INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 8

Timing the Quench to Hold mRNA Lipid Nanoparticle Size at Higher Throughput

Pages 171–195

Edited by John D. Hedengren

How to cite this chapter. "Timing the Quench to Hold mRNA Lipid Nanoparticle Size at Higher Throughput," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 171–195.

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

The app for this chapter
https://apmonitor.com/apps/mrna-lnp/

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

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

CHAPTER 8

Timing the Quench to Hold mRNA Lipid Nanoparticle Size at Higher Throughput



The controller moves 1 the feed line into the closed mixing skid

The controller reads 1 the bench analyzer that measures particle size

What disturbs it 1 samples waiting their turn, the delay before a result

The mixing, the aging loop and the quench run out of sight inside the sealed cabinet, while the size that decides release is measured on a drawn sample at the bench minutes later. Illustration, not to scale.

IN PLAIN TERMS **Stopping Small Particles at the Right Size**

The main idea. Two liquids meet in a mixer, and tiny carrier particles form in the stream at once. Those particles keep growing until a flood of fresh buffer stops them, so the time between the mixer and the flood is what sets their size.

The challenge. The only instrument that reports the size draws a sample and answers minutes later. By then the liquid it sampled has left the plant, and anything made since was made blind.

Why it matters. Size is one of the qualities the product is judged on, so the plant has to know it while the liquid is flowing rather than after the tank is full.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the history of the mixer as a flow device and the fact on parallel channels (Section 8.1), the chapter's own account of ethanol, jets and the arrest of fusion (Section 8.2), the two try-it segments and the fact from the simulator (Section 8.6) and the key points
Study	60 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 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the two open challenges

The control question The particles of this skid grow for as long as the mixed stream ages before a quench arrests them, so a loop volume and a total flow act on one quality attribute. What does timing the quench buy that ratio control on four feeds cannot, and what can a blend of channels behind one set of meters hide?

Learning objectives

- Solve the four-feed ratio block for its flows, and show why two feeds cannot hold the charge ratio when the flow ratio moves.
- Compute the aging time and the predicted particle size from the loop volume, the total flow and the fitted size law.
- Explain why the ethanol content and not the quench ratio arrests growth, and compute the encapsulation at a given charge ratio.
- Compare release on a delayed size result with release on a soft sensor, and say why the soft sensor's bias gain stays small.
- Judge a scale-out blend that meets every window while its channels do not.

Prerequisites Sections 1.4.4 (ratio station and trim), 1.4.2 (unmeasured disturbances), 1.6.2 (velocity form) and 1.6.4 (scenarios and the challenge). Chapter 7 for a ratio blend trimmed on an analyzer.

The app https://apmonitor.com/apps/mrna-lnp/

See also Volume 5, Chapter 7 for residence-time tracking and diversion on a continuous drug line, Volume 5, Chapter 3 for a composition loop through a slow analyzer, and Volume 5, Chapter 4 and Volume 5, Chapter 10 for soft sensors that late measurements correct.

8.1 Introduction

The size of a lipid nanoparticle can be moved without touching its formulation. In the simulated skid of this chapter, an aqueous stream of mRNA in acetate buffer meets an equal flow of lipid in ethanol in an impinging-jet mixer at 40 mL/min per channel, and the mixed stream spends 117 ms in a 78 µL loop of tubing before a second jet floods it with two volumes of buffer. The particles leave at 139.9 nm with a modeled size spread $\sigma(\ln d)$ of 0.349 and 92.7 % nominal encapsulation. A 293 µL loop gives the same material 440 ms and takes it to 215.8 nm at the same flows, the same charge ratio and the same mixing, while raising the quench from two volumes to 3.5 volumes moves the size by 0.46 nm.

Nanoparticle production of this kind is often approached as one mixing step, which obscures the relation between the process inputs and the particle properties [1]. The simulator is built against the form of that belief a control student arrives with, that the flow ratio and the mixing intensity decide the particle size while the total flow decides only the throughput. In this model the jets make small particles whose size follows the aqueous buffer strength, and those particles then fuse for as long as the ethanol content lets them. The total flow is therefore the clock as well as the throughput, because the aging time is the loop volume divided by it, and the flow ratio acts mostly through the ethanol content. At a 3:1 flow ratio, with the charge ratio held at 6.01 by the ratio block, the same skid makes 98.0 nm particles instead of 139.9 nm.

The class of problem is a quality attribute fixed by the age of the material rather than by the state of a vessel. A continuous drug process must follow each portion of material through a residence-time distribution, and its control objective is the satisfaction of every critical quality attribute rather than a set of steady values [2]. The vocabulary for that is as old as the description of residence times in a continuous flow system [3], and every tubular hold, aging coil or maturation line shares the shape. Volume 5, Chapter 7 follows material in volume through such a distribution to a divert valve on a flow synthesis line. However, demonstrations of timed arrest are laboratory runs at fixed recipes [1], and the control literature on continuous manufacturing treats traceability and quality attributes in general terms [2], so neither supplies a plant on which the age distribution, each channel's true particle size and a delayed assay can be read side by side while a controller decides.

This chapter asks whether timing the quench buys anything that ordinary ratio control does not already buy. Every policy runs on one simulated formulation skid with the same instruments, feeds and faults, and each is judged on particle-quality variability and accepted productivity across flow perturbations and channel counts. The methods are all established ones: the size law, the tanks-in-series distribution, ratio control, a bias-corrected soft sensor and a constrained grid search. The results are measurements on a synthetic skid. Over six paired seeds, coordinated ratio control on the 78 µL loop runs out of buffer strength at 441.7 mL/min of formulated bulk, while a supervisor free to change the loop as well reaches 600.0 mL/min, a margin of 158 mL/min. In the passing run 13 500 mL of the 15 409 mL accepted was made while at least one channel sat outside its own window, although the blend that reached the tank met every criterion. The model behind these numbers closes its volume, RNA and lipid ledgers to 10^{-14} of throughput, its size map is checked against withheld tables, and a measurement boundary keeps every live decision on declared instruments.

Section 8.2 follows the formulation method from ethanol injection to the timed quench. Section 8.3 sets out the skid: the ratio block on four feeds, the volume clock and its age distribution, the size law, the collection line, the parameters and the boundary between what is measured and what is not. Section 8.4 records how the size law was calibrated and where its domain ends. Section 8.5 gives the flow loops, the size trim and its soft sensor, the three policies and the release rule. Section 8.6 runs the nine cases and the challenge, Section 8.7 sets the simulated instruments beside the real ones, Section 8.8 states what the model cannot support, and Section 8.9 collects the numbers. The exercises of Section 8.10 work through the ratio block, the size law, three of the cases and the challenge.

> ⧖ HISTORY **The Mixer Becomes a Flow Device, 1999 Onward**
>
> Particles of this kind were first made by hand, by pipetting or vortexing one phase into the other. In 1999 a T-shaped junction was introduced as an alternative to that macroscopic mixing for lipid-based particles, and the preparation became a flow process in a channel [4]. Industrial scale went a different way. A confined impinging-jet mixer collides two high-velocity streams in a chamber, and it relies on rapid mixing at Reynolds numbers between 100 and 2000, where inertia dominates and the flow is turbulent [4]. Throughput is then raised by running mixers in parallel rather than by building a larger one. One practice of that industry is the subject of this chapter: a dilution placed in line just downstream of the mixing zone floods the stream with aqueous buffer, which drops the ethanol content to 20 percent by volume or less and stabilizes the particle size [4].

8.2 Ethanol, Jets and the Arrest of Fusion

Lipid vesicles entered the laboratory in 1965, when phospholipids swollen in water were shown to form lamellar structures across which univalent ions diffuse slowly [5]. Ethanol was the solvent that made them a process rather than a preparation. Injecting an ethanolic lipid solution into an aqueous phase produced single-bilayer vesicles without sonication in 1973 [6], and the step from that injection to a manufacturing route was taken when plasmid DNA was encapsulated by combining a lipid-ethanol stream with an aqueous DNA stream in a T-shaped chamber, a stepwise dilution that dropped the ethanol below the concentration at which the lipid stays soluble and gave particles under 200 nm with encapsulation above 80 % [7]. The chemistry that made the route useful arrived separately. Ionizable cationic lipids, designed against a proposed mechanism of action, carried small interfering RNA at doses as low as 0.01 mg/kg in rodents [8], and the same particle class later carried mRNA into clinical use [9, 10].

Mixing time became the design variable once the vesicles were made in flowing streams. A confined impinging-jet mixer was characterized through a competitive reaction set as a chemical ruler, over characteristic times from 320 ms down to 5 ms [11]. Microfluidic mixing in milliseconds then produced limit-size particles, the smallest aggregates compatible with the packing of their constituents, with diameters engineered from 20 to 80 nm by composition [12], and the same technique gave siRNA particles of 20 nm and larger with essentially complete encapsulation and polydispersity indexes as

low as 0.02 [13]. Cryogenic microscopy and X-ray evidence then described how the particles form, with ionizable lipid not complexed to cargo promoting small vesicular structures that coalesce into larger particles with dense cores [14]. Coalescence is the mechanism the present model exploits, because a process that can start it can also stop it.

> **DID YOU KNOW? A Hundred Channels at Once**
>
> Scale-out is real practice. An array of 128 mixing channels operating simultaneously gave more than 100 times the production rate of a single channel, with no loss of the particles' physical properties or of their biological potency [4].

The step this simulator is built on is the deliberate separation of the two stages. Mixing at a symmetric flow ratio, with the higher ethanol content that implies, holds the particles in a fusion-prone condition, and an aqueous buffer injected after a defined residence time arrests them at the size they have reached, with the size and the morphology controlled independently and without giving up the other quality attributes [1]. That source is the anchor of the application. Its reported mean sizes, tabulated over four residence times and five buffer strengths for several cargoes, are the data the simulator's size law was fitted to, and its qualitative statements about ethanol, about PEG-lipid saturation as what finally arrests fusion, and about a conductivity probe as a cheap handle on the buffer strength are what the application's three declared extension factors stand on. The flowsheet modeled here was assembled for teaching, out of evidence that is public. It is not a reproduction of any commercial process, no product or company is named, and the quality limits in the application are teaching values rather than a specification. Real limits for a particle product come from the product's own specification and from regulatory guidance, which for liposome drug products asks for particle size and distribution profile among the characteristics of the drug product specification [15].

8.3 The Skid as Modeled

The skid is drawn in Figure 8.1. Four pumps feed the first impinging-jet mixer IJM-1 of each channel: P-101 carries the mRNA stock at concentration c_P of counted phosphate groups in buffer S_1, P-102 a buffer concentrate S_2, P-103 water for injection, and P-104 the lipid solution in ethanol at c_N of counted ionizable nitrogen. Each has its own flow meter and its own flow loop, and a conductivity probe CT-101 reads the aqueous blend. The mixed stream crosses one of five aging loops between 3.3 µL and 667 µL, reaches IJM-2 where pump P-105 injects R_q volumes of quench buffer, and flows through a collection line of four well-mixed cells to the divert valve XV-401, the hold tank and a buffer exchange that removes the ethanol and hands over the formulated bulk. One, two or four identical channels run off the same manifold.

Table 8.1 sorts the variables by what they do. Four operator demands act through one ratio block, four flow loops and the aging loop in service, and three quality attributes and a rate are what the skid is judged on. Units are SI inside the model, which displays mL/min, µL, nanometers, mol/m^3 and bar.

8.3.1 Four Demands on Four Feeds

The operator asks for a total flow per channel, a flow ratio, a charge ratio and an aqueous buffer strength. The ratio block FY-100 solves the four feed flows that realize all four at once, as shown in Equation (8.1). It is the ratio station of Section 1.4.4 with three ratios held instead of one, and each of its outputs is the setpoint of a flow loop.

$$
\begin{aligned}
Q_o &= \frac{\mathrm{TFR}}{1+\mathrm{FRR}}, \qquad Q_a = \mathrm{TFR} - Q_o = Q_1 + Q_2 + Q_3 \\
Q_1 &= \frac{Q_o\, c_N}{(N/P)\, c_P}, \qquad Q_2 = \frac{S\,Q_a - Q_1 S_1}{S_2} \\
Q_3 &= Q_a - Q_1 - Q_2, \qquad Q_q = R_q\,\mathrm{TFR}
\end{aligned}
\tag{8.1}
$$

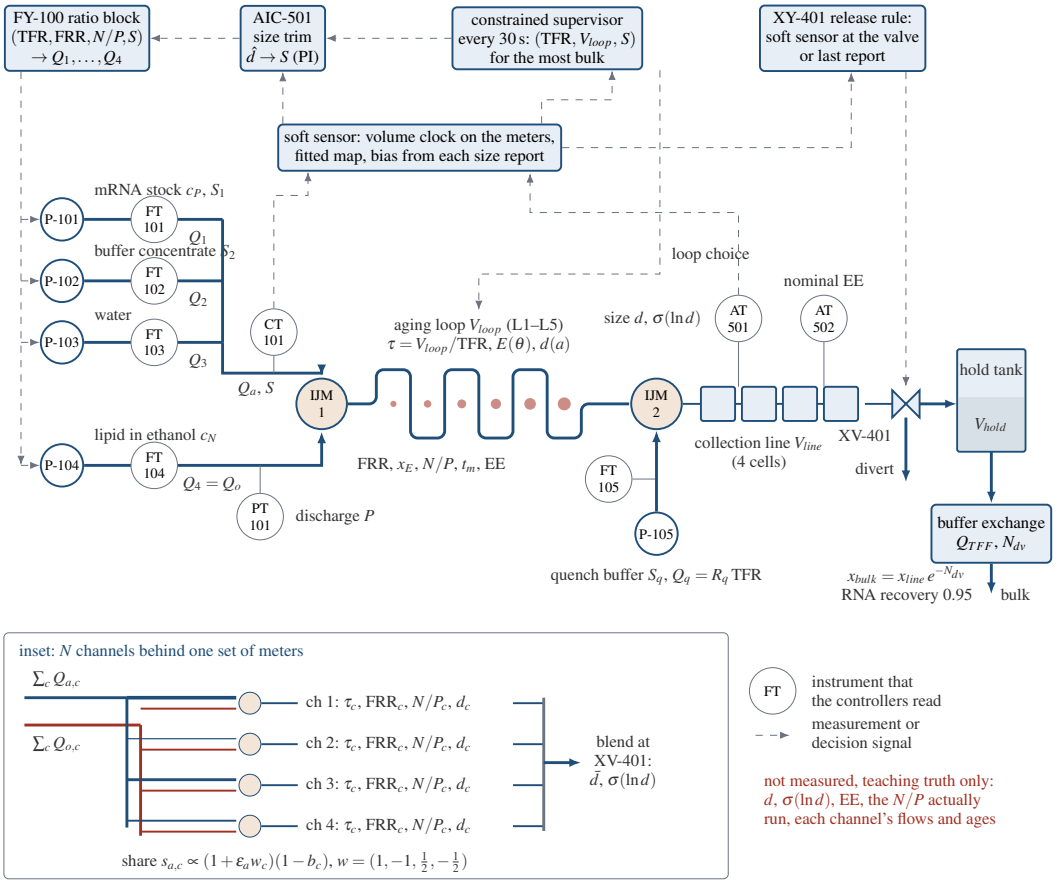


Figure 8.1 One formulation channel of the simulated skid, with the manifold of the parallel channels as an inset. Four metered feeds set the total flow TFR, the flow ratio FRR, the charge ratio N/P and the aqueous buffer strength S at IJM-1, which also fixes the ethanol fraction x_E, the mixing time t_m and the nominal encapsulation EE. The aging loop V_{loop} is the only place the particles grow, over the age distribution $E(\theta)$ of its volume clock. IJM-2 arrests them with Q_q, the collection line carries the eight transported quantities to XV-401, and the buffer exchange takes the residual ethanol out. Instruments the controllers read are circled, signals are dashed, and the legend lists the states of the model that no instrument on the skid reports. In the inset the organic header is the lower line of each pair. Process model only, not a piping diagram. Not to scale.

where c_N and c_P are molar concentrations of counted ionizable nitrogen and counted RNA phosphate rather than of lipid and RNA molecules. With an mRNA stock and a lipid solution alone the first and third relations share Q_1, so tripling the flow ratio at fixed feed compositions divides the charge ratio by three. Measured on the application with two feeds, a 1:1 to 3:1 step takes N/P from 3.00 to 1.00, while the four-feed block holds 6.000 to 6.000. The buffer concentrate and the water are therefore the two extra degrees of freedom. The block keeps Q_2 between zero and $Q_a - Q_1$ and Q_3 at or above zero, and it reports when either limit binds: at FRR $= 0.50$ the mRNA stock alone fills the aqueous stream and the realized strength is 26 mol/m^3 of the 100 mol/m^3 asked. The quench dilutes both components in the same ratio, so it cannot change N/P.

At the design point the declared feed concentrations give $c_N = 5\,\mathrm{g/L} \times 0.50/0.648\,\mathrm{kg/mol} = 3.858\,\mathrm{mol/m^3}$ and $c_P = 0.425\,\mathrm{g/L}/0.330\,\mathrm{kg/mol} = 1.288\,\mathrm{mol/m^3}$, and Equation (8.1) returns 9.986, 3.501, 6.514 and 20.000 mL/min with 80 mL/min of quench.

Table 8.1 Variable Classification and the Design-Point Values

Variable	Role	Design	Range
d	CV, particle size (true)	139.9 nm	target ±15 nm
$\sigma(\ln d)$	CV, modeled size spread	0.349	≤ 0.45
EE	CV, nominal encapsulation	92.7 %	≥80 %
S	MV of the size trim	100 mol/m^3	25–300
V_{loop}	MV of the supervisor	78 µL	3.3–667
TFR	MV of the supervisor	40 mL/min	10–80
$Q_1 \dots Q_4$	MV, feed flows	9.99/3.50/6.51/20.0	0–160 mL/min
Q_q	MV, quench flow	80 mL/min	0–1200 mL/min
$u_1 \dots u_5$	MV, pump speeds	25 %	0–100 %
XV-401	MV, accept or divert	accept	—
FRR	demand, flow ratio	1.0	0.5–5
N/P	demand, charge ratio	6.0	3–10
R_q	demand, quench ratio	2.0	0.5–4
ε_a, ε_o	DV, maldistribution	0	±30 %
b_c	DV, channel restriction	0	0–80 %
δ_L	DV, lipid lot and drift	0	±30 %
δ_S	DV, concentrate lot	0	±30 %
A, f	DV, pump pulsation	0	0–60 %, 0–5 Hz
ϕ	state, lipid volume fraction	2.49×10^{-3}	—
W, V_2	states, log-size moments	—	—
a	state, material age	6.1 s at the valve	—
c_R, e_R	states, total and held RNA	—	—
x_E	state, ethanol fraction	16.6 % in line	—
salt	state, buffer in the stream	—	—
Q_i	states, pump deliveries	as above	—
P	output, discharge pressure	2.51 bar	alarm 6 bar

Pause. Try Exercise 8.2 before reading on.

8.3.2 The Volume Clock and Its Age Distribution

The aging loop is between 5 ms and 1000 ms long while the physics step is 50 ms, so the loop is not integrated. It is a transport clock on cumulative volume with a normalized residence-time distribution, as shown in Equation (8.2).

$$\begin{aligned} \Theta(t) &= \int_0^t \mathrm{TFR}(t')\,\mathrm{d}t', \qquad \theta = \frac{\Theta(t) - \Theta(t-a)}{V_{loop}} \\ E(\theta) &= \frac{n^n \theta^{n-1} \mathrm{e}^{-n\theta}}{(n-1)!}, \qquad \int_0^\infty E(\theta)\,\mathrm{d}\theta = 1 \end{aligned} \tag{8.2}$$

where n is the number of equal stirred volumes the loop is modeled as, a teaching value of 24 whose distribution has $\sigma/\tau = 0.204$. A packet's age a is found by inverting the cumulative volume, never by taking the instantaneous V_{loop}/TFR. The application evaluates E at nine equal-probability nodes, whose dimensionless throughputs run from 0.699 to 1.347 and whose mean is renormalized to exactly one, so the material reaching IJM-2 is a volume average over a distribution of ages. At a constant flow the clock is exact and $\tau = V_{loop}/\mathrm{TFR}$, which is 117.0 ms for the 78 µL loop at 40 mL/min.

Pause. Try Exercise 8.1 before reading on.

8.3.3 The Size Law

Each node carries the fitted map of Equation (8.3), which is the model form of the anchor: an initial self-assembly size that follows a modified Gompertz curve in the aqueous buffer strength, plus a fusion-driven growth term that is a power law in the age.

$$\begin{aligned} d(a) &= d_0(S) + R(S)\,a^p, \qquad u(S) = \mathrm{e}^{-\mathrm{e}^{-D(S-C)}} \\ d_0(S) &= a_0 + a_1 u(S), \qquad R(S) = r_0 + r_1 u(S) \end{aligned} \tag{8.3}$$

where a is in milliseconds and d in nanometers. The seven coefficients were fitted inside the application, as Section 8.4 records, because the anchor's own coefficients are in supplementary material that could not be read.

Growth is fusion, so it conserves particle volume. The model transports the lipid volume fraction ϕ with no source term anywhere and derives the particle number from it and the diameter, as shown in Equation (8.4).

$$\begin{aligned} \frac{\mathrm{d}d}{\mathrm{d}a} &= p\,R(S)\,a^{p-1} f_E(x_E)\, f_d(\phi)\, f_P(d) \\ f_E &= \left(\frac{x_E}{0.5}\right)^{6}, \quad f_d = \frac{\phi}{\phi_{ref}}, \quad f_P = \min\!\left(1, \frac{900-d}{500}\right)^{+} \\ N &= \frac{6\phi}{\pi d^3}, \qquad \left.\frac{\mathrm{d}\phi}{\mathrm{d}t}\right|_{\text{growth}} = 0 \end{aligned} \tag{8.4}$$

where x_E is the ethanol volume fraction and ϕ_{ref} the lipid volume fraction of the design mixed stream. The three factors are declared teaching closures, each resting on one qualitative statement of the anchor rather than on data: ethanol raises membrane fluidity and lowers the barrier to fusion, fusion is a collision process whose rate falls with dilution, and PEG-lipid accumulating on the shrinking surface eventually arrests it. Their measured weights differ by three orders of magnitude. Two volumes of quench multiply f_E by 0.00137 and f_d by 0.333, so the arrest is the ethanol's, and f_P is exactly one below 400 nm and is never engaged by any shipped case. The quadrature of the growth integral runs in $u = a^p$ rather than in a, because the integrand carries an a^{p-1} singularity at zero. In that variable the integrand is bounded and the plant reproduces Equation (8.3) to round-off whenever the ethanol content is constant along a packet's path.

8.3.4 What the Collection Line Carries

Everything downstream of the quench junction is a state. Each channel's collection line is four well-mixed cells carrying eight quantities that mix linearly, as shown in Equation (8.5).

$$\begin{aligned} \frac{\mathrm{d}c_k}{\mathrm{d}t} &= \frac{Q_{line}}{V_{cell}}\,(c_{k-1} - c_k) + s_k, \qquad V_{cell} = \frac{V_{line}}{4} \\ d &= d_{ref}\,\mathrm{e}^{W/\phi}, \qquad \sigma(\ln d)^2 = \frac{V_2}{\phi} - \left(\frac{W}{\phi}\right)^2 \end{aligned} \tag{8.5}$$

where the eight are ϕ, $W = \phi \ln(d/d_{ref})$, $V_2 = \phi[(\ln d/d_{ref})^2 + v_0]$, the age a with source term one, total and encapsulated RNA c_R and e_R, the ethanol fraction x_E and the buffer in the whole stream. The size distribution is modeled as lognormal in d, so two streams mix as two lognormals do and the reported size is the geometric mean diameter. The growth source terms are $s_W = \phi g/d$ and $s_{V_2} = 2\phi \ln(d/d_{ref})\, g/d$ with g from Equation (8.4), which is why fusion moves W and V_2 and never ϕ. The quantity $\sigma(\ln d)$ is the modeled width of that lognormal. It is not the cumulants polydispersity index a light-scattering instrument reports [16], the transform between the two is not modeled, and the application never prints a polydispersity index for its own output. The distinction is the one a measurement study of dynamic light scattering makes when it warns that central values taken incautiously from a scattering distribution lead to interpretation errors [17].

8.3.5 Encapsulation, Pressure and the Exchange

Nominal encapsulation is set at the mixer and then carried along, as shown in Equation (8.6).

$$\begin{aligned} \mathrm{EE} &= \mathrm{EE}_{max}\,\frac{1}{1+(t_m/t_{m,1/2})^2}\cdot\frac{1}{1+(2.6\,P/N)^4} \\ t_m &= t_{m,0}\left(\frac{Q_0}{\mathrm{TFR}}\right)^{3/2}, \quad P = 10^5 + k_{res}(1+3b_c)Q_{ch}^{1.8} \\ \frac{x_{bulk}}{x_{line}} &= \mathrm{e}^{-N_{dv}}, \qquad N_{dv} = \ln\!\left(\frac{x_{line}}{x_{max}}\right) + m_{dv} \end{aligned} \tag{8.6}$$

where t_m is the jet mixing time, scaled from its design value by engulfment so that it falls as $\mathrm{TFR}^{-3/2}$, Q_{ch} is the channel's whole flow including the quench, b_c a channel restriction and x_{max} the teaching limit on residual ethanol. The mixing term is fitted to the single degraded operating point the anchor quotes, 60 % encapsulation at 0.4 mL/min against 82 % at 8 mL/min, with the exponent declared rather than fitted because one end point cannot identify both. The charge term collapses below about four nitrogens per phosphate. This is nominal encapsulation, the fraction an exclusion assay cannot reach, not intact functional RNA, and no statement about potency or transfection is made or implied anywhere in the application. The diafiltration relation is exact for a constant-volume exchange, and at the design point 16.62 % ethanol in the line needs 3.50 diavolumes to reach the 0.5 % teaching limit, against the 3.81 the controller runs with its margin.

The manifold splits each feed among the channels in shares that always add up to one, so it can neither create nor lose flow. The aqueous share of channel c is proportional to $(1+\varepsilon_a w_c)(1-b_c)$ with the pattern $w=(1,-1,\frac{1}{2},-\frac{1}{2})$, and the organic share follows ε_o in the same way. The total flow meters see nothing of either.

8.3.6 Parameters, Evidence, Initialization and Numerics

Table 8.2 gives every parameter with its value and the label its evidence earns. Sizes, delays and limits that no source fixes are teaching values chosen for the lesson, and they are labeled as such in the application's Model dialog and on its Instructions page as well.

Every case opens lined out on its own settings. Equation (8.1) is solved open loop, the loop and the line are filled with material of the age that flow would give it, the meters and the analyzers open on that state, and the soft sensor's history is primed back over twice the longest loop. Only a start-up case opens with its pumps stopped. States advance by fourth-order Runge-Kutta at a fixed 0.05 s, a step set by the pump drives ($\tau = 0.8$ s) and one collection-line cell (3 mL at 120 mL/min, 1.5 s), and the controllers execute every 0.2 s. Halving the step moves the particle size of the quench case by 0.00 nm, of the pulsation case by 0.05 nm and of the challenge by 0.18 nm, and the accepted bulk by less than 10^{-4} relative. The challenge's 1500 s needs 1.8 s of computing headless and 37.5 s of wall time in a browser at the top of the speed ladder of Section 1.6.1, where the full 40-fold rate is realized.

8.3.7 Measurements and the Truth Boundary

The controllers, the soft sensor, the supervisor and the release rule read one packet and nothing else: five flow meters with a 0.5 s lag and 0.4 % noise, their manifold total, the conductivity probe with 1.5 mol/m^3 of noise and whatever drift a case applies, the discharge pressure transmitter, and two sampled analyzers. The size analyzer draws a sample every 90 s and reports 60 s later, so a result is between 60 and 150 s old, and the encapsulation analyzer reports 420 s after its own sample. Noise is drawn once per scheduled measurement, never inside the derivative function and never while drawing. The true particle size, the true spread, the true encapsulation, the age distribution, the charge ratio actually run and every per-channel flow are states of the model with no instrument behind them. They appear in the scene as dashed teaching truth, in the scorecard, and in the exported file under names prefixed 'truth_', and one of the application's own checks fails when a controller touches any of them.

8.3.8 Conservation and Numerical Checks

The application's own test suite prints the residual of every balance rather than clipping it. Over the disturbed four-channel case the volume, RNA and lipid ledgers close to 3.9×10^{-15}, 7.4×10^{-15} and 2.0×10^{-15} of what was fed. Because the aging loop is evaluated rather than integrated, its own hold-up is carried as two explicit inventory states per channel. Without them the RNA and lipid balances closed only to 2.4×10^{-4} during a transient, which was the loop's hold-up and not a solver defect. The residence-time distribution integrates to one at every node count tested, growth has no source term in ϕ, and the priming execution feeds nothing, advances no clock, draws no sample and updates no bias.

Table 8.2 Parameters and the Evidence Each One Carries

Symbol	Meaning	Value	Source
a_0, a_1	assembly size, offset and span, nm	61.84, 65.98	fitted
r_0, r_1	growth, the same, nm/ms$^{\mathrm{p}}$	0.154, 1.653	fitted
C, D	Gompertz center mol/m^3, slope	81.03, 0.0256	fitted
p	growth exponent in the age	0.7785	fitted
x_{fit}	ethanol fraction the map was fitted at	0.50	source
n_E	ethanol exponent in f_E	6	teaching
ϕ_{ref}	design lipid volume fraction	2.50×10^{-3}	derived
d_{knee}, d_{stop}	PEG guard, nm	400, 900	source, teaching
n, M	tanks in series, quadrature nodes	24, 9	teaching
v_0	spread one aging leaves, $(\ln d)^2$	0.12	teaching
EE_{max}	encapsulation plateau	0.96	teaching
$t_{m,0}$	jet mixing time at 40 mL/min, ms	0.30	teaching
$t_{m,1/2}$, n_{mix}	mixing half point ms, exponent	406, 2	fitted, declared
c_{RNA}	mRNA stock, g/L	0.425	derived
c_L	lipid solution, g/L	5.0	source
S_1, S_2, S_q	stock, concentrate, quench mol/m^3	25, 500, 100	source, teaching
V_{loop}	five loops, µL	3.3–667	source times
V_{line}	collection line per channel, mL	12	teaching
Q_{max}, $Q_{max,q}$	pump capacities, mL/min	160, 1200	teaching
τ_{pump}, τ_{FT}	drive lag, meter lag, s	0.8, 0.5	teaching
k_{res}	channel resistance (SI)	2.7×10^{15}	teaching
P_{alarm}, P_{trip}	pressure alarm and trip, bar	6, 8	teaching
V_{hold}	hold tank, L	2.0	teaching
Q_{TFF}	exchange throughput, mL/min	700	teaching
rec_{TFF}	exchange RNA recovery	0.95	teaching
t_{SR}, t_{cyc}	size result, sampling, s	60, 90	source, teaching
t_{EE}	encapsulation result, s	420	teaching
σ_{FT}, σ_{cond}	meter, probe noise	0.4 %, 1.5 mol/m^3	teaching
σ_{size}	size assay noise	2.5 %	teaching
Δd	size window, nm	15	teaching
σ_{max}	modeled spread limit	0.45	teaching
EE_{min}	encapsulation minimum	80 %	source standard
x_{max}	residual ethanol, % v/v	0.5	teaching
N/P window	charge ratio, low and high	5.0, 7.0	teaching

Table 8.3 Verification of the Implementation

Check	Result
Volume ledger, fed minus line change minus out	3.9×10^{-15}
RNA ledger, the same	7.4×10^{-15}
Lipid ledger, the same	2.0×10^{-15}
Lipid volume fraction during growth, $\lvert \mathrm{d}\phi/\mathrm{d}t \rvert$	0 exactly
Residence-time distribution, mean throughput	1.000000000000
Charge ratio across a 1:1 to 3:1 step, four feeds	$< 10^{-9}$
Nine nodes against 2001, mean size at 440 ms	0.75 per mille
Step halved, quench case size	0.00 nm
Step halved, pulsation case size	0.05 nm
Step halved, challenge size | bulk	0.18 nm | $< 10^{-4}$
Playback speed, one run against thirty chunks	bit-identical
Each case loaded three ways	identical
Priming step: volume fed, samples, bias updates	0, 0, 0
Wall time, 1500 s challenge, headless	1.8 s
Wall time, the same in a browser at 40×	37.5 s

8.3.9 Assumptions

The size map is fitted to one published table: one formulation family of four lipids, one cargo type, pH 5.5, one mixer geometry, buffer strength 25 to 300 mol/m^3 and aging 5 to 440 ms. The application draws that box and says when the operating point has left it. Fixed pH is a real restriction, because the anchor shows pH to be an independent handle on size and morphology, and its own pH 4.5 table is the

one this map does not reproduce. Morphology is not modeled at all, and no bleb fraction, transfection or potency is predicted. Encapsulation is nominal and is set at the mixer rather than evolving. The loop's age distribution is that of 24 stirred volumes in series, whereas a real coil carries a parabolic velocity profile and secondary flow, whose dispersion is a different and usually wider shape around the same mean [18]. The buffer exchange is a declared recovery map with an exact diafiltration law, not a membrane model, and the retention of a 100 nm particle by a real device is outside it. Every parallel channel uses the same per-channel model with its own actual flows, so more throughput across more channels is not evidence that a larger channel would behave the same way. All quality limits are teaching values for this simulation, and a simulated pass is not product release.

8.4 Calibration and Its Domain

The anchor publishes the model form of Equation (8.3) and its own five coefficients in supplementary material whose endpoint returned HTTP 403, so the coefficients could not be read. The seven coefficients of Table 8.2 were therefore fitted inside the application, by least squares on $\ln d$, to the twenty mean particle sizes the anchor's Figure 2b5 tabulates for one cargo at pH 5.5 over four residence times and five buffer strengths [1]. The fit reproduces that table to 5.11 % rms with a worst cell of 10.7 %. A second derivation by damped Gauss-Newton from 200 random starts, written independently of the first, lands on the same coefficients to six significant figures and the same sum of squares, 5.2764×10^{-2}.

Agreement with the source is a fit inside its domain plus two withheld checks, not a validation. The two cargo tables that were not fitted, a second mRNA and a yeast RNA at the same pH, are reproduced at 8.76 % and 7.96 % rms with no refitting, as shown in Figure 8.2. The fourth table, taken without cargo at pH 4.5, is not reproduced at all (28.15 % rms, worst cell 79.7 %), which is why pH 5.5 is a declared boundary of the domain rather than an untested extrapolation.

Two further checks were run for this chapter and are not in the application. The first withholds one buffer row at a time and refits: the interior rows are predicted at 4.8 to 6.4 % rms, while the 25 mol/m^3 and 300 mol/m^3 edges are predicted at 20.5 and 16.5 %, because each edge is held by its own row and is not interpolated. The second withholds one residence-time column at a time. The 50, 117 and 440 ms columns are predicted at 9.6, 10.9 and 90.6 % rms, and the 5 ms column at 34.8 %. The power law cannot be extrapolated to the longest time from the shorter ones, and 90.6 % is the quantitative case for the window guard that stops the supervisor from leaving the box. One independent check of the reading itself: the anchor's Figure 4 quotes five case studies at pH 5.5 by their conditions and sizes as 70.7, 67.6, 113, 100 and 78 nm, and the corresponding table cells read 70.7, 67.7, 112.8, 100.7 and 78.6 nm.

The three extension factors of Equation (8.4) carry the map to conditions the table does not cover, and their weights were measured by changing each one for the plant, the soft sensor and the supervisor together and rerunning every case. The dilution factor carries almost nothing. Removing it entirely moves the quench case's 215.8 nm by 0.4 nm and tripling it by 0.7 nm, and no verdict changes. The PEG guard is inert in every shipped case, since none reaches 400 nm even with the window guard off. The ethanol exponent carries the rest. At $n_E = 3$ and 12 the direction of every lesson survives and the numbers do not: the 3:1 size moves from 98.0 to 103.3 and 97.4 nm, the separation between the richest and leanest channel of the four-channel case from 36.5 to 29.6 and 72.8 nm, and at $n_E = 12$ the challenge's suggested settings no longer pass, making 358.6 mL/min with 7.2 L of off-specification material accepted. Every quantitative statement this chapter makes about a flow-ratio effect rests on that one teaching exponent.

8.5 Implemented Control System

Every loop executes at $\Delta t_c = 0.2$ s, which is also the instrument scan. Within one physics step the order is the controllers, then the integrator, then the plant's own bookkeeping, and last the outputs and the historian.

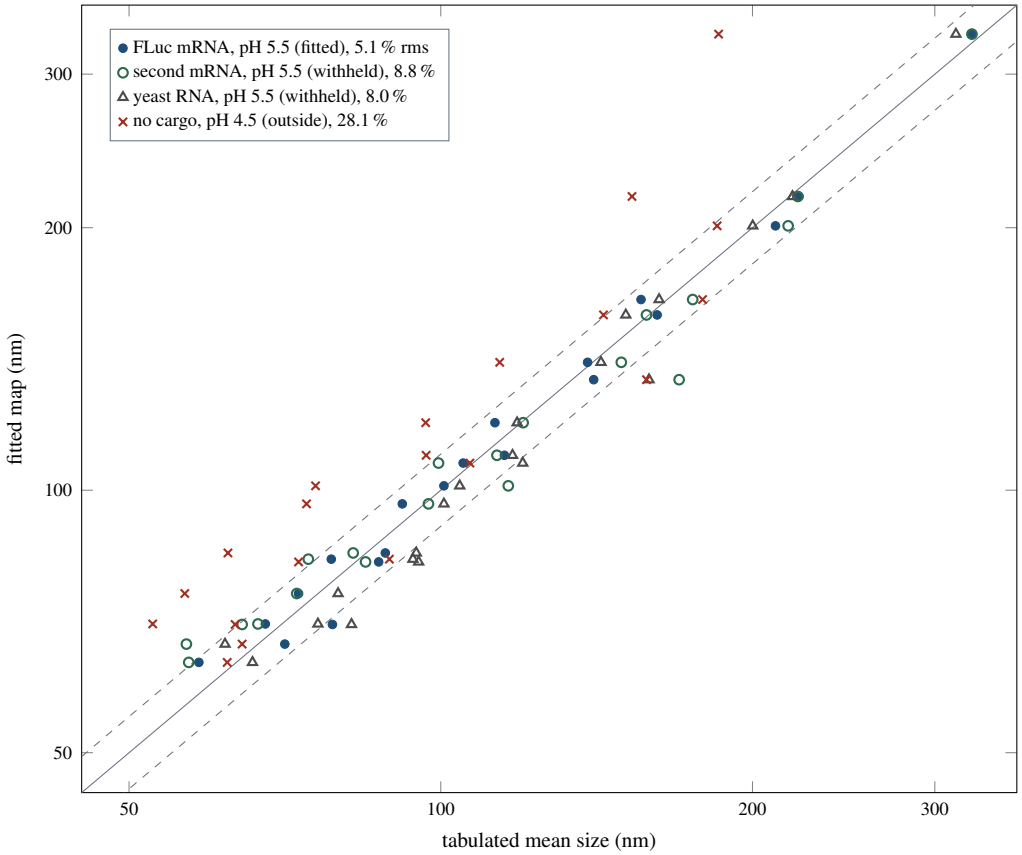


Figure 8.2 The fitted map against the tabulated mean sizes it was fitted to and the three tables withheld from the fit: filled circles for the fitted table, open circles and triangles for the two withheld cargo tables, and crosses for the table without cargo. Dashed lines are $\pm 10\,\%$. The two withheld cargo tables at pH 5.5 are reproduced at 8.8 and 8.0 % rms with no refitting, while the pH 4.5 table lies off the diagonal at 28.1 %, which fixes pH 5.5 as a boundary of the declared domain.

8.5.1 The Flow Loops

Each of the five pumps runs on the velocity-form PI algorithm of Section 1.6.2, with the error taken on the metered flow in mL/min, the output in percent of pump speed clamped to 0 and 100, back-calculation into the integral term over a tracking time of τ_I, and bumpless transfer from MANUAL. The four feed loops share $K_c = 1.2\,\%$ per mL/min and $\tau_I = 3.0\,\mathrm{s}$, and the quench loop runs at 0.25 and 4.0 s, with no derivative action anywhere. The ratio block's demand passes through a rate limit on the total flow, a declared 120 mL/min per minute, which is what makes the start-up case a lesson rather than a five-second event. A pump asked for more than its capacity reports the request beside what was applied.

8.5.2 The Size Trim and the Soft Sensor

AIC-501 is a slow PI on the same algorithm that moves the aqueous buffer strength to hold the size, with $K_c = 1.2\,\mathrm{mol/m^3}$ per nanometer, $\tau_I = 90\,\mathrm{s}$ and output limits of 25 and 300 $\mathrm{mol/m^3}$. That knob was chosen because the ratio block can move it without touching the flow ratio or the charge ratio, which is the whole reason for the two extra feed pumps. It reads either the last reported size or the

soft sensor. The same shape, a ratio blend trimmed on a composition analyzer, holds the sweetness of a soft drink in Chapter 7.

The soft sensor rebuilds, from the meters and the probe alone, the ratios on the label concentrations, its own volume clock with the same normalized distribution the plant uses, the fitted map at each of the nine ages, and the encapsulation closure. Each size result then updates a bias on the logarithm of the diameter, as shown in Equation (8.7).

$$b \leftarrow b + \lambda \left[\ln d_{lab} - (\ln \hat{d}_s + b) \right], \qquad \hat{d} = \hat{d}_{raw}\, \mathrm{e}^{b} \tag{8.7}$$

where $\hat{d}_s$ is the sensor's own prediction for the sample that was drawn, recovered from a ring of past predictions, rather than its prediction now. The same rule corrects the spread and the encapsulation on their own channels, bounded at 0.5, 0.4 and 0.15. The gain λ is 0.15 per result and was chosen by measurement, not by data: the result is 60 s old and arrives every 90 s while the trim acts in one scan, so the loop is delay-dominated, and at $\lambda = 0.4$ the challenge's accepted off-specification volume rises from a mean of 43.1 mL to 197.4 mL over the same six paired seeds at the same bulk rate. A soft sensor whose bias a late laboratory corrects also carries the moisture loop of Volume 5, Chapter 10, and Volume 5, Chapter 4 reconciles a residence-time predictor with an analyzer that reports hours after the action.

🕮 **Pause.** Try Exercise 8.3 before reading on.

8.5.3 Three Policies and a Release Rule

The policies read the same packet and the same hardware. FIXED RATIOS sends the operator's numbers to FY-100 and corrects nothing. COORDINATED FLOW AND QUENCH adds AIC-501 and automatic diavolumes. CONSTRAINED SUPERVISOR adds the choice of operating point: every 30 s it evaluates nine total flows from 20 to 80 mL/min crossed with the five aging loops, bisects the buffer strength that puts the predicted size on target, and takes the point of greatest bulk rate subject to the predicted size, the modeled spread, the encapsulation, the pressure alarm, the exchange throughput, and, with the window guard on, the box the map was fitted in, as shown in Equation (8.8).

$$\begin{aligned} \max_{\mathrm{TFR},\, V_{loop},\, S} \quad & \min\left(Q_{line}, Q_{TFF}\right) rec_{TFF} \\ \text{s.t.} \quad & |\hat{d} - d_{SP}| \le m_s \Delta d \\ & \hat{\sigma} \le \sigma_{max}, \quad \hat{\mathrm{EE}} \ge \mathrm{EE}_{min} \\ & P(\mathrm{TFR}) \le P_{alarm}, \quad Q_{line} \le Q_{TFF} \\ & 5\,\mathrm{ms} \le \tau \le 440\,\mathrm{ms} \\ & 25 \le S \le 300\,\mathrm{mol/m^3} \end{aligned} \tag{8.8}$$

where $m_s = 0.6$ is a margin on the size window and $\hat{\sigma}$ adds the variance of $\ln d$ over the age nodes to the within-node spread. Keeping the present loop is worth 8 % more bulk and no loop change may follow another within 150 s, because moving the loop puts a transient through the line. When nothing is feasible the search takes the point that misses the size by least from inside the box and names *size* as the binding constraint. It is a constrained steady-state search on a cadence, not a receding-horizon optimization, and nothing in the application is called model predictive control.

The release rule at XV-401 is a separate decision with three settings: accept everything, accept while the last reported size is inside the window, or accept while the soft sensor says the material now at the valve is inside it. The third asks the sensor what it predicted one line volume ago, computed from the meters, because the sensor predicts the quench junction while the valve is 12 mL per channel downstream. Each rule keeps 20 % of every window in hand by default, diverts while the pressure trip is latched, and holds the valve shut until the line has turned over twice since the feeds started.

Table 8.4 The Nine Teaching Cases and the Challenge

#	Case and events	Teaching point
1	Flow ratio 1:1 to 3:1 at 240 s, charge ratio held	ethanol, not mixing, sets the window
2	Total flow to 20 at 200 and 72 mL/min at 500 s	the total flow is the clock
3	Lipid lot 22 % under its label at 150 s	a feed nothing meters
4	Aging loop 78 to 293 µL at 240 s, then quench 3.5 volumes	the opening experiment
5	P-104 head pulsates ±40 % at 180 s	a spread the meter cannot see
6	Release on the reported size, flow to 21 mL/min at 200 s	a result a minute old
7	Four channels, manifold 12 % out at 180, channel 2 restricted at 480 s	one meter, four channels
8	Start-up from stopped pumps	material with another history
9	Supervisor asked for 380 nm	a target outside the window
10	Challenge: four channels, four faults, 1500 s	the whole problem

8.5.4 Plant, Controller and Internal Models

The application keeps three models apart. The plant is everything in Section 8.3. The controller's model is the fitted map on label concentrations with its learned bias, which does not know a lipid lot that is not its label, a concentrate that is weak, a maldistribution or a restriction. The identified model is the relation a learner fits by hand between the aging time and the size, which is what cases 2 and 4 are for. The supervisor evaluates the controller's model only, including for the discharge pressure, so a hidden restriction never reaches it.

8.6 Nine Cases and a Challenge, Measured

The cases appear in Table 8.4. One reference configuration underlies them, the design point of Table 8.1 run on two channels under fixed ratios with the soft-sensor release rule, and each row names only what it changes. Measurement was done on the engine alone, at its default seed, on the step the browser uses, after the last behavior change. Loading a case first, picking it from the dropdown after another one and restarting it give the same trajectory to the last bit. Each case credits one decision, an action taken and a state the skid reached because of it, as Section 1.6.4 describes.

8.6.1 What Sets the Size

Case 4 is the opening experiment. At 200 s on the 78 µL loop the skid makes 139.9 nm with $\sigma(\ln d) = 0.349$ and 92.7 % encapsulation, and the loop change at 240 s takes it to 215.8 nm with $\sigma(\ln d) = 0.356$ and the same encapsulation by 420 s, as shown in Figure 8.3. That is 3.76 times the aging time for 1.54 times the size, because only the growth term grows: at 100 mol/m^3 of buffer the map assembles particles at 97.5 nm and adds 42.7 and 119.7 nm over the two ages, a factor of 2.81 that is $3.76^{0.779}$ to three figures.

> **TRY IT IN THE APP** **4 · Quench Timing · the Opening Experiment**
>
> Load the scenario and press Run. Watch the true particle size alone. It sits near 139.9 nm while the skid runs on the 78 µL loop, and when the loop goes to 293 µL at 240 s it climbs to 215.8 nm by 420 s. No flow, no ratio and no charge ratio moved.

Raising the quench ratio from 2.0 to 3.5 volumes on the same loop moves the size by 0.46 nm, because the ethanol is already below 17 % and the growth it allows is spent. Case 2 reaches the same conclusion along the other path: the 78 µL loop gives 234 ms at 20 mL/min and 65 ms at 72 mL/min, and the size follows at 171.1 nm and 124.5 nm rather than in proportion. Case 1 moves the flow ratio instead. At 3:1 the ethanol content falls from 50 to 25 %, f_E falls to 0.0156, and the size falls to 98.0 nm while the ratio block holds the charge ratio at 6.01 and every one of the four feed pumps moves.

8.6.2 Faults That No Flow Meter Sees

Three cases turn on measurements that do not exist. In case 3 a lipid lot 22 % under its label takes the true charge ratio from 6.00 to 4.62 and the true encapsulation to 87.6 % with every flow meter exactly on setpoint. The first encapsulation result that shows it arrives 11.5 min after the lot went on, and the size itself moves by less than 0.1 nm. In case 5 a pump head pulsating ±40 % at 1.2 Hz leaves its average delivery unchanged, so the meter sees almost nothing through its own lag, while the modeled spread rises from 0.349 to 0.400 and the mean from 140.5 to 150.4 nm, because the jets get a different flow ratio every half second and the loop carries every one of them forward. In case 6 the release rule reads the last reported size and the flow drops to 21 mL/min, which takes the aging to 223 ms and the size to 168 nm: 239.5 mL of off-specification material reaches the hold tank by 420 s, against 1.2 mL when the same rule reads the soft sensor.

Case 8 shows what a start-up makes. With the pumps ramping together the first material reaches 191.8 nm at 10 s, because the loop holds a fixed volume while the flow is still a fraction of its design value, and the valve first accepts at 21.9 s and keeps all but 0.8 mL of that material out of the tank. Case 9 asks for a size the fitted map cannot deliver. The largest in-window size is 333.4 nm, at the corner of the box, and a logged sweep of 5992 points over flows, loops and buffer strengths reaches 380 nm at none of them, missing by 46.6 nm at best. The supervisor saturates at 300 mol/m^3 on the 293 µL loop, reports *size* as the binding constraint and stays inside the box. With the guard off it picks the 1000 ms loop, predicts 371.6 nm and leaves 89 % of its age distribution outside the fitted window.

8.6.3 Scale-Out and What the Blend Hides

Four channels behind one set of meters is the application's second axis. In case 7 a 12 % aqueous maldistribution leaves the manifold total at 159.6 mL/min, inside the noise of its own meters, while the four channels run at flow ratios of 1.13, 0.89, 1.07 and 0.95 and make 126.6, 163.6, 132.9 and 150.8 nm. The blend is 142.8 nm with $\sigma(\ln d) = 0.364$ and meets every window. The only online sign of a restriction is pressure: channel 2 closing to 40 % of its share takes PT-101 from 2.49 to 4.29 bar against its 6 bar alarm, and the channels separate further, to 125.3 and 190.1 nm.

> **TRY IT IN THE APP** **7 · Four Channels, One Flow Meter**
>
> Run the scenario to 480 s and compare the blend with the channels. The manifold total stays at 159.6 mL/min while the four channels make 126.6, 163.6, 132.9 and 150.8 nm, and the blend is 142.8 nm inside its window. Then run past 480 s and watch PT-101 rise from 2.49 to 4.29 bar.

Figure 8.4 sweeps that fault from an even split to 30 %. The worst channel leaves the 15 nm window near 8 % and is 89.7 nm off target at 30 %, while the blend of four channels is 14.6 nm off and the blend of two is 22.9 nm off at the same fault. The blend's spread crosses its 0.45 limit between 24 and 27 % with two channels and reaches only 0.446 at 30 % with four. More channels therefore make the blend look better at the same fault: at 24 % the four-channel blend is 9.8 nm off target with three of its four channels outside their windows, against 14.8 nm off with both channels outside when there are two, and the two-channel blend is the only one that ever leaves the size window itself, which it does at 27 %. Averaging more streams is not the same as making them alike, which is the reason the application scores the blend and reports, separately, how much of the accepted volume was made while a channel was outside its window.

8.6.4 The Challenge and the Policy Comparison

The challenge runs four channels at a 200 nm target through 1500 s and four faults in sequence: the aqueous manifold drifts 12 % out at 180 s, a lipid lot 8 % under its label goes on at 420 s, a buffer concentrate 10 % under its label follows at 620 s, the lipid strength then drifts at −12 % per hour, and the lipid pump starts pulsating at 1080 s. The target is 560 mL/min or more of formulated bulk, at most 150 mL of off-specification material accepted, residual ethanol under 0.5 % and no overflow.

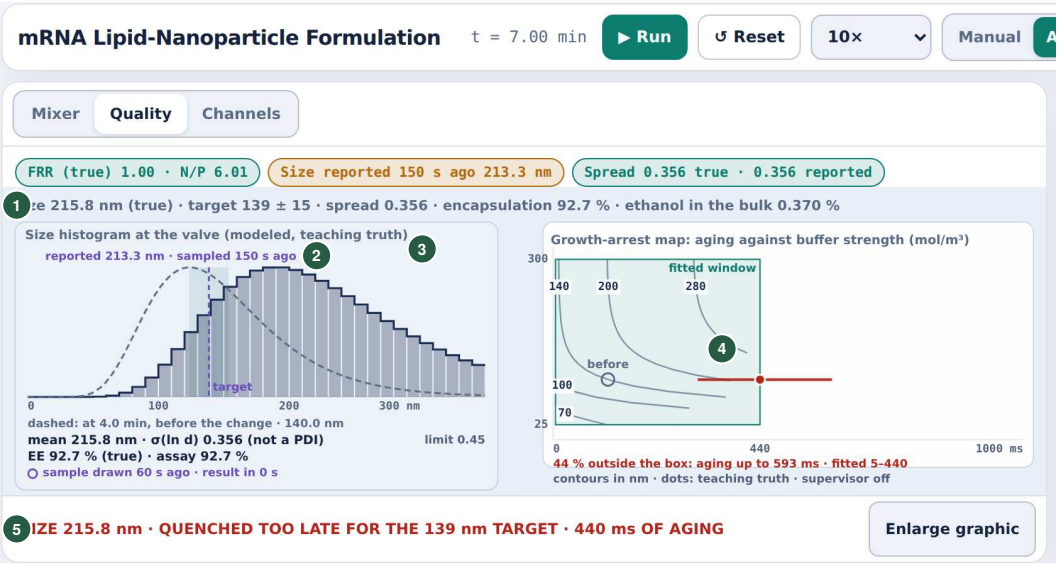


Figure 8.3 Case 4 at $t = 420$ s, three minutes after the aging loop went from 78 to 293 µL, on the Quality view in AUTO with fixed ratios and the soft-sensor release rule. The histogram at the valve is centered on 215.8 nm with $\sigma(\ln d) = 0.356$, and the dashed curve beside it is the distribution the same skid made at 240 s on the shorter loop, 140.0 nm. The at-line result, the filled circle, reads 213.3 nm from a sample drawn 150 s earlier, and its next sample is already 60 s old. The growth-arrest map puts the operating point on the 440 ms edge of the fitted box with 44 % of the age distribution beyond it. Encapsulation is 92.7 % true against an assay reporting the same value seven minutes after its own sample, a delay Table 8.2 prints, and the status line names the cause of the miss against the 139 nm target the case ships with. The picture is the application's window cropped to the header and the Quality view, with the Control panel, the trends and the scorecard cut. Callouts: (1) the summary line with the true size against the target and its window, (2) the at-line result and the age of its sample, (3) the size histogram at the valve with the dashed distribution from before the loop change, (4) the operating point on the growth-arrest map at the edge of the fitted box, (5) the status line.

The allowance is three line volumes, because when an unmeasured feed steps the material already between the mixer and the valve was made at the old condition and no release rule can reach it.

Table 8.5 compares every policy on the same plant, the same faults, the same instruments and six paired seeds. As shipped, fixed ratios at 32 mL/min per channel with the valve accepting everything make 372.5 mL/min and fail on rate at every seed, for the physical reason the case teaches: the quench triples the volume, so four channels at 32 mL/min present 384 mL/min to an exchange that passes 700, and the rest of the ceiling is unused. Raising the flow is what costs the size. Coordinated ratio control buys it back with the buffer strength and reaches 441.7 mL/min at 44 mL/min per channel, with the trim already at 273 mol/m^3 of the 300 mol/m^3 it has. At 48 mL/min it falls back to 415.7 mL/min, because the trim saturates and the product leaves the window. The supervisor, free to lengthen the loop as well, moves to the 293 µL loop at 56 mL/min per channel and reaches 600.0 mL/min, where the exchange and not the chemistry is the limit. **The timed quench is therefore worth 158 mL/min of bulk that ordinary ratio control cannot buy**, with the same hardware, feeds, faults and information. That margin is the answer to the question the application was built to ask, and it is a measurement on this plant and this formulation family, not a general claim.

The ablations in Table 8.5 say which layer does what. The size trim is what makes the rate: without it the supervisor's own point gives 368.8 mL/min and leaks 522.4 mL, because nothing corrects the

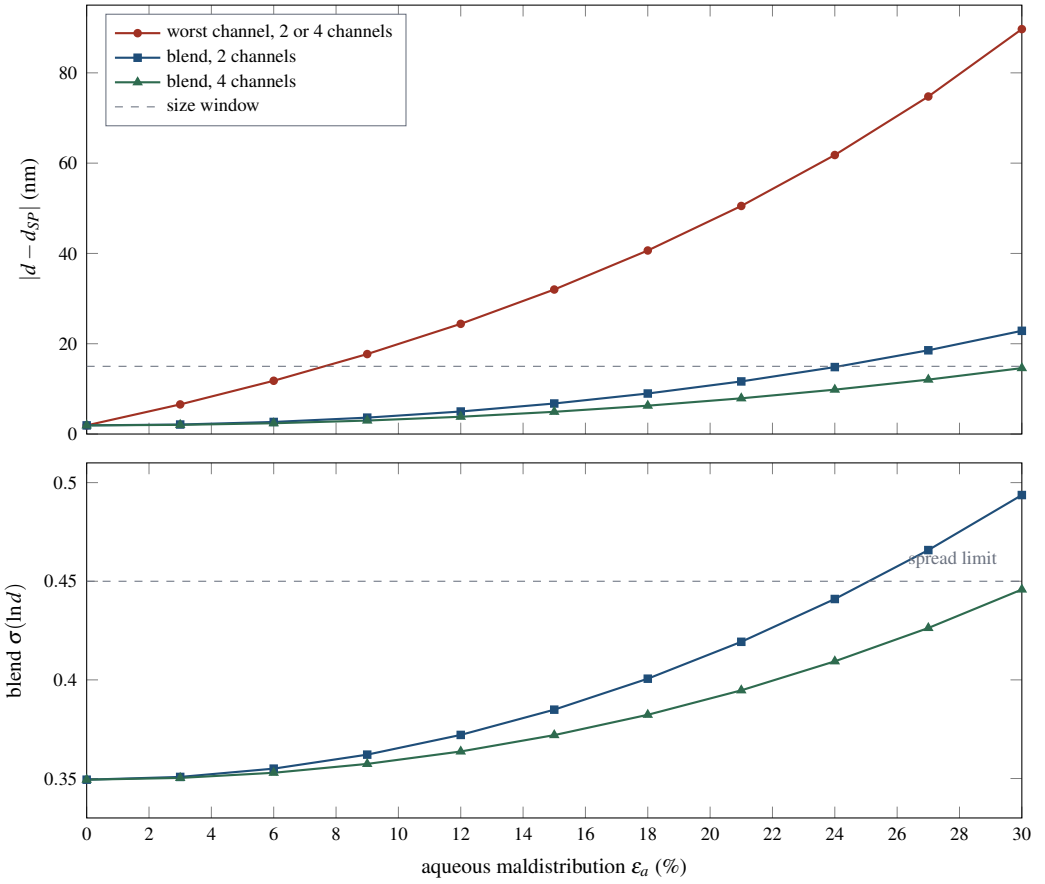


Figure 8.4 A steady aqueous maldistribution at the design recipe, read at 170 s. Above, the largest departure of any single channel from the target and the departure of the blend that leaves the valve, for two and for four channels. Below, the blend's modeled spread. Circles mark the worst channel, squares the blend of two channels and triangles the blend of four, and the dashed lines are the limits. The worst channel leaves the 15 nm window near 8 % while the blend stays inside it, and the blend's spread catches the fault only when it is severe.

concentrate lot. The conductivity probe is what makes the quality: the same supervisor on the metered blend runs 35.7 mL/min faster and leaks 776.7 mL, passing none of the six seeds, because no flow meter sees a drum that is 10 % weak while the probe on the blend line does. The window guard is worth more than either: without it the supervisor extrapolates, the trim chases a map that is wrong there, and the rate collapses to 200.4 mL/min with 1928.5 mL leaked. Figure 8.5 puts five of those rows against the challenge's two targets.

Table 8.5 The Challenge Over 1500 s, Six Paired Seeds. Rate Is Accepted Bulk per Minute and Off-Spec the Accepted Volume Outside a Window

Policy	Pass	Rate (mL/min)	Off-spec (mL)	Spread worst
fixed ratios, 32 mL/min, no gate (shipped)	0/6	372.5	0.0	0.441
with the soft-sensor gate	0/6	314.9	0.0	0.441
with the gate on the reported size	0/6	288.6	0.0	0.441
coordinated, 32 mL/min	0/6	314.4	526.8	0.446
coordinated, 40 mL/min	0/6	377.0	65.7	0.439
coordinated, 44 mL/min	0/6	441.7	0.0	0.434
coordinated, 48 mL/min	0/6	415.7	0.0	0.430
the supervisor's point, held by hand	0/6	619.2	1297.2	0.423
supervisor on the probe (suggested)	**6/6**	**600.0**	**43.1**	0.423
on the metered blend instead	0/6	635.7	776.7	0.491
without the size trim	0/6	368.8	522.4	0.411
without the bias update	0/6	600.3	349.7	0.422
with the bias update at 0.4	4/6	598.5	197.4	0.425
without the window guard	0/6	200.4	1928.5	0.647
with the gate on the reported size	5/6	581.0	51.5	0.423
with no release margin	6/6	628.1	47.9	0.423
with the valve accepting everything	6/6	644.8	63.9	0.423

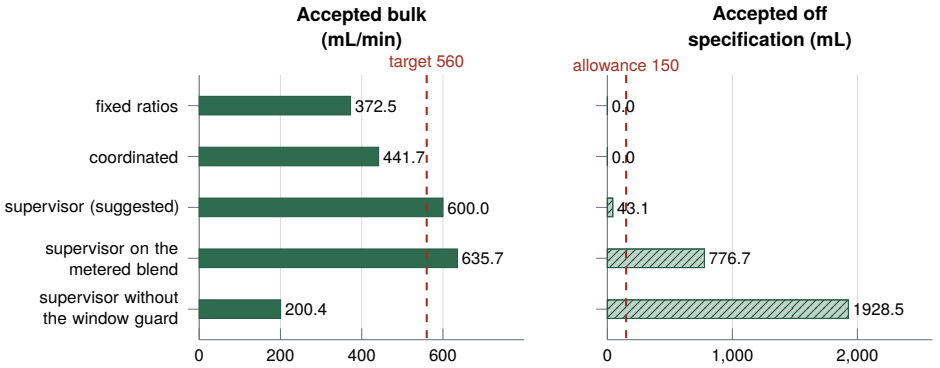


Figure 8.5 Table 8.5 redrawn: five of its policies on the same plant, the same four faults and six paired seeds. Solid bars are the accepted bulk rate, hatched bars the accepted volume made outside a window, and the dashed lines are the challenge's rate target and its off-specification allowance. Only the suggested supervisor meets both, and it passes six of six seeds. The two rows below it run faster or further and fail on quality, which is why the conductivity probe and the window guard are each necessary.

> ✦ **DID YOU KNOW?** **The Gate That Buys Nothing**
>
> From the simulator: with the supervisor in charge, a valve that accepts everything makes 644.8 mL/min and still passes six of six seeds, against 600.0 mL/min when the gate is in service (Table 8.5). Here the diversion rule costs rate and buys no quality.

Two results in that table are negatives and are reported as such. The release gate adds nothing here: with the same supervisor and the valve accepting everything the challenge makes 644.8 mL/min and still passes six of six, so the gate costs about 45 mL/min of rate and buys no quality. Its value appears where the control layer is blind, in cases 6 and 8 and in the metered-blend row, not where the control layer already holds the product inside its window. And the release margin, which the suggested settings keep at 20 % of every window, is worth 4.8 mL of off-specification material for 28 mL/min of rate, which is close to a wash.

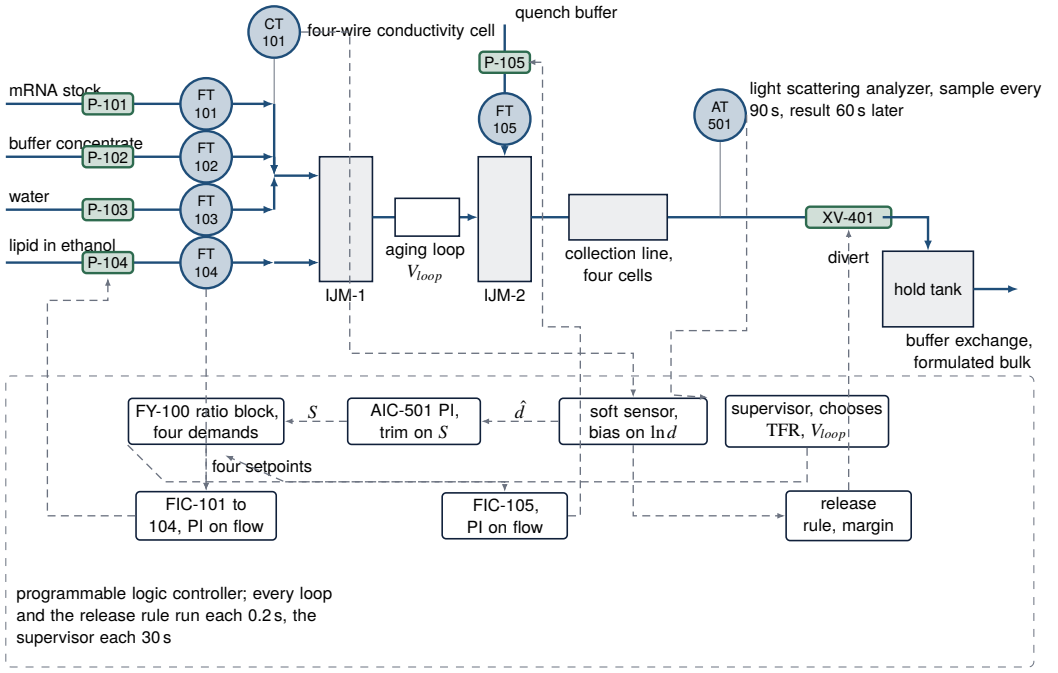


Figure 8.6 One channel of the skid on real hardware. Four metered feeds set the four demands at once through the ratio block FY-100, whose outputs are the setpoints of the four flow loops, and the quench pump follows the total flow in ratio. Only the buffer strength S is trimmed, by AIC-501 on the soft sensor, which reads the conductivity cell and the flow meters and is corrected each time the sampled analyzer answers. The supervisor sets the total flow and the aging loop, and the release rule drives the divert valve. Solid arrows are process streams and dashed gray arrows are signals. Notice that the only instrument that sees the product itself sits downstream of everything the controller moves.

Robustness was checked on twelve seeds rather than six. The shipped settings fail at all twelve. The suggested settings meet the target at eleven, with a mean rate of 598.1 mL/min and a mean of 65.4 mL accepted off-specification. The twelfth is a documented miss rather than noise: after the pulsation at 1080 s the true mean creeps to 215 to 219 nm against a 185 to 215 nm window while the soft sensor reads 193 to 202 nm and the last reported size 193 nm, so neither instrument asks the valve to divert and 295.2 mL is accepted against the 150 mL allowance. It is the blindness of case 5 arriving at the end of a long run, and every retuning that removed it cost rate at other seeds. Over the 31 seeds of the final validation the suggested settings pass 27 at both steps, and the four misses accept 166 to 466 mL.

The scoring deserves its own sentence, because the passing run is not as clean as its scorecard. At 1500 s the suggested run's four channels are at 171.0, 253.5, 185.3 and 225.4 nm against the 200 nm target, the blend is 206.3 nm with $\sigma(\ln d) = 0.412$, and 13 500 mL of the 15 409 mL accepted was made while at least one channel was outside its window. The criterion stays on the blend, because the blend is what leaves the valve and reaches the tank, and a light-scattering measurement of that tank would report the blend's mean and width. The application therefore scores the blend and reports the other figure beside it as teaching truth, in a scorecard tile and in the exported file.

8.7 Technology in Practice

Figure 8.6 puts the structure of Section 8.5 on real hardware: four metered feeds under one ratio block, the quench in ratio to the total flow, the size trim and the release decision at the valve.

What Measures It

Three kinds of instrument carry the loops. The flow meters make the ratios possible, because a ratio block can only hold what something measures. The conductivity probe CT-101 stands in for the aqueous buffer strength, and a probe of that kind serves concentration analysis over limited ranges only [19]. Its service problem is fouling: ions bonding to the electrodes of a two-electrode cell read as a liquid less conductive than it is, while a four-wire cell measures the voltage across the liquid alone and ignores any coating on its inner electrodes [19]. The drift a case applies to CT-101 is that error. The third instrument is the analyzer that decides the product.

> **TECHNOLOGY** **Sensor: Light Scattering on a Drawn Sample**
>
> The common industrial size measurement for these particles is dynamic light scattering. A laser lights a dilute sample, and the speed at which the scattered intensity fluctuates gives a diffusion coefficient and from it the hydrodynamic diameter, over a range of 0.5 to 1000 nm [4]. The number that reaches the controller depends on how the measurement is analyzed. The default output of a common inversion method departs from the cumulant z-average typically by more than 15 percent, the width descriptors are unreliable, and the governing standard ISO 22412:2017 defines the size only as the central value of the distribution [17]. The method is also biased toward larger particles and sensitive to aggregates and contaminants [4], so one stray result sits high on a trend and the next sample does not repeat it.

What Moves It

Five pumps are the final elements, and nothing else on the skid moves except the divert valve. Each pump is a metered feed whose delivery follows a command signal over a range set at the pump [19], modeled here as a first-order drive with $\tau = 0.8\,\mathrm{s}$ that reports when a request passes its capacity. The aging loop is the supervisor's second actuator, and the valves that select it are not modeled, which is why a loop change is given a 150 s lockout rather than a stroke time. XV-401 is the last element in the chain and the only one whose action is a decision rather than a flow.

> **TECHNOLOGY** **Final Element: Pumps That Meter Their Own Flow**
>
> A metering pump carries its pump mechanism, its motor and its drive electronics in one package, and a 4 to 20 mA signal sets its delivery from zero to the full flow of a range that knobs on the pump establish [19]. That is what makes a feed a controlled variable instead of a hand setting. The cheaper kind drives a plunger with a solenoid at variable intervals, so its delivery is a train of discrete pulses rather than a continuous flow, which is simple and reliable wherever non-continuous flow is permissible [19]. A formulation skid is not such a place. A head that pulses at constant average delivery leaves the meter reading its setpoint while the jets see a different ratio every half second, which is the fault of case 5.

Where the Controller Runs

A skid like this one runs on a programmable logic controller, which serves either as a local controller under a supervisory scheme or as the primary controller of a smaller system, with control cycles in the range of milliseconds to minutes [20]. In the application every flow loop, the size trim and the release rule execute each 0.2 s, while the supervisor evaluates its grid every 30 s. The structure maps onto standard blocks. A signal multiplying relay, or ratio station, multiplies a wild flow signal by a constant and sends the product to a captive flow controller as its setpoint [19]. FY-100 is that block solved for four demands at once, and AIC-501 is an ordinary slow loop whose output is one of those demands rather than a valve position.

IN PRACTICE **What a Divert Valve Has to Be Able to Say**

Guidance for continuous manufacturing asks that the residence-time distribution be characterized, because it fixes the time available for transport and transformation and so lets each portion of material be tracked [21]. Diversion follows from that tracking. Criteria should be established that trigger the start and the end of a diversion period and the restart of collection, and how much material is diverted depends on the process dynamics, that distribution, the control strategy, the severity of the disturbance and where the sampling and diversion points sit [21]. The release rule at XV-401 is exactly such a criterion, and its margin and its two line volumes of hold-off are the designer's answers to the last two factors.

What Is Modernizing It

Three developments bear on this skid. The first is better size metrology. Multi-angle dynamic light scattering represents the size distribution more accurately and is routinely used in formulation development and quality control [4]. The second is prediction between samples. Modeling, digital twins, automation and real-time monitoring are the opportunities named for this industry [4], and the soft sensor of Section 8.5 is the plain version of one: a fitted map on the meters, corrected by a bias each time the laboratory answers. The third is the network under all of it. Computing at the edge acts on data where it is collected rather than shipping it to a data center [20], which is what would let a per-channel measurement act inside one line volume.

8.8 Limitations

The size law is a published model form fitted here to one published table, so every size the application reports is an interpolation inside a box of two axes at one pH, one cargo type, one mixer geometry and one formulation family, and the agreement reported in Section 8.4 is a fit plus two withheld cargo tables rather than a validation. Outside that box the three extension factors carry the prediction, and the ethanol exponent that carries the most is a teaching value fitted to no data: at an exponent of 12 the suggested challenge settings no longer pass. The quantitative size of every flow-ratio result in this chapter therefore depends on one number that a measurement could change, while the direction of each result does not.

The plant is lumped where the lesson does not need detail. The aging loop is evaluated on a volume clock rather than integrated, and that clock is first-order accurate in the step under a fast-swinging flow: against an exact integration the age of a packet under the 1.2 Hz pulsation is 8.6 % low at 0.05 s and 4.5 % low at half the step, where the instantaneous V_{loop}/TFR the model avoids would be 19.9 % low and would not converge at all. The error does not reach a lesson, because the size is a power law of the age averaged over nine nodes and then over a collection line, and halving the step moves the pulsation case by 0.05 nm. The nine-node quadrature is likewise 0.75 per mille above a 2001-node evaluation in the mean size at 440 ms, but it resolves the share of the distribution inside the box only in ninths, which overstates that share's resolution by about eleven points.

OPEN CHALLENGE **Channels That Drift Apart**

Parallel mixing lines are how throughput is won, and they are also where quality is lost. Scaling out risks altered mixing regimes and mixer-to-mixer variability, with inconsistent flow rates, pressure imbalance between the channels or a channel clogging, and the result is non-uniform product [4]. The hard part is that the blend hides it, as Figure 8.4 shows. Progress would be a measurement per channel, or a flow loop per channel rather than one manifold, so that a drifting line can be found before its material is averaged into the tank.

Encapsulation is nominal and set at the mixer, morphology is absent, and nothing in the application speaks to intact RNA, transfection or potency. The buffer exchange is a declared recovery map. Faults enter on a schedule as steps and ramps, not with the texture a plant's own upsets have. The pressure trip is an educational guard rather than certified protection, and every window, limit and delay is

a teaching value. A simulated pass is not batch release, and real limits come from the product's own specification, a pharmacopeial chapter or regulatory guidance [15]. No learning study has been run, so no claim is made that the application improves learning. What the cases establish is which experiments a learner can carry out and which misconception each one addresses.

OPEN CHALLENGE **Running for Weeks Instead of Hours**

Deposits on channel walls change the mixing and the residence time, and a surface treatment that resists them has held particle size and an encapsulation above 85 percent throughout a run [22]. The demonstration was one chip over a few hours. Production will need either much longer continuous runs or cleaning and reuse that are reliable and reproducible across hundreds of parallel units [22]. Progress would be a documented campaign length at which the size law and the age distribution are still the ones the controller assumes.

8.9 Conclusion

A formulation skid whose particle size is set by the age of the material has two actuators for one quality attribute, a loop volume and a flow, and they are not equivalent. The simulator of this chapter carries a published size law along a normalized age distribution on a volume clock, with the true size, the age distribution and each channel's own product logged beside the instruments that the controllers read. Moving the aging loop from 78 to 293 µL takes the product from 139.9 nm to 215.8 nm at the same flows and the same charge ratio, while 3.5 volumes of quench instead of 2.0 move it by 0.46 nm, so the arrest is the ethanol's and the quench ratio is a cost rather than a quality knob. On the challenge, coordinated ratio control runs out of buffer strength at 441.7 mL/min while a supervisor that may also choose the loop reaches 600.0 mL/min, a margin of 158 mL/min with the same instruments and faults. The release gate adds nothing once that control layer is in place, and the conductivity probe, the size trim and the window guard are each necessary. A 12 % maldistribution separates four channels by 37 nm while their one flow meter reads its setpoint, and 13 500 mL of the 15 409 mL the passing run accepted was made with a channel outside its own window. The map is fitted inside one box and the factors that carry it outside are teaching values. A fuller model would fit the ethanol dependence to measured data, add pH as an axis, model the morphology the anchor controls alongside the size, and replace the exchange map with a membrane model. The simulator shows what a quality attribute looks like when its controlled variable is a time.

KEY POINTS

- The size of these particles is set by how long they travel before the quench floods them, so the clock and not the mixing is the thing to control.
- The ratio block holds the flow ratio, the charge ratio and the buffer strength at once, which leaves the buffer strength free as the trim knob.
- A supervisor free to choose the aging loop as well reaches 600.0 mL/min against 441.7 mL/min for coordinated ratio control.
- One flow meter in front of four channels reads its setpoint while the channels separate by 37 nm, and the blend still meets every window.

Where else this applies. Volume 5, Chapter 7 tracks material through a residence-time distribution to a divert valve, Chapter 7 trims a ratio blend on a slow analyzer, and Volume 5, Chapter 10 runs a soft sensor that a late measurement corrects.

References

[1] C. Devos, A. Udepurkar, P. Sagmeister, A. S. Hodlewsky, J. Chen, A. Hatas, N. Ostrovsky, M. Al-Jazrawe, J. I. Ren, A. Y. Liu, R. D. Braatz, and A. S. Myerson, "Manufacturing mRNA-loaded lipid nanoparticles with precise size and morphology control," *ACS Nano*, vol. 19, no. 38, pp. 33 991–34 002, 2025.

[2] A. S. Myerson, M. Krumme, M. Nasr, H. Thomas, and R. D. Braatz, "Control systems engineering in continuous pharmaceutical manufacturing May 20–21, 2014 continuous manufacturing symposium," *Journal of Pharmaceutical Sciences*, vol. 104, no. 3, pp. 832–839, 2015.

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

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

[5] A. D. Bangham, M. M. Standish, and J. C. Watkins, "Diffusion of univalent ions across the lamellae of swollen phospholipids," *Journal of Molecular Biology*, vol. 13, no. 1, p. 238, 1965.

[6] S. Batzri and E. D. Korn, "Single bilayer liposomes prepared without sonication," *Biochimica et Biophysica Acta (BBA) - Biomembranes*, vol. 298, no. 4, pp. 1015–1019, 1973.

[7] L. B. Jeffs, L. R. Palmer, E. G. Ambegia, C. Giesbrecht, S. Ewanick, and I. MacLachlan, "A scalable, extrusion-free method for efficient liposomal encapsulation of plasmid DNA," *Pharmaceutical Research*, vol. 22, no. 3, pp. 362–372, 2005.

[8] S. C. Semple, A. Akinc, J. Chen, A. P. Sandhu, B. L. Mui, C. K. Cho, D. W. Y. Sah, D. Stebbing, E. J. Crosley, E. Yaworski, I. M. Hafez, J. R. Dorkin, J. Qin, K. Lam, K. G. Rajeev, K. F. Wong, L. B. Jeffs, L. Nechev, M. L. Eisenhardt, M. Jayaraman, M. Kazem, M. A. Maier, M. Srinivasulu, M. J. Weinstein, Q. Chen, R. Alvarez, S. A. Barros, S. De, S. K. Klimuk, T. Borland, V. Kosovrasti, W. L. Cantley, Y. K. Tam, M. Manoharan, M. A. Ciufolini, M. A. Tracy, A. de Fougerolles, I. MacLachlan, P. R. Cullis, T. D. Madden, and M. J. Hope, "Rational design of cationic lipids for siRNA delivery," *Nature Biotechnology*, vol. 28, no. 2, pp. 172–176, 2010.

[9] P. R. Cullis and M. J. Hope, "Lipid nanoparticle systems for enabling gene therapies," *Molecular Therapy*, vol. 25, no. 7, pp. 1467–1475, 2017.

[10] X. Hou, T. Zaks, R. Langer, and Y. Dong, "Lipid nanoparticles for mRNA delivery," *Nature Reviews Materials*, vol. 6, no. 12, pp. 1078–1094, 2021.

[11] B. K. Johnson and R. K. Prud'homme, "Chemical processing and micromixing in confined impinging jets," *AIChE Journal*, vol. 49, no. 9, pp. 2264–2282, 2003.

[12] I. V. Zhigaltsev, N. Belliveau, I. Hafez, A. K. K. Leung, J. Huft, C. Hansen, and P. R. Cullis, "Bottom-up design and synthesis of limit size lipid nanoparticle systems with aqueous and triglyceride cores using millisecond microfluidic mixing," *Langmuir*, vol. 28, no. 7, pp. 3633–3640, 2012.

[13] N. M. Belliveau, J. Huft, P. J. C. Lin, S. Chen, A. K. K. Leung, T. J. Leaver, A. W. Wild, J. B. Lee, R. J. Taylor, Y. K. Tam, C. L. Hansen, and P. R. Cullis, "Microfluidic synthesis of highly potent limit-size lipid nanoparticles for in vivo delivery of siRNA," *Molecular Therapy - Nucleic Acids*, vol. 1, p. e37, 2012.

[14] J. A. Kulkarni, M. M. Darjuan, J. E. Mercer, S. Chen, R. van der Meel, J. L. Thewalt, Y. Y. C. Tam, and P. R. Cullis, "On the formation and morphology of lipid nanoparticles containing ionizable cationic lipids and siRNA," *ACS Nano*, vol. 12, no. 5, pp. 4787–4795, 2018.

[15] U.S. Food and Drug Administration, Center for Drug Evaluation and Research, "Liposome drug products: Chemistry, manufacturing, and controls; human pharmacokinetics and bioavailability; and labeling documentation. Guidance for industry," Guidance document, https://www.fda.gov/media/70837/download, Apr. 2018.

[16] D. E. Koppel, "Analysis of macromolecular polydispersity in intensity correlation spectroscopy: The method of cumulants," *The Journal of Chemical Physics*, vol. 57, no. 11, pp. 4814–4820, 1972.

[17] N. Farkas and J. A. Kramar, "Dynamic light scattering distributions by any means," *Journal of Nanoparticle Research*, vol. 23, no. 5, p. 120, 2021.

[18] G. I. Taylor, "Dispersion of soluble matter in solvent flowing slowly through a tube," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 219, no. 1137, pp. 186–203, 1953.

[19] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[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] U.S. Food and Drug Administration, Center for Drug Evaluation and Research and Center for Biologics Evaluation and Research, "Q13 continuous manufacturing of drug substances and drug products. Guidance for industry," ICH harmonised guideline, https://www.fda.gov/media/165775/download, Mar. 2023.

[22] O. Milić, S. M. Savić, M. Zurković, B. Stanojević, and S. Savić, "Lipid nanoparticles for gene therapy: Unresolved challenges in manufacturing, transdermal delivery, machine learning, endosomal escape, and the protein corona," *Pharmaceutics*, vol. 18, no. 8, p. 991, 2026.

8.10 Exercises

8.1 Concept check · 2 min

Which operator demand of this skid sets how long the material ages before the second junction arrests it?

(a) The flow ratio, because it fixes the mixing time in the jets

(b) The quench ratio, because it fixes how far the stream is diluted

(c) The total flow, because the age is the loop volume divided by it

(d) The charge ratio, because it fixes how much lipid is uncomplexed

8.2 Concept check · 2 min

A two-feed version of this skid carries only the mRNA stock and the lipid solution, at fixed feed concentrations. The operator steps the flow ratio from 1:1 to 3:1. What happens to the charge ratio N/P?

(a) It holds, because the quench dilutes both components equally

(b) It holds, because the ratio block solves for it

(c) It rises by a factor of three, since three volumes of aqueous arrive

(d) It falls by a factor of three, from 3.00 to 1.00

8.3 Concept check · 1 min

AIC-501 holds the particle size by moving the aqueous buffer strength rather than the flow ratio or the total flow. Why?

(a) The ratio block can move it while the flow ratio and the charge ratio stay on their demands

(b) It is the only input the conductivity probe can measure

(c) Moving the total flow would change the throughput, and the buffer strength does not

(d) The size law is linear in the buffer strength and nonlinear in the others

8.4 Theory and calculation · 7 min

The ratio block of Equation (8.1) is asked for TFR $= 40\,\mathrm{mL/min}$, FRR $= 1.0$, $N/P = 6.0$ and $S = 100\,\mathrm{mol/m^3}$, with $c_N = 3.858\,\mathrm{mol/m^3}$, $c_P = 1.288\,\mathrm{mol/m^3}$, $S_1 = 25\,\mathrm{mol/m^3}$, $S_2 = 500\,\mathrm{mol/m^3}$ and $R_q = 2.0$.

(a) Find the four feed flows and the quench flow. *Checkpoint:* 9.99, 3.50, 6.51 and 20.0 mL/min, with 80 mL/min of quench.

(b) The same demands are asked at FRR $= 3.0$. Find the new feed flows, and say which feed takes up the change. *Checkpoint:* 4.99, 5.75, 19.26 and 10.0 mL/min.

(c) At FRR $= 0.50$ the block cannot deliver $S = 100\,\mathrm{mol/m^3}$. It keeps Q_1 and the aqueous total Q_a, gives the concentrate only what the stock leaves, $Q_2 = Q_a - Q_1$, and sets Q_3 to zero. Find the strength it delivers. *Checkpoint:* About $26\,\mathrm{mol/m^3}$.

8.5 Theory and calculation · 8 min

Use the fitted map of Equation (8.3) with $a_0 = 61.84\,\mathrm{nm}$, $a_1 = 65.98\,\mathrm{nm}$, $r_0 = 0.154$ and $r_1 = 1.653\,\mathrm{nm/ms^p}$, $C = 81.03\,\mathrm{mol/m^3}$, $D = 0.0256\,\mathrm{m^3/mol}$ and $p = 0.7785$, at $S = 100\,\mathrm{mol/m^3}$.

(a) Find the assembly size $d_0(S)$ and the growth coefficient $R(S)$. *Checkpoint:* 97.5 nm and $1.047\,\mathrm{nm/ms^p}$.

(b) Find the aging time of the 78 µL and the 293 µL loop at 40 mL/min, and the size the map gives at each. *Checkpoint:* 117 and 440 ms, giving 140.2 and 217.2 nm.

(c) The age rises by a factor of 3.76 and the size by a factor of 1.54 on the application. Reconcile the two. *Checkpoint:* $3.76^{0.779} = 2.80$ on the growth term alone.

8.6 Theory and calculation · 5 min

At the design point the mixed stream is 50 % ethanol by volume before the quench. The quench adds R_q volumes of aqueous buffer per volume of mixed stream, and the exchange must bring the line's ethanol to the 0.5 % teaching limit.

(a) Find the ethanol fraction after 2.0 volumes of quench, and the factor f_E of Equation (8.4) before and after it. *Checkpoint:* 16.7 %, with 1.000 and 0.00137.

(b) Find the same at 3.5 volumes, and say what the difference is worth. *Checkpoint:* 11.1 % and 1.2×10^{-4}.

(c) Find the diavolumes needed from 16.62 % to the limit, and the residual the controller's extra 0.30 leaves. *Checkpoint:* 3.50 diavolumes, leaving 0.37 %.

8.7 App challenge · 8 min

Load scenario 4, Quench timing, the opening experiment. The aging loop goes from 78 to 293 µL at 240 s. Use the default seed and read the exported columns `truth_d_nm` and `tau_ms`.

(a) Run to 200 s, then on to 420 s, and read the true size and the aging time at each. *Checkpoint:* 139.9 nm at 117 ms, then 215.8 nm at 440 ms.

(b) Without resetting, set the quench ratio to 3.5 volumes and run to 700 s. Compare with the same run left at 2.0 volumes. *Checkpoint:* 215.5 nm against 216.0 nm.

(c) Say what part (b) costs, reading the discharge pressure in each run. *Checkpoint:* 2.49 against 4.10 bar.

Accept a size within 0.5 nm and a pressure within 0.05 bar.

8.8 App challenge · 8 min

Load scenario 6, The size result is a minute old. The release rule starts on the last reported size, and the total flow steps to 21 mL/min at 200 s. Use the default seed and read the Off-specification accepted tile.

(a) Run to 420 s and read the tile, the true size and the aging time. *Checkpoint:* 239.5 mL, with 168 nm at 223 ms.

(b) Reload the scenario, set the release rule to Soft sensor, and run the same step to 420 s. *Checkpoint:* 1.2 mL.

(c) Explain the factor of 200, using the age of the reported size.

Accept a volume within 2 mL.

8.9 App challenge · 8 min

Load scenario 7, Four channels, one flow meter. A 12 % aqueous maldistribution arrives at 180 s and channel 2 restricts at 480 s. Use the default seed and read the exported columns `truth_ch1_d_nm` to `truth_ch4_d_nm`, `truth_d_nm`, `truth_spread` and `P_bar`.

(a) Run to 470 s and read the four channel sizes, the blend and its spread. *Checkpoint:* 126.6, 163.6, 132.9 and 150.8 nm, blending to 142.8 nm at 0.364.

(b) Does the blend meet the 15 nm size window and the 0.45 spread limit? Say what an instrument at the valve would report. *Checkpoint:* Yes to both.

(c) Run on to 600 s and read the pressure and the channel sizes again. *Checkpoint:* 4.29 bar, with 125.3, 190.1, 131.2 and 148.1 nm.

Accept a size within 0.5 nm and a pressure within 0.05 bar.

8.10 Application area · 4 min

Section 8.4 reports the fitted map at 5.11 % rms on the table it was fitted to, 8.76 and 7.96 % on two withheld cargo tables at the same pH, and 28.15 % on a table taken at pH 4.5.

(a) Say what those four figures together establish, and what they do not.

(b) A leave-one-column-out refit predicts the 440 ms column at 90.6 % rms. Name the feature of the application that this number justifies, and say what it protects against.

8.11 Application area · 3 min

The challenge target asks for 560 mL/min of bulk with at most 150 mL accepted outside a window, and the suggested settings pass it at six of six paired seeds.

(a) State two reasons from the chapter why that pass is not a statement about a real product.

(b) The passing run made 13 500 mL of its 15 409 mL while at least one channel sat outside its own window. Say why the application still scores the blend.

8.12 Review the agent · 5 min

An AI assistant is asked how to raise throughput on this skid and replies as follows. “Hold the 78 µL loop and raise the total flow to 56 mL/min per channel. That shortens the age, so recover the size by taking the quench ratio from 2.0 up to 3.5 volumes: a longer quench leaves the particles less diluted at the junction and holds them larger. The flow loops and the exchange have the capacity.” Find the flaw and say how to check it.