INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 12 Control in Petrochemicals and Polymers

CHAPTER 11

A Butyl Rubber Reactor Simulator for Fouling, Actuator Margin and Wash Policy

Pages 257–281

Edited by John D. Hedengren

How to cite this chapter. "A Butyl Rubber Reactor Simulator for Fouling, Actuator Margin and Wash Policy," 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. 11, pp. 257–281.

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

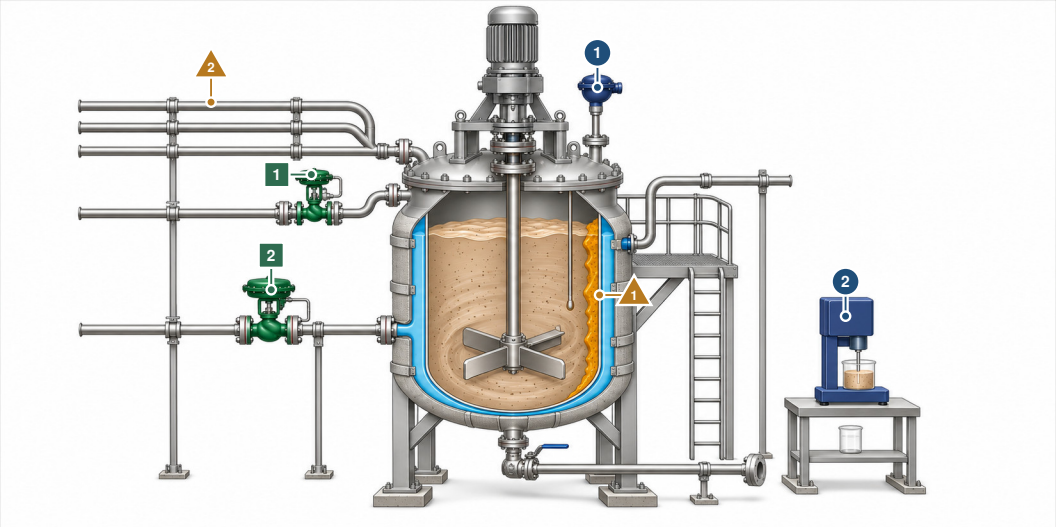
The app for this chapter
https://apmonitor.com/apps/butyl-rubber/

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

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

CHAPTER 11

A Butyl Rubber Reactor Simulator for Fouling, Actuator Margin and Wash Policy



The controller moves 1 catalyst valve 2 refrigerant valve that sets the jacket temperature

The controller reads 1 reactor temperature 2 Mooney viscosity of the rubber, a laboratory test every two hours

What disturbs it 1 polymer coating the cold wall 2 isobutylene feed rate

Notice the amber layer between the slurry and the cold jacket: it insulates the wall, so the jacket must run ever colder to remove the same heat until it reaches its floor. Illustration, not to scale.

IN PLAIN TERMS A Cold Reactor That Coats Its Own Walls

The main idea. Butyl rubber is made in a stirred tank kept colder than any freezer. As it forms, some rubber sticks to the cold walls and blankets them, like frost on a freezer coil.

The challenge. The controls hide this by running the cooling ever colder, until it can go no colder and the plant must stop to wash. Running colder slows the coating but spends the cooling sooner, so how cold to run and when to wash are hard choices.

Why it matters. The warning to watch is how hard the cooling works, not the tank temperature the controls keep flat.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the two History boxes and the Did you know? box in Section 11.2, the two Try it boxes and the second Did you know? box in Section 11.6, and Key points at the end of the Conclusion. No equations.
Study	50 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 11.6 with Figure 11.3, Section 11.7, and Section 11.8 with its Open challenge boxes.

The control question A butyl rubber reactor fouls its own cooling surface, and its temperature loop holds the setpoint while the jacket walks toward a refrigeration floor it cannot pass. Which trend shows how much of the run is left, what reactor temperature gives the most saleable rubber per elapsed hour, and what should decide when to wash?

Learning objectives

- Compute the design duty, the jacket temperature and the cooling margin, and read the run length from the closed-form countdown.
- Explain why a flat controlled variable and a drifting controller output together mean that the loop is spending a finite resource.
- Explain why colder operation shortens the run, and find the grade window that bounds the temperature optimum.
- Compare a fixed wash schedule with a wash on the cooling margin by the yield per elapsed hour, and name the measurement that ends a run.
- Judge when a constraint on the cooling margin, by override or by predictive control, pays on a reactor whose end of run creeps.

Prerequisites Sections 1.4.2 (run length against a limit), 1.4.4 (capacity set by heat removal), 1.6.2 (velocity-form PID), 1.4.6 (laboratory bias update) and 1.6.3 (the predictive option). Chapter 2 for the same countdown in coke.

The app https://apmonitor.com/apps/butyl-rubber/

See also Volume 4, Chapter 2 for a protective override and Volume 4, Chapter 8 for a run against a cash floor. Volume 3, Chapter 2 for matching a cascade to where the disturbance enters and Volume 3, Chapter 3 for the speed of an inner loop. Volume 8, Chapter 9 for corrections from late, sampled results.

11.1 Introduction

A control course teaches that a well tuned loop holding its setpoint is a loop doing its job. On the reactor modeled here it is a loop counting down. Butyl rubber is polymerized at −95 °C in a methyl chloride slurry, and the polymer deposits on every cooled surface from the first minute of a run. The overall heat transfer coefficient falls from 2400 to 976 W/(m^2 K) over forty-eight hours. A temperature controller absorbs that almost perfectly. Measured, the reactor temperature never leaves its setpoint by more than 0.01 °C while the jacket walks from −99.4 °C to −105.9 °C at hour 48 and reaches its −106 °C floor at hour 48.9, the temperature difference the duty requires grows from 4.4 K to 11.0 K, and the fraction of the jacket's capability still spare falls from 0.598 at hour zero to zero at hour 49.0. The controlled variable is a flat line for two days and then the loop has nothing left.

The misconception the simulator corrects is that the controlled variable is where the information is. On any plant with a slowly dying gain, the controlled variable stays put by construction and the manipulated variable is the diagnostic. Heat exchanger fouling, catalyst decay, valve wear and filter blinding all have this shape, and in every case the trend worth an alarm is the controller output rather than its input. The coke in a cracking furnace coil is the same countdown, read in tube metal temperature. Section 1.4.2 gives its run length in closed form, and Chapter 2 works it on a four-pass furnace. The class is large enough that scheduling cleaning against a fouling model is a research subject in its own right [1], and the reactor here adds the feature that makes it a control problem rather than a maintenance one: the actuator has a hard floor, so the end of the run is a saturation event, abrupt in the controller even though its cause was gradual.

Two further results in this application are not the intuitive ones and both are measured. Colder operation gives a higher molecular weight and slower fouling and a shorter run, because the refrigeration floor does not move when the setpoint does, and the measured optimum in tonnes of rubber per hour of elapsed time is near −94 °C rather than at the coldest reachable temperature. And a wash triggered on the cooling margin, the obvious feedback replacement for the fixed schedule most plants

run, yields less than the schedule when the feed fouls faster than the schedule assumed, because the margin runs out about fifteen hours before the product leaves grade.

This chapter works with an interactive butyl rubber reactor whose fouling rate, temperature acceleration and refrigeration floor are hidden, so that the countdown must be read from the actuator, the temperature optimum found by experiment, and a wash policy judged by the yield it delivers rather than by whether it uses feedback. The polymerization kinetics, the fouling model and cascade control are standard, and the app exports its complete time history (Section 1.6.4). Section 11.2 gives the history of the product and of its control. Section 11.3 writes out the slurry reactor with its fouling dynamics and its instrument chain in the detail needed to rebuild it. Section 11.4 derives the run length in closed form and the behavior at the refrigeration floor. Section 11.5 describes the four control architectures, the two wash policies and the predictive option. Section 11.6 reports the scenario sequence, whose temperature and scheduling conclusions contradict the obvious answers and are measured rather than argued, and compares the predictive controller with the shipped structures. Section 11.7 names the hardware and the practice behind the model, and Sections 11.8 to 11.10 give the limitations, the lessons that transfer to other plants and the conclusions.

11.2 Background

Butyl rubber is a copolymer of isobutylene with a small amount of isoprene, made by cationic polymerization at very low temperature, and the History box below tells how it was found [2, 3]. The account of its development given by one of its inventors thirty years later [4] makes clear that the low temperature was not a design preference but a requirement: the chain transfer reaction that limits molecular weight is strongly temperature dependent, so useful molecular weight is only reachable below about −90 °C. The diluent is methyl chloride because the polymer is insoluble in it and precipitates as a slurry rather than thickening the reactor [5].

> ⧖ HISTORY **A Rubber With Just Enough Double Bonds, 1937 to 1944**
>
> In 1941 a team of five chemists described butyl rubber as a new hydrocarbon product. Their work on polymers of simple olefins had then run for ten years. Long chains made from the butylenes behaved much like rubber, but a fully saturated polymer could not be vulcanized in the usual way [2]. The answer, in an application first filed in December 1937 by Robert M. Thomas and William J. Sparks, was to copolymerize isobutylene with a small amount of a diolefin such as isoprene, at low temperature, with aluminum chloride dissolved in ethyl or methyl chloride as the catalyst. The polymer kept a low, controlled unsaturation that sulfur could cure, and the patent was granted in 1944 [6]. Because polymers of isobutylene are so impermeable, they are widely used in inner tubes and in layers of tires [7].

That slurry is the source of the control problem. Polymer particles adhere to the cooled surfaces, and the deposit is an insulator between the reaction and the refrigerant. Butyl plants have therefore always run in cycles, with a scheduled shutdown for hydroblasting, and the run length between washes is one of the numbers a plant is managed by.

> ⧖ HISTORY **Washing the Reactor Clean, 1946**
>
> A patent application filed in 1946 describes the continuous butyl reactor: a central draft tube with an agitator, and return tubes around it, all cooled by liquid ethylene. A run proceeded smoothly for 10 to 150 hours while a layer of polymer built up on the vessel, the stirrer and the shaft, until the layer interfered seriously with the transfer of heat. The tough, rubbery coating was very hard to remove by hand, and light naphtha dissolved too little of it. The inventors drained the refrigerant, warmed the reactor with ethylene gas and filled it with warm liquid isobutylene. That softened the polymer and loosened its hold enough for the stirrer to break it away [8].

Control of the reactor developed in the order the instruments allowed. The first requirement is temperature, because molecular weight follows it directly, represented here by −1.40 Mooney units per degree. A jacket temperature loop is straightforward and the refrigeration system is the constraint underneath it, which is why split range and model-based algorithms for jacketed reactor temperature remain a subject even on much simpler services [9]. Product quality entered as a laboratory Mooney viscosity with a cycle measured in hours, which supports a cascade that corrects over shifts rather than minutes. Inferential estimation of polymer properties from reactor conditions became common later and carries the maintenance problems any soft sensor has [10]. Constraint handling above all of it, which is what decides whether a conversion loop should be allowed to recover production on a reactor that is nearly out of cooling, is the layer a modern plant delegates to a multivariable controller [11].

> + DID YOU KNOW? **Seven Percent of the Synthetic Rubber**
>
> In the world statistics for 2002, isobutylene-isoprene rubber, butyl and halobutyl together, made up 7 % of the production volume of the major synthetic rubbers, the same share as polyisoprene. Emulsion styrene-butadiene rubber led with 28 % and polybutadiene followed with 24 % [12].

The safety literature on polymerization reactors concentrates on thermal runaway, and the mechanisms are well catalogued [13]. This reactor is not a runaway problem, and the distinction is worth making explicitly. Its conversion is limited by catalyst rather than by temperature, so the heat release does not accelerate with temperature the way an adiabatic exotherm does. What accelerates with temperature here is the fouling, at 0.1386 per degree, so a warming reactor fouls faster and reaches its stop sooner. The failure is economic rather than energetic, and it is the more common of the two on operating plants.

11.3 Process Model

Reactor R-201 is a 30 m^3 slurry reactor with 190 m^2 of cooled surface, charged with 6.0 t/h of isobutylene, 18 t/h of methyl chloride and 0.12 t/h of isoprene. At 85 % conversion it makes 5.10 t/h of polymer at 21.1 % slurry solids, releasing 2.018 MW. Figure 11.1 shows the arrangement and every symbol used in the equations, and Section 1.4.4 places the reactor among the polymer reactors whose capacity is set by heat removal.

11.3.1 Reaction and Energy Balance

Conversion is set by catalyst and by temperature, Equation (11.1), and the heat released follows the polymer made, Equation (11.2).

$$X^{*} = \frac{k\tau}{1 + k\tau}, \qquad k = k_p \, a \, F_{cat} \, e^{\,b(T - T_{ref})} \tag{11.1}$$

$$Q = \Delta H_p \, F_{IB} \, X + \dot{m} \, c_p \, (T_{feed} - T) \tag{11.2}$$

where a is the catalyst activity, $k_p = 1.35$ per hour per kg/h of catalyst, τ is the residence time of 0.871 h, $b = -0.012/^{\circ}\mathrm{C}$ and $T_{ref} = -95\,^{\circ}\mathrm{C}$. The conversion X approaches X^{*} with the residence

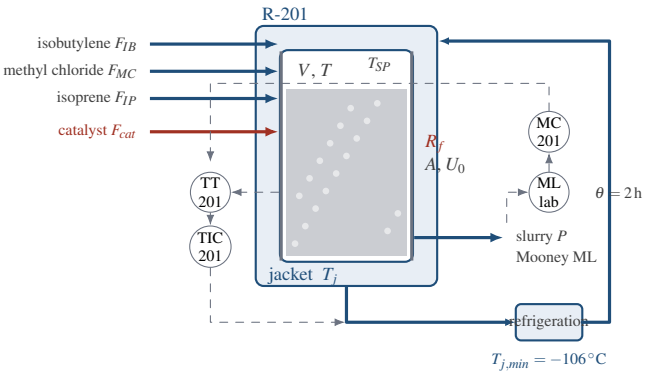


Figure 11.1 Butyl rubber slurry reactor R-201. Isobutylene F_{IB}, methyl chloride F_{MC} and isoprene F_{IP} enter with catalyst F_{cat} and polymerize at temperature T in volume V. Heat leaves through cooled surface A of clean coefficient U_0, carrying a growing fouling resistance R_f, to a jacket at T_j whose refrigeration cannot pass $T_{j,min}$. The slurry product P is characterized by its Mooney viscosity ML, reported by a laboratory on a cycle θ. TIC-201 manipulates the jacket temperature and MC-201 trims the reactor setpoint T_{SP} from the laboratory result.

Table 11.1 Classification of Variables at the Start of a Clean Run

Symbol	Role	Nominal	Range
T (TT-201)	CV, continuous	−95 °C	−102 to −86 °C SP
ML	CV, laboratory	51.0	46–56 spec
X conversion	CV	0.85	0.50–0.95 SP
T_j	MV	−99.42 °C	≥ -106 °C
F_{cat}	MV	4.82 kg/h	0–8 kg/h
F_{IB}	DV, measured	6.0 t/h	3–9 t/h
catalyst activity	DV, unmeasured	1×	0.4–1.3×
fouling tendency	DV, unmeasured	1×	0.4–3×
R_f	state	0	growing
U	derived	2400 W/(m^2 K)	falling
margin	derived	0.598	0 ends the run

time as its time constant, and the catalyst reaches the reactor through a 0.08 h lag. The heat of polymerization is $\Delta H_p = 930\,\mathrm{kJ/kg}$, and the second term chills $\dot{m} = 24.1\,\mathrm{t/h}$ of feed from $T_{feed} = -40\,^{\circ}\mathrm{C}$ at $c_p = 1.9\,\mathrm{kJ/(kg\,K)}$, so the design duty is 1.32 MW of reaction and 0.70 MW of feed chilling. The temperature dependence of conversion is deliberately weak, which is what makes this a fouling problem rather than a runaway problem.

11.3.2 Fouling, and the Actuator Margin

The fouling resistance accumulates and accelerates with temperature, Equation (11.3), and the overall coefficient follows from it, Equation (11.4).

$$\frac{\mathrm{d}R_f}{\mathrm{d}t} = k_f \frac{P}{P_0} \phi \exp\left[c_f (T - T_{ref})\right] \tag{11.3}$$

$$U = \frac{1}{1/U_0 + R_f} \tag{11.4}$$

with $k_f = 1.29 \times 10^{-5}\,\mathrm{m^2\,K/(W\,h)}$, $c_f = 0.1386/^{\circ}\mathrm{C}$, which is a doubling of the fouling rate for every five degrees of warming, ϕ the fouling tendency of the feed, and P the polymer rate against a reference $P_0 = 5.2\,\mathrm{t/h}$, so the design production of 5.10 t/h fouls at 0.98 of k_f.

Table 11.2 Model Parameters

Symbol	Meaning	Value	Class
V	reactor volume	30 m^3	adjustable
A	cooled surface	190 m^2	adjustable
U_0	clean coefficient	2400 W/(m^2 K)	adjustable
k_f	fouling rate	1.29×10^{-5}	adjustable
c_f	fouling temperature rise	0.1386/°C	adjustable
$T_{j,min}$	refrigeration floor	−106 °C	adjustable
k_{ML}	Mooney against temperature	−1.40/°C	adjustable
θ	Mooney laboratory cycle	2.0 h	adjustable
τ_{MW}	product mixing time	0.60 h	adjustable
τ_j	refrigerant loop lag	0.035 h	adjustable
wash time	shutdown and restart	8 h	adjustable
ML band	specification	46 to 56	visible

The jacket temperature the duty requires is Equation (11.5), and the cooling margin is the fraction of the available temperature difference still unused, Equation (11.6).

$$T_j = T - \frac{Q}{UA} \tag{11.5}$$

$$m = 1 - \frac{Q/(UA)}{T - T_{j,min}} \tag{11.6}$$

The margin is the clock. It is readable at any instant, it falls smoothly, and it reaches zero at the moment the jacket reaches its floor. At the design condition it starts at 0.598 and reaches zero at hour 49.0.

11.3.3 Product Quality

Mooney viscosity follows the reactor temperature directly, Equation (11.7), with the mixing time of the reactor between the temperature and the product.

$$\tau_{MW} \frac{\mathrm{dML}}{\mathrm{d}t} = \mathrm{ML}_{ref} + k_{ML}\,(T - T_{ref}) - \mathrm{ML} \tag{11.7}$$

with $k_{ML} = -1.40$ Mooney units per degree Celsius and $\tau_{MW} = 0.60$ h. The specification band is 46 to 56, so the reactor temperature has about seven degrees of usable range before the product leaves grade at one end or the other.

11.3.4 Assumptions

The reactor is a single well mixed volume. A real butyl reactor has an axial circulation driven by a large impeller and the fouling is not uniform, so the coefficient in this model is an average of something that varies over the surface.

The fouling resistance grows linearly in time at a given temperature and production rate and is removed completely by a wash. Real deposits have an induction period, an asymptotic component and a removal term, and a hydroblast does not always return the surface to its clean condition.

Conversion is a function of catalyst and temperature with no molecular weight distribution. Mooney viscosity is a correlation on reactor temperature rather than a moment calculation, so nothing about the breadth of the distribution can be studied.

Isoprene incorporation and the unsaturation specification are represented by a linear relationship to the feed ratio, which is a simplification of a copolymerization that depends on temperature as well.

The refrigeration floor is a hard limit on jacket temperature with no dynamics of its own. A real refrigeration system loses capacity as its own condensers foul and as ambient temperature rises.

The wash takes a fixed eight hours and restores the surface exactly. Both are convenient rather than true.

11.4 Dynamic Analysis

11.4.1 The Countdown Is Readable in Closed Form

With the loop holding the reactor and the duty roughly constant, the required temperature difference is $\Delta T = Q/(UA)$, and substituting Equation (11.4) gives Equation (11.8).

$$\begin{aligned} \Delta T(t) &= \frac{Q}{A}\left(\frac{1}{U_0} + R_f(t)\right) \\ &= \Delta T_0 + \frac{Q}{A}\, k_f\, \frac{P}{P_0}\, \phi\, e^{c_f (T - T_{ref})}\, t \end{aligned} \tag{11.8}$$

At a fixed reactor temperature the required difference grows linearly in time, which is why the jacket trend in scenario 2 is a straight line and why its slope is a direct reading of the fouling rate. Setting $\Delta T = T - T_{j,min}$ gives the run length, Equation (11.9).

$$t_{run} = \frac{A\, P_0}{Q\, k_f\, P\, \phi}\, e^{-c_f (T - T_{ref})} \left[(T - T_{j,min}) - \frac{Q}{A U_0} \right] \tag{11.9}$$

This is the constant-rate limit of Equation (1.2): the fouling rate here does not grow during a run held at one temperature. At the design condition Equation (11.9) gives 48.9 h, and the simulated run of scenario 2 spends its margin at 49.0 h. Both temperature effects are visible in that expression and they pull opposite ways. The exponential prefactor lengthens the run as the reactor is cooled, because the fouling slows. The bracket shortens it, because the distance to the refrigeration floor is what the run is spending, and it goes to zero at $T = T_{j,min} + Q/(AU_0)$, which is −101.6 °C at the design duty and −101.4 °C once the larger feed-chilling duty at that temperature is included. From a clean surface held at −101 °C the expression gives a run of 6.7 h. The measured run at −101 °C in Table 11.3 is shorter still, 2.9 h, because that setpoint arrives after two hours of fouling at −95 °C, and the wash follows 0.8 h after the setpoint change. Both are that bracket almost closed.

The maximum of t_{run} is not the operating optimum, because a longer run at a lower rate is not necessarily more rubber. The quantity to maximize is on-specification production per hour of elapsed time including the wash, which for a run that stays in grade is Equation (11.10), the on-stream rate of Equation (1.3) with the wash as the downtime.

$$Y = \frac{P\, t_{run}}{t_{run} + t_{wash}} \tag{11.10}$$

which is why the eight-hour wash appears in every economic comparison in Section 11.6 and why the answer to scenario 5 depends on it.

11.4.2 Behavior at the Refrigeration Floor

Nothing in Equation (11.8) is discontinuous. The discontinuity is the actuator. While $T_j > T_{j,min}$ the loop is closed and the reactor is held. At the floor the loop is open and the reactor obeys Equation (11.11).

$$\rho V c_p \frac{\mathrm{d}T}{\mathrm{d}t} = Q(T) - U(t)\, A\, (T - T_{j,min}) \tag{11.11}$$

The reactor holds 39.9 MJ/K, so a sustained shortfall of 0.1 MW would warm it 9 °C per hour, but no such shortfall builds up. Each degree of warming lowers the feed-chilling duty and raises the duty the floor can remove by UA, 0.18 MW/K when the floor is first reached, so within a few minutes the reactor settles onto the temperature at which the two terms balance and then creeps with that temperature as the fouling grows. Warming raises $\mathrm{d}R_f/\mathrm{d}t$ by Equation (11.3), which lowers U, which

moves the balance point faster, so the creep accelerates. The loss of control is abrupt and the loss of product is slow, and the measured consequence in Section 11.6 is that the margin gives about a day of warning before off-grade rubber and the reactor temperature gives sixteen hours once it has moved a degree.

11.5 Controller Design

11.5.1 Velocity Form PID

Every loop uses the shared velocity-form PID of Section 1.6.2 with the derivative on the measurement, written in this chapter's symbols as Equations (11.12) to (11.16).

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{11.12}$$

$$P_k = K_c e_k, \qquad D_k = -K_c \tau_D\, a\, \frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{11.13}$$

$$s_k = s_{k-1} + \big(P_k - P_{k-1}\big) + \big(D_k - D_{k-1}\big)$$

$$c_k = u_{bias} + s_k + I_{k-1} \tag{11.14}$$

$$u_k = \min\big(\max(c_k, u_{min}), u_{max}\big) \tag{11.15}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\, e_k + \frac{u_k - c_k}{\tau_t}\right]\Delta t \tag{11.16}$$

where a is the controller direction, s_k is the never-clipped running sum of proportional and derivative increments (S_k of Equation (1.15)), c_k is the demanded output (v_k of Equation (1.16)), and the integral state I is seeded with the initial output so that $u_{bias} = 0$. Anti-windup is the back-calculation term of Equation (11.16), with tracking time $\tau_t = \tau_I$, so only the integral is corrected at a limit [14]. The derivative acts on the measurement through a 0.05 h filter, and $\tau_D = 0$ on every loop as shipped. In manual, and for any loop not in service, the integral tracks the value applied, so the transfer to automatic is bumpless. The controller scan Δt is 0.02 h (72 s) for all three loops and integration is fixed-step fourth-order Runge-Kutta at 0.005 h.

The low limit of the temperature loop is the refrigeration floor itself, $u_{min} = T_{j,min}$, and the high limit is −80 °C, so the back-calculation of Equation (11.16) holds the integral at the floor. Once the jacket reaches −106 °C the command stays there rather than winding on, and a warmer setpoint or a rate cut acts on the jacket at the next scan. In manual the operator may ask for as low as −120 °C, and the floor is imposed on the applied jacket temperature after the station.

11.5.2 Four Architectures and Two Wash Policies

Off leaves the jacket in manual. *Temperature* closes TIC-201 from the reactor temperature to the jacket temperature with $a = +1$, because the process gain is positive and a reactor above its setpoint calls for a colder jacket, with $K_c = 4.0$ degrees of jacket per degree and $\tau_I = 0.30$ h. *Cascade* adds MC-201 from the two-hour Mooney laboratory result to the reactor temperature setpoint, with $a = -1$ because a warmer reactor makes a lower Mooney, $K_c = 0.30$ °C per Mooney unit, $\tau_I = 8$ h, and the setpoint confined to −102 °C to −86 °C. The laboratory result is held for its two-hour cycle and passed through a 0.05 h filter. The fourth architecture, *MPC*, replaces MC-201 and the conversion loop with a receding-horizon plan and is described in Section 11.5.3.

The outer integral time is long on purpose. A Mooney result up to two hours old, with a laboratory repeatability of 0.4 units (1σ) in scenario 6, drives a setpoint through a plant gain of 1.40 units per degree, so $K_c = 0.30$ °C per unit is a loop gain of 0.42. At that gain a 2 h integral time only makes the setpoint busier, 0.75 °C peak to peak over two days against 0.70 °C, and no cycle appears. Raising K_c to 1.0 °C per unit as well, a loop gain of 1.4, gives a sustained cycle with a 4.0 h period, twice the laboratory cycle, in which the setpoint swings between its −102 °C limit and −86.8 °C and the Mooney between 42.8 and 56.4. The sample age is what sets the outer loop speed, not the tuning

preference [10]. The quality cascade of Chapter 8 and the per-head correction of Volume 8, Chapter 9 also act on results that describe product made some time before.

A conversion loop manipulates the catalyst rate between 0 and 8 kg/h to hold conversion, with $K_c = 12$ kg/h per unit of conversion and $\tau_I = 1.0$ h. Whether it should be in automatic is the subject of scenario 7.

The wash policy is an operator decision, a fixed elapsed time, or a threshold on the cooling margin. The two automatic policies are implemented as a shutdown, an eight-hour wash and a restart in which the catalyst is ramped in over 1.5 h and the quality trim is held for 3.0 h. Whatever the policy, a high-temperature trip on the reactor transmitter at −80 °C kills the reaction and starts the same wash, so a run that nobody stops ends in a trip. The threshold is an illustrative simulation choice, eight degrees above the warm-slurry alarm, rather than a limit derived from a plant safety study. This model cannot establish protection adequacy.

11.5.3 Model Predictive Control

The predictive option treats the cooling margin of Equation (11.6) as a constraint on a prediction rather than as a wash trigger, in the receding-horizon form surveyed by Mayne et al. [15] and set up in general in Section 1.6.3. Every $T_s = 0.1$ h, five controller scans, it moves two handles: the reactor temperature setpoint T_{SP}, which TIC-201 follows exactly as it follows the output of MC-201, and the catalyst command u_c, which is applied as the conversion loop's command is. MC-201 and the conversion PI follow the commands actually applied, scan by scan and without a preload, so a hand-back to the cascade is bumpless. With the conversion loop in manual the catalyst stays with the operator and the plan moves the setpoint alone. TIC-201, the refrigeration floor, the wash policy, the catalyst ramp after a restart and the high-temperature trip remain in force beneath and above the plan, and the plan never decides a wash.

Controller model. The internal model uses only the declared defaults of Table 11.2 and the design feeds, never from the run's parameters, the plant state or the disturbances, Equations (11.17) to (11.22).

$$\tau_T \frac{\mathrm{d}T}{\mathrm{d}t} = T_{SP} + b_T - T \tag{11.17}$$

$$\tau_{MW} \frac{\mathrm{d}M}{\mathrm{d}t} = \mathrm{ML}_{ref} + k_{ML}(T - T_{ref}) - M \tag{11.18}$$

$$\tau_c \frac{\mathrm{d}c}{\mathrm{d}t} = u_c - c, \qquad \tau \frac{\mathrm{d}X}{\mathrm{d}t} = X^*(\hat{a}\,c, T) - X \tag{11.19}$$

$$R_f(t+s) = \hat{R}_f + \hat{r}\,s + z(s) \tag{11.20}$$

$$\frac{\mathrm{d}z}{\mathrm{d}s} = \hat{r}\left(c_f\,\delta T + \frac{\delta X}{\bar{X}}\right) \tag{11.21}$$

$$T_{j,need} = T - \frac{Q(X,T)}{A}\left(\frac{1}{U_0} + R_f\right) \tag{11.22}$$

Here τ_T is the closed-loop lag of TIC-201, b_T its offset filtered over 0.5 h, M the Mooney of the polymer being made, corrected to $\widehat{\mathrm{ML}} = M + b_M$ by the laboratory bias b_M, c the delivered catalyst with the 0.08 h lag τ_c, X^* the conversion law of Equation (11.1) with the activity estimate $\hat{a}$, Q the declared duty of Equation (11.2), $\hat{R}_f$ and $\hat{r}$ the inferred fouling resistance and its trend, δ a deviation from the present point and $\bar{X}$ the present conversion. The model is linearized at the measured point at every decision (the resistance taken at mid-horizon for the slopes) and sampled exactly at T_s, giving six states and four outputs: the Mooney, the margin row, the conversion and the temperature. The margin row

$$T_{j,need} - m\,T \;\geq\; (1-m)\,T_{j,min} \tag{11.23}$$

is Equation (11.6) rearranged into a form linear in the states, with m the margin the plan keeps, 5 % as shipped, and $T_{j,min}$ read as TIC-201 reads it.

The lag τ_T was identified with TIC-201 closed at its shipped tuning, because an open-loop step test removes the loop the plan relies on. Setpoint steps of $\pm 1\,^{\circ}\mathrm{C}$ at the design point with noise off give a gain of 1.00 in both directions and a ringing response, a peak of 1.56 at 0.055 h and a period of 0.12 h, that no low-order model fits to better than a 4 to 9 percent residual. The transmitter's mean delay by the area method is 0.080 h and a first-order fit at the decision instants gives 0.085 h, so $\tau_T = 0.08\,\mathrm{h}$. Through the 0.60 h mixing lag the model predicts −1.095 Mooney units per degree at one hour against a measured −1.099 and −1.342 against −1.346 at two hours, but only −0.094 against −0.215 at 0.1 h, which is the ring the model leaves out. At the design point the conversion law gives $\partial X/\partial c = 0.0265$ per kg/h, the production is $6.0\,\mathrm{t/h} \times X$, and the duty rises by 1.55 MW per unit of conversion.

Fouling as a trend of its own future. The resistance is never read. At every scan the controller computes $1/U = A(T - T_j)/Q(X,T)$ from the transmitter, the jacket and the declared duty at the measured conversion, filters it over 0.1 h, keeps a sample only while the reactor is within 0.1 °C of its setpoint and 0.3 h after any setpoint move of 0.05 °C or more, and regresses the last 4 h for the trend $\hat{r}$. The ratio of $\hat{r}$ to the declared rate law of Equation (11.3) at the measured temperature and production is a tendency $\hat{f}$, bounded between 0.2 and 5, which supplies the trend when less than an hour of samples exists after a start or a wash. Equation (11.20) extrapolates that trend in a straight line over the horizon and adds the first-order effect of the plan's own temperature and production on the rate. Against the simulated truth, sampled every half hour outside washes, the inferred resistance has a mean relative error of 0.1, 0.3 and 0.6 % under the predictive controller on cases 7, 8 and 9, and the trend a mean relative error of 2.2, 2.7 and 4.0 %. The worst tendency errors, 0.56 on case 8 and 0.69 on case 9, fall in the hours after a wash while the window refills. Under the shipped loop the jacket on its floor leaves no quiet samples, and on case 7 the tendency freezes at 0.90 against a true 1.00.

Estimators. The design is offset-free in the sense of Pannocchia and Rawlings [16], with one disturbance estimate per channel. The catalyst activity multiplier comes from a steady-state Kalman filter on the sampled pair $[X, \ln \hat{a}]$ with process covariances 10^{-8} and 2×10^{-5} and measurement covariance 10^{-6} per scan, whose gain is $L = [0.156,\ 4.11]^{\mathsf{T}}$. It updates only while the delivered catalyst exceeds 1 kg/h and the conversion exceeds 0.3, within bounds of 0.3 and 1.6, and in scenario 10 it reads 0.650 two hours after the poisoning to 0.65. The Mooney bias takes half of each laboratory residual, within ±5 units, and the residual is formed against the model's Mooney at the instant the sample was taken, held in a 20 h history, and read 0.2 h after the sample once the laboratory filter has settled. Results during the 3.0 h restart hold are skipped, the rule the shipped trim follows. On the challenge, with a laboratory repeatability of 0.5 units, the bias stays within 0.69 units of zero.

Objective and constraints. The horizon is $N = 80$ steps (8 h) in ten move blocks, their lengths running 1, 1, 2, 3, 5, 8, 10, 10, 20 and 20 steps, so the quadratic program has 20 move variables and three slacks. The objective, Equation (11.24), prices production as a linear cost in the manner of economic MPC [17].

$$
\begin{aligned}
J = \sum_{k=1}^{N} \Big[& - w_P\,(6.0\,\mathrm{t/h})\,X_k T_s + q_M \big(\widehat{\mathrm{ML}}_k - \mathrm{ML}_{SP}\big)^2 \Big] \\
& + \sum_{j} \Big[s_T\,\Delta T_{SP,j}^2 + s_c\,\Delta u_{c,j}^2 \Big] + \sum_{i=1}^{3} \rho_i \big(\varepsilon_i^2 + \varepsilon_i\big)
\end{aligned}
\tag{11.24}
$$

with $w_P = 1$ per tonne, $q_M = 0.005$ per Mooney unit squared, $s_T = 2$ per $^{\circ}\mathrm{C}^2$ and $s_c = 0.5$ per $(\mathrm{kg/h})^2$ at each block boundary. Three outputs are soft, judged at every second step, each with one slack $\varepsilon_i \geq 0$ over the whole horizon: the Mooney within $\mathrm{ML}_{SP} \pm 2.5$ at $\rho_1 = 100$, the margin row of Equation (11.23) at $\rho_2 = 10^4$ per kelvin, and the conversion at or below its target at $\rho_3 = 10^3$, so the plan never runs the reactor above the rate the operator commissioned. Where the present point already violates a limit, the limit re-enters through a funnel with a 1 h time constant. The hard limits

are the ranges the shipped loops clamp to, $T_{SP} \in [-102, -86]\,°\mathrm{C}$ and $u_c \in [0, 8]\,\mathrm{kg/h}$, with rate limits of 0.3 °C and 0.8 kg/h per decision.

Acceptance and fallback. A plan stands only while the solver's status is optimal, every planned command lies inside its range, the plan's first command matches the move being applied, and the cost, the slacks and every published prediction are finite. Otherwise the historian is marked, MC-201 and the conversion loop take the two handles from the commands they have been tracking, and the plan is retried at every scan. Six forced failures injected at hour 20 of scenario 10 with noise off were each rejected, the setpoint moved −0.0175 °C and the catalyst 0.0019 kg/h at the hand-over, and the next plan was accepted one scan later. A plan costs a mean of 0.57 to 0.59 ms in Chrome over the 480 plans of scenario 10, the largest 7.3 ms for the first. That load is 0.6 % of a single core when one simulated hour passes each second and 35 % at the top speed of sixty.

Three models kept apart. The plant is Equations (11.1) to (11.7) with the true activity, fouling tendency and exponential temperature acceleration, the refrigerant lag and the loop's ring. The identified model is the closed-loop lag τ_T from the step tests above. The controller's model is Equations (11.17) to (11.22) at the declared defaults, linearized at each decision, with the inferred fouling trend, the estimated activity and the two biases. The controller reads the transmitter, the jacket temperature, the conversion, the laboratory result, the commands applied and $T_{j,min}$, and nothing else, so a parameter changed by the learner is model mismatch rather than knowledge.

11.6 Results

Eleven scenarios run from watching the coefficient fall by hand to a graded four-day challenge. Table 11.4 lists them with the point each one teaches and what completes it. The app runs at 0.5 to 60 simulated hours per second, 5 h per second by default, so a two-day cycle takes about ten seconds. Reset restarts a case from time zero and keeps the reader's own settings, which is how the comparisons in scenarios 4 to 8 are made (Section 1.6.4).

11.6.1 The Controlled Variable Is Not Where the Information Is

With the temperature loop closed in scenario 2 the reactor stays within 0.010 °C of its setpoint for forty-eight hours. That residual is the steady error of a PI loop following a ramp, the jacket's 0.134 °C per hour divided by K_c/τ_I, and otherwise the loop is working perfectly. Over the same period the jacket walks from −99.4 °C to −105.9 °C, the required temperature difference grows from 4.4 K to 10.9 K, and the cooling margin falls in a straight line, 0.0122 per hour, from 0.598 to 0.012. The margin reaches zero, and the jacket its floor, at hour 49.0. With the loop open in scenario 1 and the jacket held at −99.4 °C, the reactor warms instead, to −83.3 °C at hour 48, and the coefficient falls further, to 600 $\mathrm{W/(m^2\,K)}$, because warmer slurry fouls faster. The Mooney follows at 1.40 units per degree, leaves the grade at hour 22.7 and reaches 35.1 by hour 48.

> **TRY IT IN THE APP 2 · The Temperature Loop, and What It Is Spending**
>
> Load the scenario, press Run and watch the jacket trace walk toward the dashed refrigeration floor while the reactor trace stays flat. Over 48 hours the reactor stays within 0.010 °C of its setpoint, the jacket walks from −99.4 °C to −105.9 °C, and the cooling margin falls from 0.598 to 0.012.

The end of a run is abrupt in the controller and slow in the process. Scenario 3 starts the reactor with the fouling of about 43 hours at the design rate and the jacket at the −105.2 °C that the fouled surface needs, so the reactor holds its setpoint within 0.01 °C until the jacket reaches its floor and the margin is spent, both at 6.3 h. The reactor is 1 °C above its setpoint at 13.8 h, the Mooney leaves the grade at 29.9 h, and the laboratory result confirms it at 30.2 h. The creep accelerates with the fouling, from 0.13 °C per hour just after the floor to 0.29 °C per hour by hour 43, and the reactor passes −88 °C at 43.6 h. The margin trend gives about a day of warning before off-grade rubber, the



Figure 11.2 The butyl rubber application at hour 48 of scenario 3, the end of a run, cropped to the plant view, with the temperature loop in automatic and the reactor started at the fouling resistance of about 43 hours at the design rate. The jacket and the temperature controller output have been at the −106 °C floor since hour 6.3, as (1) the badges show with the margin at zero, while (2) the reactor has crept to −86.6 °C and the Mooney viscosity has fallen to 39.5, with the laboratory result at 40.4. (3) The fouling layer lines the cooled wall, (4) the bar beside the vessel is the temperature difference the duty needs against the floor, and (5) the cooling margin through the run has reached the point marked no control. The run has made 153 t on specification and 18.1 h of off-grade rubber.

reactor temperature gives sixteen hours once it has moved a degree, and the laboratory gives none. Figure 11.2 shows the case at hour 48.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.6.2 The Temperature Optimum Is Warm

Table 11.3 reports run length, Mooney viscosity and yield at seven reactor setpoints, measured in scenario 4. Each setpoint is entered 0.1 h after the scripted step at hour 2, the wash is triggered at a cooling margin of 0.03, and the run, the Mooney and the tonnage are read when the wash starts. The two yields divide the tonnes made, and the tonnes on specification, by the run plus its eight-hour wash.

> **TRY IT IN THE APP** **4 · Temperature Does Three Things, and They Disagree**
>
> Run to the first automatic wash after the step to −99 °C at hour 2: the run lasts 27.0 h and the Mooney ends at 56.6, above the grade. Press Reset, run past the step, set the Reactor temperature setpoint to −93 °C and run to the wash again. The run now lasts 47.6 h, with the Mooney at 48.2.

Colder is shorter, and the arithmetic is not close. The fouling does slow as the reactor gets colder and it is beaten by the collapse of the temperature difference. At −101 °C there are only five degrees between the reactor and the refrigeration floor, the fouling laid down in the first two hours nearly exhausts them, and the wash follows 0.8 h after the setpoint change, before the product has lined out at its Mooney of 59.4.

Table 11.3 Reactor Setpoint Against Run Length, Product and Yield, at 85 % Conversion

Setpoint (°C)	Run (h)	Mooney at wash	Made (t/h)	On spec (t/h)
−101	2.9	57.1	1.37	1.25
−99	27.0	56.6	3.94	0.51
−97	40.6	53.8	4.26	4.26
−95	46.5	51.0	4.35	4.35
−93	47.6	48.2	4.37	4.37
−91	45.8	45.4	4.34	0.34
−89	42.4	42.6	4.29	0.28

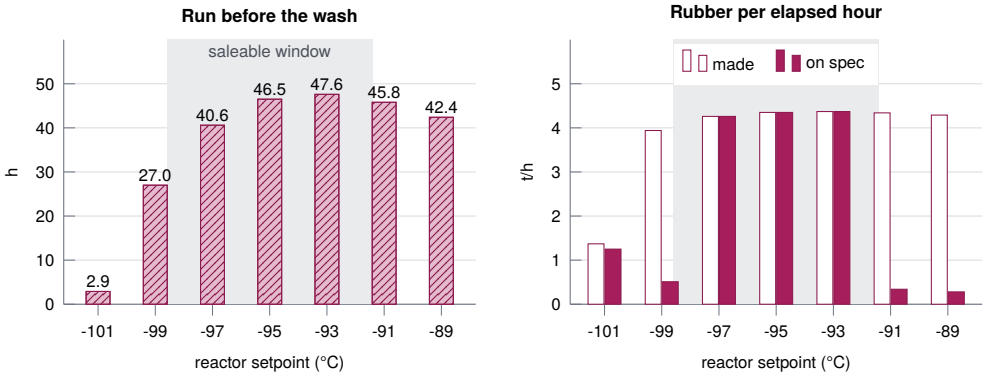


Figure 11.3 The run and yield columns of Table 11.3 against the reactor setpoint. Left, the run before the wash. Right, the tonnes made (open bars) and on specification (solid bars) per elapsed hour, counting the eight-hour wash. The shaded band is the saleable window of −98.6 °C to −91.4 °C. Notice that the run collapses on the cold side, that the tonnes made are nearly flat from −97 °C to −89 °C, and that only the setpoints inside the band keep their rubber on specification.

Pause. Try Exercise 11.2 before reading on.

The yield optimum is near −94 °C and the peak is very flat, 4.35 t/h at −95 °C, 4.37 t/h at −93 °C and 4.34 t/h at −91 °C counting every tonne made. What stops the plant running warmer is not the physics but the grade. At −91 °C the Mooney is 45.4, below its 46 minimum, and only 18.3 t of the 233.7 t made is on specification. The grade binds on the cold side as well, where the Mooney of 56.6 at −99 °C is above the 56 maximum. The saleable window is −98.6 °C to −91.4 °C, so the economic optimum is a constrained one and the binding constraint is a number on a contract. Figure 11.3 draws the run and yield columns of Table 11.3 against the setpoint.

11.6.3 The Case for a Constraint Override

Scenario 7 starts the reactor with the fouling of about 38 hours at the design rate and poisons the catalyst to an activity of 0.65 at hour 2, with the conversion loop in automatic. The loop does exactly what it was asked to do and raises the catalyst from 4.82 kg/h to 7.2 kg/h by hour 12 to recover the production rate. Recovering the production recovers the heat load, which on a reactor with a nearly saturated jacket is the resource the run is spending. The jacket reaches its floor at 11.6 h against 16.2 h with the loop in manual, and the Mooney leaves grade at 35.0 h against 43.7 h. That is the standard argument for a constraint override, and nothing in a conventional cascade knows about it. The protective override of Volume 4, Chapter 2 is the same idea where the limit is reached in seconds.

On this reactor the arithmetic goes the other way. Manual makes 4.72 t/h of polymer instead of 5.07 t/h, and because the end of a run creeps rather than runs away, the hours that automatic loses cost

less than the tonnes that manual loses. At hour 36 the yield is 4.92 t/h in automatic against 4.74 t/h in manual, and stopping each run when its rubber leaves grade, with an eight-hour wash, gives 4.12 t/h against 4.00 t/h. An override that gave way to the cooling margin has to be justified by that arithmetic, and it pays only where the end of a run is abrupt or the wash is long.

11.6.4 A Schedule Is a Controller With No Feedback

Scenario 8 runs a fixed wash at hour 44 with the feed fouling 1.6 times faster than the schedule was written for. The margin is gone at hour 31.3, so the jacket sits on its floor for the last 12.7 hours of the first run. The reactor creeps rather than runs away, that run stays in grade, and over four days the Mooney is out of grade for 0.9 h, at the end of the second run. The schedule yields 4.54 t/h. A wash on a margin of 0.10 makes no off-grade rubber and yields 3.99 t/h, because it washes 26.2 hours into every run and pays for three washes in four days. Lower thresholds of 0.05 and 0.02 both give 4.09 t/h. The schedule is lucky rather than right. Moved to hour 60 it yields 3.87 t/h with 13.6 h off grade, because the rubber of the first run leaves the grade at hour 46.4.

The general form transfers to other plants. A schedule is right on the day it is set and wrong every day the plant is different, and the cure is to find the measurement the schedule was standing in for and use that instead [1]. On this reactor that measurement is not the cooling margin, which runs out about fifteen hours before the product leaves grade, but the quantity that ends the run, the Mooney viscosity or the reactor temperature that sets it.

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

11.6.5 The Graded Challenge

The challenge runs ninety-six hours with the fouling already 1.35 times the design rate, a loss of catalyst activity to 0.78 at hour 18, a rate increase to 6.8 t/h at hour 40 and a dirtier feed at 2.1 times the design fouling at hour 66. It starts with cautious tuning ($K_c = 1.0$, $\tau_I = 1.2$ h on the temperature loop), the conversion loop in manual and the wash left to the operator, and it is judged by the criterion in the last row of Table 11.4.

As shipped, with no wash chosen, the jacket reaches its floor at 38.6 h, the reactor passes −88 °C at 58.2 h, and once the feed fouls at 2.1 times the design rate the creep of the floor turns into a runaway that the −80 °C trip ends at 68.5 h with an eight-hour wash. The catalyst station keeps its 4.82 kg/h through that wash, so the restart at hour 76.5 adds 100 t of rubber within grade and the yield is 3.65 t/h, which the 15.00 °C excursion still fails. The suggested configuration, the Mooney cascade with the tuned loops, the conversion loop in automatic and a wash on a margin of 0.10, passes with a worst error of 0.83 °C, all of it at start-ups, a lined-out error of 0.21 °C, and a yield of 3.83 t/h from three washes and 367 t on specification. The cascade settles the setpoint near −93.5 °C, on the flat top of Table 11.3. Holding −97 °C with the temperature loop alone makes the same 367 t with a worst error of 3.4 °C. With the conversion loop switched to manual at hour 66 the same configuration still passes, at 3.86 t/h and a worst error of 1.07 °C, and with the conversion loop in manual from the start it passes at 3.90 t/h from two washes instead of three. The same configuration on the 44-hour clock yields 4.09 t/h but lets the reactor run 22 °C above its setpoint, with 11.5 h off grade, until the trip at hour 93.9. Passing needs the margin wash and the tighter temperature tuning, without which a lined-out run sits 1.1 °C off setpoint, but neither the conversion loop in automatic nor the cascade: the temperature loop alone at −95 °C passes at 1.46 °C.

11.6.6 The Predictive Controller Against the Shipped Structures

Table 11.5 runs cases 7, 8 and 9 and the challenge twice, once per structure, from identical seeds, durations, instruments and noise. The shipped structure keeps its pinned tuning, with the suggested configuration on the challenge, and the predictive controller runs at the settings of scenario 10 on every case with no retuning. The application has no preview, so neither structure receives one. Three further rows are measured beside them. The best fixed setting of the shipped structure is the highest yield with no off-grade hour found in hindsight over the conversion target (0.60 to 0.90) under the

Table 11.4 Scenarios and Completion Requirements

#	Scenario	Kind	Teaching point, and what is required to complete
1	Watch the coefficient fall	Demonstration	*A gain that dies.* With the jacket held in manual at –99.4 °C, run until the Mooney leaves the 46 to 56 grade (hour 22.7), then on to hour 48 with nothing changed.
2	The temperature loop	Experiment	*The output is the diagnostic.* With the temperature loop in automatic, run until the cooling margin falls below 0.10 (hour 40.8), then on to hour 48.
3	The end of a run	Demonstration	*Saturation is a discontinuity.* From the fouling of about 43 hours with the temperature loop in automatic, run until the jacket reaches its floor after the first hour, the reactor rises 1 °C above setpoint, and the laboratory Mooney falls below 46 (hour 30.2).
4	Temperature does three things	Experiment	*A constrained optimum.* Run to the first margin-triggered wash after the step to –99 °C at hour 2, then Reset, set –93 °C after the step and run to a wash more than 40 hours into the run.
5	Rate against run length	Experiment	*Fixed downtime arithmetic.* Run to the first wash after the isobutylene push to 7.2 t/h at hour 2, then Reset, raise isoprene to 0.14 t/h to 0.15 t/h after the push and run to a wash with more than 100 t on specification.
6	The Mooney cascade	Experiment	*Sample age sets the speed.* Run 48 hours on the cascade with $\tau_{I,M} \geq 6$ h, then Reset, set $\tau_{I,M} \leq 3$ h before the activity drop at hour 3, and run 48 hours again.
7	The loop that helps then hurts	Experiment	*The arithmetic of an override.* Run 12 hours with the conversion loop in automatic, then Reset, put it in manual before the poisoning at hour 2, and run 12 hours again.
8	The clock against the margin	Experiment	*Which measurement ends a run.* Run 96 hours with the wash on a 44-hour clock, then Reset, switch to the margin policy before the dirtier feed at hour 2, and run 96 hours again.
9	A dirtier feed	Exploration	*Early warning.* Run past the fouling step to 2.2 times the design rate at hour 1 until the margin policy at 0.08 starts a wash (hour 19.8), and confirm that the Mooney, the reactor temperature and the conversion stayed on target.
10	The margin as a constraint	Experiment	*What a margin costs.* From case 7's fouled reactor and poisoning, run 48 hours under the MPC at a 5 % margin, then Reset under the temperature loop alone and run 48 hours, then return to the MPC at 2 %, Reset and run 48 hours again.
11	Challenge	Challenge	*Operating, not only controlling.* Over 96 hours, hold the reactor within ±1.5 °C of setpoint from hour 63.4, start-ups included, and within ±0.3 °C once a run is more than four hours old, and finish with a yield above 3.4 t/h, with the model parameters and disturbances as given.[a]

[a] With the shipped configuration (temperature loop at $K_c = 1.0$ and $\tau_I = 1.2$ h, conversion loop in manual, no wash chosen) the attempt fails by a wide margin: the reactor runs away after the jacket reaches its floor until the high-temperature trip shuts it down at hour 68.5, and the reported worst error is 15.00 °C. The yield, 3.65 t/h, clears the target only because the restart at hour 76.5 resumes at the operator's catalyst rate. With the suggested configuration (Mooney cascade, conversion loop in automatic, wash on a margin of 0.10, tuning of Section 11.5) it passes with a worst error of 0.83 °C, 0.67 °C inside the band, a lined-out error of 0.21 °C, 0.09 °C inside, and a yield of 3.83 t/h, 0.43 t/h above the target.

case's own wash policy and over a grid of margin washes. The margin override on case 7 is a classical structure built in the measurement script only: a PI loop on the app's cooling margin with a setpoint of 5 %, low-selected against the conversion loop, with the reactor setpoint raised after the first step to –93.21 °C, the warm edge of the plan's Mooney band. The last row of each case removes the fouling trend from the plan by holding the resistance at its present value over the horizon.

Scenario 10 is case 7's reactor, started at about 38 hours of fouling with the cooling margin at 13.4 % and poisoned at hour 2, with the predictive controller in charge, the conversion loop in automatic and the wash left to the operator. The plan recovers the conversion to its 0.85 target with the catalyst at 7.59 kg/h by hour 4.3, then buys margin cheaply first by walking the reactor up its Mooney band, reaching –93.25 °C at hour 27.0. From about hour 13, when the margin has fallen to 5 %, it brings the catalyst down, 7.49, 4.01, 2.27 and 1.55 kg/h at hours 12, 24, 36 and 48, and the margin then stays between 4.79 and 5.03 %, as shown in Figure 11.4. The jacket never reaches its

Table 11.5 The Predictive Controller Against the Shipped Structures, Measured on the Engine

Case	Structure	Yield (t/h)	First wash (h)	Washes	Floor (h)	Off grade (h)	Mooney IAE	Catalyst travel (kg/h)	Plans, fallbacks
7, 48 h	temperature loop, conversion loop automatic	3.689	none	0	36.36	13.01	144.4	3.11	–
	MPC	4.412	none	0	0	0	85.6	8.90	480, 0
	best fixed: conversion 0.75	4.484	none	0	27.64	0	59.2	1.53	–
	margin override at 5 %, warm setpoint	4.319	none	0	0	0	118.6	8.70	–
	MPC without the fouling trend	4.485	none	0	0	0	83.4	9.06	480, 0
8, 96 h	cascade, wash on a 44-hour clock	4.545	44.01	1	26.24	0.95	104.2	10.35	–
	MPC	4.473	44.01	1	0	0	58.9	13.65	880, 0
	best fixed: conversion 0.84	4.542	44.01	1	24.82	0	97.0	9.73	–
	MPC without the fouling trend	4.529	44.01	1	0	0	50.4	13.71	880, 0
9, 48 h	cascade, wash on a margin of 0.08	4.085	19.83	2	0	0	34.1	9.93	–
	MPC	4.065	21.78	1	0	0	8.2	10.64	400, 0
	best fixed: conversion 0.90, wash at 0.03	4.206	19.25	2	0	0	33.8	15.23	–
	MPC without the fouling trend	4.067	19.83	1	0	0	3.0	10.54	401, 0
Challenge, 96 h	suggested cascade, wash on a margin of 0.10	3.827	30.28	3	0	0	74.6	38.40	–
	MPC with the suggested regulatory settings	4.139	33.04	3	0	0	24.5	31.51	765, 0
	best fixed: conversion 0.85, wash at 0.03	4.143	33.18	3	0.02	0	80.3	26.89	–
	MPC without the fouling trend	3.805	30.30	3	0	0	22.0	37.39	721, 0

Yield is tonnes on specification per elapsed hour including washes. Washes counts the washes started. Floor is the hours the jacket spends at −106 °C outside washes. Mooney IAE accumulates the absolute error of the true Mooney against its target outside washes, and catalyst travel sums the absolute changes of the catalyst command, both by the scorecard's definitions. A best fixed setting that allows off-grade rubber yields more on two cases, 4.517 at a conversion target of 0.76 on case 7 (0.28 h off grade) and 4.548 at 0.87 on case 8 (2.77 h off grade).

floor, the reactor never leaves its setpoint by more than 0.02 °C, the Mooney never falls below 48.51, and all 211.8 t are on specification, a yield of 4.412 t/h. The loop of case 7 makes 243.3 t but only 177.1 t on specification, 3.689 t/h, after 36.36 h on the floor with the reactor up to −88.20 °C. The comparison has a crossover. The yield at hours 12, 24, 36 and 48 is 5.08, 4.96, 4.71 and 4.41 t/h under the plan and 5.00, 5.04, 4.92 and 3.69 t/h under the loop, so the loop leads from some point between hours 12 and 24 until after its rubber leaves grade. The margin the plan keeps is a price: at 2 % the plan yields 4.530 t/h and at 8 % 4.284 t/h, with no hour on the floor and no off-grade hour at either.

Most of the case 7 win belongs to the margin row, not to the plan. The classical override at the same 5 % with the warm setpoint yields 4.319 t/h, so it carries 0.63 of the 0.72 t/h gained, and the prediction and coordination of the two handles add 0.09 t/h. At a 2 % margin the two yield 4.423 and 4.530 t/h. Neither ingredient works alone: the override at −95 °C yields 3.970 t/h, and the warm setpoint without an override yields 3.266 t/h, less than the shipped loop, because its floor comes

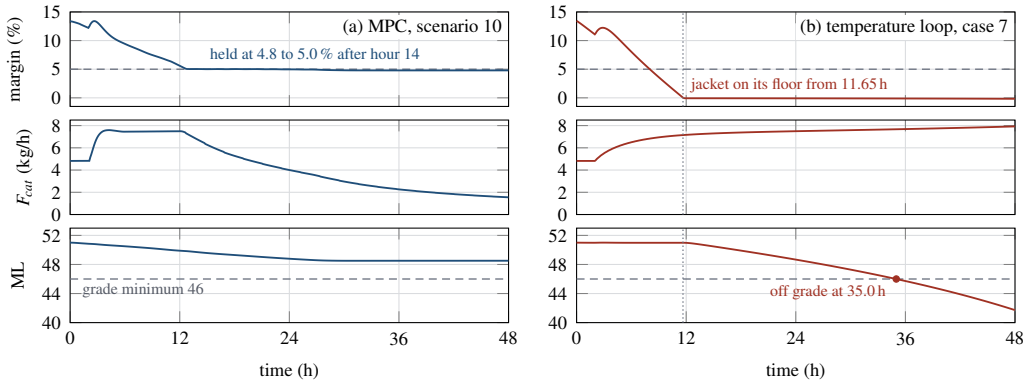


Figure 11.4 The same fouled reactor, poisoned to an activity of 0.65 at hour 2, over two days under (a) the predictive controller of scenario 10 and (b) the temperature loop with the conversion loop in automatic, case 7. Top, the cooling margin with the 5 % the plan keeps. Middle, the catalyst command. Bottom, the Mooney viscosity of the polymer being made against the grade minimum. The plan cuts catalyst from about hour 13 and the jacket never reaches its floor, while the loop holds production until the jacket stops at hour 11.65 and its rubber leaves grade at hour 35.0.

later but its reactor then creeps further. A fixed conversion target of 0.75 with no wash beats the plan at 4.484 t/h, all of it on specification, by sitting 27.64 h on the refrigeration floor and ending the two days with a Mooney of 46.31, a third of a unit inside the grade. Counted over a fixed 48 hours, rubber made after the loop's product leaves grade weighs against the loop, which is why these margin structures pay here while the per-run arithmetic of scenario 7 favors holding production.

The predictive controller loses cases 8 and 9, and the mechanism is the one scenario 8 already teaches. The end of a run on this reactor creeps, so the tonnes the plan gives up to protect the margin are worth more than the grade they protect. On case 8 the shipped cascade spends 26.24 h on its floor before the wash at hour 44 and loses only 0.95 h of grade, while the plan has cut the catalyst to 1.75 kg/h by hour 44 and yields 4.473 t/h against 4.545 t/h. On case 9 the margin wash at 0.08 ends the run in either structure, and the plan's cut only stretches the first run from 19.83 to 21.78 h at a lower rate, 4.065 t/h against 4.085 t/h. The best fixed setting beats the plan on both, 4.542 and 4.206 t/h. On the challenge the plan yields 4.139 t/h against the suggested cascade's 3.827 t/h, 397.3 t against 367.4 t, because its cut delays the first margin wash from hour 30.28 to 33.04, and the best fixed setting of the cascade in hindsight ties it at 4.143 t/h. On the variant of case 8 with the wash moved to hour 60, which scenario 8 itself measures as a mistake, the plan yields 4.248 t/h against 3.872 t/h with 13.61 h off grade.

The fouling trend is what keeps the margin the operator set. Without it the plan holds 3.38 and 2.59 % at its lowest on cases 7 and 8 instead of 5 % and makes more, 4.485 and 4.529 t/h, but on the challenge it yields 3.805 t/h, below the suggested cascade, because it no longer stretches the first run, which reaches its wash at hour 30.30. In no case does the predictive controller reach the floor or make off-grade rubber, and no plan was rejected in 2525 solves across the four cases.

DID YOU KNOW? A Tighter Grade, Less Rubber

From the simulator: on case 9 of Table 11.5 the predictive controller holds the Mooney to an IAE of 8.2, against 34.1 for the cascade with a margin wash. Yet it yields 4.065 t/h against 4.085 t/h. Holding the grade more tightly did not make more saleable rubber.

11.7 Technology in Practice

What Measures It

Two measurements carry the chapter's loops. TT-201 reads the reactor temperature for TIC-201 on every 72 s scan, and in the setpoint steps of Section 11.5.3 its mean delay is 0.080 h. The thermocouple and the resistance thermometer that real reactors use are described in Chapter 10 and Chapter 4. The quality loop MC-201 reads the Mooney viscosity, the measurement that decides whether a run's rubber can be sold, and the box below describes how that number is made.

TECHNOLOGY Sensor: The Mooney Viscometer

Mooney viscosity is the resistance of the uncured rubber to flow. The standard test, ISO 289 or ASTM D1646, preheats a sample for one minute, then turns a large rotor in it at 2 revolutions per minute for four minutes and records the torque. At a test temperature of 100 °C the result is reported as ML(1+4) 100 °C [18, 19]. It is a laboratory test, not an online instrument. MC-201 of the chapter reads its result as a staircase, each step held for the two-hour cycle and describing rubber made earlier, with a repeatability of 0.4 units in scenario 6.

What Moves It

TIC-201 moves the jacket temperature, and the conversion loop moves the catalyst rate between 0 and 8 kg/h. The first is the handle the fouling consumes. In a real reactor it is the temperature of a refrigerant boiling against the cooled surface, as the box below explains.

TECHNOLOGY Final Element: Boiling Refrigerant Against the Wall

The butyl reactor of a 1946 patent was cooled by liquid ethylene, circulated around a central draft tube and its return tubes or held in a refrigerating jacket [8]. A boiling liquid holds its temperature at its boiling point until it has boiled away, and that boiling point rises and falls with the pressure on the liquid [20]. The refrigerant side of the wall can therefore be held at a steady temperature set by pressure, and what changes over a run is the polymer layer on the slurry side. TIC-201 of the chapter moves no valve directly: its output is the jacket temperature itself, reached through a 0.035 h refrigerant lag and stopped at the −106 °C floor that stands for the coldest the refrigeration can deliver.

Where the Controller Runs

Figure 11.5 draws the loops of Section 11.5 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 [21]. In the app the three PI blocks run in velocity form every 72 s, MC-201 writes the setpoint of TIC-201 as a cascade, and the laboratory result enters MC-201 held for its two-hour cycle. The MPC of Section 11.5.3 replaces MC-201 and the conversion loop with a plan every 0.1 h, five scans, and the wash policy and the −80 °C trip stay in force above every layer. The lag of TIC-201 that the plan uses was identified with the loop closed at its shipped tuning, and the box below describes how industry tests a plant for such a controller.

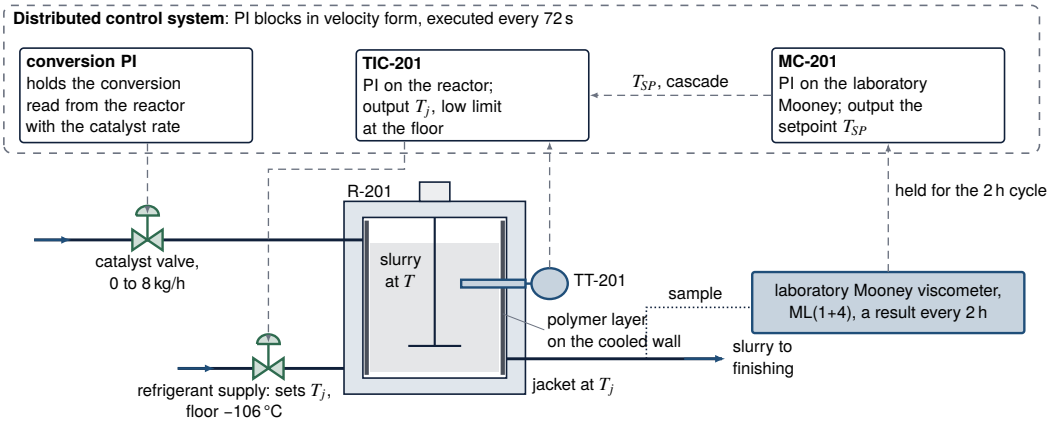


Figure 11.5 The loops of R-201 with their hardware. Blue marks are sensing devices: the temperature element and transmitter TT-201 and the laboratory Mooney viscometer. Green symbols are final elements: the refrigerant supply that sets the jacket temperature and the catalyst valve. Dashed arrows are signals. Notice that the only measurement of product quality reaches MC-201 from a laboratory, held for two hours, and that MC-201 acts only by moving the setpoint of TIC-201, whose output walks toward the −106 °C floor as the wall fouls. The feeds, the trip and the MPC are left out. Not to scale, and not a piping diagram.

> **IN PRACTICE** **Commissioning a Predictive Controller**
>
> An industrial predictive controller gets its model from a test on the plant. A pre-test steps each manipulated variable to adjust the instruments and PID controllers beneath it and to find each output's settling time. The plant test runs around the clock with engineers watching. Each handle is stepped 8 to 15 times, with an output signal-to-noise ratio of at least six, and the test can take 5 to 15 days. No PID configuration or tuning may change during the test, and operators may step in to avoid critical situations but make no correlated moves. If the PID tuning later changes much, the model may need rebuilding [11]. The chapter's closed-loop step tests on TIC-201 follow the same rule.

What Is Modernizing It

Three developments bear on this reactor. Soft sensors trained on data from the distributed control system and the laboratory estimate Mooney viscosity between laboratory results and can cut how often the laboratory must measure it [22]. In-line fouling sensors, applied in dairy plants, detect the deposit as it forms and help trigger cleaning on the actual buildup rather than on fixed schedules [23]. A 2003 survey counted polymers among the areas with the most nonlinear predictive control applications [11]. The cooling margin of this chapter is an inferred measurement of the deposit, and the MPC writes it into its plan as a constraint rather than a wash trigger. How far each development can be trusted is the subject of the Open challenge boxes in Section 11.8.

11.8 Limitations

The reactor is a single well mixed volume with a uniform fouling layer. Real fouling is not uniform, and the practical consequence is that a real plant loses circulation in one region before the average coefficient predicts it.

The fouling model is linear in time with an exponential temperature acceleration and no induction period, no asymptote and no removal term. That shape makes the countdown clean and it is simpler than the fouling literature supports.

Molecular weight is a correlation on reactor temperature rather than a calculation from chain transfer kinetics, so the distribution breadth, the gel content and the branching that a real butyl plant manages are all absent [3].

> **OPEN CHALLENGE** **Mooney Viscosity Without the Laboratory**
>
> In industrial rubber mixing, no reliable online instrument measures Mooney viscosity. The laboratory reports it every 4 to 6 hours, against batches that take 2 to 5 minutes, so an abnormal batch is found long after it was made. Soft sensors trained on data from the distributed control system could fill the gap, but laboratory results are too costly and slow to give them enough labeled data. Semi-supervised methods that also learn from the unlabeled process data are one answer under study [22]. Progress would be an estimate that a quality loop like MC-201 could close on between samples, and that stays right as the plant drifts.

The refrigeration floor is fixed. A real system loses capacity in summer and as its own surfaces foul, which means the countdown on a real plant has a second term this model does not carry.

The wash is a fixed eight hours restoring the surface exactly. A real hydroblast varies, and the fraction of the original coefficient recovered is one of the numbers a plant argues about.

> **OPEN CHALLENGE** **Washing on the Deposit, Not the Clock**
>
> Fouling of heated and cooled surfaces is unsolved well beyond rubber. A 2025 review of fouling in dairy heat treatment finds a notable gap in evaluating in-line measurement and prediction of fouling, in the laboratory and in plants. Computational flow models are still far from simulating, optimizing or scheduling whole heat exchangers over complete cleaning cycles, and the learned models that predict fouling well are hard to interpret, which slows their industrial use [23]. Progress would be a deposit sensor or model trusted enough to trigger a wash on the actual buildup rather than on a fixed schedule, the decision that scenario 8 shows the cooling margin cannot make alone.

Finally, the shipped structures deliberately do not include the multivariable layer that a real plant uses to arbitrate between production rate, quality and cooling margin [11]. Scenario 7 poses the question an override has to answer, and on this reactor the measured arithmetic favors holding production. No classical override is implemented, and the predictive option of Section 11.5.3 carries the margin as a constraint instead.

The predictive controller's fouling model is a trend of its own future: a straight line at the trend inferred over the last four hours, with only the first-order effect of the plan's temperature and production on the rate. It leaves out the exponential acceleration of Equation (11.3) along the plan, it freezes while the jacket is on its floor and refills after a wash, and it infers the coefficient from the duty at the design feeds, so a changed feed lands in the coefficient. The margin settles at 4.8 % rather than the 5 % asked for, 4.79 on scenario 10 and 4.78 on case 8, because each soft output carries one slack over the whole horizon and the funnel re-grants the present shortfall at every decision. The margin prediction is conservative while the plan cuts catalyst: a plan applied open loop for eight hours realizes up to 3.06 points more margin than it predicted, while its Mooney is within 0.04 units and its temperature within 0.16 °C. The plan does not decide the wash, and it loses cases 8 and 9 to the shipped structure.

11.9 Discussion

Three results generalize past butyl rubber.

The first is that a monotonically drifting controller output is a measurement. On this reactor the controlled variable is flat for two days while the plant runs out of cooling, and the only trend that says so is the jacket temperature. Trending controller outputs and alarming on their rate of change is cheap, and it is the difference between planning a shutdown and being surprised by one.

The second is that a fixed limit turns a monotone trade into an optimum. Colder gives a higher molecular weight and slower fouling, and the run length collapses anyway because the distance to the refrigeration floor is what the run is spending. Any time a decision variable improves two things at once, it is worth asking what fixed limit it is also consuming.

The third is that a schedule is an open-loop controller, and that replacing it with feedback helps only when the feedback measures what ends the run. The fixed wash at hour 44 is right for one fouling rate and lucky or wrong for every other. The cooling margin is already measured, but on this reactor it runs out about fifteen hours before the product leaves grade, so a wash triggered on it comes early and gives up 0.55 t/h against the schedule. Replacing a schedule with the measurement of the binding constraint it approximates is among the cheapest improvements available on a cyclic plant.

11.10 Conclusion

A butyl rubber reactor insulates itself with the product it is making. The simulated reactor of this chapter holds −95 °C to within 0.01 °C for two days while its overall coefficient falls from 2400 to 964 W/(m^2 K), its jacket walks from −99.4 °C to its −106 °C floor, and its cooling margin falls from 0.598 to zero at hour 49, so the run ends at an actuator stop rather than at a measurement. Colder operation is better for the molecular weight and for the fouling rate and worse for the plant: measured yields of on-specification rubber per hour of elapsed time are 1.25 t/h at −101 °C and 4.37 t/h at −93 °C, with the optimum near −94 °C and the binding constraint the Mooney grade of 46 to 56 rather than anything in the equipment. A wash scheduled at a fixed hour is correct only for the fouling rate it was written for, and a wash triggered on the cooling margin is not the cure on this reactor, because the margin runs out long before the product leaves grade. A fuller model would add non-uniform fouling, a refrigeration system with capacity of its own, and a molecular weight distribution model. A model predictive controller that keeps a predicted cooling margin of 5 % never reaches the floor or makes off-grade rubber in any case measured, yet it wins on yield only on the poisoned reactor and the challenge, loses the scheduled wash and the dirtier feed by at most 0.08 t/h, and owes most of its poisoned-reactor win to the margin row that a classical override also carries. The reactor is a place to learn to read a plant from its controller outputs.

KEY POINTS

- A butyl rubber reactor insulates itself with its own product, so its run ends when the jacket reaches its floor.
- The controlled variable stays put and the manipulated variable is the diagnostic: the cooling margin is the clock.
- Colder operation shortens the run, because the floor does not move with the setpoint. The optimum is near −94 °C, bounded by the Mooney grade.
- A fixed wash schedule suits only the fouling rate it was written for, and a margin wash comes about fifteen hours early.
- A predictive controller keeping a 5 % margin never reaches the floor, yet wins on yield only on the poisoned reactor and the challenge.

Where else this applies. A run spent against a limit is also the coke countdown of Chapter 2 and the fouling preheat train of Volume 11, Chapter 3.

References

[1] J. Tian, Y. Wang, and X. Feng, "Simultaneous optimization of flow velocity and cleaning schedule for mitigating fouling in refinery heat exchanger networks," *Energy*, vol. 109, pp. 1118–1129, 2016.

[2] R. M. Thomas, I. E. Lightbown, W. J. Sparks, P. K. Frolich, and E. V. Murphree, "Butyl rubber. a new hydrocarbon product," *Rubber Chemistry and Technology*, vol. 14, no. 1, pp. 177–195, 1941.

[3] J. L. White, T. D. Shaffer, C. J. Ruff, and J. P. Cross, "Incorporation of isoprene in isobutylene/isoprene copolymers: Nmr identification of branching in butyl rubber," *Macromolecules*, vol. 28, no. 9, pp. 3290–3300, 1995.

[4] R. M. Thomas, “Early history of butyl rubber. charles goodyear medal address—1969,” *Rubber Chemistry and Technology*, vol. 42, no. 4, pp. G90–G96, 1969.

[5] O. Sofronova, E. Markina, S. Chelnokova, and A. Sakhabutdinov, “The effect of alkyl halides on the stability of a butyl rubber suspension in methyl chloride,” *International Polymer Science and Technology*, vol. 40, no. 3, pp. 11–14, 2013.

[6] R. M. Thomas and W. J. Sparks, “Mixed olefinic polymerization process and product,” U.S. Patent 2,356,128, filed 20 October 1939 (continuation in part of an application filed 29 December 1937), granted 22 August 1944, 1944.

[7] I. E. Nifant’ev, S. A. Korchagina, M. S. Chinova, and A. N. Tavtorkin, “Polyisobutylenes with controlled molecular weight and chain-end structure: Synthesis and actual applications,” *Polymers*, vol. 15, no. 16, p. 3415, 2023.

[8] L. R. Gould and G. G. Wanless, “Method for washing butyl reactors with isobutylene,” U.S. Patent 2,580,019, filed 9 October 1946, granted 25 December 1951, 1951.

[9] M. G. Balaton, L. Nagy, and F. Szeifert, “Model-based split-range algorithm for the temperature control of a batch reactor,” *Engineering*, vol. 04, no. 09, pp. 515–525, 2012.

[10] P. Kadlec, B. Gabrys, and S. Strandt, “Data-driven soft sensors in the process industry,” *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[11] 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.

[12] 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.

[13] A. Kummer and T. Varga, “What do we know already about reactor runaway? a review,” *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

[14] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[15] 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.

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

[17] M. Ellis, H. Durand, and P. D. Christofides, “A tutorial review of economic model predictive control methods,” *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[18] A. Aggarwal, F. Grunert, S. Ilisch, and A. Blume, “A quantitative analysis of influencing parameters for cure torque in silica-filled SSBR compounds,” *Polymers*, vol. 18, no. 20, p. 2596, 2026.

[19] H. Yangthong, P. Nun-Anan, A. Krainoi, B. Chaisrikhwun, S. Karrila, and S. Limhengha, “Hybrid alumina–silica filler for thermally conductive epoxidized natural rubber,” *Polymers*, vol. 16, no. 23, p. 3362, 2024.

[20] 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.

[21] 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.

[22] Y. Zhang, H. Jin, H. Liu, B. Yang, and S. Dong, “Deep semi-supervised just-in-time learning based soft sensor for Mooney viscosity estimation in industrial rubber mixing process,” *Polymers*, vol. 14, no. 5, p. 1018, 2022.

[23] J. Tarapata, T. R. Murphy, E. W. Finnegan, T. F. O’Callaghan, and J. A. O’Mahony, “Approaches for measuring and predicting fouling during thermal processing of dairy solutions,” *Comprehensive Reviews in Food Science and Food Safety*, vol. 24, no. 4, p. e70209, 2025.

11.11 Exercises

11.1 ○ Concept check · 2 min

In scenario 2 the temperature loop holds the reactor within 0.01 °C of −95 °C for two days while the surface fouls. Which reading tells the operator how many hours of the run are left?

(a) The jacket temperature against the −106 °C floor, or the cooling margin computed from it

(b) The reactor temperature deviation, which grows in step with the fouling resistance

(c) The laboratory Mooney result, because the molecular weight follows the heat transfer coefficient

(d) The integrated absolute error on the scorecard, which accumulates the fouling over the run

11.2 Concept check · 1 min

Fouling slows by half for every five degrees of cooling, yet in Table 11.3 the run at −99 °C is much shorter than the run at −95 °C. Why?

(a) Colder polymerization is faster, so the reaction duty rises enough to shorten the run

(b) The run spends the distance to a fixed refrigeration floor, and cooling takes that distance away faster than it slows the fouling

(c) The Mooney viscosity rises above its grade limit and the wash policy stops the run early

(d) The temperature loop saturates at its high limit and stops correcting the jacket temperature

11.3 Concept check · 2 min

In scenario 8 the feed fouls 1.6 times faster than the 44-hour wash schedule assumed. A wash on a cooling margin of 0.10 yields 3.99 t/h against the schedule's 4.54 t/h. What is the reason?

(a) The margin is computed from noisy measurements, so it triggers washes at random times

(b) The schedule washes later and makes no off-grade rubber, while the margin policy makes some

(c) The margin runs out well before the rubber leaves grade, so the margin policy washes early and pays for an extra wash

(d) The schedule happens to wash at the moment the margin reaches zero, which is the optimum

11.4 Theory and calculation · 6 min

Reactor R-201 is fed 6.0 t/h of isobutylene, 18 t/h of methyl chloride and 0.12 t/h of isoprene at $T_{feed} = -40\,^{\circ}\mathrm{C}$, and holds $T = -95\,^{\circ}\mathrm{C}$ at a conversion of 0.85. Use $\Delta H_p = 930\,\mathrm{kJ/kg}$ of polymer, $c_p = 1.9\,\mathrm{kJ/(kg\,K)}$, $A = 190\,\mathrm{m^2}$, $U_0 = 2400\,\mathrm{W/(m^2\,K)}$ and $T_{j,min} = -106\,^{\circ}\mathrm{C}$.

(a) Compute the heat of reaction, the feed-chilling duty and the total duty Q of Equation (11.2), in MW to three decimals. *Checkpoint:* 1.318 MW, 0.700 MW and 2.018 MW.

(b) Compute the temperature difference the duty requires on a clean surface, the jacket temperature of Equation (11.5) and the cooling margin of Equation (11.6). *Checkpoint:* 4.42 K, −99.42 °C and 0.598.

(c) Compute the overall coefficient at which the margin reaches zero with the same duty, and compare it with the 964 W/(m^2 K) the chapter reports when the jacket reaches its floor. *Checkpoint:* about 965 W/(m^2 K).

11.5 Theory and calculation · 7 min

Use Equation (11.9) with the values of the previous exercise, $k_f = 1.29 \times 10^{-5}\,\mathrm{m^2\,K/(W\,h)}$, $c_f = 0.1386/^{\circ}\mathrm{C}$, $P = 5.10\,\mathrm{t/h}$, $P_0 = 5.2\,\mathrm{t/h}$, $\phi = 1$ and $T_{ref} = -95\,^{\circ}\mathrm{C}$. The run starts from a clean surface and the reactor is held at one temperature throughout.

(a) Compute the prefactor $AP_0/(Qk_fP\phi)$ and the run length at −95 °C. *Checkpoint:* 7.44 h per kelvin and 48.9 h.

(b) Repeat at −93 °C and at −101 °C, with the duty recomputed for the feed chilled to each temperature. *Checkpoint:* 1.992 MW and 49.3 h, then 2.094 MW and 6.7 h.

(c) Holding Q at its design value, find the temperature that maximizes t_{run}, and check it against the Mooney grade window of 46 to 56, with ML = 51.0 at −95 °C and $k_{ML} = -1.40$ per degree. *Checkpoint:* −94.4 °C, inside the window −98.6 to −91.4 °C.

11.6 Theory and calculation · 6 min

The reactor holds 30 m^3 of slurry at 700 kg/m^3 with $c_p = 1.9\,\mathrm{kJ/(kg\,K)}$. When the jacket reaches its floor the coefficient is 964 W/(m^2 K) on $A = 190\,m^2$, and 24.12 t/h of feed is chilled from −40 °C.

(a) Compute the heat capacity of the contents and the rate at which a sustained shortfall of 0.1 MW would warm them. *Checkpoint:* 39.9 MJ/K and 9.0 K per hour.

(b) Each degree of warming raises the duty the floor removes by UA and lowers the feed-chilling duty. Find how far the reactor warms before a 0.1 MW shortfall is absorbed, and the time constant of that approach. *Checkpoint:* 0.51 K and about 3.4 min.

(c) Before the floor, TIC-201 with $K_c = 4.0$ and $\tau_I = 0.30\,h$ moves the jacket down a ramp of 0.134 K per hour. Compute the steady error of the PI loop on that ramp. *Checkpoint:* 0.010 °C.

11.7 App challenge · 8 min

Load scenario 2, The temperature loop, and what it is spending, and change no setting. Run a little past hour 48 at speed 20 with the default seed and export the CSV file. The time column `time_h` is in hours and a row is logged every 0.035 h.

(a) Read column `margin` in the first row and in the last row at or before hour 48 (at 47.985 h), and compute its mean slope per hour. *Checkpoint:* 0.598 and 0.012, a slope of −0.0122 per hour within 0.0003.

(b) Find the first row in which `margin` is below 0.10, and extrapolate the straight line of part (a) to zero. *Checkpoint:* hour 40.8 within 0.1 h, and about hour 49.0.

(c) Find the largest absolute difference between columns `temp_C` and `temp_sp_C`, and explain its size. *Checkpoint:* 0.010 °C within 0.002 °C.

11.8 App challenge · 9 min

Load scenario 4, Temperature does three things, and they disagree. A script steps the reactor setpoint from −95 to −99 °C at hour 2, and the wash starts when the cooling margin falls below 0.03. Use speed 5 and the default seed, and export the CSV file after each run.

(a) Change no setting and run until the wash has started. In the last CSV row before column `washing` becomes 1, read `run_hours`, `metric_onspec_t` and `mooney`. Compute the tonnes on specification per elapsed hour, counting the run and its eight-hour wash. *Checkpoint:* 27.0 h, 17.8 t, a Mooney of 56.6, and 0.51 t/h within 0.02.

(b) Press Reset. Run past hour 2 until the step to −99 °C has appeared on the trend, pause before hour 3, set Reactor temperature setpoint to −93.0 °C, and run until the wash has started. Read the same three columns and compute the same yield. *Checkpoint:* a run of 47.6 to 48.2 h, 242.7 to 245.6 t, a Mooney of 48.2, and 4.37 t/h within 0.01.

Explain why the colder run is shorter and why almost none of its rubber is on specification.

11.9 App challenge · 8 min

Load scenario 8, The clock against the margin. The fouling tendency steps to 1.6 at hour 2, and the wash is on the clock at 44 h. Use speed 20 and the default seed. Run each case a little past hour 96, export the CSV file, and read it in the last row at or before hour 96 (the row at 95.97 h). The yield is the tonnes on specification divided by the 96 elapsed hours.

(a) Change no setting. Read columns `metric_onspec_t`, `metric_offspec_h` and `downtime_h`, compute the yield, and find the first hour at which column `margin` is zero or below. *Checkpoint:* 436.3 t, 0.9 h, 8.0 h, 4.54 t/h within 0.01, and hour 31.3 to 31.4.

(b) Press Reset, set When to wash to On margin before hour 2, keeping the threshold at 0.10, and run again. Read the same three columns, compute the yield, and find the hours at which column

`washing` turns to 1. *Checkpoint:* 383.4 t, 0.0 h, 17.8 h, 3.99 t/h within 0.01, washes at about 26.2, 60.2 and 94.2 h.

State which measurement the schedule was standing in for.

11.10 Application area · 3 min

The Limitations section notes that a real refrigeration system loses capacity in summer. Suppose the floor rises from −106 to −105 °C while everything else stays at the design point of the earlier exercises.

(a) Use the prefactor of 7.44 h per kelvin to find the new run length. *Checkpoint:* 41.5 h, 7.4 h shorter.

(b) Use Equation (11.10) with the eight-hour wash to compare the yield before and after, for runs that stay in grade. *Checkpoint:* 4.38 against 4.28 t/h.

11.11 Application area · 3 min

Section 11.2 argues that this reactor is not a runaway problem, and Section 11.5 calls the −80 °C trip an illustrative threshold eight degrees above the warm-slurry alarm at −88 °C.

(a) State what does not accelerate with temperature on this reactor, and what does. *Checkpoint:* the heat release does not, the fouling does, doubling for every five degrees.

(b) In scenario 3 the creep reaches 0.29 °C per hour. At that rate, how long would the reactor take from the alarm to the trip? On the challenge as shipped, the reactor passes −88 °C at 58.2 h and trips at 68.5 h. What does the difference say about warning time? *Checkpoint:* about 28 h by extrapolation against 10.3 h measured.

11.12 Review the agent · 5 min

A learner asks an AI assistant to make the Mooney cascade of scenario 6, The Mooney cascade, respond faster to the two-hour laboratory result. The assistant replies: “The result is two hours old, so the outer loop must act harder on each one. Raise Mooney loop gain from 0.30 to 1.0 and cut Mooney integral time from 8 h to 2 h. The Mooney will then be corrected within one laboratory cycle.” Find the flaw, and say how you would check the proposal in the app before using it.