diafiltration. Filtration, wash and drying share one 60 min clock. At the discharge the solid is acceptable only if its peptide purity is at least 98.5 %, its acetonitrile at most 0.010 %, its total solvent at most 2.5 % and the clock did not run out. These limits are teaching values. Real ones come from the product's specification, a pharmacopeia and the regulator's assessment, and passing them here certifies no batch. The scorecard leads with the acceptable mass and its rate per campaign hour, both read at the discharge.

🕮 **Pause.** Try Exercise 9.3 before reading on.

9.3.5 Parameters, Evidence and Instruments

Every parameter appears in Table 9.2 with its evidence. Four are anchored in read sources. The crude purity of 70 % sits inside the 60 to 70 % of [12], and the transmembrane pressure and diavolumes inside the 2 to 9 bar and 4 to 8 diavolumes of [13]. Everything else is a teaching value, and the membrane, precipitation and cake values carry no data at all. No controller sees the plant except through an instrument. FT-201 reads the flow through a 3 s lag with 0.5 % colored noise, AT-202 the inlet modifier through 6 s, and AT-203 the absorbance $a_{UV} = \ell \sum_i \varepsilon_i c_i + a_0$ of the detector cell through 2 s with 0.004 AU of noise, a mixed signal in which product and variants absorb nearly alike

(1.00, 0.98 and 1.06 AU per g/L per cm). An inline refractive index gives the modifier at the detector. AT-401 reads the pH through 25 s with 0.01 of noise, PT-501 and PT-601 the pressures with 0.02 bar of noise, and AT-503 the permeate's absorbance. The pool assay draws the pool once, at the transfer, and reports its purity and late-variant fraction 15 min later with standard deviations of 0.06 and 0.03 points. The time step is 1 s, the controllers run every 5 s, one row is logged every 20 s, and the seed is 20261020.

9.3.6 Assumptions

The crude is a mass and a composition, with three impurity classes and one downstream aggregate standing in for a related-substance profile that in practice has dozens of members. The isotherm is a candidate effective law with fixed temperature, pH and ion pairing, and one preparative pass replaces the two or three of the patents, so the trade-offs appear at a lower purity than a multipass route reaches. Solvent volumes are additive, the apparent pH of organic-rich liquor is an offset, and the moments assume size-independent growth with no agglomeration, breakage or dissolution. The cake is washed as one mixed stage, which understates a real displacement wash, and its resistance is evaluated on its current composition, so the wash lowers the resistance of a formed cake by 16 % in the reference case. The scorecard therefore reports the value at the end of the filtration. The precipitation vessel and the conditioning vessel share one inventory in the model although the scene draws R-701 apart.

9.4 The Recipe, the Loops and the Three Cuts

A supervisor runs thirteen stages, load, column wash, gradient, strip, conditioning, concentration, diafiltration, final concentration, precipitation, filtration, cake wash, dewatering and discharge. Each ends on a measurement or a time, a loaded volume, 1.5 column volumes of wash, the gradient time, 6 min of strip at $\varphi_1 + 0.15$, the pH held within 0.12 of setpoint for 2 min after at least 3 min, a vessel volume of 3 L, the diavolume count, a peptide concentration from the soft sensor's inventory, the antisolvent fraction plus a 10 min hold, and the filter volumes. Each transition fires once, after the integration of the step in which its condition is met. In the reference case the stages take 4.7, 60.0, 6.0, 5.3, 44.8, 114.6, 37.1, 24.5, 15.3, 5.2 and 20.2 min from the wash to the drying, a campaign of 5.66 h.

The five loops are the velocity-form PID of the framework, Equations (1.11) to (1.14) of Section 1.6.2, with the derivative off, executed every $\Delta t_c = 5\,\mathrm{s}$. The bias u_b is the output the loop is initialized at, and the anti-windup is back-calculation on the integral with a tracking time $T_t = \tau_I$. FIC-201 drives the pump (3 L/min at full speed, 8 s lag) at $K_c = 18\,\%$ per L/min and $\tau_I = 25\,\mathrm{s}$. AIC-202 drives the B valve on the gradient program at 2.2 % per percent and 30 s, and idles with setpoint tracking when the column stops. AIC-401 drives a split range with a 0.04 deadband at 70 % per pH and 150 s, dosing only during conditioning. PIC-501 sets a retentate valve whose demand is $8\,\mathrm{bar}\,(u/100)^{1.2}$ through a 20 s lag at 22 % per bar and 70 s, and PIC-601 a gas regulator of 5 bar span through 15 s at 35 % per bar and 70 s. Every one has a positive gain from its output to its measurement, and each tracks its operator station in Manual, so transfers are bumpless (Section 1.6.2).

The cut is the learner's decision, and three rules make it. The time rule collects while $t_1 \le t_{el} \le t_2$ on the gradient clock. The UV rule collects while $a_{UV} \ge 0.60\,\mathrm{AU}$ and, once it has shut, stays shut. Its stop threshold of 0.35 AU has no effect as implemented, because the opening threshold is tested at every scan (a defect of the app as shipped). The downstream-aware rule reads a soft sensor that turns the UV into a peptide concentration with one average absorptivity and splits it by a calibration against the modifier at the detector, as shown in Equation (9.7).

$$\begin{aligned} \hat{f}_L &= \frac{f_{L,max}}{1 + \mathrm{e}^{-(\varphi - \varphi_{L50})/w_L}}, \qquad \hat{f}_E = \frac{1}{1 + \mathrm{e}^{(\varphi - \varphi_{E50})/w_E}} \\ \hat{x}_L &= 100\,\frac{\int Q\hat{c}\,\hat{f}_L\,\mathrm{d}t}{\int Q\hat{c}\,\mathrm{d}t}, \qquad \hat{c} = \frac{a_{UV} - a_0}{\ell\bar{\varepsilon}} \end{aligned} \tag{9.7}$$

Table 9.3 Code Verification and the Checks a Source Allows

Check	Result
Code verification	
product ledger, charge to cake, 45 campaigns	$\le 2 \times 10^{-11}$ relative
solid against $\rho_s k_v M_3$	$\le 3.7 \times 10^{-15}$
column with integrated fluxes, 30–480 cells	4×10^{-15}–10^{-14}
tracer mean against V_{liq}/Q	−0.000 %
band exits against linear solvent strength	0.033, 0.046, 0.034 min
linear pulse, mean and spread at 60 cells	0.12, 2.2 %
diafiltration $c/c_0 = \mathrm{e}^{-sN_D}$, $s_P = 0.010$	0.43 %
$s = 0$ retention	no product passes
Ruth filtration time, reference cake	14.2 against 15.3 min[a]
drying time by formula and engine	20.2, 20.2 min
grid, 60 against 480 cells at 0.25 s	0.12 %, 0.04 points
Δt 1 s against 0.25 s	0.12 %, 0.13 points
Δt halved, challenge (suggested)	59.91 and 59.89 g
playback at 1800 times real time	bit-identical
Against read sources (consistency only)	
91 % net yield over 10 diavolumes [11]	$s_P \le 0.0094$, model 0.010
crude purity 60–70 % [12]	model 70 %
TMP 2–9 bar, 4–8 diavolumes [13]	model 2 bar, 4
cake resistance of a peptide precipitate	no source

[a] The rest is PIC-601's pressure ramp, 1.22 bar after 1 min.

where both signals are carried forward by the declared V_c in volume, so the sensor describes what is at the valve. The supervisor of Volume 5, Chapter 7 follows material in volume in the same way. It opens when the valve stream is above 0.60 AU and its predicted purity $100(1 - \hat{f}_E - \hat{f}_L)$ reaches 80 %, and it shuts for good when $\hat{x}_L$ reaches $\hat{x}_{L,max}$ or either opening condition fails. The calibration constants were fitted inside the app to seven runs of this column with crudes of 65 to 75 %, with a weighted root-mean-square error of 6.23 points of mass percent on those runs. No loop or rule reads a true composition. Four stage ends read a model volume standing in for a level, a totalizer, a sight glass and a dew point, and the application keeps three models apart: the plant of Section 9.3, the purity and recovery curve that trends let a learner draw, and the soft sensor, which knows only the calibration and the declared holdup.

9.5 Code Verification, and What the Sources Can Check

Code verification, agreement with sources and educational evaluation are kept apart. Whether the program solves the equations it states is the subject of Table 9.3, and every entry passes. The column's inventory, with its inlet and outlet fluxes integrated in the same stages, closes to 10^{-14} at 30 to 480 cells. A tracer leaves after $V_{liq}/Q = 188.5$ s to round-off, and at vanishing load each band leaves the gradient within 0.05 min of linear-solvent-strength theory [10]. With the step held at 0.25 s, 60 cells put the outlet peak within 0.12 % and the pooled recovery within 0.04 points of a 480-cell reference, while first-order upwinding is 1.58 % and 0.91 points out. The shipped 1 s step adds 0.12 % and 0.13 points. Halving it over a whole campaign moves the isolated mass of case 2 by 0.033 %, and no verdict among 31 lesson variants changes at 120 cells, at half the step or at 30 cells.

Agreement with sources is thin, and the table says how thin. A net yield of 91 % over ten water diavolumes in [11] would correspond to an observed product sieving of $-\ln(0.91)/10 = 0.0094$ if every loss were sieving, which is consistent with the declared 0.010 without fixing it. The crude purity and the membrane operating window lie inside the disclosed ranges. No source consulted for this chapter measures the retention of these peptides on this column, a membrane's sieving for this solvent, or the filterability of a peptide precipitate, so the isotherm, the growth inhibition and the cake are teaching closures. Validation runs of the load, the gradient and the filtration against withheld data cannot be made without such data. Educational evaluation is not attempted, and no learning gain is claimed.

Table 9.4 The Nine Cases and the Challenge

#	Case and departure	Teaching point
1	crude at 60 %	a cut belongs to its batch
2	reference, cut 23–31 min	purity against recovery
3	200 g charged	capacity and displacement
4	collector holdup +1.2 L, unmeasured	volume is delay
5	conditioning at pH 4.5	aggregate the filter feels
6	membrane sieving ten times specified	each diavolume costs
7	antisolvent at 0.40 L/min, cut to 31.25 min	fines
8	filter pressure and wash	compressible cake
9	cut shut at 31.75 min, UV rule infeasible	the pool that passes
C	cut to 32, 2 diavolumes, 0.25 L/min, 3 bar	the whole campaign

Table 9.5 Measured Outcomes of the Cases (Default Seed, at the Discharge)

#	Measured
1	23–31: pool 96.27 %, solid 96.88 %, 53.0 g rejected. Re-cut 24–30.5: pool 98.47 %, 45.3 g acceptable. Downstream aware: pool 94.43 %, rejected.
2	24–30: 47.7 g pooled, 99.48 %, 29.7 min on the filter, 43.4 g, 9.1 g/h. 23–31: 61.7 g, 98.92 %, 11.6 µm, $\alpha = 12.6 \times 10^{12}$ m/kg, 40.7 min, 56.2 g, 9.9 g/h. 23–32: 98.11 %, 1.235 % late, 8.3 µm, out of time filtering.
3	200 g: recovery 42.7 % (73.4 at 120 g), 2.79 % late, $\alpha = 1.6 \times 10^{14}$ m/kg, out of time filtering.
4	pool 97.84 %, solid 98.26 %, rejected. Both cut times +0.8 min: 55.6 g acceptable.
5	pH 4.5: 2.21 g aggregate, solid 95.41 %, 66.4 min. pH 6.0: 0.25 g, acceptable.
6	33.9 g through the membrane, 27.0 g isolated. 2 diavolumes: 0.0127 % acetonitrile, rejected. 8: 17.8 g.
7	0.40 L/min: 8.1 µm, 68.3 min, rejected. 0.15, 0.06, 0.02: 9.3, 10.7, 12.5 µm in 56.2, 47.8, 41.5 min.
8	3.5 against 1.5 bar: 24.8 against 40.7 min. 0.6 bar: drying unfinished, 3.11 % solvent.
9	pool 98.39 % (assay 98.35), 0.939 % late, 9.0 µm, $\alpha = 3.4 \times 10^{13}$ m/kg, out of time drying, solid 98.83 %. Rescued by $t_2 = 31$ or 3 bar. UV rule: 0 of 24 isolation settings rescue it.
C	shipped: 53.5 g at 98.22 %, 0.0273 % acetonitrile, out of time. Suggested: 59.9 g, 10.0 g/h.

9.6 Nine Cases, a Challenge and Three Cut Policies

The cases are listed in Table 9.4. The figures quoted here come from the app's own engine, run without a display at seed 20261020 with the browser's step sequence and with each setting typed in before the clock starts, and every case loads identically when opened first, chosen after another case or rerun from Reset. The outcomes of the cases and their main variants are listed in Table 9.5. Credit for a case requires both a change the learner made and the plant state it produced, and all 11 progress checks pass the automated walk-through.

9.6.1 The Column's Own Trade

The product band is above 0.1 g/L at the valve from 20.0 to 36.0 min of the gradient, with a flat top near 5.5 g/L from 25.9 to 28.5 min, the early variant peaking at 20.4 min on its front and the late variant at 33.9 min on its tail. That ordering is the overloaded displacement of Equation (9.2), and it is why the cut is a decision. In case 3 a 200 g charge exceeds the bed's 157 g and breaks through during the load, the recovery falls from 73.4 to 42.7 % and the pool carries 2.79 % of the late variant. In case 4 an unmeasured 1.2 L extension of the collector line delays the valve by 48 s at 1.5 L/min, and the same cut collects the front shoulder, so the pool falls from 98.92 to 97.84 %.

🕮 **Pause.** Try Exercise 9.2 before reading on.

Table 9.6 Cut Policies on the Challenge Case With Equal Downstream Settings (4 Diavolumes, 0.06 L/min, 1.5 bar, 0.30 L Wash), Default Seed

Policy	Window (min)	Rec. (%)	Pool (%)	Late (%)	Memb. (g)	Size (μm)	α (10^{12} m/kg)	Filter (min)	Isol. (g)	Acc. (g)	Rate (g/h)	Solv. (L/g)
purity-maximizing	24.05–30.03	56.8	99.48	0.086	3.70	14.4	5.7	30.7	43.4	43.4	9.02	2.91
reference	23.05–31.03	73.4	98.92	0.358	4.67	11.6	12.6	41.8	56.2	56.2	9.89	2.25
recovery-maximizing at the pool limit	23.05–31.53	77.0	98.62	0.698	4.88	9.8	23.9	62.0[a]	59.0	0	0	—
UV threshold	20.30–32.78	95.0	91.64	2.069	5.96	7.3	91.2	80.2[b]	38.9	0	0	—
downstream aware, $\hat{x}_{L,max} = 0.80\,\%$	22.55–31.20	78.2	98.51	0.430	4.95	11.2	14.4	46.8	59.9	59.9	10.03	2.11
fixed at that window	22.63–31.28	78.3	98.53	0.485	4.95	10.9	16.3	49.4	60.0	60.0	9.96	2.11
best fixed window in the sweep	22.55–31.03	77.0	98.58	0.340	4.88	11.9	11.8	42.4	58.9	58.9	10.08	2.14

Rec.: recovery into the pool. Isol.: isolated. Acc.: acceptable. Window: the collection times realized on the gradient clock (the rule scans every 5 s). Memb.: product through the membrane plus product adsorbed on it. Solv.: all water and organic solvent used per gram of acceptable peptide, 80.9 L of water and 45.5 L of organic solvent per campaign (37.0 L of it acetonitrile in the mobile phase) for every policy here. [a]Out of filter time while drying. [b]Out of filter time while filtering.

> **TRY IT IN THE APP** **4 · A Longer Line to the Collector**
>
> Run the scenario and compare the UV trace with the shaded cut window on the Column view. The pool falls from 98.92 to 97.84 % and the solid is rejected. Press Reset, move both Cut opens and Cut shuts 0.8 min later on the XV-203 card, and run again: 55.6 g is now acceptable.

9.6.2 Purity, Recovery and the Filter

The comparison of cut policies is run on the challenge case with every policy given the same downstream recipe, four diavolumes, 0.06 L/min of antisolvent, 1.5 bar and a 0.30 L wash, as listed in Table 9.6. A purity-maximizing cut (24 to 30 min) pools 99.48 % and isolates every gram it pools, but only 56.8 % of the product reaches the pool, and the rate is 9.02 g/h. A recovery-maximizing cut that holds the pool at the solid's 98.5 % limit shuts at 31.5 min, pools 98.62 % at 77.0 % recovery and is lost on the filter clock at 62.0 min, still drying. The UV rule recovers 95.0 % into a 91.64 % pool and never finishes filtering. The downstream-aware rule at a 0.80 % limit opens at 22.55 min, shuts at 31.20 min and isolates 59.9 g at 10.03 g/h. Figure 9.2 draws the recovery and the acceptable mass of every policy side by side.

The comparison over the whole cut grid is shown in Figure 9.3. Of 44 windows from $t_1 = 22.5$ to 24 min and $t_2 = 29$ to 33 min, the 20 that failed all ran out of filter time, and the dividing line is the pool's late variant, at most 0.568 % in the 24 that passed and at least 0.665 % in those that failed. A failing pool reached a purity of 98.86 % while a passing one held 98.46 %, so the boundary in the purity and recovery plane is not a purity contour. The last acceptable shutting time is 31.25 min for every opening time tried, and the acceptable mass rises with recovery up to that edge and drops to zero past it, as shown in Figure 9.3b. Membrane loss rises with the pooled mass, from 3.70 to 5.96 g, because sieving takes a fixed share of what the membrane sees. Solvent use is nearly the same for every policy, about 126 L per campaign, because the column program and the diafiltration volume do not depend on the cut, so solvent per gram follows the acceptable mass, from 2.91 to 2.11 L/g.

Pause. Try Exercise 9.1 before reading on.

> **TRY IT IN THE APP** **9 · The Wide Pool That Passes Its Assay**
>
> Load the scenario, press Run, set the speed to 1800× and follow the campaign to the Filter view. The pool assays at 98.35 %, inside the limit. Watch the filter clock: the cake needs 79.3 min of the 60 it is allowed. Press Reset, set Cut shuts to 31 min and run again, and the batch is saved.

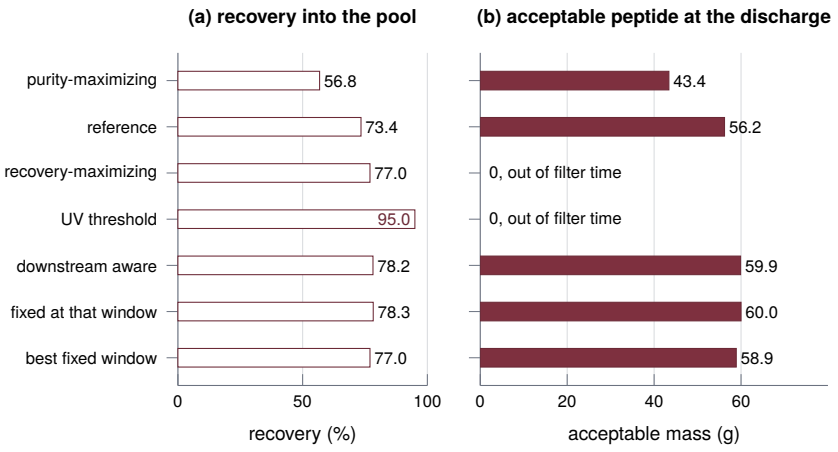


Figure 9.2 The cut policies of Table 9.6 on the challenge case with equal downstream settings: (a) recovery into the pool, (b) acceptable peptide at the discharge. Notice that the UV threshold rule recovers the most and isolates nothing acceptable, like the cut that holds the pool at its purity limit, and that the two fixed windows match the downstream-aware rule.

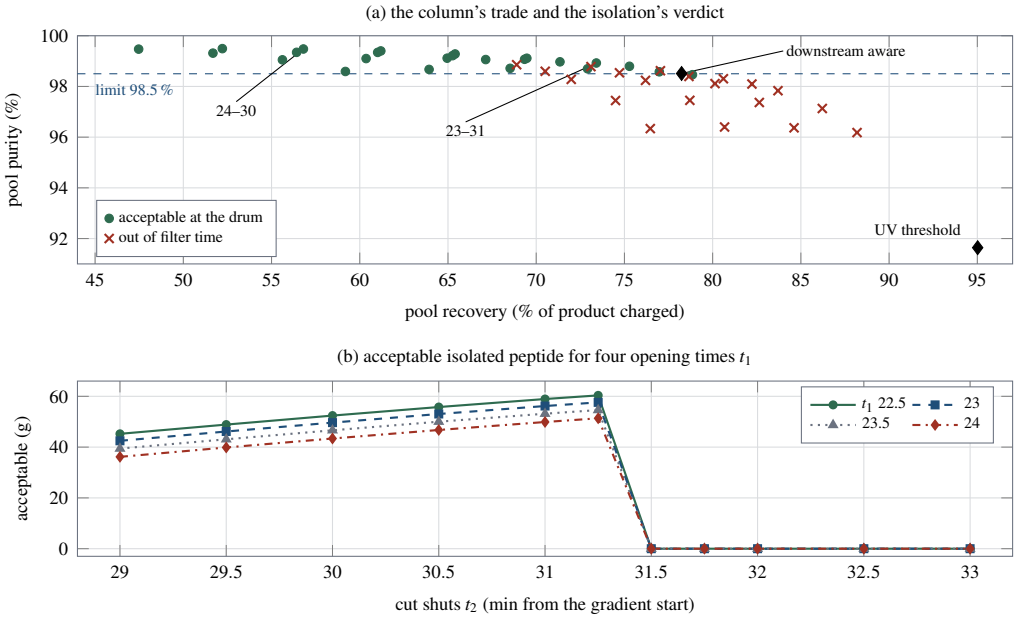


Figure 9.3 Cut sweep on the challenge case with the suggested downstream settings. (a) Pool purity against recovery for 44 windows, marked by the isolation's verdict, with the UV and downstream-aware rules. (b) Acceptable isolated peptide against the shutting time for four opening times, from the earliest to the latest solid with dots, dashed with squares, dotted with triangles and dash-dotted with diamonds. Every window that fails runs out of filter time, and none shutting after 31.25 min survives.

9.6.3 What the Soft Sensor Contributes

The downstream-aware rule's margin over the reference cut, 59.9 g against 56.2 g, comes from its window rather than from what it knows about the pool. A fixed cut set to the window it realized, 22.55

Table 9.7 The Soft Sensor Against the Crude (Reference Case, Cut 23–31 min), and the Downstream-Aware Rule on the Same Crude

Crude (%)	True pool purity	True pool late	Predicted purity	Predicted late	Ratio late	Rule, acceptable (g)
60	96.27	0.484	97.26	0.786	1.62	0, pool 94.43 %
65	98.11	0.416	97.21	0.778	1.87	0, pool 97.14 %
70	98.92	0.358	97.19	0.776	2.17	59.9
75	99.28	0.304	97.18	0.775	2.55	60.9
80	99.47	0.251	97.18	0.776	3.10	61.6

to 31.20 min, isolates 60.0 g at 9.96 g/h, and the best fixed window of the sweep, 22.5 to 31 min, reaches 10.08 g/h, slightly above the rule's 10.03 g/h. The reason is the calibration. The soft sensor's prediction is a function of the modifier at the valve and of the UV, and it hardly moves with the crude, as listed in Table 9.7. On the 23 to 31 min cut it predicts 0.78 % of the late variant for every crude from 60 to 80 %, while the truth falls from 0.484 to 0.251 %, and the rule's window stays at 22.55 to 31.20 min for crudes from 65 to 80 %. Over six paired seeds the suggested configuration isolates 59.9 to 64.3 g and passes the challenge every time, the fixed 23 to 31 min cut 55.9 to 56.7 g and never, and a fixed 23 to 31.5 min cut filtered at 3 bar 58.6 to 59.5 g and every time. Over the 31 noise seeds of the final validation the suggested configuration passes 18 (19 at half the step). At six seeds the rule's cut shuts within half a minute of opening and is not reopened (the stop-threshold defect of Section 9.4), at five the filtration, wash and drying need 62.7 to 68.9 min, and at two the solid is 98.06 to 98.39 % pure. The rule's advantage disappears on a weaker crude. At 65 % it pools 97.14 % and is rejected while the fixed cut isolates 55.0 g.

> ✦ DID YOU KNOW? **A Soft Sensor That Does Not See the Crude**
>
> From the simulator (Table 9.7). On the reference cut the soft sensor predicts between 0.775 and 0.786 % of the late variant for every crude from 60 to 80 %, while the true pool holds from 0.484 down to 0.251 %. Its estimate follows the modifier at the valve, not the crude.

9.6.4 Sensitivity to Selectivity, Sieving and Cake Resistance

The plant was then perturbed while the soft sensor kept its calibration, as listed in Table 9.8. Selectivity matters most. With the late variant's affinity 10 % lower (b_L/b_P from 1.73 to 1.56) its band moves into the product, the reference cut pools 0.96 % of it and every policy but the purity-maximizing cut runs out of filter time, while the narrow cut still isolates 43.5 g. With 10 % more affinity every policy gains filter margin and the ranking of the declared plant holds. Sieving costs mass in proportion. Five times the specified sieving takes the downstream-aware rule from 59.9 to 44.0 g with 21.1 g lost to the membrane, and the ranking is unchanged. A higher cake resistance costs time and not mass until the clock binds. At 1.5 times α_{ref} the rule's cake takes 59.0 min of the 60, the closest of the four to failing, because its pool carries the most late variant.

> ✦ DID YOU KNOW? **A Weaker Selectivity Reverses the Ranking**
>
> From the simulator (Table 9.8). On the declared plant the narrow 24 to 30 min cut isolates the least acceptable peptide, 43.4 g. With the late variant's affinity 10 % lower, the three wider policies all run out of filter time, and that narrow cut, at 43.5 g, is the only one left.

9.6.5 The Isolation Cases

Cases 5 to 8 break the isolation one lever at a time. At pH 4.5 the loop overshoots to pH 3.0 on the way down, 2.21 g of aggregate forms, mostly during the diafiltration, and the solid falls to 95.41 %

Table 9.8 Acceptable Peptide (g) With the Plant Perturbed and the Soft Sensor's Calibration Kept. Filter: Out of Filter Time

Plant	24–30	23–31	22.5–31	Downstream
declared	43.4	56.2	58.9	59.9
b_L −10 %	43.5	filter	filter	filter
b_L +10 %	43.3	56.0	58.7	59.7
sieving ×2	40.2	52.1	54.7	55.6
sieving ×5	31.8	41.2	43.3	44.0
α_{ref} ×0.8	43.4	56.2	58.9	59.9
α_{ref} ×1.25	43.4	56.2	58.9	59.9
α_{ref} ×1.5	43.4	56.2	58.9	59.9[a]

[a]59.0 min on the filter, against 50.9 for 22.5–31.

because the aggregate is insoluble. A membrane sieving ten times its specification passes 33.9 g, and the permeate's own UV cell shows it within minutes. Dosing the antisolvent at 0.40 L/min instead of 0.06 makes 8.1 µm particles from the same 57.6 g, and the cake needs 68.3 min. On the filter, 2.33 times the pressure is 1.64 times faster, because the cake's resistance rises from 1.3×10^{13} to 2.1×10^{13} m/kg as it compresses. The challenge as shipped fails twice, on the clock and on 0.0273 % acetonitrile left by two diavolumes in the occluded liquor. A headless campaign of 5.66 h takes 0.82 s on the build host, 24 719 steps per second, and in a browser a campaign plays in 11 to 15 s at 1800 times real time. The shoulder experiment of Figure 9.4, the reference cut and then the cut at 32 min, took 27.6 s of wall time at 1800 times real time in a usability review.

9.7 Technology in Practice

What Measures It

The cut is decided on one optical signal. AT-203 reads the absorbance of the stream in a UV cell at the column outlet. In the model it carries a 2 s lag and 0.004 AU of noise, and it sees the product and its variants nearly alike. Real purification meets the same limit, because one optical method alone offers limited selectivity between a product and contaminants of closely similar structure [18]. The variable of each loop comes from a common transmitter: a flowmeter, an inline analyzer of the modifier, a pH electrode and two pressure transmitters, each with the lag and noise of Section 9.3. Composition beyond the UV signal comes only from the pool assay, 15 min after the transfer, the role that analytical chromatography of typically 30 min or more plays in practice [18].

> **TECHNOLOGY** **Sensor: The UV Absorbance Detector**
>
> The most popular detectors for liquid chromatography use the absorption of ultraviolet or visible light by the analytes. The stream passes a flow cell, which in an analytical instrument holds 1 to 10 µL over a path of 0.2 to 1 cm, and the chromatogram is absorbance against elution time. The mobile phase must not absorb at the chosen wavelength, and the choice of wavelength changes each analyte's signal. Gas released as the pressure falls across the column can form bubbles that interfere with the detector's response [19]. One optical method alone also has limited selectivity. In protein purification, contaminants with subtle structural differences from the product are found by analytical chromatography that typically takes 30 min or more [18].

What Moves It

The pump of FIC-201 delivers 3 L/min at full speed through an 8 s lag, and AIC-202 sets the share of solvent B with a valve, the arrangement of a proportioning valve ahead of a single pump [19]. The flow measurement carries 0.5 % colored noise and no stroke pulses. The cut valve XV-203 switches at once in the model, but the line between the detector and the valve holds 1.2 L, and case 4 shows that 1.2 L more delays the cut by 48 s at 1.5 L/min. Figure 9.5 sets the loops on such hardware.



Figure 9.4 The shoulder collected, case 2 with the cut shut at 32 min, on the Column view at $t = 2443$ s (40.7 min), 34.0 min into the gradient. The figure is a composite of the app's header and its Column view, with the control panel, the trends and the Learn panel cut. Callouts: (1) the soft sensor's badges, (2) the UV cell AT-203, (3) the cut valve XV-203, (4) the pool with the soft sensor's estimate, (5) the shaded collection. The shaded collection from 23.0 to 32.0 min is complete and shut, 2.3 g remain on the bed, and the late-variant band peaks at the valve, about 1.4 g/L against under 1 g/L of product on the app's trend panel. The pool holds 13.5 L and 68.9 g by the soft sensor, which predicts 96.4 % purity and 1.78 % late variant against a true 1.24 % that the trend panel shows, and 80.1 % of the charged product is in the pool. The status reads ELUTING TO WASTE at 0.40 AU and 33.8 % organic. This pool will assay at 98.07 % and run out of filter time.

TECHNOLOGY **Final Element: The Gradient Pump**

A liquid chromatograph draws its mobile phase from one or more reservoirs. In one design two piston pumps feed a mixing chamber: their combined speed sets the flow and their relative speed the composition. In another, a solvent proportioning valve ahead of a single pump sets how much of each solvent is drawn in each pump cycle. The back-and-forth stroke of a reciprocating pump gives a pulsed flow that adds noise to the chromatogram. A second cylinder that delivers while the first refills, or a pulse damper with a compressible fluid behind a flexible diaphragm, smooths the flow [19].

Where the Controller Runs

In a plant, loops like these would run in a distributed control system, whose supervisory level oversees localized controllers. Programmable logic controllers often serve as those local controllers, and the loops run with cycle times from milliseconds to minutes [20]. The app executes the five PI loops every 5 s on a 1 s integration step. Above them sits the plantwide structure: the supervisor of thirteen stages, each ended on a measurement or a time, and the cut rule, the one decision that couples the column to the filter. ICH Q11 lists in-process controls and controls implicit in the design of the process, such as the sequence of purification steps, among the parts of a control strategy [21], the two parts that the loops and the stage sequence play in the model.

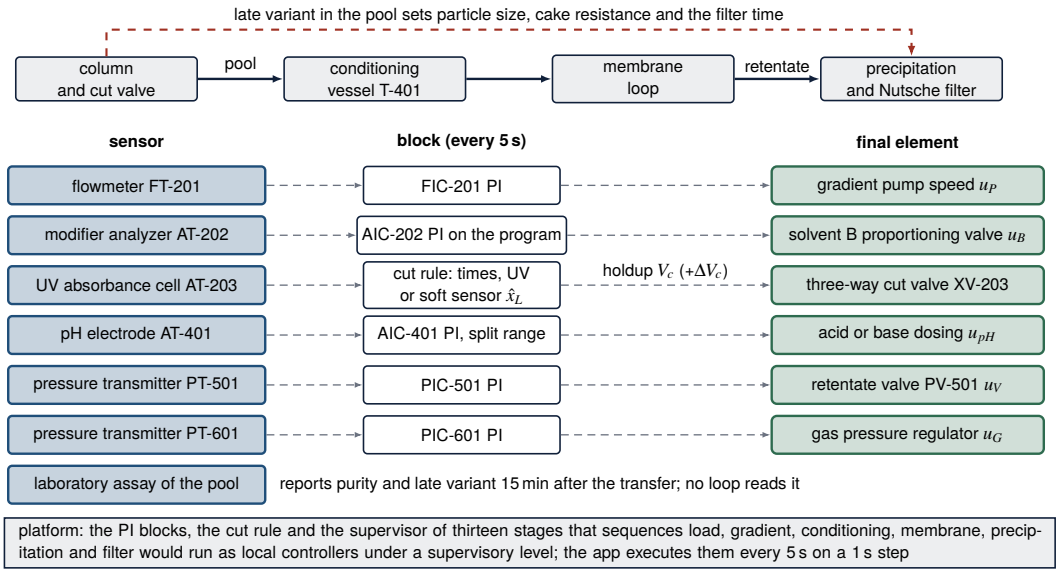


Figure 9.5 The loops of Section 9.4 on real hardware, sensor to block to final element, below the four blocks of the train. Notice that no loop closes from the filter back to the cut: the coupling drawn as the dashed line across the top is left to the cut rule and its holdup, and the laboratory assay arrives after the cut is over.

> **IN PRACTICE** **One Design Space Across Several Unit Operations**
>
> The international guideline on the development and manufacture of drug substances, ICH Q11, asks that the formation, fate and purge of impurities be understood, and that controls be set for impurities as they progress through multiple process operations. A design space may be set for one unit operation or for a combination of selected unit operations. These are chosen for their impact on the critical quality attributes and need not be sequential, and the linkages between steps are evaluated so that the cumulative generation and removal of impurities is controlled. A design space that spans multiple unit operations can give more operational flexibility [21].

What Is Modernizing It

The first development is in-line spectroscopy with multivariate calibration in place of a late laboratory result. In the purification of protein drugs, spectroscopic methods detect rapidly and automatically with no sample preparation, conditioning or destruction, and several sensors combined measure better than one. Where sensors sit at different points of a stream, the dispersion between them must be accounted for [18], which is the lesson of case 4 in another form. The second is model predictive control. Once a fast dynamic model of the critical quality attributes exists, it can be used for control, and a predictive controller respects constraints and can tailor production toward real-time release [18]. In this chapter nothing optimizes the cut, the diavolumes or the dosing, and a predictive or batch-to-batch policy is not implemented (Section 9.8).

9.8 Limitations

The route is a teaching assembly of disclosed options, and patents establish that the options exist, not that the train works. The isotherm, the growth inhibition, the membrane and every cake coefficient are teaching values, and the coupling of this chapter rests on K_{inh}, a number chosen so the lesson appears at a realistic cut. The soft sensor is a modifier calibration biased high by a factor of 1.6 to 3.1 on the late variant and blind to the crude, so the downstream-aware rule acts as a fixed window in

disguise, and the best fixed window found with hindsight matches it. The pool assay arrives after the cut is over, so it cannot correct the cut it describes, unlike the laboratory result that corrects the soft sensor of Volume 5, Chapter 10. The weaker crude of case 1 arrives 10 points below the calibration in one step, which is not how a synthesis drifts, but it isolates the calibration's blind spot. Nothing optimizes the cut, the diavolumes or the dosing, and a model predictive or batch-to-batch policy is not implemented. The precipitation has no agglomeration, the wash is one mixed stage, and the drying is a relaxation. Every limit and the 60 min clock are teaching values.

OPEN CHALLENGE Telling the Product From Its Variants on Line

Ideally a process measurement would tell the product, process-related and product-related contaminants apart in real time. For protein drugs this is hard: the contaminants are chemically close to the product, and one optical method alone has limited selectivity. Several sensors combined by multivariate data analysis do better, but a calibration model must be tested on an external data set before productive use, because internal validation is not sufficient. In-line soft sensors have reached the needed detection limits for aggregates and fragments at laboratory scale, and future work may close the loop on the steps of interest [18]. The soft sensor of this chapter, fitted inside the app to seven runs, shows the gap: it is blind to the crude.

9.9 Conclusion

This chapter studied a GLP-1 purification and isolation train as a simulator in which the chromatography cut, the membrane exchange and the filter cake share one material, with a product ledger that closes to 2×10^{-11} of the charge. A cut that adds 4.4 g to the pool and keeps it inside its purity limit loses the batch, because the late variant it collects shrinks the particles from 11.6 to 9.0 µm and the cake misses its hour. On the challenge, a purity-maximizing cut isolates 43.4 g, cuts that maximize recovery fail on the filter, and the downstream-aware rule isolates 59.9 g, but a fixed cut at the window it realizes does as well, because its soft sensor predicts 0.78 % late variant whatever the crude. Lower selectivity leaves only the narrow cut standing. The coupling, the cake and the isotherm are teaching closures without data. A fuller model would add measured retention, sieving and cake resistance for one peptide, a sensor that sees the crude, and an optimized schedule of cut and isolation. The simulator shows that a fraction cut is a decision about the filter as much as about the pool.

KEY POINTS

- A pool can pass its purity assay and still lose the batch: the late variant it carries shrinks the particles, and the cake misses its hour.
- Over 44 cut windows the dividing line is the pool's late variant, not its purity: at most 0.568 % where the batch passed, at least 0.665 % where it failed.
- An unmeasured 1.2 L between the detector and the valve delays the cut by 48 s at 1.5 L/min and drops the pool from 98.92 to 97.84 %.
- The downstream-aware rule isolates 59.9 g, but a fixed cut at its window does as well: its soft sensor predicts 0.78 % late variant whatever the crude.

Where else this applies. A product stream cut on an analyzer's running average returns on the batch column of Volume 8, Chapter 3.

References

[1] S. Murugesan, D. M. Hallow, J. P. Vernille, J. W. Tom, and J. E. Tabora, "Lean filtration: Approaches for the estimation of cake properties," *Organic Process Research & Development*, vol. 16, no. 1, pp. 42–48, 2012.

[2] L. Aumann and M. Morbidelli, "A continuous multicolumn countercurrent solvent gradient purification (MCSGP) process," *Biotechnology and Bioengineering*, vol. 98, no. 5, pp. 1043–1055, 2007.

[3] N. Kubota and J. W. Mullin, "A kinetic model for crystal growth from aqueous solution in the presence of impurity," *Journal of Crystal Growth*, vol. 152, no. 3, pp. 203–208, 1995.

[4] G. Guiochon, “Preparative liquid chromatography,” *Journal of Chromatography A*, vol. 965, no. 1–2, pp. 129–161, 2002.

[5] N. Marchetti, F. Dondi, A. Felinger, R. Guerrini, S. Salvadori, and A. Cavazzini, “Modeling of overloaded gradient elution of nociceptin/orphanin FQ in reversed-phase liquid chromatography,” *Journal of Chromatography A*, vol. 1079, no. 1–2, pp. 162–172, 2005.

[6] R. B. Merrifield, “Solid phase peptide synthesis. I. the synthesis of a tetrapeptide,” *Journal of the American Chemical Society*, vol. 85, no. 14, pp. 2149–2154, 1963.

[7] J. Lau, P. Bloch, L. Schäffer, I. Pettersson, J. Spetzler, J. Kofoed, K. Madsen, L. B. Knudsen, J. McGuire, D. B. Steensgaard, H. M. Strauss, D. X. Gram, S. M. Knudsen, F. S. Nielsen, P. Thygesen, S. Reedtz-Runge, and T. Kruse, “Discovery of the once-weekly glucagon-like peptide-1 (GLP-1) analogue semaglutide,” *Journal of Medicinal Chemistry*, vol. 58, no. 18, pp. 7370–7380, 2015.

[8] G. I. Bell, R. F. Santerre, and G. T. Mullenbach, “Hamster preproglucagon contains the sequence of glucagon and two related peptides,” *Nature*, vol. 302, no. 5910, pp. 716–718, 1983.

[9] D. J. Drucker, J. F. Habener, and J. J. Holst, “Discovery, characterization, and clinical development of the glucagon-like peptides,” *Journal of Clinical Investigation*, vol. 127, no. 12, pp. 4217–4227, 2017.

[10] L. R. Snyder, J. W. Dolan, and J. R. Gant, “Gradient elution in high-performance liquid chromatography,” *Journal of Chromatography A*, vol. 165, no. 1, pp. 3–30, 1979.

[11] A. Stadelmaier, R. O. Schönleber, D. Samson, and F. Dettner, “Purification of glucagon-like peptide 1 analogs,” European Patent 3 433 268 B1, Jul. 28, 2021.

[12] C. T. Dewar *et al.*, “Improved purification process of semaglutide,” European Patent 4 181 946 A1, May 24, 2023.

[13] S. Krishnan *et al.*, “Synthesis of an intermediate and its use for the preparation of GLP-1R/GIPR agonists,” U.S. Patent Application Patent 2026/0 116 848 A1, Apr. 30, 2026.

[14] A. Isidro-Llobet, M. N. Kenworthy, S. Mukherjee, M. E. Kopach, K. Wegner, F. Gallou, A. G. Smith, and F. Roschangar, “Sustainability challenges in peptide synthesis and purification: From R&D to production,” *The Journal of Organic Chemistry*, vol. 84, no. 8, pp. 4615–4628, 2019.

[15] B. F. Ruth, “Studies in filtration III. derivation of general filtration equations,” *Industrial & Engineering Chemistry*, vol. 27, no. 6, pp. 708–723, 1935.

[16] R. van Reis and A. Zydney, “Bioprocess membrane technology,” *Journal of Membrane Science*, vol. 297, no. 1–2, pp. 16–50, 2007.

[17] H. M. Hulburt and S. Katz, “Some problems in particle technology,” *Chemical Engineering Science*, vol. 19, no. 8, pp. 555–574, 1964.

[18] 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.

[19] D. Harvey, “High-performance liquid chromatography,” Section 12.5 of *Analytical Chemistry 2.1*, Chemistry LibreTexts open textbook, Creative Commons BY-NC-SA 4.0 license, 2021.

[20] 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.

[21] International Conference on Harmonisation, “Development and manufacture of drug substances (chemical entities and biotechnological/biological entities) Q11,” ICH harmonised tripartite guideline, Step 4 version, https://database.ich.org, May 2012.

9.10 Exercises

9.1 ⌕ Concept check · 2 min

In case 9 the cut shuts at 31.75 min instead of 31 min. The pool assays at 98.35 %, and the isolated solid is 98.83 % pure with 0.0019 % acetonitrile. Why is the batch lost?

(a) The isolated solid fails the 98.5 % purity limit once the assay error is allowed for

(b) The longer collection leaves too much acetonitrile in the liquor occluded in the cake

(c) The late variant slows particle growth, the cake resists more, and the filter misses its hour

(d) The larger pool passes more peptide through the membrane, and the cake is too small to wash

9.2 Concept check · 2 min

In case 4 the line from the detector to the cut valve is 1.2 L longer than the recipe allows for, and nobody has measured it. The cut is still set at 23 to 31 min of the gradient clock at 1.5 L/min. What does the pool collect, and how should the cut be corrected?

(a) Material that eluted later than planned, so both cut times should move earlier by about 0.8 min

(b) The same material as before, because the detector trace is unchanged and the cut follows the clock

(c) The same material as before, because the soft sensor carries the signal forward to the valve

(d) Material that eluted about 48 s earlier than planned, so both cut times should move later

9.3 Concept check · 2 min

Case 8 raises the Nutsche filter pressure from 1.5 to 3.5 bar, a factor of 2.33. The cake is compressible with $s = 0.60$. What does Equation (9.6) predict for the filtrate rate once the cake dominates the resistance?

(a) It rises by about 1.4 times, because the specific resistance grows as $\Delta P^{0.6}$

(b) It rises by about 2.3 times, because the flow through a cake is proportional to the pressure

(c) It stays the same, because the specific resistance rises in proportion to the pressure

(d) It falls, because the compressed cake closes up faster than the pressure pushes liquor through

9.4 Theory and calculation · 6 min

Use Equation (9.2) and Table 9.2: $q_{max} = 50\,\mathrm{kg/m^3}$, $\varepsilon = 0.60$, $b_P = 0.30$ and $b_L = 0.52\,\mathrm{m^3/kg}$ at $\varphi_{ref} = 0.30$, and $S_P = 60$, $S_L = 72$.

(a) Find the product's retention factor at vanishing load at 30 and at 22 % organic. *Checkpoint:* 10.0 and 1215.

(b) Find the selectivity b_L/b_P at 22, 30 and 33 % organic. *Checkpoint:* 4.53, 1.73 and 1.21.

(c) The modifier analyzer AT-202 reads 0.5 points high, and AIC-202 holds its measurement on the program. By what factor do the retention factors of the product and of the late variant change? *Checkpoint:* 1.35 and 1.43, both higher.

9.5 Theory and calculation · 6 min

Diafiltration at constant volume keeps a fraction e^{-sN_D} of a species whose observed sieving is s. Acetonitrile passes freely ($s = 1$). In case 6 the membrane sieves ten times its specification ($s_P = 0.10$), four diavolumes isolate 27.0 g, and two diavolumes leave 0.0127 % acetonitrile in the solid against the 0.010 % limit.

(a) Find the product lost during four diavolumes at the specified $s_P = 0.010$ and at $s_P = 0.10$. *Checkpoint:* 3.92 % and 33.0 %.

(b) Scale the 27.0 g of case 6 to two and to eight diavolumes. *Checkpoint:* 33.0 g and 18.1 g.

(c) Find the smallest setting on the app's 0.5 diavolume grid that meets the acetonitrile limit in case 6, and the mass it isolates. *Checkpoint:* 2.5 diavolumes, 0.0077 % and 31.4 g.

9.6 Theory and calculation · 6 min

Use Equation (9.6) with $\rho_s = 1150\,\mathrm{kg/m^3}$, $\varepsilon_c = 0.60$, $A_f = 0.05\,\mathrm{m^2}$ and $s = 0.60$, and the scorecard values of Table 9.5.

(a) Find the depth of the 56.6 g cake of the reference case. *Checkpoint:* 2.46 mm.

(b) Case 2 gives 11.6 µm and 12.6×10^{12} m/kg, and case 9 gives 9.0 µm and 3.4×10^{13} m/kg, both at 1.5 bar. Split the rise in α into the factor from particle size and the rest. *Checkpoint:* 2.70 in all: 1.66 from size and 1.62 from the rest.

(c) Predict α for case 9 filtered at 3 bar, and the factor on the cake-dominated filtrate rate. *Checkpoint:* 5.15×10^{13} m/kg and 1.32.

9.7 App challenge · 8 min

Load scenario 4, A longer line to the collector. The collector line holds an unmeasured 1.2 L beyond the recipe. Use the default seed and run each case at 1800 × to the discharge.

(a) Find the extra delay at 1.5 L/min. Run the scenario as shipped and read the pool purity and the Isolated purity tile. *Checkpoint:* 48 s; 97.84 % and 98.26 %, nothing acceptable.

(b) Press Reset, set Cut opens to 23.5 and Cut shuts to 31.5 min, and run again. Read the same tiles and the acceptable mass. *Checkpoint:* 98.40 % and 98.81 %, 55.4 g.

(c) Press the reload button, set the Collection rule to Downstream aware, and run again. *Checkpoint:* 96.90 % and 97.40 %, nothing acceptable.

Accept purities within 0.01 points and masses within 0.1 g. Explain why the soft sensor does not correct the delay.

9.8 App challenge · 8 min

Load scenario 9, The wide pool that passes its assay, with the cut shut at 31.75 min. Use the default seed and run each case at 1800 × to the discharge.

(a) Run the scenario as shipped. Read Late variant in the pool (true), Particle size · cake resistance and Time on the filter. *Checkpoint:* 0.939 %; 9.0 µm and 34.0×10^{12} m/kg; out of time while drying, off at 79.3 min.

(b) Press Reset, set the Filter pressure to 3.0 bar, and run again. Read the same tiles and the acceptable mass. *Checkpoint:* 51.5×10^{12} m/kg, 54.9 min, 60.2 g at 9.8 g/h.

(c) Press the reload button, set the Collection rule to UV threshold, and run again. Read the pool tile and Time on the filter. *Checkpoint:* 95.0 % and 91.64 %; out of time while filtering.

Accept resistances within 0.2×10^{12} m/kg and times within 0.1 min. Compare (b) with the prediction of $\alpha \propto \Delta P^{0.6}$.

9.9 App challenge · 7 min

Load the challenge, the most acceptable peptide per hour. Use the default seed and run each case at 1800 × to the discharge.

(a) Run the challenge as shipped. Read the Challenge target tile, Residual solvent and Particle size · cake resistance. *Checkpoint:* cake not off the filter in time; 0.0273 % acetonitrile; 6.7 µm.

(b) Press the reload button, press Reveal suggested tuning, and run again. Read the acceptable mass, the pool tiles and the Challenge target tile. *Checkpoint:* 59.9 g at 10.0 g/h; 78.2 and 98.51 %, 0.430 %; met.

(c) Press Reset, set Crude purity from the synthesis to 65 %, and run again. This is an experiment, not an attempt. *Checkpoint:* pool 97.14 %, isolated purity 97.70 %, target missed on purity.

Accept masses within 0.1 g and purities within 0.01 points.

9.10 Application area · 4 min

Every policy of Table 9.6 uses 80.9 L of water and 45.5 L of organic solvent per campaign, 37.0 L of it acetonitrile.

(a) Find the total solvent, the acetonitrile share, and the solvent per gram of acceptable peptide for the reference cut (56.2 g) and the downstream-aware rule (59.9 g). *Checkpoint:* 126.4 L, 29.3 %, 2.25 and 2.11 L/g.

(b) Explain why the solvent per gram follows the acceptable mass, and say what the recovery-maximizing cut costs in solvent. *Checkpoint:* 6.2 % less per gram; all 126.4 L for no acceptable product.

9.11 Application area · 4 min

The pool assay reports purity and late-variant fraction 15 min after the transfer, with standard deviations of 0.06 and 0.03 points. In the reference case the stages from conditioning to precipitation take 5.3, 44.8, 114.6, 37.1 and 24.5 min.

(a) Find how long before the filtration starts the assay result is available. *Checkpoint:* 211 min.

(b) In the cut sweep of Figure 9.3, passing pools carry at most 0.568 % of the late variant and failing pools at least 0.665 %. Express the gap in assay standard deviations, and propose a use of the result. *Checkpoint:* 3.2 standard deviations.

9.12 Review the agent · 5 min

An AI assistant reviews the challenge and writes: “The downstream-aware rule isolates 59.9 g against 56.2 g for the fixed 23 to 31 min cut, because its soft sensor predicts the late variant that the filter answers to. The 3.7 g is the value of that prediction. Use the rule on every campaign. When a weaker crude arrives, the sensor will see the extra late variant and shut the cut earlier.” Find the flaw and say how you would check it.