INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 12 Control in Petrochemicals and Polymers

CHAPTER 9

Polypropylene Loop Simulator to Teach Quality Control Against a Cavitation Limit

Pages 194–225

Edited by John D. Hedengren

How to cite this chapter. "Polypropylene Loop Simulator to Teach Quality Control Against a Cavitation Limit," in *Control in Petrochemicals and Polymers*, John D. Hedengren, Ed., vol. 12 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 194–225.

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

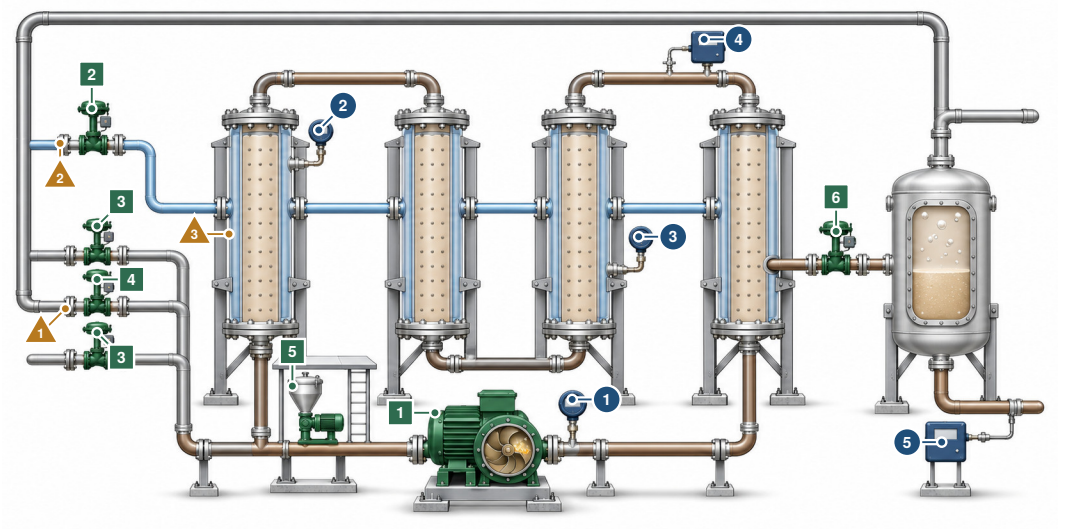
The app for this chapter
https://apmonitor.com/apps/polypropylene-loop/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v12-petrochemicals-polymers/#polypropylene-loop

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

CHAPTER 9

Polypropylene Loop Simulator to Teach Quality Control Against a Cavitation Limit



The controller moves 1 circulation pump speed, set by the operator 2 cold-water valve on the jackets 3 hydrogen and ethylene valves 4 fresh propylene valve, which holds the pressure 5 catalyst feeder 6 slurry take-off valve

The controller reads 1 reactor pressure 2 reactor temperature 3 slurry density 4 loop chromatograph: dissolved hydrogen, ethylene and inerts 5 melt flow rate and density analyzer

What disturbs it 1 light inerts in the propylene feed 2 cooling-water temperature 3 jacket fouling

Notice the bubbles in the pump's impeller: the hydrogen that valves 3 add for the grade also raises the liquid's bubble point toward the pressure the pump needs, and the chromatograph sees that hydrogen within minutes while the analyzer under the flash shows the grade about an hour later. Illustration, not to scale.

IN PLAIN TERMS **One Gas Sets the Product and Threatens the Pump**

The main idea. Plastic for cups and fibers grows as tiny grains in a pipe loop full of liquid. A large pump keeps the liquid moving, and a little dissolved hydrogen sets how easily the melted plastic flows.

The challenge. The same hydrogen brings the liquid closer to boiling, and bubbles in the pump stop it from pushing. The pump feels a change in minutes and the plastic only after about an hour, so chasing the plastic can wreck the pump.

Why it matters. A guard controller that limits the hydrogen saves the pump, but on its own it leaves the plant short of the plastic that was ordered.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes in Section 9.2, the Did you know? boxes in Sections 9.3 and 9.6, the two Try it boxes in Section 9.6, and Key points at the end of the Conclusion. No equations.
Study	65 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 9.6 from Table 9.5 with Figure 9.3, Section 9.7, and Section 9.8 with its Open challenge box.

The control question In a liquid-full polypropylene loop the hydrogen that sets melt flow rate also raises the bubble point of the liquid toward the pressure the circulation pump needs. How should a quality controller and a margin controller share the one hydrogen setpoint, and what does each arrangement cost in pump damage and in off-specification polymer?

Learning objectives

- Compute the bubble point of the loop liquid and its gradients, and turn a grade request into a pressure requirement.
- Explain why one hydrogen move reaches the margin in minutes, the product in more than an hour and the pump in seconds.
- Describe the three details that make a low selector work: tracking of the loser, a range that lets the constraint win and parking of the idle controller.
- Compare the override and the predictive controller on one grade transition, and attribute the result to the plan or to the temperature.
- Plan a grade transition that buys margin with pressure before spending it on hydrogen.

Prerequisites Sections 1.4.4 (what spends the margin), 1.4.5 (inventory-averaged quality), 1.4.8 (one-directional guards), 1.6.2 (velocity-form PID) and 1.6.3 (the predictive option). Chapter 8 for a polyolefin grade change against a cooling limit.

The app https://apmonitor.com/apps/polypropylene-loop/

See also Volume 4, Chapter 2 for a protective override against surge and Volume 3, Chapter 4 for selectors that cross-limit fuel and air. Volume 9, Chapter 4 for offset-free predictive control with an estimated bias, and Volume 3, Chapter 2 for matching a cascade to where the disturbance enters.

9.1 Introduction

A constraint is usually drawn as a horizontal line on a trend and taught as a region to stay out of. A bulk polypropylene loop reactor refuses that picture. Its product specification is melt flow rate, the only effective handle on melt flow rate is dissolved hydrogen, and dissolved hydrogen is what raises the bubble point of the circulating liquid toward the pressure that keeps the circulation pump from cavitating. The line on the trend is therefore not a place the operator avoids. It is another controller wanting the same valve, in the opposite direction, and one of the two has to lose. A block diagram of the melt-flow cascade contains no term in which that collision can appear, because the hydrogen valve enters the quality transfer function once and the pump does not enter it at all. The misconception this simulator corrects is the belief that a constraint is tuned away from rather than arbitrated with.

The shape of the problem is general. Wherever the only effective actuator for a quality or throughput objective also consumes a mechanical or thermodynamic margin, a constraint controller and a quality controller share one manipulated variable and a selector decides which of them writes it. A centrifugal compressor raises throughput with the same guide vanes and speed that carry it toward surge, and the anti-surge controller is a genuine controller with its own measurement and its own authority [1], as Volume 4, Chapter 2 shows. A distillation column buys purity with reflux and pays for it in vapor loading toward the flooding limit. A pump on a hot suction gains flow at the cost of net positive suction head [2]. Each is arranged as a ratio, override or selective control problem rather than as a single loop [3], and each fails in the same way when the losing controller is allowed to wind up or is parked where it cannot act quickly. Polymer plants meet the pattern more often than most, because product quality there is inferred from a slow sampled analyzer and the mechanical constraints are fast [4, 5].

Loop reactor models in the literature answer polymerization questions. Population-balance treatments of the liquid-phase loop resolve the particle size distribution, the multiple site types and the recycle structure [6, 7], and commercial bulk processes have been matched at steady state and in transient [8]. The thermodynamic state of the circulating liquid receives far less attention even though it is what decides whether the pump survives a grade change. However, a student handed the melt-flow

transfer function and asked for proportional and integral settings never encounters the reason the resulting controller wrecks an impeller, because the pump is not in the model that was handed over. This chapter teaches override control on a bulk polypropylene loop reactor whose bubble point is computed from the same liquid composition the quality controllers manipulate, so that a well-tuned melt-flow loop and a wrecked circulation pump become one event. The kinetics, the bubble-point correlation and the selector are standard. What the app adds is that all of them act on one plant, in a browser, with every variable logged.

Section 9.2 traces how the bulk loop and its quality specifications came about. Section 9.3 carries every equation and every constant of the plant: nineteen states, a re-verified vapor-pressure surrogate, a Henry-law bubble point, an impeller-eye depression, an algebraic head-breakdown curve and a deposition-velocity settling limit. Section 9.4 prices each dissolved species in bubble point, turns a quality request into a pressure requirement and separates the three clocks of the plant. Section 9.5 describes the ten loops, the low selector, the interlocks and the predictive controller. Section 9.6 runs one grade transition twice, once to destroy the pump and once to abandon the grade, measures what the selector reports and what slow disturbances cost, and shows that the graded challenge is won only by raising the operating pressure and slowing the melt-flow controller together. It ends with the predictive controller on the same transition. Section 9.7 names the hardware and the practice behind the model, and Sections 9.8 and 9.9 state what the model cannot show and what a plant would add.

9.2 Polypropylene Process Technology

The catalysts of the bulk loop descend from Ziegler's low-pressure ethylene process [9] and from Natta's crystalline, stereoregular polymers of propylene [10], whose contested origin the first History box below tells. That stereoregularity is what made polypropylene a material rather than a curiosity, and it is the reason a polypropylene reactor has to control a variable an ethylene reactor does not.

> ⧖ HISTORY **A Crystalline Plastic From a Gas, 1951 to 1954**
>
> Near the end of 1953 Karl Ziegler's group found that ethylene polymerizes even at normal pressure on catalysts made by mixing an aluminum alkyl with titanium tetrachloride, and the catalysts became known to industry around the turn of the year [11]. Early in 1954 Giulio Natta's group polymerized propylene with them. Some of the product was crystalline, and Natta traced the crystallinity to a regular steric arrangement along the chain rather than to a higher molecular weight [12]. The two shared the 1963 Nobel Prize in Chemistry [11]. Yet in 1951 J. Paul Hogan and Robert L. Banks, trying to turn propylene into gasoline over a catalyst carrying chromium oxide, had already made crystalline polypropylene. Three laboratories filed patent applications on the discovery between 1951 and 1953, and the contest lasted three decades, until an appellate court ruled for Hogan and Banks in 1983 [13].

The process technology went through three stages, driven by the catalyst generations that the second History box below traces, until the reactor was almost the whole plant. The third stage removed the diluent. In a bulk or liquid-pool process the monomer is the reaction medium, which raises the propylene concentration at the active site, removes the solvent recovery train, and takes the reactor from a stirred vessel to a pipe loop that a single axial pump keeps in circulation. The Spheripol configuration pairs that liquid loop with a gas-phase reactor downstream for impact copolymer, and it is the arrangement most of the world's polypropylene is made in [14]. Metallocene catalysis arrived in the 1980s and became a commercial option in the 1990s, because a chiral single-site metallocene fixes stereochemistry at one well-defined center and therefore produces a narrow molecular weight distribution and a controllable comonomer sequence [15]. Single-site catalysts changed the grade slate rather than the reactor, and the bulk loop remained the vessel.

HISTORY From Washed Powder to Liquid Monomer, 1960s Onward

The first polypropylene catalysts, used in slurry processes from the 1960s, made about one tonne of polymer per kilogram of catalyst and 5 to 10 percent of atactic polymer, so the catalyst had to be washed out of the product and the atactic part removed. The catalysts of the 1970s reached about ten tonnes per kilogram, and supports with electron donors raised the isotacticity. The fourth generation, the present standard, sits on a spherical support made for liquid monomer, yields 30 to 50 tonnes per kilogram and made both cleaning steps obsolete. Modern suspension processes therefore use liquid propylene instead of a solvent, which the industry calls a bulk process [16].

The controlled quality variables followed from what a converter can measure. Molecular weight itself is a laboratory determination, and average molecular weight enters processing through melt viscosity, so the industry adopted melt flow rate, the mass of polymer extruded through a standard capillary in ten minutes under a standard load, as the operating specification. Hydrogen sets it. Hydrogen terminates a growing chain by transfer to the metal alkyl, and increasing the hydrogen concentration therefore shortens the chains and raises melt flow rate, with a response that depends on the catalyst and donor system and is a fitted quantity in every plant model that exists [17]. Density is the second specification and it follows ethylene incorporation, because the comonomer units interrupt the isotactic sequences and reduce crystallinity. Inference of both from composition and temperature, rather than waiting for the analyzer, became the standard route to closed-loop quality control in polyolefin plants [4], and grade transitions on a liquid-propylene loop are now posed as constrained optimal control problems [5]. What that literature treats sparingly is the arrangement below the quality loop, where the loop is liquid-full and the same dissolved hydrogen that sets the specification is also the largest avoidable term in the bubble point of the circulating liquid.

9.3 Process Model

Reactor R-101 is a 65 m^3 liquid-full loop of 0.65 m bore and 196 m developed length, built as four vertical legs of 40 m joined by 180° bends, with cooling jackets over 326.7 m^2 of the wall. Axial pump P-101 sits on the horizontal return and circulates the slurry at 5.50 m/s against the loop resistance. Catalyst and hydrogen and ethylene enter at separate points, fresh propylene make-up enters on pressure control, and product leaves through take-off valve XV-101 to a flash at 18 bar where the monomer is recovered and returned. The circuit, its instrument tags and each symbol carried by the equations of this section are drawn in Figure 9.1.

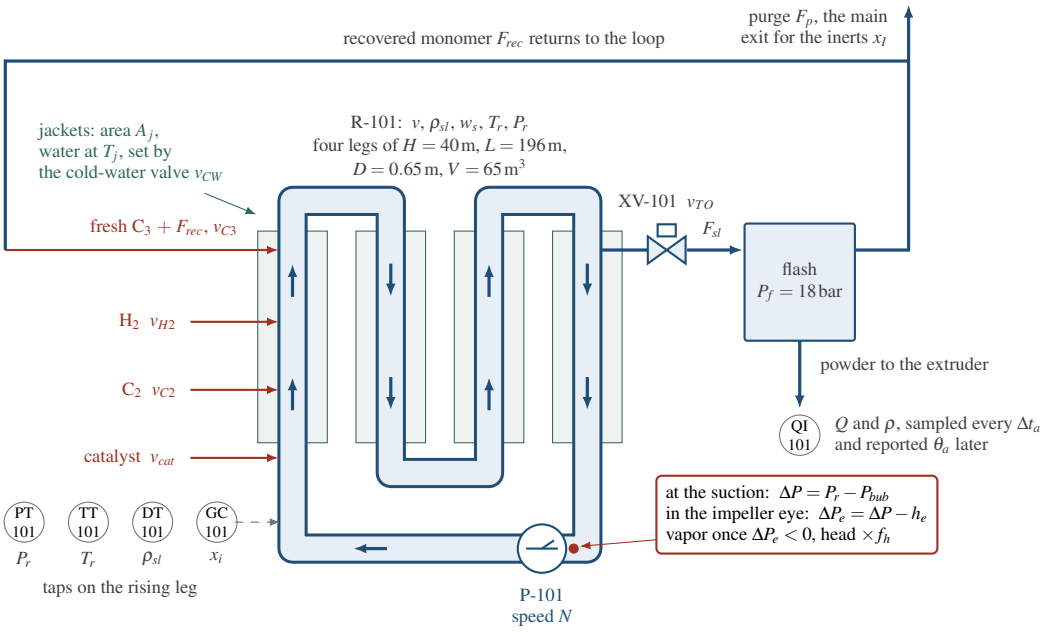


Figure 9.1 Loop reactor R-101 with its four jacketed legs, the axial circulation pump, the four injection points and the product take-off. Liquid propylene circulates at velocity v round a developed length L of bore D and volume V, carrying solids fraction w_s at slurry density ρ_{sl}, temperature T_r and pressure P_r. Catalyst, ethylene, hydrogen and fresh propylene enter separately on the leg leaving the pump, and the four instruments that the control system reads are PT-101, TT-101, DT-101 and the loop chromatograph GC-101 reporting the dissolved mole fractions $x_i = \{x_{H2}, x_{C2}, x_I\}$. The margin the application is about is measured at the pump suction as $\Delta P = P_r - P_{bub}$ and inside the impeller eye as ΔP_e, one eye depression h_e lower, where vapor appears and multiplies the delivered head by f_h. Slurry F_{sl} leaves through XV-101 to the flash at P_f, whose recovered monomer F_{rec} returns with the fresh propylene. The purge F_p and the 15 % of the flashed inerts that recovery does not return are the only exits for the accumulated inerts. Analyzer QI-101 reports melt flow rate Q and density ρ on a sample interval Δt_a with dead time θ_a.

> ✦ **DID YOU KNOW? Pipe Loops Up to Fifty Meters Tall**
>
> Loop reactors are long jacketed pipes arranged as four or six vertical legs up to 50 m tall, and an in-line axial pump drives the slurry of a polyethylene loop round at 6 to 10 m/s. Liquid-propylene loops run near 70 °C and 4 MPa with about 40 wt% of polymer [16].

Everything the control system touches appears in Table 9.1, where melt flow rate is in g/10min, density in g/cm^3, slurry density in kg/m^3 and every valve position in %. Nineteen of those quantities are integrated. Units throughout are the ones on a control-room faceplate, temperature in °C, pressure in bar, mass in kg, amount in mol and velocity in m/s, and the clock is **in minutes**. Every rate below is therefore per minute, and every time constant is in minutes unless it is quoted in seconds.

Table 9.1 Classification of Variables. Nominal Values Are Those of the Solved Design Point at Which Every Scenario Opens.

Symbol	Role	Nominal	Range
Q melt flow rate	CV, sampled	3.00	0.5–40 SP
ρ density	CV, sampled	0.90550	0.888–0.907 SP
R production	CV, inferred	22.0	8–34 t/h SP
ΔP margin	CV, inferred	3.43	0.5–6 bar SP
T_r	CV, measured	70.0	58–84 °C SP
P_r	CV, measured	38.00	30–43 bar SP
ρ_{sl}	CV, measured	591	480–650 SP
r_H H_2/C_3	CV, measured	0.900	0.02–5.0
r_E C_2/C_3	CV, measured	0	0–6.0
v_{H2} hydrogen	MV	11.9	0–100 %
v_{C2} ethylene	MV	0.0	0–100 %
v_{C3} propylene	MV	43.3	0–100 %
v_{CW} cold water	MV	45.0	0–100 %
v_{TO} take-off	MV	45.0	0–100 %
v_{cat} catalyst	MV	44.7	0–100 %
N pump speed	MV, operator	100	40–110 %
T_{cw} water supply	DV, measured	22	4–34 °C
y_I feed lights	DV, measured	1.47	0–5 mol%
purge setting	DV, measured	100	0–160 %
header pressure	DV, measured	100	70–110 %
b_a batch activity	DV, unmeasured	0	±35 %
z poison	DV, unmeasured	0	0–0.40
jacket fouling	DV, unmeasured	0	0–60 %
analyzer bias	DV, unmeasured	0	±15 %
n_{H2} (mol)	state	394.7	≥ 0.01
n_{C2} (mol)	state	0	≥ 0
n_I (mol)	state	22280	≥ 1
M_s (kg)	state	19203	500–40950
A_c (—)	state	0.994	≥ 0
T_r, T_j (°C)	states	70.0, 41.2	20–94, 0–94
P_r (bar)	state	38.00	8–60
v (m/s)	state	5.50	0–14
$\bar{q}$ (—), $\bar{e}$ (wt%)	states	1.099, 0	—
ϕ settled layer	state	0	0–0.50
seven valve lags (%)	states	as above	0–110

9.3.1 Inventories and Composition

Solids and the three dissolved lights are separate inventories, as shown in Equations (9.1) to (9.4).

$$\frac{\mathrm{d}M_s}{\mathrm{d}t} = r_p - \dot{m}_{\text{out}} \tag{9.1}$$

$$\frac{\mathrm{d}n_{H2}}{\mathrm{d}t} = F_{H2} - \zeta\, \dot{n}_{C3}\, r_H - o_{H2} \tag{9.2}$$

$$\frac{\mathrm{d}n_{C2}}{\mathrm{d}t} = F_{C2} - r_{C2} - o_{C2} \tag{9.3}$$

$$\frac{\mathrm{d}n_I}{\mathrm{d}t} = \dot{n}_{\text{feed}}\, y_I - o_I \tag{9.4}$$

where M_s is the polymer inventory in kg, r_p is the polymerization rate in kg/min, $\dot{m}_{\text{out}}$ is polymer leaving with the take-off, n_i are dissolved amounts in mol, F_{H2} and F_{C2} are the gas feeds in mol/min,

$\zeta \dot{n}_{C3} r_H$ is the hydrogen consumed by chain transfer, r_{C2} is the ethylene built into chains, y_I is the light-inert fraction of the fresh propylene, and o_i is the loss of species i with the liquid drawn off and with the recycle purge. Equation (9.4) has no reaction term at all, which is the reason inerts accumulate until the purge and the unrecovered part of the flashed liquid remove them, the purge taking 53 % at the design point.

The liquid volume is what the solids leave behind, and the liquid composition follows from the mass that is not one of the three lights, as shown in Equation (9.5).

$$\begin{aligned} V_L &= V - M_s/\rho_p, \qquad M_L = \rho_L V_L \\ n_{C3} &= \frac{1000}{\mathrm{MW}_{C3}} \left(M_L - \tfrac{1}{1000} \textstyle\sum_i n_i \mathrm{MW}_i \right) \\ x_i &= n_i / n_L, \quad r_H = 1000 \, \frac{x_{H2}}{x_{C3}}, \quad r_E = 100 \, \frac{x_{C2}}{x_{C3}} \end{aligned} \tag{9.5}$$

where V is the loop volume, ρ_p is the particle density, ρ_L is the liquid density, n_L is the total liquid amount, r_H is the composition ratio in mol H_2 per kmol C_3 and r_E is mol C_2 per 100 mol C_3. Slurry density is $\rho_{sl} = (M_s + M_L)/V$ and the solids weight fraction is $w_s = M_s/(M_s + M_L)$. Liquid density is taken linear about the design point, $\rho_L = \rho_{L,0}\,[1 + \kappa(P_r - 38) - \alpha(T_r - 70)]$.

9.3.2 The Bubble Point and the Cavitation Margin

The constraint is one algebraic equation, and it is the reason the application exists. Raoult's law supplies the solvent term and Henry's law supplies the three solutes, as shown in Equation (9.6).

$$P_{bub} = x_{C3} P^{\mathrm{sat}}_{C3}(T_r) + \sum_{i \in \{H2, C2, I\}} x_i H_i(T_r) \tag{9.6}$$

where P^{sat}_{C3} is the propylene saturation pressure and H_i is the Henry constant of solute i. The saturation pressure is a three-constant Antoine surrogate, $\log_{10}(P^{\mathrm{sat}}/\mathrm{bar}) = A - B/(T[\mathrm{K}] + C)$ with $A = 4.23311$, $B = 932.239$ and $C = -4.262$, fitted by least squares to seven propylene saturation pressures between 0 °C and 90 °C. Re-checking that fit against all seven reference points gives a worst relative residual of 0.45 %, at 90 °C, and no more than 0.27 % anywhere from 0 °C to 80 °C, against the 0.5 % tolerance the application's own verification suite applies. Each Henry constant is fitted at 70 °C and slides with a power of the solvent's own vapor pressure, $H_i(T) = H_{i,0}\,[P^{\mathrm{sat}}_{C3}(T)/P^{\mathrm{sat}}_{C3}(70)]^{n_i}$, so that every solute becomes less volatile relative to the solvent as propylene approaches its critical point. The margin the pump lives on and the margin the impeller actually sees are Equation (9.7).

$$\Delta P = P_r - P_{bub}, \qquad \Delta P_e = \Delta P - h_e \tag{9.7}$$

$$h_e = \frac{\rho_{sl}}{2 \times 10^5} \left[\sigma_v v^2 + \sigma_N \left(\frac{N v_r}{100} \right)^2 \right] \tag{9.8}$$

where h_e is the impeller-eye depression in bar, a two-term net positive suction head surrogate carrying a flow term and a peripheral-speed term [2], N is pump speed in percent and v_r is the rated velocity. At the design point $h_e = 0.36$ bar, so the transmitter reports 3.43 bar of margin while the impeller has 3.08 bar. Vapor appears when ΔP_e goes negative, and the void fraction and the head it costs are Equation (9.9).

$$\alpha_v = \begin{cases} 0, & \Delta P_e \geq 0 \\ \alpha_{\max} \left(1 - \mathrm{e}^{\Delta P_e / \delta} \right), & \Delta P_e < 0 \end{cases} \tag{9.9}$$

$$f_h = \frac{1}{1 + (\alpha_v / \alpha_{50})^2} \tag{9.10}$$

where δ is the void-growth pressure scale and α_{50} is the void fraction at which the pump delivers half its head. The form reproduces the shape of a real head breakdown, which is nothing and then everything over a fraction of a bar, and nothing else about cavitation.

9.3.3 Circulation Momentum and the Settling Limit

The circulating slurry is a 196 m column of fluid between the pump and the loop resistance, so velocity obeys a momentum balance rather than an algebraic pump-curve intersection, as shown in Equation (9.11).

$$\rho_{sl} L \frac{\mathrm{d}v}{\mathrm{d}t} = 60\left(\Delta P_{\mathrm{pump}} f_h - \Delta P_{\mathrm{fric}}\right) \tag{9.11}$$

$$\begin{aligned} \Delta P_{\mathrm{pump}} &= \Delta P_{\mathrm{rated}}\left[R_{sh}\left(\tfrac{N}{100}\right)^2 - (R_{sh}-1)\left(\tfrac{v}{v_r}\right)^2\right] \\ \Delta P_{\mathrm{fric}} &= \tfrac{1}{2} K_f (1 + k_\phi \phi_S) \rho_{sl} v^2 (1 + 2.2\phi) \end{aligned} \tag{9.12}$$

where $K_f = fL/D + \sum K$ is the loop resistance coefficient, ϕ_S is the solids volume fraction, ϕ is the settled-layer state, R_{sh} is the pump shut-off head ratio, and the factor 60 converts an acceleration in m/s^2 to the per-minute clock the rest of the model keeps. Linearizing (9.11) about the design point gives a restoring slope of 79.7 kPa per m/s against 115.8 t/m^2 of inertia, so the circulation time constant is $\tau_v = 1.45$ s, confirmed by kicking the velocity state 0.30 m/s low and measuring 1.48 s to recover 63 %. That figure is the reason the application has two time scales, and it is the one place where reading the clock as seconds rather than minutes changes the answer by a factor of sixty.

Below a critical velocity the solids stop being carried and slide along the bottom of the horizontal runs [18]. The application uses a Durand form for that velocity and grows a settled layer whenever the circulation falls under it, as shown in Equation (9.13).

$$v_{\mathrm{dep}} = F_L \sqrt{2 g D \left(\rho_p / \rho_L - 1\right)} \tag{9.13}$$

$$\begin{aligned} \frac{\mathrm{d}\phi}{\mathrm{d}t} &= k_f \max\!\left(0, 1 - \tfrac{v}{v_{\mathrm{dep}}}\right) \frac{\phi_S}{0.328}\left(1 - \frac{\phi}{\phi_{\max}}\right) \\ &\quad - 6 \times 10^{-4} \phi \max\!\left(0, \tfrac{v}{v_{\mathrm{dep}}} - 1.25\right) \end{aligned} \tag{9.14}$$

Evaluating (9.13) at the design point gives 3.65 m/s for 900 kg/m^3 powder in 440 kg/m^3 liquid propylene through a 0.65 m bore, which is a ratio of 1.51 against the 5.50 m/s circulation. The erosion term in (9.14) removes a settled layer with a time constant of 108 h at the design velocity and still 68 h at 110 % pump speed, so a layer built during a transient does not come off again within any run the application makes.

9.3.4 Kinetics, Heat and Pressure

One lumped active-site inventory carries the kinetics, with the propagation rate first order in site concentration and in monomer, as shown in Equations (9.15) and (9.16).

$$r_p = k_p A_c C_{C3} f_H f_E f_z \left(1 + b_a/100\right) \tag{9.15}$$

$$\frac{\mathrm{d}A_c}{\mathrm{d}t} = k_a v_{cat} - \frac{A_c \dot{m}_{\mathrm{out}}}{M_s} - k_d A_c \tag{9.16}$$

with $k_p = k_{p,\mathrm{ref}} \exp[-(E_p/R)(1/T_r - 1/343.15)]$, $f_H = 1 + k_H r_H/(K_H + r_H)$, $f_E = 1 + k_E r_E$, $f_z = 1/(1 + K_z z)$ and $k_d = k_{d,0} \exp[\beta_d (T_r - 70)]$, where C_{C3} is the liquid monomer concentration in mol/m^3, A_c is the active-site inventory, v_{cat} is the feeder setting in percent, b_a is the batch activity bias and z is the poison level. Equation (9.16) has a decay time constant of 100 min at 70 °C, which makes production the slowest regulatory loop on the plant.

Reaction heat leaves through the jacket and with the cold feed, and the jacket coefficient carries a velocity exponent, as shown in Equation (9.17).

$$\begin{aligned} C_r \frac{\mathrm{d}T_r}{\mathrm{d}t} &= Q_{rxn} + Q_{\mathrm{feed}} - Q_j - Q_{\mathrm{loss}} \\ M_j c_{pw} \frac{\mathrm{d}T_j}{\mathrm{d}t} &= Q_j - F_{cw} c_{pw} (T_j - T_{cw}) \\ Q_j &= U_0 A_j \left(v/v_r\right)^{0.80} (1-\varepsilon)\,(T_r - T_j) \end{aligned} \tag{9.17}$$

where $Q_{rxn} = (-\Delta H_p) r_p$, $C_r = M_s c_{pS} + M_L c_{pL} + M_{st} c_{p,st}$ is the total thermal capacitance including the loop steel, ε is the combined external and settled-layer fouling, and F_{cw} is the cold-water flow set by valve v_{CW}. The exponent 0.80 is what converts a cavitating pump into a thermal runaway, because the velocity and the heat-removal coefficient fall together.

The loop is liquid-full, so pressure is not a level. It is what the liquid's own compressibility says about the mass that will not fit, as shown in Equation (9.18).

$$\frac{\mathrm{d}P_r}{\mathrm{d}t} = \frac{\dot{M}_L + (\rho_L/\rho_p)\,\frac{\mathrm{d}M_s}{\mathrm{d}t}}{V_L\,\rho_{L,0}\,\kappa} + \frac{\alpha}{\kappa}\,\frac{\mathrm{d}T_r}{\mathrm{d}t} \tag{9.18}$$

where $\dot{M}_L$ is the net liquid mass rate into the loop and κ and α are the isothermal compressibility and the thermal expansivity. The second term is worth 4.6 bar per degree in a vessel with nowhere to go, which is why the pressure controller cannot be left in manual. The first term carries a sign that surprises people, because polypropylene at 900 kg/m^3 is denser than liquid propylene at 440 kg/m^3, so polymerizing shrinks the contents.

9.3.5 Product Quality and the Analyzer

Melt flow rate and comonomer incorporation are correlations for the polymer being made at this instant, as shown in Equation (9.19).

$$\begin{aligned} \ln Q_{\mathrm{inst}} &= \ln K_Q + a_T\left(\tfrac{1}{343.15} - \tfrac{1}{T_r}\right) \\ &\quad + m \ln(b_0 + r_H + b_E r_E - 0.6 z) \\ w_{C2} &= \frac{a_{E1} r_E}{1 + a_{E2} r_E} \\ \rho &= \rho_h - k_\rho w_{C2} + k_{\rho Q} \ln(Q/Q_{\mathrm{ref}}) \end{aligned} \tag{9.19}$$

where $m = 2$, so melt flow rate goes as the square of the chain-transfer group and the process gain is $\mathrm{d}Q/\mathrm{d}r_H = 2Q/(b_0 + r_H)$. Evaluating that derivative by re-solving the steady state gives 6.25 g/10min per unit of ratio at the design grade and 12.50 g/10min per unit at 12 g/10min, so the process gain doubles across one grade transition. What leaves the loop is not Q_{inst} but a first-order average over the polymer inventory, as shown in Equation (9.20).

$$\frac{\mathrm{d}\bar{q}}{\mathrm{d}t} = \frac{r_p\,(\ln Q_{\mathrm{inst}} - \bar{q})}{M_s}, \qquad \frac{\mathrm{d}\bar{e}}{\mathrm{d}t} = \frac{r_p\,(w_{C2} - \bar{e})}{M_s} \tag{9.20}$$

whose time constant is $M_s/\dot{m}_{\mathrm{out}} = 52.4$ min at the design point. Analyzer QI-101 samples that product every $\Delta t_a = 6.0$ min and reports $\theta_a = 4.0$ min later, so the effective dead time in front of the quality loop is about 7 min, being half the sample interval plus the transport.

9.3.6 Parameters and Initialization

Table 9.2 lists the constants that set the behavior reported in this chapter. The application exposes 95 parameters in all, of which 20 are hidden from the model panel because they are quantities a plant would have to infer rather than read.

Table 9.2 Model Parameters. The Seven Groups Below the Lower Rule Are Hidden From the Application's Model Panel Because They Are Quantities a Plant Would Have to Infer. The Rest Are Adjustable.

Symbol	Meaning	Value	Unit
V	loop volume	65.0	m^3
D	pipe bore	0.65	m
L	developed length	196	m
H	vertical leg height	40	m
A_j	jacketed area	326.7	m^2
ρ_p	particle density	900	kg/m^3
$\rho_{L,0}$	liquid density	440	kg/m^3
κ	compressibility	1.3×10^{-3}	bar
α	thermal expansivity	0.0060	1/K
$H_{H2,0}$	Henry, H_2 at 70 °C	2600	bar
$H_{C2,0}$	Henry, C_2 at 70 °C	130	bar
$H_{I,0}$	Henry, inerts at 70 °C	72	bar
v_r	rated velocity	5.5	m/s
K_f	loop resistance	10.22	—
R_{sh}	shut-off head ratio	1.90	—
$\alpha_{\max}$	maximum void fraction	0.25	—
α_{50}	void at half head	0.055	—
δ	void-growth scale	0.35	bar
F_L	Durand factor	1.00	—
k_f	settling growth rate	0.0040	1/min
E_p	propagation activation	45000	J/mol
$k_{d,0}$	site decay at 70 °C	0.0100	1/min
β_d	decay temperature term	0.045	1/K
$-\Delta H_p$	heat of polymerization	2046	kJ/kg
K_Q	melt-flow constant	3.2552	—
m	exponent on transfer group	2.00	—
b_0	spontaneous transfer	0.060	—
a_T	Arrhenius term on Q	4711	K
k_ρ	density per wt% C_2	0.00170	—
U_0	jacket coefficient	69.0	kJ/min m^2K
M_j	jacket water inventory	12000	kg
P_f	flash pressure	18	bar
$T_{\max}$	swelling limit	80.0	°C
P_{hh}	relief set pressure	44.0	bar
Δt_a	analyzer interval	6.0	min
θ_a	analyzer dead time	4.0	min
σ_v, σ_N	eye depression terms	2.0, 2.0	—
n_{H2}, n_{C2}, n_I	Henry exponents	0.35, 0.75, 0.85	—
$k_{p,\mathrm{ref}}$	propagation, 70 °C	0.0299873	kg m^3/mol min
k_a	catalyst activation	6.4646×10^{-4}	1/% min
ζ	H_2 chain-transfer use	7.13×10^{-4}	—
M_{st}	loop steel mass	78540	kg
μ_L	liquid viscosity	7.5×10^{-5}	Pa s

No scenario hard-codes a state. Every scenario solves for the state that makes its own setpoints an exact equilibrium, using three scalar solves for the inert inventory that balances its own feed against its own purge, the velocity where the pump curve crosses the loop resistance, and the jacket temperature that carries the reaction heat away. Evaluating all nineteen derivatives at the solved design point gives a worst relative residual of 2.7×10^{-15}, which is the numerical noise floor of the arithmetic that produced it. The design point is 70.0 °C and 38.00 bar with 591 kg/m^3 of slurry, which is 50.0 wt% solids and 19.20 t of powder in the 65 m^3 loop, making 22.0 t/h of a 3.00 g/10min homopolymer of density 0.90550 g/cm^3 at $r_H = 0.900\,\mathrm{mol/kmol}$. The residence time is 52.4 min on 9.1 g/min of catalyst, the circulation is 5.50 m/s at 263 kW of pump shaft power and a 35.6 s lap time, and reaction heat is 12.50 MW of which the cold monomer feed carries 1.46 MW and the jacket removes 10.83 MW with water at 41.2 °C through a 45.0 % cold-water valve.

The engine of Section 1.6.1 integrates the plant by fourth-order Runge-Kutta at a fixed step of

0.01 min (0.6 s), which places 2.4 integration steps inside one circulation time constant, well inside the stability limit of the method. Controller execution is every 0.05 min (3 s) and the historian keeps one sample every 0.17 min (10.2 s), every 17th integration step. The noise seed is 20260914.

A setting entered before Run gives the same run as the same setting carried through Reset. The start-up solves the steady state for the chosen setpoints and lines up only the stations the reader has not edited, and a loop switched to Manual or Auto stays where it was put. A setting the plant cannot hold, such as 84 °C or a negative margin, starts the run at the scenario's own steady state and acts on it as a step. A helper button beside the manual stations, "What does this grade cost in margin?", re-solves the steady state for the setpoints entered and reports the melt flow rate, the bubble point, the margin and the reactor pressure that the minimum margin would require.

9.3.7 Assumptions

Modeling effort went into the thermodynamic constraint and was withheld from almost everything else. Every simplification that resulted breaks down somewhere, and the paragraphs below name where.

The bubble point is a surrogate, not an equation of state. Equation (9.6) is linear in each solute mole fraction, carries no binary interaction parameters and treats the inert slate as one lumped species of molecular weight 34. A design calculation would use a cubic or associating equation of state and would find a genuinely nonideal liquid phase [19]. The Antoine fit is claimed only from 0 °C to 90 °C, and above 90 °C it extrapolates toward a critical point it does not represent, propylene's being 91.1 °C and 46.0 bar. The application clamps its temperature slider at 84 °C for that reason, and no number the model produces above 90 °C means anything. The Henry constants and their temperature exponents are fitted, and 2600 bar for hydrogen in a light hydrocarbon at 70 °C is the right order and nothing more.

One well-mixed volume. A real loop is a plug-flow circuit with the 35.6 s lap time the model itself reports, a measurable temperature rise across the pump, and composition gradients between the injection points and the take-off, which is precisely why catalyst enters at the pump discharge and hydrogen enters on a return leg. Residence-time distribution effects on the molecular weight distribution are the subject of a literature this model cannot enter [7].

One lumped active-site type with first-order decay. Ziegler-Natta catalysts have several site families with different hydrogen and comonomer responses, and that multiplicity is the origin of the broad molecular weight distribution that makes polypropylene processable [17]. Hydrogen activation of the catalyst appears in (9.15) as a fitted saturating multiplier rather than as a mechanism.

Quality is two correlations, not a chain-length distribution. There is no polydispersity, no xylene solubles, no isotacticity index and no bimodality in (9.19), so the model says nothing about the properties a converter cares about beyond melt flow rate and density.

Cavitation is a lumped void fraction with an algebraic breakdown curve. Equations (9.9) and (9.10) capture the shape of a real head breakdown and none of the acoustics, the damage rate, or the difference between incipient and fully developed cavitation. The eye depression in (9.8) is a two-term surrogate with fitted coefficients rather than a computed suction-specific speed.

Fouling is one scalar. Equation (9.14) represents settled powder in the horizontal runs. Real loop fouling includes a wall polymer film, grows preferentially at the bends, and is removed in a shutdown rather than by running faster.

The flash and the recycle are fixed recovery fractions. Inerts accumulate because recovery returns 85 % of what flashes and only the purge and the remaining 15 % remove them, which is the right mechanism on the right timescale, but the purge composition is assumed through enrichment factors rather than computed from a separation.

The state clamps bound a runaway rather than model it. Reactor temperature is clipped at 94 °C after every step, and the post-trip excursions reported in Section 9.6 run into both that clamp and the 90 °C validity limit of the Antoine fit. Those numbers are correct as statements about the simulator and meaningless as statements about a reactor.

9.4 The Constraint and Its Gradients

9.4.1 What Each Dissolved Species Costs

At the design point the 34.57 bar bubble point divides into 28.86 bar of propylene vapor pressure, 3.48 bar of accumulated light inerts at 4.83 mol% of the liquid, 2.23 bar of hydrogen at 856 mol ppm, and nothing from ethylene because the design grade is a homopolymer. The order of those two middle terms is the first surprise the application delivers, because the inerts nobody set on purpose are worth more bubble point than the hydrogen the quality controller is spending.

🕮 **Pause.** Try Exercise 9.1 before reading on.

Four gradients follow, each measured by re-solving the steady state on either side of the design point rather than by differentiating an expression. Raising r_H by one unit raises the bubble point by 2.44 bar, and the figure is almost exactly linear, at 2.439 bar per unit taken as a secant from 0.60 to 1.50 and 2.440 bar per unit taken locally. Raising r_E by one unit costs 0.87 bar locally at the homopolymer point and 0.86 bar taken as a secant out to $r_E = 2.0$. Each degree of reactor temperature is worth 0.629 bar. Reactor pressure is the only handle that buys margin, and it buys it at 1.004 bar per bar, because raising the pressure moves the bubble point itself by only −0.004 bar per bar through the liquid density.

The application reports the constraint as the sum Equation (9.6) makes it rather than as a single number. All four terms are trended separately on the pump view, and the hydrogen and ethylene terms appear again on the faceplate of the loop that moves them, beside the dissolved concentration in mol ppm and the valve position that produced it. Raising $r_{H,\mathrm{SP}}$ is therefore watched three ways at once: the hydrogen term of the sum rises, the margin gives up very nearly the same amount within the minute, and the melt flow rate the setpoint was asked for arrives an hour later. That is the whole collision, in three traces a student can read off one screen.

9.4.2 The Grade Envelope

Those numbers turn a quality request into a pressure requirement, as shown in Figure 9.2. The bubble point rises linearly with dissolved hydrogen while melt flow rate rises as the square of it, so the two lines that matter are the reactor pressure and the reactor pressure less the minimum margin. At 38.00 bar with a 2.00 bar minimum margin the bubble point may reach 36.00 bar, which is $r_H = 1.488$ and a melt flow rate of 7.80 g/10min. A 6 g/10min grade still leaves 2.46 bar. The 12 g/10min grade needs $r_H = 1.860$ and a bubble point of 36.91 bar, so it leaves 1.09 bar of margin at 38.00 bar and 4.11 bar at 41.00 bar. Raising the pressure to 41.00 bar moves the ceiling from 7.80 g/10min all the way to 25.2 g/10min, which is the whole 25 g/10min fiber grade, because melt flow rate is quadratic in the group that the linear margin constrains.

9.4.3 Three Time Scales in One Plant

The reason the collision cannot be tuned away is that the two controllers who want the hydrogen valve read clocks that differ by a factor of five to eight, and the equipment they are arguing about reads a third clock faster than either. Stepping r_H from 0.900 to 1.500 with the inventory loops closed and the quality loops open gives a liquid ratio that reaches 63 % of its change in 13.4 min, a bubble point that reaches 63 % of its change in 13.8 min, and a product melt flow rate that reaches 63 % of its change in 74.5 min. The margin is that bubble point read against a reactor pressure that wanders by a few tenths of a bar with the instrument noise, so its own 63 % crossing comes at 8.9 min in the shipped run and as late as 11.6 min on five other noise seeds. The same step moves melt flow rate from 3.00 to 7.92 g/10min and the steady-state margin from 3.43 to 1.97 bar. Against the 1.45 s circulation time constant of (9.11), the quality loop is roughly 3000 times slower than the equipment it can destroy. That ratio, and not a tuning constant, is why the constraint needs its own controller with its own measurement and its own authority.

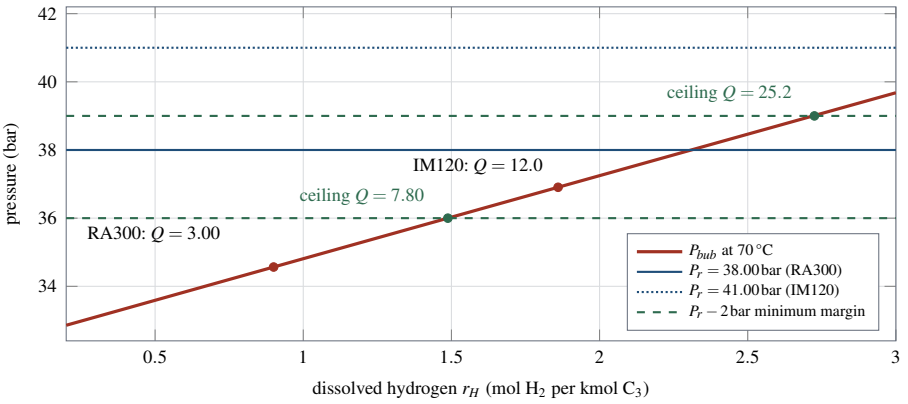


Figure 9.2 Bubble point of the loop liquid against dissolved hydrogen, with the reactor pressure and the minimum-margin line for two grade recipes. Every point is a separately solved steady state at 70 °C and the design inert level. The bubble point is the thick solid curve, the reactor pressure is a solid line at 38.00 bar and a dotted line at 41.00 bar, and the minimum-margin lines are dashed. Where the thick curve crosses a dashed line is the largest melt flow rate that pressure can support, 7.80 g/10min at 38.00 bar and 25.2 g/10min at 41.00 bar. The 12 g/10min grade sits above the 38 bar ceiling and comfortably below the 41 bar one.

9.5 Control Architecture

9.5.1 The Implemented PID

All ten loops run the velocity-form PID of Section 1.6.2 on a 0.05 min execution period, with the derivative on a measurement PV^f that a 0.4 min first-order filter smooths, advanced by its exact zero-order-hold recursion. Its limiting step carries two features of this app, as shown in Equation (9.21).

$$\begin{aligned} c_k &= u_{b,k} + s_k + I_{k-1} \\ u_k &= \mathrm{slew}\big(\mathrm{clamp}(c_k, u_{\min}, u_{\max,k}),\, u_{k-1},\, S\big) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right] \Delta t_c \end{aligned} \tag{9.21}$$

where e_k is the error of Equation (1.14), s_k is the never-clipped running sum of proportional and derivative increments of Equation (1.15), c_k is the demanded output (v_k in Equation (1.16)), $u_{b,k}$ is the output bias, $u_{\max,k}$ is an upper clamp that may move from scan to scan, S is the output slew rate, Δt_c is the execution period and $T_t = \tau_I$ is the tracking time, since no loop sets its own. The direction a of Table 9.3 is $+1$ on a loop whose process gain is positive and -1 where it is negative. Because the $(u_k - c_k)/T_t$ term acts on the difference between what was applied and what was demanded, output clamping and slew limiting both bleed the integral down over one integral time instead of resisting it. Manual operation, an interlock that overrides a loop, and a melt-flow controller that loses the low selector all set I_k directly so that c_k equals the applied output, which makes every transfer bumpless without any extra logic. The bias $u_{b,k}$ of each controller is preloaded with the actuator position that the solved equilibrium requires of that loop, which leaves both the error and the integral at zero when a scenario starts, and it stays constant on every loop except the constraint controller of Section 9.5.3.

9.5.2 Ten Loops in Three Tiers

Table 9.3 lists the loops, their default tunings and the response time each one has to live with. Three regulatory loops are fast, three are medium and two of the three quality loops carry the analyzer dead time.

Table 9.3 Implemented Loops With Their Default Tunings. Direction $a = +1$ Where the Process Gain Is Positive and -1 Where It Is Negative.

Loop	Actuator	K_c	τ_I (min)	a	Output range
Jacket water T_j	v_{CW}	8	0.6	-1	0–100 %
Reactor T_r	$T_{j,\mathrm{SP}}$	8	8	$+1$	8–62 °C
Pressure P_r	v_{C3}	4	2.5	$+1$	0–100 %
Slurry ρ_{sl}	v_{TO}	0.8	12	-1	0–100 %
Production R	v_{cat}	1.2	30	$+1$	0–100 %
Ratio r_H	v_{H2}	15	10	$+1$	0–100 %
Ratio r_E	v_{C2}	20	8	$+1$	0–100 %
Melt flow Q	$r_{H,\mathrm{SP}}$	0.55	45	$+1$	0.02–5.0
Density ρ	$r_{E,\mathrm{SP}}$	700	50	-1	0–6.0
Margin ΔP	$r_{H,\max}$	1.5	8	-1	0.02–clamp

The melt-flow and density controllers write composition setpoints rather than valve positions, and both are rate-limited at 0.05 units of ratio per minute. That limit is not a tuning constant in disguise. Hydrogen entering the loop leaves again only through a twenty-minute inventory, so a quality controller allowed to slam its setpoint spends bubble-point margin it cannot get back, and the rate limit is also what gives the constraint controller time to be a controller rather than a trip. Five architectures are offered as one selection. Manual leaves every valve to the operator. Inventory closes temperature, pressure, slurry density, production and both composition ratios and leaves melt flow rate and density to come out as they will, which is how loop reactors were run for years and is not unreasonable, because the ratios are what the catalyst sees. Quality adds the two outer quality cascades. Override adds the margin controller and the selector. The fifth keeps the override structure whole and places the predictive controller of Section 9.5.5 over it.

Each loop in Table 9.3 is presented as its own faceplate, carrying a MAN/AUTO switch, the setpoint it reads, the manual output station it writes and its two tuning constants, so every manipulated variable in the table is reachable directly and opening a cascade at any tier hands the student the signal that controller had been writing. The two composition ratios are the ones that matter pedagogically. Both $r_{H,\mathrm{SP}}$ and $r_{E,\mathrm{SP}}$ are outputs of the quality tier and setpoints to the regulatory tier, and taking the melt-flow controller to manual therefore turns $r_{H,\mathrm{SP}}$ into a direct operator handle on dissolved hydrogen, which is the same variable Equation (9.6) charges for. That is not a simulator convenience. A polyolefin plant, liquid loop or gas-phase fluidized bed (Chapter 8) alike, is operated through composition ratios rather than through feed valves, because the ratio is what the catalyst sees while the polymer property it sets is an inventory away. The ratio setpoint is the handle an operator is given and the flow controller beneath it is the handle the ratio controller is given.

9.5.3 The Low Selector

The melt-flow controller and the bubble-point controller both want the hydrogen setpoint and they want it in opposite directions, so a low selector decides, as shown in Equation (9.22).

$$r_{H,\mathrm{SP}} = \min\left(r_H^{Q},\; r_H^{B}\right) \tag{9.22}$$

where r_H^{Q} is the melt-flow controller's demand and r_H^{B} is the bubble-point controller's allowance. The acetylene converter of Chapter 3 takes its reagent away from a quality loop through the same kind of selector. Three details decide whether such an arrangement works [3]. The loser must not wind up, so whichever controller did not write the setpoint is told what was actually applied and its integral is tracked to it through the same track operation manual mode uses. The constraint controller must be able to win, so its output range reaches 0.02 mol/kmol, which is essentially no hydrogen at all.

And the idle controller must be parked where it can act, so its upper clamp and bias are set every execution to $\min(5.0, \max(r_{H,\mathrm{SP}}, r_H) + 0.5)$ rather than to the top of its range, with back-calculation holding it there. Section 9.6 measures what that third detail is worth. The implementation carries two separate flags out of the arrangement, one for the strict comparison and one for whether the constraint controller has left that clamp, because Section 9.6 shows that the two answer different questions and that only the second is fit to drive an operator indicator.

9.5.4 Hard Interlocks

Underneath every architecture are seven interlocks that no switch turns off, because each prevents an outcome that ends the run rather than degrading it. Hydrogen is slammed shut below 0.20 bar of margin at the pipe wall. Catalyst is limited to 15 % when the void fraction exceeds 2 % and to 30 % when the circulation falls below the deposition velocity. Catalyst is cut and the cold-water valve driven wide open above the 80.0 °C swelling limit, and above 82.0 °C the catalyst and every gas feed are latched shut until the unit is reset, because a loop full of tacky agglomerate is emptied with a hydro-jet rather than with a valve. The take-off is forced to at least 75 % above 60 wt% solids, and the propylene make-up is limited to 3 % above 43.5 bar. What matters pedagogically is which of those protects what. The hydrogen interlock defends the vessel against flashing its whole liquid-full contents. It is not a pump protection, and Section 9.6 shows it firing 2.15 min after the impeller has already been full of vapor.

9.5.5 Model Predictive Control

The fifth architecture puts a linear model predictive controller over the override structure [20, 21], an instance of the option of Section 1.6.3. It is never the application's default, and only scenario 11 loads it. Every $T_s = 2\,\mathrm{min}$, or 40 controller executions, it plans the three supervisory setpoints $u = [r_{H,\mathrm{SP}}, r_{E,\mathrm{SP}}, T_{\mathrm{SP}}]$ over $N = 90$ steps (180 min) in twelve move blocks of $\{1,1,1,1,2,2,3,4,6,9,15,45\}$ steps and applies the first move until the next plan, the first plan being solved at $t = 0$. The regulatory loops of Table 9.3 and the bubble-point controller stay closed beneath it. Its hydrogen request takes the place of r_H^Q in the low selector of (9.22), so the bubble-point controller and the seven interlocks keep the authority they have without it, and the displaced melt-flow and density controllers track the applied setpoints on every execution with no preload. The plan acts only while the header and the melt-flow, density, both ratio and reactor-temperature loops are in Auto.

The controller's model is its own. It sees the plant of Section 9.3 only through the instruments, the setpoints in force, the declared parameter values and its own state, so a parameter changed during a run is mismatch it is not told about. Its eight states are the two ratios with one intermediate lag each, the reactor temperature, the product log melt flow rate q, the product ethylene content e and a constant, as shown in Equation (9.23).

$$\begin{aligned} \tau_{H2}\dot{h}_H &= r_{H,\mathrm{SP}} - h_H, \qquad & \tau_{H1}\dot{r}_H &= h_H - r_H \\ \tau_{E2}\dot{h}_E &= r_{E,\mathrm{SP}} - h_E, \qquad & \tau_{E1}\dot{r}_E &= h_E - r_E \\ \tau_T \dot{T}_r &= T_{\mathrm{SP}} - T_r \\ \hat{\tau}\,\dot{q} &= \ln Q_{\mathrm{inst}}(r_H, r_E, T_r) - q, \qquad & \hat{\tau}\,\dot{e} &= w_{C2}(r_E) - e \end{aligned} \tag{9.23}$$

Here $\tau_{H1} = 9.5$, $\tau_{H2} = 2.35$, $\tau_{E1} = 16.3$, $\tau_{E2} = 1.7$ and $\tau_T = 1.5\,\mathrm{min}$, Q_{inst} and w_{C2} are the correlations of (9.19) at the declared constants with no poison, and the residence time $\hat{\tau}$ (52.4 min at design) is inferred from the measured slurry density and production. The outputs are $q + b_Q$, the density $\rho = \rho_h - k_\rho e + k_{\rho Q}(q + b_Q - \ln Q_{\mathrm{ref}}) + b_\rho$ and the margin $\Delta P = P_{r,\mathrm{SP}} - P_{bub}(r_H, r_E, T_r) + b_m$ from (9.6) at the design inert fraction of 4.83 mol%. The correlations and the bubble point are relinearized at each decision at the end of the previous plan, with the affine remainder kept, and discretized by an exact zero-order hold over T_s.

The ratio lags come from setpoint steps on scenario 5 with the inventory loops closed and no noise. They reproduce those responses with root-mean-square residuals of 3.7 % for hydrogen and 1.7 % for ethylene, and an independent two-lag fit gives 6.0 and 5.8 min for hydrogen, a different split of nearly the same total lag (11.9 against 11.85 min). The temperature loop reaches a setpoint step within about a minute and overshoots it by about 40 %, and its 1.5 min lag (residual 7.2 %) is slow on purpose because the plan moves T_{SP} at only 0.05 °C/min. At the design point the model gives 4.08 % of melt flow rate and −0.609 bar of margin per degree (against the plant's 0.629 bar of bubble point) and −2.44 bar per unit of r_H. The fixed inert fraction puts the model's margin 0.07 bar below the plant's at 66 °C, and over six decisions of scenario 11 the prediction 60 min ahead stays within 10 % of the melt flow rate and 0.36 bar of the margin.

The estimator runs on every execution whatever architecture is selected, integrating q and e from the measured ratios and temperature and running the lag states on the applied setpoints. The biases b_Q and b_ρ move by half of each analyzer residual, taken against the estimate held when the sample was drawn, and b_m filters over 5 min the difference between the inferred margin the override reads and the model's margin, which makes the prediction offset free [22], as the Kalman estimator of Volume 9, Chapter 4 does with a fuel bias. No event is shown to the plan in advance. Each plan minimizes Equation (9.24) subject to Equation (9.25).

$$
\begin{aligned}
J = {} & \sum_{k=1}^{N} w_k \left[\left(\frac{q_k + b_Q - \ln Q_{\mathrm{SP}}}{s_Q} \right)^2 + \left(\frac{\rho_k - \rho_{\mathrm{SP}}}{s_\rho} \right)^2 \right] \\
& + \sum_{j=1}^{12} \Delta u_j^{\top} S \, \Delta u_j + r_T \sum_{k=1}^{N} \left(T_{\mathrm{SP},k} - T_g \right)^2 + \sum_i \rho_i \varepsilon_i^2
\end{aligned}
\tag{9.24}
$$

$$
\begin{aligned}
& 0.02 \le r_{H,\mathrm{SP}} \le 5.0, \qquad 0 \le r_{E,\mathrm{SP}} \le 6.0 \\
& \max(58, T_g - \delta_T) \le T_{\mathrm{SP}} \le \min(84, T_g + \delta_T) \\
& |\Delta u_j| \le [0.1,\ 0.1,\ 0.1\,^{\circ}\mathrm{C}], \qquad \varepsilon_i \ge 0 \\
& |q_k + b_Q - \ln Q_{\mathrm{SP}}| \le s_Q (1 + \varepsilon_Q) \\
& |\rho_k - \rho_{\mathrm{SP}}| \le s_\rho (1 + \varepsilon_\rho) \\
& \Delta P_k \ge \Delta P_{\min} + \beta - \varepsilon_m
\end{aligned}
\tag{9.25}
$$

where $s_Q = \ln 1.1$ and $s_\rho = 0.0006\,\mathrm{g/cm^3}$ are the specification half-bands, $w_k = 1$ except $w_N = 6$, $S = \mathrm{diag}(3, 3, 30)$ weighs block moves measured from the command applied now, $r_T = 0.1$ per °C^2 draws the temperature toward the grade's own T_g within the freedom $\delta_T = 4\,^{\circ}\mathrm{C}$, $\Delta P_{\min} = 2.0$ bar is the override's setpoint and $\beta = 0.7$ bar the back-off above it. The input limits are the quality controllers' clamps, and the move caps are T_s times the rates at which the applied commands travel between decisions, 0.05 per minute on each ratio and 0.05 °C/min on the temperature. The bands are soft, each with one slack for the whole horizon, judged every third step and priced at $\rho_Q = \rho_\rho = 10$ per band squared and $\rho_m = 10^4$ per bar^2. The program has 39 variables. With one block and no constraints it agrees with independently formed normal equations to 1.7×10^{-13}, and a plan took 1.75 ms on average in Node.js over the 149 decisions of a five-hour scenario 11.

The setpoints follow a plan only when the solver reports it optimal, every number it returns (moves, predicted outputs, objective, slacks) is finite, and the moves respect their bounds. Any other plan leaves the setpoints where they stand, marks the run and hands the setpoints to the tracked melt-flow and density controllers under the override, and control returns with the first later plan that passes. With a failure forced from 100 to 110 min in scenario 11 and no noise, no execution moves a ratio setpoint by more than 0.0025, one execution of its slew, and the hydrogen valve moves by at most 0.049 %.

Table 9.4 Scenario Sequence and the Mechanism Each Shows

#	Scenario	What it shows
1	Manual at the design point	two clocks from one valve move
2	r_H step 0.90 to 1.50	gradient measurement, dead time
3	Production 22 to 27 t/h	activity dynamics, jacket ceiling
4	Inventory loops only, 7 h	what ratio control cannot see
5	Quality cascades lined out	tuning against the constraint
6	Grade transition, no override	cavitation, settling, trip
7	Same transition with override	the selector, and what it costs
8	Inerts creeping over a shift	a constraint spent by nobody
9	Cold night then fouled jacket	heat path into the bubble point
10	Pump speed 100 to 80 %	eye margin against settling
11	The override as a constraint	a predictive plan over the override
12	Graded challenge, 5 h	buy the margin before spending it

9.6 Results

Table 9.4 lists twelve scenarios in the order a class would work through them. A scenario pins the architecture, the grade recipe, the ten tunings, every disturbance setting and the timed events that arrive during the run.

9.6.1 The Transition That Destroys the Pump

Scenario 6 asks for 12 g/10min at $t = 15$ min with quality cascades in service, no override, and a melt-flow controller at $K_c = 1.6$ and $\tau_I = 20$ min, which is the tuning a step test at the starting grade produces. Nothing looks wrong for half an hour. The following sequence is read from the running model.

At $t = 51.35$ min, 36.35 min after the setpoint move, the impeller-eye margin goes negative while the transmitter at the pipe wall still reads 0.354 bar of margin and the liquid ratio has reached 2.169. Within 0.01 min the void fraction passes the 0.5 % at which the application counts cavitation, with the circulation still at 5.5 m/s. At $t = 53.4$ min the pump is delivering less than half its head and the velocity has fallen to 4.48 m/s. At the scan at $t = 53.5$ min the hydrogen interlock fires on 0.20 bar of wall margin, by which time the eye margin is −0.09 bar and has been negative for 2.15 min. At $t = 55.2$ min the velocity crosses the 3.67 m/s deposition velocity and the settled layer starts to grow. At $t = 67.6$ min the reactor passes the 80 °C swelling limit, and at the scan at $t = 68.25$ min the high-temperature trip latches the unit off at 82.1 °C with the circulation down to 1.61 m/s. Four hours after the start the run has spent 44.9 min cavitating, has 15.2 % of the bore under settled powder, and has made 44.2 t of polymer of which 38.7 t is off specification. Its worst reported margin of −20.1 bar is reached at $t = 78.5$ min, well after the trip and inside the region where the state clamp and the Antoine validity limit of Section 9.3.7 both apply, so the defensible statement is that the wall margin went negative at $t = 62.6$ min and stayed negative until the tripped reactor had cooled, at $t = 95.2$ min. On five other noise seeds the trip arrives between 64.5 and 125.3 min and the pump cavitates for 42.7 to 70.1 min.

Two things in that sequence are worth separating. The steady state alone would have survived, because 12 g/10min at 38.00 bar leaves 1.09 bar of margin. What cavitates the pump is the overshoot, since the liquid ratio reaches 2.29 against the 1.86 the grade needs while the analyzer has not yet registered any progress. And the eye margin, not the wall margin, is what decides, because the two differ by 0.36 bar at design flow and every alarm and interlock in the plant is written against the wall.

Table 9.5 Melt Flow Rate Step From 3 to 8 g/10min at $t = 20$ min, Read at $t = 300$ min. Both Columns Use the Default Tuning of Table 9.3. The Architecture Switch Is the Only Difference.

	Quality cascades	Plus override
Cavitation (min)	42.3	0.0
Worst margin (bar)	−19.82	+1.07
Minimum velocity (m/s)	1.51	5.50
Peak reactor temperature (°C)	94.0	70.2
Settled layer (% of limit)	15.2	0.0
Reactor still running	no	yes
Polymer made (t)	45.6	110.6
Off specification (t)	33.4	103.3
Melt flow rate at 300 min	1.71	6.55
Pump energy (MWh)	1.16	1.32
First cavitation (min)	56.5	—
First constraint selection (min)	—	45.7

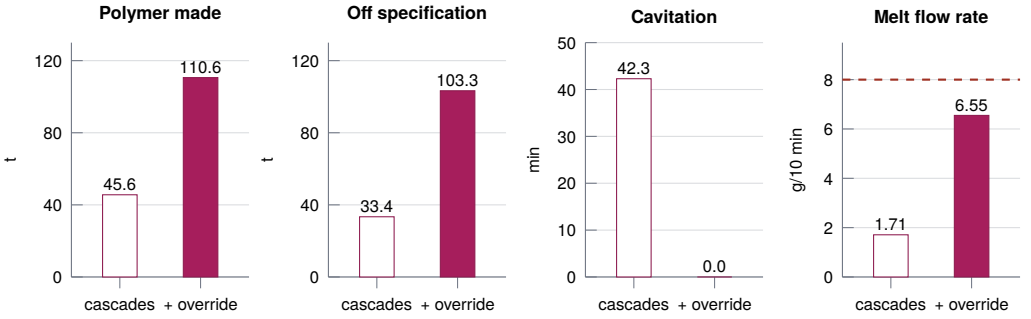


Figure 9.3 Four rows of Table 9.5 as bars, open for the quality cascades and solid for the cascades plus the override, each run read at 300 min. The dashed line in the last panel is the melt flow rate setpoint of 8. Notice that the override removes every minute of cavitation and more than doubles the polymer made, yet more of that polymer is off specification, because the margin it holds stops the melt flow rate short of the setpoint.

> **TRY IT IN THE APP** **6 · Grade Transition RA300 to IM120, No Override**
>
> Load the scenario, press Run and watch the Bubble-point margin strip on the Pump & bubble point view. After the melt-flow setpoint moves at 15 min, the dashed line at the impeller eye crosses zero at 51.35 min while the transmitter at the pipe wall still reads 0.354 bar. The circulation then slows and the reactor trips at 68.25 min.

Pause. Try Exercise 9.2 before reading on.

9.6.2 One Variable Changed

Table 9.5 compares the same run with and without the override, with the architecture switch as the only difference between the two columns. To keep the comparison free of the aggressive tuning that scenario 6 supplies, both runs start from scenario 5 at the design grade with the default melt-flow tuning of $K_c = 0.55$ and $\tau_I = 45$ min, and both receive a setpoint step from 3 to 8 g/10min at $t = 20$ min. The 8 g/10min target is chosen because its steady state at 38.00 bar does not cavitate the pump, at 1.95 bar of margin at the wall and 1.59 bar at the impeller eye, although it lies just above the 7.80 g/10min ceiling that a 2.00 bar minimum margin allows. Figure 9.3 draws four rows of the table as bars.

The override does exactly what a constraint controller is for and nothing more. It first takes the hydrogen setpoint at $t = 45.7$ min, 11 min before the unprotected run cavitates, and the pump then never leaves 5.50 m/s, the settled layer never grows, and the reactor is still running at five hours having made 110.6 t against 45.6 t. It also does not deliver the grade. Melt flow rate at 300 min is 6.55 g/10min against a setpoint of 8, which the 2.00 bar margin setpoint never allows at 38.00 bar, and 103.3 t of the 110.6 t is off specification, which is more off-specification polymer in absolute tonnes than the run that destroyed the pump produced. That is the honest trade and it is the centerpiece of the application. An override buys the equipment, not the product, and reading only the off-specification tonnage would rank the two runs the wrong way round. On five other noise seeds the protected run again never cavitates, while the unprotected run trips on four of them and on the fifth cavitates for 45.4 min without tripping.

The same comparison on scenario 6 and 7, where the melt-flow tuning is the aggressive $K_c = 1.6$ and $\tau_I = 20$ min and the target is 12 g/10min, gives 0.0 min of cavitation against 44.9 min, a worst margin of 1.21 bar against −20.1 bar, no settled layer against 15.2 %, and 88.6 t made against 44.2 t over four hours. Melt flow rate reaches only 7.49 g/10min at four hours against a setpoint of 12, and 83.2 t of the 88.6 t is off specification.

> ✦ DID YOU KNOW? **The Run That Saved the Pump Made More Off-Spec**
>
> From the simulator: in Table 9.5 the run without the override destroyed its pump and made 45.6 t of polymer, 33.4 t of it off specification. The run with the override kept the pump and made 110.6 t, of which 103.3 t was off specification, because the margin it protected never allowed the grade.

9.6.3 The Selector Chatters, and It Matters

A low selector whose loser is tracked to the winner's output has a diagnostic problem, and the application is where the problem and its remedy are both visible. Once the melt-flow controller's integral has been dragged down to the allowance, its own output is approximately equal to the allowance, so a strict inequality test reports that the constraint is not selected even while the constraint is what determines the answer. Counting the fraction of each of the first five twenty-minute windows after selection begins at $t = 41.8$ min in scenario 7 in which the constraint controller strictly wins gives 33.7 %, 19.8 %, 7.0 %, 19.0 % and 16.8 %, and the total is 39.9 min of the 198 min that follow. Over $70 \le t \le 240$ min the margin averages 1.99 bar against its 2.00 bar setpoint, and over the last hour of the run the melt flow rate is pinned between 7.2 and 7.5 g/10min against a setpoint of 12, so the constraint clearly governs the outcome for the whole period. An indicator built on the strict test alone therefore under-reports by a factor of about four and flickers at the execution rate.

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

The remedy is to ask a different question. Rather than testing which controller wrote this execution's setpoint, the application tests whether the constraint controller has come off the parking clamp of Section 9.5.3, which back-calculation holds it against for exactly as long as the margin is comfortable. That test reads true for 167.2 min of the same 198 min, and it is what colors the selector card, drives the badge and shades the hydrogen setpoint strip. The strict test is kept as a third state of the same panel and as its own exported column, so a student can see the difference between "the constraint wrote this setpoint" and "the constraint is holding the hydrogen down" rather than having one silently substituted for the other. The general lesson survives the fix. A selector indicator has to report authority and not arbitration, and the trend to read alongside either state is the separation between what the quality controller asked for and what the constraint allowed.

> **TRY IT IN THE APP** **7 · The Same Transition, With the Bubble-Point Override**
>
> Press Run on the Quality view and watch the low selector panel. Selection begins at 41.8 min. Over the 198 min that follow, the strict test of which controller wrote the setpoint is true for only 39.9 min, while the panel reports the constraint holding the ceiling for 167.2 min. The pump never cavitates.

The residual handover also leaves a real cycle. Averaging the inferred margin in three-minute blocks over $70 \le t \le 240$ min in scenario 7 gives a range of 1.70 to 2.37 bar, a peak-to-peak swing of 0.68 bar, and about 31 min between like extrema (0.36 to 0.64 bar and 23 to 31 min on five other noise seeds), which is the same order as the hydrogen inventory time constant the constraint loop is working against. Detuning the constraint controller from $\tau_I = 8$ min to 20 min cuts the swing to 0.29 bar and changes the worst margin only from 1.21 to 1.17 bar, while raising its gain from 1.5 to 3.0 raises the swing to 1.35 bar and costs 0.91 g/10min of melt flow rate averaged over the last hour. The shipped tuning is therefore a defensible compromise rather than a settled loop, and the application says so in the scenario note rather than claiming a settled margin.

The third detail of Section 9.5.3 is worth a number of its own. Forcing the constraint controller's clamp and bias back to the top of its 5.0 mol/kmol range on every execution, which is where an unconsidered implementation leaves it, changes the worst margin in scenario 7 from 1.21 to 0.48 bar, a loss of 0.74 bar out of a 2.00 bar setpoint, because the idle controller has to travel its whole range downward before it can act. That version does not cavitate in this scenario, and its ripple of 0.94 bar peak-to-peak is of the same order as the shipped 0.68 bar, so the cost of getting the parking wrong here is measured in margin rather than in a new oscillation.

9.6.4 Constraints Spent by Nobody in Particular

Scenario 4 runs the inventory loops alone for seven hours through a weak catalyst batch, a lighter feed specification and a trace of oxygenate. Melt flow rate holds 3.00 through the weak batch, reading 2.978 at 120 min and 3.003 at 240 min, while production dips from 22.0 to 16.9 t/h, because catalyst activity sets how fast chains are made and hydrogen sets how long they are. The oxygenate then moves melt flow rate to 2.65 g/10min, 11.7 % off target, with the H_2/C_3 ratio still at 0.901. Meanwhile the light inerts climb from 4.83 to 7.58 mol% and take the margin from 3.43 to 2.30 bar with every controlled variable on setpoint. Ratio control is an excellent inner loop and a poor outer one, which is the argument for the cascade, and the inert creep is invisible from every instrument except the loop chromatograph.

Scenario 8 makes that the whole story. A cheaper propylene cargo and a purge trimmed from 100 % to 30 % at $t = 20$ min take the inerts from 4.83 to 10.62 mol% by $t = 240$ min, raise their share of the bubble point from 3.48 to 7.66 bar, and walk the margin down onto the override setpoint at $t = 105$ min. The override takes its first hydrogen 91 min after the purge cut, which is the window a monitoring calculation on the loop chromatograph would have to flag the accumulation. The constraint controller then begins taking hydrogen away for a reason that has nothing to do with hydrogen. The liquid ratio falls from 0.900 to 0.36 at $t = 233$ min and the product melt flow rate to 1.30 g/10min at $t = 273$ min against a setpoint of 3.00, although the purge is restored to 130 % at $t = 240$ min, and 63.7 t is off specification by $t = 360$ min. Nothing alarms until the margin does.

Scenario 9 supplies the thermal route to the same place. Fourteen degrees of free cooling water takes the cold-water valve from 45.7 % to about 26 % and the jacket-limited production ceiling from 27.4 to 34.5 t/h. Removing 40 % of the jacket coefficient then drives the jacket water from 41.2 to a transient 17.5 °C and an average of 21.9 °C, with the cold-water valve reaching 100 % transiently and then working around 62 %, overshoots the reactor to 72.43 °C on the way, and leaves a ceiling of 23.8 t/h with 93 % of it in use on average. At 0.629 bar of bubble point per degree, the 2.43 °C overshoot alone is 1.53 bar of margin, and the worst margin over the run is 0.70 bar. The route from a fouled jacket to a cavitating pump is four steps long and there is no alarm on the first three.

Scenario 10 is the pump itself. Trimming the speed from 100 to 80 % saves 128 kW of shaft power, from 263 down to 135 kW, and improves the impeller-eye depression from 0.36 to 0.23 bar so that the steady-state eye margin actually rises from 3.08 to 3.21 bar. Everything else gets worse. The velocity falls from 5.50 to 4.40 m/s, which is 1.21 times the deposition velocity instead of 1.51, the jacket coefficient falls with $v^{0.80}$ so the heat-removal capacity in use climbs from 81 to 92 %, and the production ceiling falls from 27.4 to 24.2 t/h. Solving the steady state across the speed range puts the settling threshold at 66.4 % speed, so that at 66 % the circulation is already down to 3.63 m/s against the 3.65 m/s deposition velocity.

9.6.5 The Challenge

The challenge asks for 12 g/10min held within ±1.2 g/10min with the margin never below 2.00 bar and no cavitation at all, over five hours, starting from quality cascades with no override and a melt-flow controller at $K_c = 1.8$ and $\tau_I = 14$ min. Left alone that run cavitates the pump from $t = 44.3$ min for 115.5 min in all, takes the worst margin to −0.33 bar, puts 7.6 % of the bore under settled powder, and makes 72.8 t of which 47.2 t is off specification. It ends at 12.47 g/10min, inside the band, and it does not trip in the shipped run, but on five other noise seeds the same run trips the reactor between 59 and 109 min.

The answer is not a tuning constant alone. Raising the reactor pressure to 41.00 bar first is what the IM120 grade recipe in the application's own grade selector specifies, because on a liquid-full loop the operating pressure is part of the recipe in the same way the temperature is. Putting the override in service costs nothing when it is not needed. And slowing the melt-flow controller to $K_c = 0.15$ and $\tau_I = 80$ min recognizes both that the process gain doubles between 3 and 12 g/10min and that hydrogen leaves the loop only through a twenty-minute inventory, so an overshoot is paid for in bubble point long before it appears in melt flow rate. Done that way the five hours reach 11.88 g/10min, hold a minimum margin of 2.14 bar, never cavitate, and make 111.0 t of which 30.25 t is off specification. The override never acts in that run, and neither of the other two changes is enough alone, since 41.00 bar with the given tuning still trips the reactor after 45.2 min of cavitation and the slow tuning at 38.00 bar cavitates the pump for 66.9 min. The suggested settings give the same result whether they are entered before Run or carried through Reset. The general shape of that answer is to move the operating point until the constraint is not binding, keep the override as insurance, and stop asking a slow loop to move quickly.

The application is shown during that transition in Figure 9.4, at an instant when the strict selector comparison has momentarily handed the pen back while the constraint is still holding the ceiling. Each number quoted in that caption was read from the captured run and agrees with the same model stepped to $t = 51.01$ min outside the browser.

9.6.6 Comparison With Model Predictive Control

Scenario 11 repeats the transition of scenario 7, melt-flow tuning included, with the predictive architecture selected, and its own comparison is the same case reloaded with the override structure, which reproduces scenario 7 to the last digit. Table 9.6 runs the shipped structure at its shipped tuning and the plan at the settings of scenario 11 on that transition and on four other cases. Each pair shares its seed, instruments, noise, events and duration, the plan is not retuned between cases, and neither structure is shown anything in advance. Beside the headline the table carries the counterfactual runs that decide what the headline means.

Over four hours the plan never lets the override act, against 39.9 min of strict selection and 169.7 min off the parking clamp under the override structure, and the true margin never falls below 2.0 bar, with a worst value of 2.21 bar against 1.21 bar and 97.8 min below 2.0 bar. The product enters the 10.8 to 13.2 g/10min band at 136.8 min and stays, and 43.7 t of the 87.1 t made is off specification against 83.2 t of 88.6 t, as shown in Figure 9.5.

The comparison is not like-for-like, because the plan holds a temperature handle the override structure lacks. It lowers T_{SP} to 66.0 °C by 97.0 min, the floor of its freedom and 4 °C below the

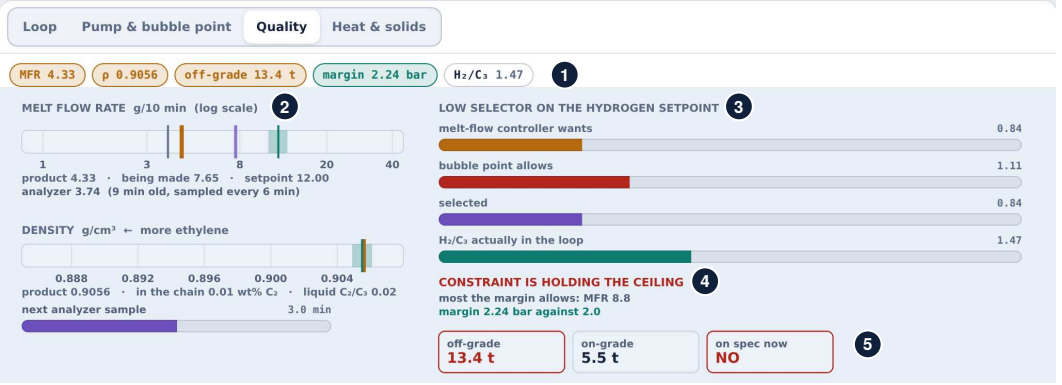


Figure 9.4 The Quality view of scenario 7 at $t = 51.01$ min, 36 min after a melt flow rate setpoint step from 3 to 12 g/10min with the bubble-point override in service. (1) The badges carry the product melt flow rate, the off-specification tonnage and the margin. (2) The melt flow rate scale marks the product, the polymer being made, the setpoint and the analyzer. (3) The low selector panel is the point of the view. The melt-flow controller is asking for 0.84, the bubble point allows 1.11, 0.84 is selected, and the liquid still holds 1.47 because the constraint took the pen at 41.8 min and hydrogen leaves only through a twenty-minute inventory. (4) The panel reads "constraint is holding the ceiling" rather than "constraint has the pen", because the strict comparison of the two signals is momentarily false at this instant while the constraint controller is still well off its parking clamp of 1.97. Product melt flow rate is 4.33 against 7.65 g/10min being made and an analyzer reading of 3.74 that is 9 min old, and the margin is 2.24 bar at the wall. (5) The reactor has made 13.43 t off specification against 5.49 t on grade, and the pump has not cavitated.

recipe, and spends the bubble point thereby freed on hydrogen, with $r_{H,\mathrm{SP}}$ peaking at 2.38 against the 1.86 that the grade needs at 70 °C. At 70 °C and 38.00 bar the grade leaves 1.09 bar of margin at steady state (Figure 9.2), so no structure can hold it there with 2.0 bar to spare. With no temperature freedom the plan makes 82.9 t off specification and ends at 6.17 g/10min, level with the override, although it still keeps the override idle and spends only 0.1 min below 2.0 bar.

The matched rows separate the temperature from the plan. Given the same ramp to 66 °C by hand, the override structure makes 62.4 t off specification and spends 54.5 min below 2.0 bar. The best such ramp, to 68 °C, makes 56.6 t (80.9 min below), and started at $t = 0$ with hindsight it makes 67.0 t. The plan held to the same path makes 43.8 t without a minute below 2.0 bar. Lowering the setpoint toward 66 °C by hand in steps of about 0.5 °C every 10 min instead of along the smooth ramp makes 68.3 t off specification under the override structure, with a worst margin of 0.82 bar and 40.3 min below 2.0 bar, against 1.21 bar for the smooth ramp, so each half-degree step evidently rings the margin through the pressure loop. On one temperature path, writing the margin as a constraint saves 18.6 t over the override and keeps the margin, while the colder reactor is what makes the grade reachable at all. Removing the back-off lowers the off-specification polymer to 36.4 t but lets the override select for 12.2 min, and pricing the margin at 1 instead of 10^4 per bar^2 returns the plan to the melt-flow controller's behavior, 83.7 t with 191.8 min of selection.

The colder reactor has a cost the controller's model does not see. Propagation is 17 % slower at 66 °C, so at steady state the production controller needs 49.1 % of the catalyst feeder against 42.9 % at 70 °C, about 14 % more, and the four hours make 1.6 t less polymer. Isotacticity, xylene solubles and the molecular weight distribution, which a real recipe protects with its temperature, are absent from the plant model as well, and the recipe for this grade asks for 70 °C at 41.0 bar. The plan has found an operating point the simulated plant permits, not one a recipe owner has approved.

The other cases divide the same way. Through the inert accumulation of scenario 8 the plan makes

Table 9.6 The Override Structure Against Model Predictive Control, Seed 20260914. Polymer in t, Margin in bar, Times in min

Case	Structure	Off spec	Q at end	Worst ΔP	Below 2.0 bar	Override selects, holds	$r_{H,\mathrm{SP}}$ travel	Lowest T_r (°C)	Mean v_{cat} (%)
11, 240 min	override structure	83.2	7.49	1.21	97.8	39.9, 169.7	17.89	69.8	43.55
	slow melt-flow tuning	83.0	6.65	1.21	51.9	11.4, 65.5	6.55	69.8	43.78
	hand ramp to 66 °C	62.4	9.76	1.21	54.5	20.2, 84.0	15.12	65.8	46.34
	hand ramp to 68 °C	56.6	11.23	1.32	80.9	31.5, 134.5	16.35	67.8	44.83
	the same from $t = 0$	67.0	10.27	1.24	67.4	28.0, 101.1	16.50	67.9	45.02
	predictive	43.7	11.89	2.21	0.0	0.0, 0.0	7.91	65.9	46.13
	no temperature freedom	82.9	6.17	2.00	0.1	0.0, 0.0	9.30	69.8	43.99
	the same, T_g ramped to 66 °C	43.8	12.00	2.32	0.0	0.0, 0.0	5.52	65.9	46.37
	no back-off	36.4	11.93	1.55	85.4	12.2, 155.5	9.11	66.5	45.18
	margin price 1	83.7	4.36	1.31	108.7	191.8, 207.7	95.00	69.9	42.54
8, 360 min	override structure	63.7	1.68	1.42	124.9	19.8, 184.0	7.58	69.8	46.54
	predictive	0.0	3.01	2.17	0.0	0.0, 0.0	7.72	67.5	47.36
	no temperature freedom	83.2	1.54	1.56	82.0	32.8, 111.6	16.48	69.8	46.60
9, 300 min	override structure	37.3	2.49	0.70	5.5	2.1, 6.2	4.24	69.8	44.89
	predictive	0.0	3.00	0.70	5.4	4.7, 6.2	3.16	69.8	44.64
5, 300 min	quality cascades	0.0	3.00	2.92	0.0	—	0.07	69.8	44.73
	predictive	0.0	3.00	2.94	0.0	0.0, 0.0	1.06	69.8	44.73
12, 300 min	quality cascades, as loaded	47.2	12.47	−0.33	220.2	—	6.11	66.7	20.27
	predictive	43.2	11.88	2.08	0.0	0.0, 0.0	9.86	65.9	46.30
	override structure, suggested	30.3	11.88	2.14	0.0	0.0, 0.0	3.30	69.8	42.85
	predictive, suggested	30.0	12.00	2.42	0.0	0.0, 0.0	4.94	69.5	42.67

Worst ΔP and the minutes below 2.0 bar use the true margin at the pipe wall. Selects counts the minutes of a true strict selector test and holds the minutes off the parking clamp, as in Section 9.6. Travel sums the absolute changes of the selected hydrogen setpoint. Hand ramps move T_{SP} at 0.05 °C/min from the 15 min setpoint step, and T_g ramped to 66 °C moves the plan's grade temperature along the same path. The slow tuning is the challenge's $K_c = 0.15$ and $\tau_I = 80$ min, and the suggested challenge runs at 41.0 bar. Only the challenge as loaded cavitates (115.5 min), no run trips, and the plan solves 120, 150 and 180 times in 240, 300 and 360 min without a rejected plan.

no off-specification polymer against 63.7 t by taking the reactor to 67.5 °C, and with its temperature held it loses, at 83.2 t. In the jacket case of scenario 9 it holds the grade but the override selects for 4.7 against 2.1 min, likely because the 40 % fouling step arrives faster than a plan made every 2 min can answer. On the lined-out scenarios 5 and 10 nothing is gained, and the plan moves the hydrogen setpoint 1.06 and 1.07 units over five hours against 0.07 and 0.12 under the shipped structures, most likely because it answers analyzer and margin noise that the slow melt-flow controller filters. The challenge as loaded is met by the plan at 11.88 g/10min and a worst margin of 2.08 bar, but with the reactor as cold as 65.9 °C, which the grade recipe does not specify, and with the suggested settings the two structures tie at 30.0 and 30.3 t. On five other noise seeds the five-hour transition gives 43.3 to 43.5 t under the plan against 105.1 to 105.2 t, and halving the integration step leaves scenario 11 at 43.7 t.

9.7 Technology in Practice

What Measures It

The variable the chapter is about has no sensor. As Section 9.9 says, the margin is computed from a pressure transmitter, a thermowell and the loop chromatograph, so an error in any of the three becomes an error in the margin. The temperature term alone is worth 0.629 bar per degree, and the chromatograph is the only instrument that sees the inerts creep of scenarios 4 and 8. The pressure

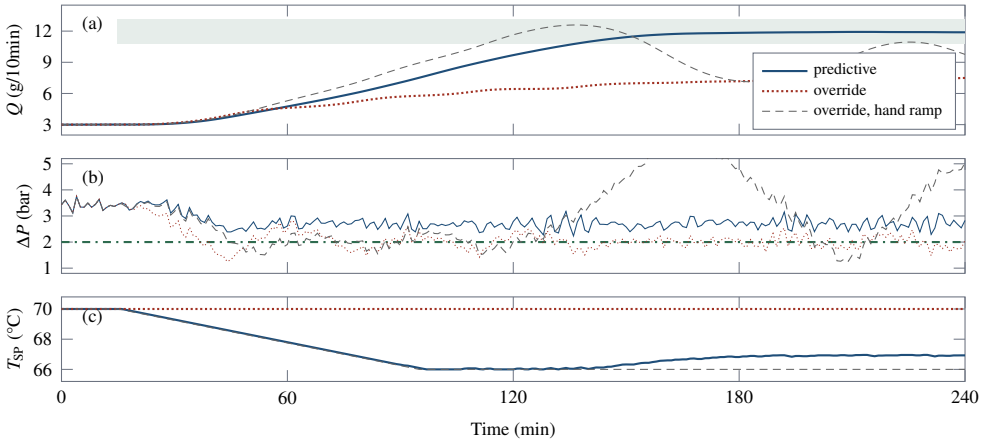


Figure 9.5 Scenario 11 over four hours under the predictive controller (solid), under the override structure (dotted), and under the override structure with the temperature setpoint ramped to 66 °C by hand (dashed). (a) Product melt flow rate and the 10.8 to 13.2 g/10min band (shaded) about the setpoint of 12 entered at 15 min. (b) True margin at the pipe wall and the 2.0 bar override setpoint (dash-dotted). (c) Reactor-temperature setpoint.

term is the subject of the box below. In the model PT-101 reads the pipe wall, where the reactor pressure wanders by a few tenths of a bar with the instrument noise, and the impeller sees 0.36 bar less than the wall.

> **TECHNOLOGY** **Sensor: A Pressure Transmitter Behind a Diaphragm Seal**
>
> Half of the margin is a pressure, and a tap on a slurry invites a plugged impulse line. A remote or chemical seal puts a flexible isolating diaphragm at the process connection and carries its pressure to the transmitter through a sealed, liquid-filled capillary, which removes the impulse line that clogs [23]. The fill must stay free of gas, because a bubble absorbs diaphragm motion and the reading falls short. A transmitter mounted below or above the tap reads high or low by the head of the fill liquid, heat near the capillary expands the fill and adds a false pressure, and the fill's viscosity in the narrow capillary slows the response [23]. On a trend these faults show as an offset or a sluggish line, and every bar of offset is a bar of margin that is not there.

What Moves It

Hydrogen reaches the loop through a valve that the ratio controller moves, so the selected ratio and not the valve is what the quality tier writes. A variable-speed motor is a good final element for a pump, because slowing the pump at low flow instead of throttling a valve saves a great deal of energy over time [23]. Scenario 10 shows the other side of that saving on a slurry loop, where the speed that saves 128 kW also brings the circulation toward the deposition velocity. In the model the pump speed is an operator setting, and its drive is the subject of the next box.

> **TECHNOLOGY** **Final Element: The Axial Circulation Pump and Its Drive**
>
> A single axial pump in each loop moves a large volume of slurry, which carries the heat to the water-cooled walls and keeps the particles from settling out of suspension [16]. A variable-frequency drive sets an induction motor's speed through the frequency of the power it supplies and its torque through voltage and current. The drive also limits speed, current, torque and the rate of acceleration, and shuts down on overcurrent or overvoltage [23]. The failure that matters is cavitation. Where liquid boils at a local pressure below its vapor pressure and the pressure then recovers, the bubbles collapse as microjets that pit any metal in time and crackle like gravel [23]. In the chapter's model, head and circulation velocity then fall together.

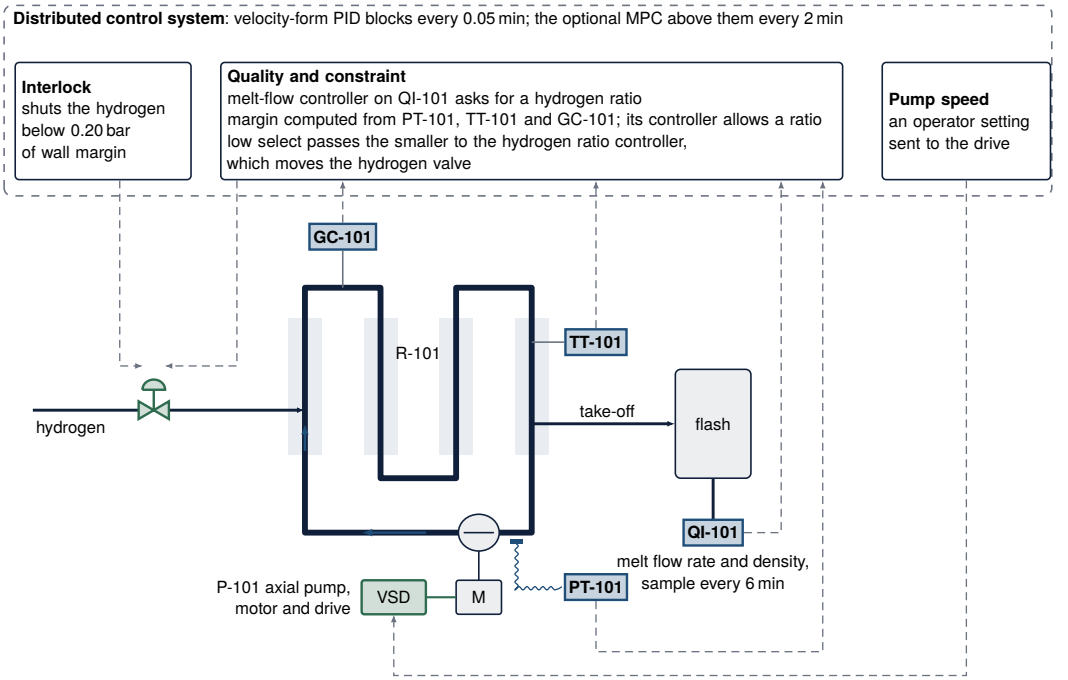


Figure 9.6 The hydrogen loops and the circulation pump with their hardware. Blue tags are sensing devices: the pressure transmitter PT-101 behind a diaphragm seal at the pump suction (its filled capillary drawn wavy), the temperature element TT-101, the loop chromatograph GC-101 and the product analyzer QI-101 under the flash. Green symbols are final elements: the hydrogen valve and the pump's variable-speed drive (VSD) with its motor (M). Dashed arrows are signals. Notice that the margin has no transmitter of its own: it is computed from three instruments, and the low selector then decides whether the melt-flow controller or the margin controller writes the hydrogen ratio setpoint. The interlock acts on the hydrogen valve directly, and the pump speed is an operator setting. The temperature, pressure, slurry-density and production loops are left out. Not to scale, and not a piping diagram.

Where the Controller Runs

Figure 9.6 draws the hydrogen loops of Section 9.5 with hardware in place of the schematic's symbols. A distributed control system holds loops like these at a supervisory level that sends set points to field controllers, with loops nested or cascading and cycle times from milliseconds to minutes [24]. An override is built there from a low-select function block between two controller outputs, and a de-selected controller must not wind up while it has no control. One remedy is the feedback input of the PID block, which stops integration once the block's output is no longer selected [23]. In the app all ten loops execute every 0.05 min, the loser of the low selector is tracked through the same operation manual mode uses, and the predictive controller is a supervisory layer that plans every 2 min.

What Is Modernizing It

Soft sensors built on industrial polypropylene data estimate quality between laboratory results, with the laboratory method kept as the certification reference [25]. A 2003 survey counted polymers among the areas with the most nonlinear predictive control applications and expected their number to grow where linear control had proven difficult, such as polymerization [20]. The EU reference document for polymer plants counts minimizing start-ups and stops as a best available technique. It

credits operating stability, assisted by computer monitoring and control systems, with fewer stops, and avoids an emergency stop by timely identification of deviating conditions followed by a controlled shutdown [16]. The override of this chapter is such an identification, made a controller. Whether a quality estimate can be trusted through a change of grade is the subject of the Open challenge box in Section 9.8.

> **IN PRACTICE** **Keep the Loop Liquid, in Two Layers**
>
> A liquid-propylene loop runs at a pressure high enough, typically 30 to 70 bar, to keep one continuous liquid phase with no separate gas. Gas in the loop can bring poor mixing, flow problems and mechanical damage to the loop pump [26]. A pump is best guarded in two layers: a soft override, where a low-select function block hands the final element to a constraint controller before the limit, and a hard shutdown behind it in case the soft layer is mistuned or fails [23]. The margin controller and the hydrogen interlock of this chapter are those two layers.

9.8 Limitations

The model is not a design tool for a loop reactor and no number in it belongs in a process design. The bubble point of (9.6) is a fit over a narrow band rather than thermodynamics, and the kinetic, quality, Henry, cavitation and fouling coefficients were chosen so that a teaching grade slate behaves like a polypropylene loop rather than to match any catalyst or any licensor's process. One well-mixed volume cannot show the composition and temperature gradients that decide where catalyst and hydrogen are actually injected [6], and one lumped site type cannot produce a molecular weight distribution. Cavitation is a lumped void fraction with an algebraic head-breakdown curve, so the application can say when a pump starts losing metal and never how fast. The post-trip excursions reported in Section 9.6 run into a 94 °C state clamp and past the 90 °C validity limit of the vapor-pressure fit, and are reported here only as the simulator's own behavior. Disturbances are entered as commanded steps and timed events, where a plant drifts. Three elements of the scenarios are stylized on purpose. No plant executes a transition from 3 to 12 g/10min by typing one setpoint, but a single step is the input whose consequences can be measured off a trend. A jacket that loses 40 % of its coefficient in one step stands for a season of operation, and it arrives as a step so that the four links from a dirty jacket to a cavitating pump can be counted. The poison and activity disturbances are separated cleanly, one moving melt flow rate and the other not, which a real feed impurity would not do. Fixing the integration step, the controller period and the random seed makes one run identical on every machine, and a distributed control system is not that repeatable. Five features are present in spite of all that trimming, the Henry-law bubble point, the eye depression, the two-time-scale split between the analyzer and the hydraulics, the nonlinear process gain and the selector, and every one of them is there because it changes the controller that has to be built. Section 9.3.7 names the condition under which every remaining assumption stops holding.

The predictive controller's model is narrower still. It fixes the inert fraction at its design value, which leaves its margin 0.07 bar below the plant's at 66 °C, an error the margin bias alone absorbs as the inerts move, and it omits the catalyst activity, the production, the poison, the pump, the hydrogen interlock and the pressure loop through which a temperature step rings the margin. Nothing in it makes 66 °C a poor temperature for the 12 g/10min grade, so the 4 °C freedom is a declared design choice rather than a property of the plant, and the extra catalyst and the polymer properties a recipe protects are charged to no one. Each output has one slack spanning all 90 steps, so a slack-free plan predicts that the line holds without promising it step by step, and in scenario 11 the margin slack is positive in 93 % of logged samples, at most 0.17 bar, which makes the 2.7 bar line a price rather than a wall. The plan loses where a disturbance outruns its 2 min decision, as in the jacket case, and on lined-out cases it moves the hydrogen setpoint nine to fifteen times as far as the shipped loops do.

OPEN CHALLENGE Quality Estimates That Survive a Change of Grade

Melt flow rate and the xylene-soluble fraction are laboratory results, the second by a gravimetric standard method (ASTM D5492), so neither gives continuous information on a reactor whose grade is moving [25]. Models that estimate them between samples exist. In a 2026 study of 425 industrial polypropylene runs, they lost accuracy when they crossed from one polymer architecture to another, and the study calls that gap between accuracy and transferability a central unresolved issue [27]. A companion study keeps the laboratory method as the certification reference and lets the model fill in between results [25]. Progress would be an estimate that holds through a transition and across homopolymer, random and impact grades.

9.9 Discussion

The organizational split around this constraint is as real as the technical one. The process engineer owns melt flow rate and density, the rotating-equipment engineer owns the circulation pump, and the bubble-point margin belongs to both and is trended by neither. That the margin is not measured at all makes it worse. It is computed from a pressure transmitter, a thermowell and a loop chromatograph through an equation-of-state routine in the control system, and the application does the same with (9.6) at its declared constants, an inert fraction reported 3 min after each 3 min sample and a 1.5 min filter. Getting that calculation trended, alarmed and put on a controller is most of the practical work, and the reason it is worth doing is that an axial pump in a 65 m^3 liquid-full loop is not a spared item.

Methods well beyond the ten implemented loops fit the same documented model. The selector of (9.22) is the regulatory answer to a problem that model predictive control states directly, by carrying the bubble-point margin as an inequality constraint on the same optimization that tracks melt flow rate and density, which removes the chatter measured in Section 9.6 and handles the doubling process gain through the model instead of through a selector [4, 5]. Grade transitions on this plant are naturally posed as optimal control problems in which the operating pressure is a decision variable rather than a recipe entry [5]. Moving horizon estimation on the same nineteen states would estimate the inert level and the catalyst activity bias from the pressure, temperature and analyzer records, so the invisible disturbance of scenario 8 becomes a soft sensor. Tuning the melt-flow controller by a rule that takes the destination gain rather than the origin gain is a small change with a measurable payoff, since the SIMC form for a 52 min lag with 7 min of effective dead time [28] returns different settings at the two ends of one transition.

9.10 Conclusion

This chapter has presented a bulk polypropylene loop reactor as an interactive laboratory for quality control against a thermodynamic constraint. The simulator integrates nineteen states and solves rather than assumes its design point, 70.0 °C and 38.00 bar with 50.0 wt% solids making 22.0 t/h of a 3.00 g/10min homopolymer, at which the bubble point of 34.57 bar divides into 28.86 bar of propylene, 3.48 bar of accumulated inerts and 2.23 bar of hydrogen, and all nineteen derivatives vanish to 2.7×10^{-15} relative. Hydrogen buys melt flow rate at a process gain that doubles from 6.25 to 12.50 g/10min per unit of ratio across one grade transition, and it spends bubble point at 2.44 bar per unit, against a design margin of 3.43 bar at the wall and 3.08 bar at the impeller eye. The plant therefore carries three clocks, 74.5 min for the product melt flow rate, 13.8 min for the bubble point and 1.45 s for the circulation. A melt flow rate step from 3 to 8 g/10min, whose steady state leaves 1.95 bar of margin, cavitates the pump for 42.3 min under quality cascades and not at all with a bubble-point override, and the override raises production from 45.6 to 110.6 t while leaving melt flow rate at 6.55 against a setpoint of 8. The application is deliberately explicit that this is a trade rather than a fix. Parking the idle constraint controller at the top of its range instead of just above the active demand costs 0.74 bar of worst-case margin, and the shipped constraint tuning leaves a 0.68 bar ripple at about a 31 min period that detuning the reset time to 20 min cuts to 0.29 bar. The predictive controller added as a fifth architecture never lets the override act on the grade change and makes 43.7 t off specification against 83.2 t, but from a reactor at 66.0 °C, and on the same temperature path

the margin written as a constraint is worth 18.6 t. Further work would add a second gas-phase reactor for impact copolymer, a plug-flow circuit with real composition gradients, a cavitation damage-rate correlation, and the operating pressure as a decision of the predictive controller. The loop reactor makes a constraint that has to become a controller into something a student can operate.

KEY POINTS

- The hydrogen that sets the grade also brings the loop liquid toward boiling.
- Hydrogen reaches the margin in minutes, the product in more than an hour and the pump in seconds, so the constraint needs its own controller.
- A low selector works when the loser cannot wind up, the constraint controller can win, and the idle controller is parked where it can act.
- An override buys the equipment, not the product: it removed 42.3 min of cavitation and left melt flow rate at 6.55 against 8.
- The challenge is won by raising the pressure before spending margin on hydrogen and by slowing the melt-flow controller.

Where else this applies. A fast limit and a slow quality loop sharing one actuator through a selector is also the anti-surge override of Volume 4, Chapter 2 and the hydrogen selector of Chapter 3.

References

[1] J. T. Gravdahl and O. Egeland, "Centrifugal compressor surge and speed control," *IEEE Transactions on Control Systems Technology*, vol. 7, no. 5, pp. 567–579, 1999.

[2] E. Grist, "Net Positive Suction Head (NPSH) and the pump operational range," in *Cavitation and the Centrifugal Pump*. Abingdon, UK: Routledge, 2023, pp. 51–69.

[3] C. A. Smith, "Ratio, override, and selective control," in *Automated Continuous Process Control*. New York, NY: John Wiley & Sons, 2002, pp. 74–126.

[4] M. Ohshima and M. Tanigaki, "Quality control of polymer production processes," *Journal of Process Control*, vol. 10, no. 2-3, pp. 135–148, 2000.

[5] M. Al-haj Ali, "Optimal grade transition control for liquid-propylene polymerization reactor," *Asian Journal of Control*, vol. 12, no. 3, pp. 413–425, 2010.

[6] J. J. Zacca and W. H. Ray, "Modelling of the liquid phase polymerization of olefins in loop reactors," *Chemical Engineering Science*, vol. 48, no. 22, pp. 3743–3765, 1993.

[7] A. S. Reginato, J. J. Zacca, and A. R. Secchi, "Modeling and simulation of propylene polymerization in nonideal loop reactors," *AIChE Journal*, vol. 49, no. 10, pp. 2642–2654, 2003.

[8] Z.-H. Luo, P.-L. Su, D.-P. Shi, and Z.-W. Zheng, "Steady-state and dynamic modeling of commercial bulk polypropylene process of Hypol technology," *Chemical Engineering Journal*, vol. 149, no. 1-3, pp. 370–382, 2009.

[9] K. Ziegler, E. Holzkamp, H. Breil, and H. Martin, "Das Mülheimer Normaldruck-Polyäthylen-Verfahren," *Angewandte Chemie*, vol. 67, no. 19-20, pp. 541–547, 1955.

[10] G. Natta, P. Pino, P. Corradini, F. Danusso, E. Mantica, G. Mazzanti, and G. Moraglio, "Crystalline high polymers of α-olefins," *Journal of the American Chemical Society*, vol. 77, no. 6, pp. 1708–1710, 1955.

[11] K. Ziegler, "Consequences and development of an invention," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/ziegler-lecture.pdf, read November 2026.

[12] G. Natta, "From the stereospecific polymerization to the asymmetric autocatalytic synthesis of macromolecules," Nobel Lecture, Stockholm, 12 December 1963, 1963, nobelprize.org/uploads/2018/06/natta-lecture.pdf, read November 2026.

[13] American Chemical Society, "Discovery of polypropylene and the development of a new high-density polyethylene," National Historic Chemical Landmark booklet, Office of Communications, Washington, DC, 1999, designated at Phillips Petroleum Company, Bartlesville, Oklahoma, 12 November 1999; acs.org/education/whatischemistry/landmarks/polypropylene.html, read November 2026.

[14] P. Galli and G. Vecellio, "Technology: driving force behind innovation and growth of polyolefins," *Progress in Polymer Science*, vol. 26, no. 8, pp. 1287–1336, 2001.

[15] H. H. Brintzinger, D. Fischer, R. Mülhaupt, B. Rieger, and R. M. Waymouth, "Stereospecific olefin polymerization with chiral metallocene catalysts," *Angewandte Chemie International Edition in English*, vol. 34, no. 11, pp. 1143–1170, 1995.

[16] European IPPC Bureau, "Reference document on best available techniques in the production of polymers," European Commission, Joint Research Centre, Institute for Prospective Technological Studies, Seville, Tech. Rep., Aug. 2007.

[17] J. B. P. Soares and A. E. Hamielec, "Kinetics of propylene polymerization with a non-supported heterogeneous Ziegler-Natta catalyst—effect of hydrogen on rate of polymerization, stereoregularity, and molecular weight distribution," *Polymer*, vol. 37, no. 20, pp. 4607–4614, 1996.

[18] M. C. Roco and C. A. Shook, "Critical deposit velocity in slurry flow," *AIChE Journal*, vol. 31, no. 8, pp. 1401–1404, 1985.

[19] D.-Y. Peng and D. B. Robinson, "A new two-constant equation of state," *Industrial & Engineering Chemistry Fundamentals*, vol. 15, no. 1, pp. 59–64, 1976.

[20] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[21] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

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

[23] 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.

[24] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[25] J. Hernández-Fernández, "Temporally validated and laboratory-reconciled soft sensing of xylene-soluble fraction in industrial polypropylene manufacturing," *Polymers*, vol. 18, no. 19, p. 2401, 2026.

[26] M. F. Bergstra, P. Denifl, M. Gahleitner, D. Jeremic, V. Kanellopoulos, D. Mileva, P. Shutov, V. Touloupidis, and C. Tranninger, "Polymerization in the Borstar polypropylene hybrid process: Combining technology and catalyst for optimized product performance," *Polymers*, vol. 14, no. 21, p. 4763, 2022.

[27] J. Hernández-Fernández and J. Lopez-Martinez, "Reactor-grade non-additive polypropylene architecture defines the physicochemical limits of melt flow index model transferability: An industrial chemometric and machine learning study," *Polymers*, vol. 18, no. 15, p. 1880, 2026.

[28] S. Skogestad, "Simple analytic rules for model reduction and PID controller tuning," *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

9.11 Exercises

9.1 ○ Concept check · 2 min

At the design point the bubble point of the loop liquid is 34.57 bar. Which statement about the three dissolved lights is right?

(a) Hydrogen is worth the most after propylene, because its Henry constant is the largest of the three

(b) The accumulated inerts are worth more than the hydrogen, although no controller sets them

(c) Ethylene is worth the most after propylene, because the density loop keeps it in the liquid

(d) The three lights together are worth less than a bar, so the propylene term sets the margin

9.2 ○ Concept check · 2 min

In scenario 6 the melt-flow setpoint steps from 3 to 12 g/10min at 38.00 bar with no override, and the pump cavitates about 36 min later. The steady state of the new grade leaves 1.09 bar of margin at the pipe wall. Why does the pump cavitate?

(a) The steady margin is smaller than the eye depression, so the new grade cannot be held at all

(b) The inerts accumulate during the transition and take the margin that the hydrogen left

(c) The controller overshoots the liquid ratio to 2.29 against the 1.86 the grade needs, before the analyzer registers progress

(d) The pump speed falls as the slurry grows denser, so the eye depression grows during the transition

9.3 Concept check · 2 min

In scenario 7 the strict selector test, "the bubble-point controller wrote this scan's setpoint", is true for only 39.9 min of the 198 min after selection begins. Over the same period the margin averages 1.99 bar against its 2.00 bar setpoint and the melt flow rate is pinned far below its setpoint. Why does the strict test read false most of the time?

(a) The override is idle most of the time, and the ratio controller holds the margin by itself

(b) The analyzer dead time delays the selector, so the test lags the constraint by one sample

(c) The constraint controller sits on its parking clamp at the top of its range for most of the period

(d) The melt-flow controller's integral is tracked to the applied setpoint, so its demand stays close to the allowance and the two signals keep crossing

9.4 Theory and calculation · 7 min

At 70 °C with no ethylene and no poison, Equation (9.19) gives the instantaneous melt flow rate as $Q = K_Q\,(b_0 + r_H)^2$ with $K_Q = 3.2552$ and $b_0 = 0.060$. At the design point $r_H = 0.900$, the margin is 3.43 bar at 38.00 bar, the bubble point rises 2.44 bar per unit of r_H, and the margin rises 1.004 bar per bar of reactor pressure.

(a) Find the ratio r_H that the 12 g/10min grade needs. *Checkpoint:* 1.860.

(b) Find the steady margin of that grade at 38.00 bar, and the reactor pressure that would leave 2.00 bar. *Checkpoint:* 1.09 bar and 38.91 bar.

(c) Find the largest melt flow rate that 38.00 bar supports with a 2.00 bar minimum margin. *Checkpoint:* 7.78 g/10min, within 0.05 g/10min.

9.5 Theory and calculation · 7 min

The melt-flow loop runs from the ratio setpoint $r_{H,\mathrm{SP}}$ to the product melt flow rate Q. Treat it as first order with the inventory time constant $\tau = 52.4$ min, an effective dead time $\theta = 7$ min, and the process gain $k = \mathrm{d}Q/\mathrm{d}r_H$ of the correlation $Q = K_Q\,(b_0 + r_H)^2$ with $b_0 = 0.060$.

(a) Show that $k = 2Q/(b_0 + r_H)$, and evaluate it at 3.00 g/10min ($r_H = 0.900$) and at 12.0 g/10min ($r_H = 1.860$). *Checkpoint:* 6.25 and 12.5 g/10min per unit of ratio.

(b) The SIMC rule with a closed-loop time constant $\tau_c = \theta$ gives $K_c = \tau/[k(\tau_c + \theta)]$ and $\tau_I = \min[\tau, 4(\tau_c + \theta)]$. Find both settings at each gain, and compare them with the default ($K_c = 0.55$, $\tau_I = 45$ min) and with the tuning the challenge hands over ($K_c = 1.8$, $\tau_I = 14$ min). *Checkpoint:* $K_c = 0.60$ and 0.30, $\tau_I = 52.4$ min at both.

9.6 Theory and calculation · 6 min

The loop holds a slurry of density $\rho_{sl} = 591\,\mathrm{kg/m^3}$ over a developed length $L = 196$ m, and linearizing the momentum balance at the design point gives a restoring slope of 79.7 kPa per m/s. The powder has $\rho_p = 900\,\mathrm{kg/m^3}$, the liquid $\rho_L = 440\,\mathrm{kg/m^3}$, the bore is $D = 0.65$ m, $F_L = 1.00$ and $g = 9.81\,\mathrm{m/s^2}$. The circulation runs at 5.50 m/s at 100 % pump speed.

(a) Find the circulation time constant in seconds, and its ratio to the 74.5 min the product melt flow rate takes to cover 63 % of a hydrogen step. *Checkpoint:* 1.45 s, a ratio of about 3100.

(b) Find the deposition velocity of Equation (9.13) and its ratio to the circulation. *Checkpoint:* 3.65 m/s, a ratio of 1.51.

(c) Show from Equation (9.12) that, with no vapor and a fixed solids fraction, the steady velocity is proportional to the pump speed. Find the speed at which the circulation reaches the deposition velocity. *Checkpoint:* 66.4 %.

9.7 App challenge · 7 min

Load scenario 2, Hydrogen step test, and change no setting. The H_2/C_3 ratio setpoint steps from 0.90 to 1.50 at 20 min. Run to 116 min at speed 30 with the default seed and export the CSV file. For a run longer than an hour the time column is headed Time (h) and is in hours. After the step the steady state has a bubble point of 36.03 bar and a product melt flow rate of 7.92 g/10min.

(a) Read column `Pbub_bar` and column `rh_actual` in the last row before the step and in the row at 115 min. Compute the bubble point gained per unit of ratio. *Checkpoint:* 34.57 and 36.02 bar, 0.902 and 1.500, about 2.43 bar per unit, within 0.03 bar.

(b) Find the time after the step at which `Pbub_bar` first covers 63.2 % of its change, and the same for the product melt flow rate in column `mfr_prod`. Give each within 0.5 min. *Checkpoint:* 13.8 min and 74.5 min.

State what the ratio of the two times means for a melt-flow controller that is tuned to make the product respond faster.

9.8 App challenge · 8 min

Scenarios 6 and 7 run the same grade change, a melt-flow setpoint step from 3 to 12 g/10min at 15 min with $K_c = 1.6$ and $\tau_I = 20$ min on the melt-flow controller. Scenario 6, Grade transition RA300 to IM120, has no override, and scenario 7, The same transition, with the bubble-point override, adds the margin controller. Load each in turn, change no setting, run past 4 h at speed 30 with the default seed and export the CSV file. In the row at 4.00 h (the last row before it, at 3.998 h) read `metric_cav_min`, `metric_worst_margin_bar`, `metric_offgrade_t` and `metric_ongrade_t`.

(a) Give the four values for scenario 6. *Checkpoint:* 44.9 min, −20.14 bar, 38.7 t and 5.5 t.

(b) Give the four values for scenario 7. *Checkpoint:* 0.0 min, 1.21 bar, 83.1 t and 5.5 t.

Accept 0.1 min, 0.02 bar and 0.2 t. Explain why the on-grade tonnage is the same in both runs, and which run is better.

9.9 App challenge · 10 min

Load the challenge, Challenge: reach 12 g/10 min without spending the margin, the last case in the list. It steps the melt-flow setpoint from 3 to 12 g/10min at 10 min and is judged at 300 min: the product melt flow rate within ±1.2 g/10min of 12, the margin never below 2.0 bar and no cavitation. For each run use the default seed, run past 5 h at speed 30, export the CSV file and, in the row at 5.00 h (the last row before it, at 4.998 h), read `mfr_prod`, `metric_worst_margin_bar`, `metric_cav_min` and `metric_offgrade_t`. Accept 0.02 g/10min, 0.02 bar, 0.1 min and 0.2 t.

(a) Run the case as loaded. *Checkpoint:* 12.47 g/10min, −0.33 bar, 115.5 min, 47.2 t.

(b) Press the reload button and, before Run, set Reactor pressure setpoint to 41.00 bar, Control architecture to + bubble-point override, and on the melt-flow controller card Proportional gain to 0.15 and Integral time to 80 min. *Checkpoint:* 11.88 g/10min, 2.14 bar, 0.0 min, 30.25 t.

(c) Press the reload button and make the same changes except the architecture, which stays at Quality cascades. *Checkpoint:* The same four values as in (b).

Say which runs pass, and what result (c) says about the override.

9.10 Application area · 4 min

Every alarm and interlock on the loop reads the margin at the pipe wall, ΔP, while vapor forms in the impeller eye, one eye depression h_e lower. Use Equation (9.8) with $\rho_{sl} = 591\,\mathrm{kg/m^3}$, $\sigma_v = \sigma_N = 2.0$ and $v_r = 5.50\,\mathrm{m/s}$.

(a) Find h_e at the design circulation of 5.50 m/s and 100 % speed, and at 4.40 m/s and 80 % speed. *Checkpoint:* 0.36 bar and 0.23 bar.

(b) The hydrogen interlock closes the valve at 0.20 bar of wall margin. Find the eye margin at that instant at design flow, and say what the interlock protects. *Checkpoint:* −0.16 bar.

9.11 Application area · 4 min

An energy audit proposes running the circulation pump at 70 % speed instead of 100 %. At full speed the pump draws 263 kW at 5.50 m/s, shaft power scales with the cube of speed, the velocity is proportional to speed, the deposition velocity is 3.65 m/s and the jacket coefficient scales with $v^{0.80}$.

(a) Find the shaft power saved, and compare it with the saving at 80 % speed. *Checkpoint:* 173 kW against 128 kW.

(b) Find the velocity, its ratio to the deposition velocity and the fraction of the jacket coefficient that remains, and give a recommendation. *Checkpoint:* 3.85 m/s, 1.05, and 0.75.

9.12 Review the agent · 5 min

An AI assistant summarizes scenario 11 as follows. “On the same transition the predictive controller makes 43.7 t of off-grade polymer against 83.2 t under the override, and the override never has to act. Writing the bubble-point margin as a constraint in the plan is therefore worth about 40 t per transition. The plant should replace the override structure with the predictive controller.” Find the flaw in the reasoning and say how it could be checked on the app.