INTERACTIVE CONTROL ENGINEERING
Principles and Applications

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

CHAPTER 6

Slowest Part, Hottest Core: Lead/Lag Control of a Composite Autoclave Cure

Pages 141–169

Edited by John D. Hedengren

How to cite this chapter. "Slowest Part, Hottest Core: Lead/Lag Control of a Composite Autoclave Cure," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 141–169.

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

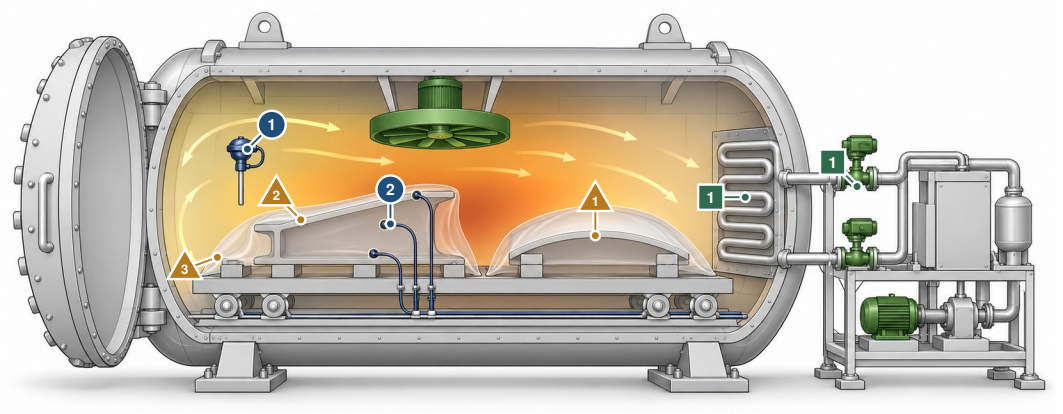
The app for this chapter
https://apmonitor.com/apps/autoclave-cure/

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

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

CHAPTER 6

Slowest Part, Hottest Core: Lead/Lag Control of a Composite Autoclave Cure



The controller moves 1 heater and cooling coil, driven by one split signal

The controller reads 1 air temperature 2 thermocouples at the bag side, middle and tool face of each part

What disturbs it 1 a part in the lee of another, in weaker air flow 2 a thicker part that heats itself as it cures 3 a torn vacuum bag

The thin parts follow the air while the thick part lags behind it and can overheat at its core, so the coldest thermocouple sets the pace and the hottest sets the limit. Illustration, not to scale.

IN PLAIN TERMS One Hot Oven, Many Different Parts

The main idea. Parts made of carbon fiber and resin are baked in a big pressurized oven until the resin sets hard. A thin part warms up in minutes, but a thick part on a heavy mold needs more than an hour.

The challenge. The schedule must wait for the slowest part, and the hottest part must not get too hot. The setting resin makes its own heat, so the center of a thick part can overheat before any sensor notices.

Why it matters. Too fast a cycle can spoil a thick part, and too slow a cycle wastes hours on every load.

THREE WAYS TO READ THIS CHAPTER

Explore	8 min	the first page, the history in Section 6.2, the first fact in Section 6.3, the second fact in Section 6.4, the try-it in Section 6.7, and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 6.7), Technology in Practice (Section 6.8), Limitations (Section 6.9) and the open challenge

The control question An autoclave cures parts of different thickness, on tools of different mass, with one actuator, the air temperature. How can the coldest of twelve thermocouples set the ramp while the hottest one, and a core that no thermocouple reaches, set the limit?

Learning objectives

- Compute the film, tool and conduction time constants of a part and name the one that sets its lag behind the air.
- Explain how a low selector with external reset runs the recipe on the lagging thermocouple while a max loop holds the hottest one at its limit.
- Explain why a limit loop on thermocouples cannot stop the exotherm of a thick part, and how a feedforward from a cure observer acts first.
- Relate the cure target to the glass transition and end a hold on the observer's estimate instead of a clock.
- Compare the cycle time and margins of lead/lag, the feedforward, an optimized recipe and a shrinking-horizon MPC.

Prerequisites Sections 1.4.2 (recipes and guards), 1.4.4 (path and endpoint constraints) and 1.6.2 (the PI of the air loop), and Chapter 5 for a batch whose hottest and slowest units differ. Section 1.6.3 is needed only for the MPC.

The app https://apmonitor.com/apps/autoclave-cure/

See also Volume 4, Chapter 3 for min and max selection over limiters whose integrators track the applied command, and Volume 4, Chapter 2 for a protective override that rests on its clamp until it is needed. Volume 9, Chapter 4 for offset-free MPC with an estimated disturbance.

6.1 Introduction

An autoclave is a heated pressure vessel in which composite parts, laid up from prepreg and sealed under vacuum bags on their tools, are taken through a recipe of ramps and holds. In the load simulated here a 3 mm skin on a thin aluminum plate, a 25 mm spar on a 450 kg invar tool, a 10 mm frame on steel and a 6 mm part in the lee of the spar share one vessel at 7 bar gauge. With the air stepped to 110 °C, the air reads 105 °C at 4.2 min and the coldest thermocouple of the skin gets there 10.1 min later, while the shadowed part needs 57.8 min and the spar 82.4 min. The standard recipe for the resin, run on the air's clock, cures every part but takes 438.3 min, and 37.5 min pass between the air reaching the hold band and the spar's tool face doing so.

A block diagram from air temperature to part temperature hides the two facts that decide this process. The first is that the controlled variable is not one measurement but the slowest and the fastest of many, so a recipe written for the part has to be run on the part that lags and limited by the part that leads. The second is that the product is a reaction. Thick thermoset sections release their heat of cure faster than conduction removes it, so the center overshoots the tool and the air, a behavior that through-thickness cure models describe [1, 2] and that cure cycles with cooling and reheating steps were later designed to suppress [3, 4]. In the simulator a 35 mm spar ramped at 3 °C/min passes 195 °C inside the laminate at 101.8 min, when its hottest thermocouple reads 187.3 °C, and peaks at 243.9 °C, although the limit loop on the hottest thermocouple has already pulled the air down. The misconception this chapter corrects is that a limit loop on the measured temperatures protects the part. It acts after the measurement moves, and a self-heating core moves first. The semi-batch reactor of Chapter 2 teaches the same lesson from the other side: there the protection is the dosing that limits the heat still to be released, not the temperature loop.

The class is broader than composites. Any batch whose many units share one heating medium, with a lagging unit that sets the schedule, a leading unit that sets the limit and a heat release the schedule can trigger, has the same shape, among them pharmaceutical shelf loads, kilns, heat-treatment furnaces and fed reactors. In the freeze-dryer of Chapter 5 the hottest vials likewise limit the shelf while the slowest decide the endpoint. Selector and override structures for such constraints, and the

windup that follows when a loop is passed over, are a subject in control theory [5], and model-based cure control of the autoclave has its own literature [6–10].

This chapter works with a four-part autoclave load rebuilt as a browser simulation, which runs the same mixed load on the air's clock and on the lagging part's clock, shows a core passing the limit between its thermocouples, and lets the recipe be shortened under the same constraints. The true temperature and degree of cure through every part are logged beside the twelve thermocouples a controller may read. On this load the chapter measures what a lead/lag supervisor, an exotherm feedforward from a cure observer and an offline recipe search do to cycle time and margins. The kinetics, the glass-transition relation and the viscosity law are published forms [11–13], the parameters come from a public report [14], lead/lag and override control are established practice, and model-based and predictive cure control have been reported [7–9]. The measured results belong to this model. Among them, the feedforward buys margin rather than minutes, and a slab bound that it acts on overestimates the core's rise while the reaction runs and underestimates it after. On the implementation side, the model's energy ledgers close to within 10^{-11} of the heater's energy over a cycle, its operator-split integrator is checked against an independent explicit solution, and the equations and constants given here suffice to rebuild it. A model predictive controller (MPC) that plans the air over the remaining recipe, a shrinking horizon like that of a 1994 miniature-autoclave study [8], is compared with the shipped, retuned and prediction-fed loops, with the safety defect of its first build and the fix's price.

Section 6.2 traces the domain from the cure kinetics to model-based cure control. Section 6.3 sets out the load as modeled, Section 6.4 explains what sets the lags of the parts and the exotherm limit, and Section 6.5 describes the control layers, with the MPC in Section 6.5.6. Section 6.6 separates verification of the code from agreement with the source of the resin data. Section 6.7 measures the nine cases and the challenge, each with a completion rule (Section 1.6.4), and the MPC in Section 6.7.6. Section 6.8 describes the hardware of a real autoclave, Section 6.9 states the limitations and Section 6.10 concludes.

6.2 Autoclaves, Cure Models and Cure Control

The thermoset chemistry came first. In 1973 an autocatalytic rate law was proposed from isothermal and scanning calorimetry of a polyester, in good agreement with measured rates and integral heats [11]. Its form, a rate proportional to $(k_1 + k_2\alpha^m)(1-\alpha)^n$, is the one extended to two terms in the parameter source used here [14]. A 1982 model of mold filling and curing for reaction injection molding predicted the pressure rise, extent of reaction and temperature in instrumented molds [13], and the viscosity law with a divergence at the gel point used here carries the names of its authors. In 1990 the DiBenedetto equation was re-evaluated as a convenient constitutive relation between the glass transition and the extent of reaction of a thermoset [12].

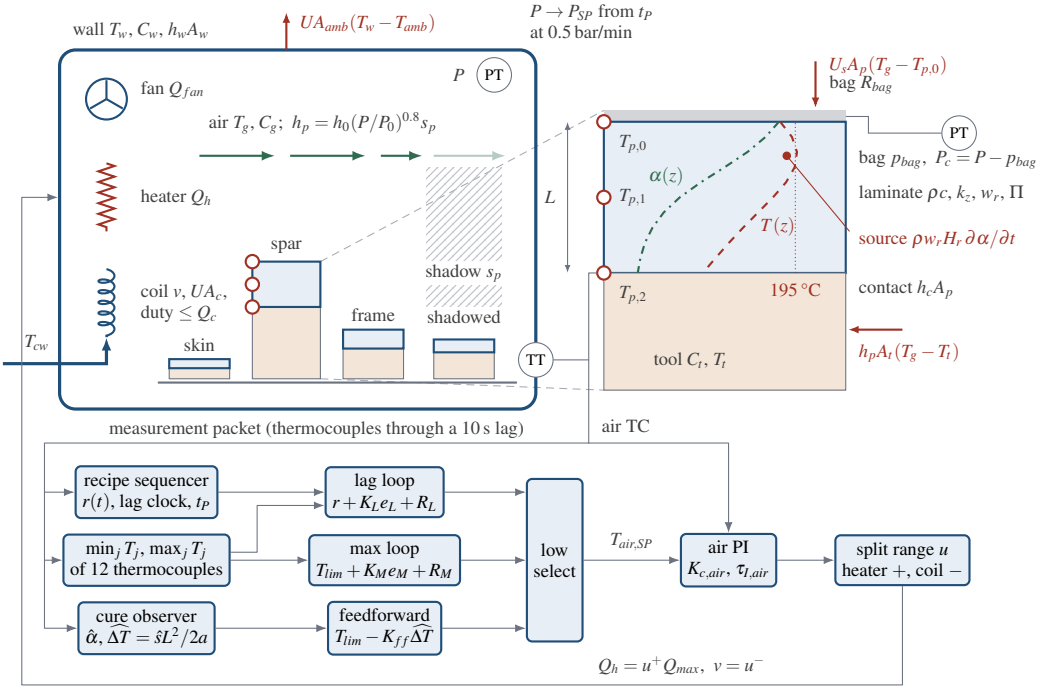
The process model followed for prepregs. In 1983 a computer code for flat-plate laminates gave, for a specified cure cycle, the temperature distribution, the degree of cure, the viscosity, the voids and the residual stress through the part [1]. Thick sections came next. A two-dimensional simulation of 1991 coupled transient anisotropic conduction to the cure kinetics with a heat generation term, predicted complex gradients in temperature and degree of cure, and showed the influence of the tool [2]. The consequence for recipes was drawn in 1997, when a cycle with cooling and reheating steps, timed by the cure rate and temperature at the midpoint, reduced the overshoot of 15 and 30 mm laminates [3], and in 2002, when measurements on a 20 mm laminate showed substantial lag and overshoot at its center and the conclusion that cycles recommended by prepreg manufacturers for thin laminates should be modified for thick ones [4].

Control of the vessel moved from the clock toward the part. A rule-based expert system for real-time control of autoclave temperature and pressure was published as a book in 1990 [6]. In 1994 model-based optimal temperature profiles were carried out through a knowledge-based system [7], and shrinking-horizon model predictive control was applied to the curing of laminates in a miniature

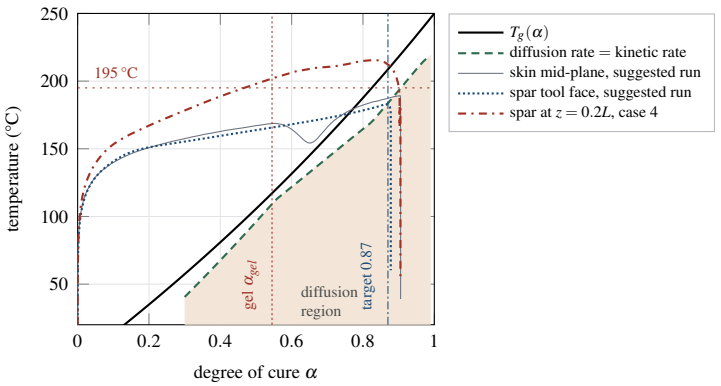
Table 6.1 Variables by Role, With Their Ranges and Initial Values

Variable	Role	Initial	Range
$T_{lag} = \min_j T_{p,j}$	CV of the lag loop (TC)	20 °C	—
$T_{max} = \max_j T_{p,j}$	CV of the max loop (TC)	20 °C	limit $\leq$ 190 °C
$T_{g,m}$	CV of the air loop (air TC)	20 °C	—
$T_{air,SP}$	MV of the supervisor	20 °C	20–225 °C
u	MV, split range	0	−100 to 100 %
P_{SP}	MV, pressure setpoint	7 bar g	0–7 bar g
s_p	DV, shadow factor	0.45	0.1–1
L_{spar}	DV, spar thickness	25 mm	10–50 mm
α_0	DV, resin lot	0	0 or 0.12 (aged)
leak, heater	DV, events	—	bag torn, half heater
T_g, T_w	states, air and wall	20 °C	—
$T_{p,i}$, $T_{t,p}$	states, 55 nodes, 5 tools	20 °C	—
$\alpha_{p,i}$, $\Pi_{p,i}$	states, cure and voids	0, 0	0–1, %
P, p_{bag}	states, vessel and bag	1.013 bar, 0.005 bar	absolute
five ledgers	states, energies	0	—

autoclave, treated as a batch reactor [8]. Time-optimal cure cycles computed by sequential quadratic programming improved on the manufacturers' recommended cycles in 1997 [15], and a 1999 study posed the shortest operating time as a constrained optimization and tracked the result with internal model control [9]. Cure monitoring in commercial carbon/epoxy composites during the autoclave cycle [16] and dynamic optimization for cure control [10] continued the line. The heat path to the parts was studied as well, and computational fluid dynamics of an autoclave run agreed best with measured part temperatures only when thermal radiation and a measured inlet velocity profile were included [17]. The lead/lag arrangement taught here is not the subject of these papers, and the simulator builds it from override control with external reset [5] on the thermocouples the parts carry.



(a)



(b)

Figure 6.1 The autoclave load as modeled and its cure map. (a) The vessel with heater, coil, fan and wall around the air node, four parts on their tools with the shadow behind the spar, and the spar enlarged through its thickness L with its three thermocouples $T_{p,j}$, the bag, the tool, the reaction source and its true temperature $T(z)$ (dashed) and degree of cure $\alpha(z)$ (dash-dotted) at 101.8 min of case 4 (teaching truth), with the 195 °C limit dotted. Below are the control layers as implemented, all fed by the measurement packet. Process model only, not to scale. (b) Temperature against degree of cure from the application's kinetics, with the glass transition (solid black), the line where the diffusion rate equals the kinetic rate (dashed, bounding the shaded diffusion region), gelation and the cure target (thin vertical lines), and three true trajectories: the skin's mid-plane (thin gray), the spar's tool face in the suggested run (dotted) and the spar at $z = 0.2L$ in case 4 (dash-dotted).

> ⧖ HISTORY **1985: The Cure Cycle Becomes a Design Variable**
>
> In 1985 a progress report to the Langley Research Center of NASA took up the optimal design of autoclave cure cycles for thick laminates. It coupled heat conduction through the laminate with the chemical kinetics of the prepreg resin, solved the two together by finite elements, and treated the surface temperature, the one the autoclave can control, as the design variable. Its goal was a temperature uniform across the thickness. A simulation of a 192-ply laminate agreed closely with the center temperature of an earlier report, but one complete cycle of 275 minutes in the autoclave took 17,056 seconds of computer time, because the equations were stiff [18].

6.3 The Load as Modeled

The vessel is 3 m long and 2 m across, as shown in Figure 6.1a. A 300 kW heater bank, a cooling coil on 15 °C water and a fan act on one well-mixed air node, a wall of 3 MJ/K exchanges heat with the air and loses it through the insulation to a 20 °C room, and four parts sit on a cart. Each part is a carbon/epoxy laminate resolved through its thickness L_p on 11 nodes, with its bag side ($z = 0$) facing the gas through the bag and breather and its far side on a lumped tool. Three thermocouples per part sit at the bag side, mid-thickness and the tool face, and an air thermocouple, a pressure transmitter and a vacuum transducer complete the instruments. An optional fifth part, a 15 mm fitting on 200 kg of steel, can be loaded before a run. The roles of the variables are listed in Table 6.1.

6.3.1 Air, Wall and Utilities

The air and the wall exchange heat with the utilities, the parts and the room, as shown in Equation (6.1).

$$\begin{aligned} C_g \frac{\mathrm{d}T_g}{\mathrm{d}t} &= Q_h - Q_c + Q_{fan} - h_w A_w (T_g - T_w) \\ &\quad - \textstyle\sum_p U_{s,p} A_p (T_g - T_{p,0}) \\ &\quad - \textstyle\sum_p h_p A_{t,p} (T_g - T_{t,p}) \\ C_w \frac{\mathrm{d}T_w}{\mathrm{d}t} &= h_w A_w (T_g - T_w) - U A_{amb} (T_w - T_{amb}) \\ Q_h &= f_h \, u^+ Q_{max}, \quad v = u^-, \quad Q_{fan} = Q_{fan,0} \, P / P_0 \\ Q_c &= v \, G_c \, (T_g - T_{cw}) \\ G_c &= \min \left(U A_c, \; Q_{c,max} / \max(1\,\mathrm{K}, \, T_g - T_{cw}) \right) \end{aligned} \tag{6.1}$$

where $u^+ = \max(0, u)$ and $u^- = \max(0, -u)$ split one signal between the heater and the coil's valve, f_h is 1 or 0.5 after a heater-bank failure, P is the absolute vessel pressure and $P_0 = 1.013\,25$ bar. The coil conductance G_c is fixed at its value at the start of each step, so the coil passes $U A_c (T_g - T_{cw})$ up to its rated duty of 300 kW. Without the cap the 6 kW/K coil removes 990 kW at 180 °C, three times the heater's authority, and the air loop reversed between heating and cooling 38 times in the cool-down of the suggested challenge run, against twice with it. Every part's coefficient follows the gas density at a fixed fan speed, as shown in Equation (6.2).

$$\begin{aligned} h_p &= h_0 \left(\frac{P}{P_0} \right)^{0.8} s_p, \qquad h_w = \frac{h_p}{s_p} \\ U_{s,p} &= \frac{1}{1/h_p + R_{bag}} \end{aligned} \tag{6.2}$$

At 7 bar gauge h is 73.2 W/(m^2 K), 5.23 times its value at one atmosphere, so a vessel that is not yet pressurized also starves the parts of heat. The shadowed part's factor of 0.45 leaves it 32.9 W/(m^2 K).

6.3.2 Parts Through Their Thickness

Each laminate conducts in z only and carries the heat of its own cure, as shown in Equation (6.3).

$$\begin{aligned}
\rho c\,\frac{\partial T}{\partial t} &= k_z\frac{\partial^2 T}{\partial z^2} + \rho\, w_r H_r \frac{\partial \alpha}{\partial t}, \qquad 0 < z < L_p \\
-k_z\frac{\partial T}{\partial z}\Big|_0 &= U_{s,p}\big(T_g - T(0)\big) \\
-k_z\frac{\partial T}{\partial z}\Big|_{L_p} &= h_c\big(T(L_p) - T_{t,p}\big) \\
C_{t,p}\frac{\mathrm{d}T_{t,p}}{\mathrm{d}t} &= h_c A_p\big(T(L_p) - T_{t,p}\big) \\
&\quad + h_p A_{t,p}(T_g - T_{t,p})
\end{aligned} \tag{6.3}$$

where $T(z)$ and $\alpha(z)$ belong to part p, discretized on 11 nodes with half cells at both faces, and $T_{p,j}$ in Equation (6.1) is the node of thermocouple j (nodes 0, 5 and 10). The source term heats a fully cured laminate adiabatically by $w_r H_r/c = 185.5\,\mathrm{K}$. Three time scales follow from the declared values. The bag-side film needs $\rho c L/U_s$, 1.5 min for the skin and 12.8 min for the spar. The tool needs $C_t/(h_p A_t)$, 29.3 min for the spar's invar and 27.3 min for the shadowed part's aluminum in its lee. Conduction across the laminate needs L^2/a with $a = k_z/\rho c = 3.45\times10^{-7}\,\mathrm{m^2/s}$, 0.4 min for the skin, 30.2 min for the 25 mm spar and 59.1 min at 35 mm. The thin parts therefore follow their surface and the spar follows its tool and its own conduction, both about half an hour.

6.3.3 Resin Kinetics, Glass Transition and Viscosity

The cure rate is the sum of two autocatalytic terms of the Kamal and Sourour type [11] placed in series with a diffusion rate, the form fitted in the public material model of the resin [14], as shown in Equation (6.4).

$$\begin{aligned}
\dot\alpha_k &= A_1 \mathrm{e}^{-E_1/RT}(1-\alpha)^{n_1}(\alpha+\alpha_{01})^{m_1} \\
&\quad + A_2 \mathrm{e}^{-E_2/RT}(1-\alpha)^{n_2}\alpha^{m_2} \\
\dot\alpha_d &= K_d\,\mathrm{e}^{-B/f}, \qquad f = a_d\big(T - T_g(\alpha)\big) + b(T_g) \\
\frac{\partial\alpha}{\partial t} &= \Big(\frac{1}{\dot\alpha_k} + \frac{1}{\dot\alpha_d}\Big)^{-1}
\end{aligned} \tag{6.4}$$

where T is in kelvin, b rises linearly from 0.021 to 0.031 as T_g goes from 120 to 195 °C and is constant outside, and f is floored at 10^{-30}. The glass transition follows the DiBenedetto relation in the form re-evaluated by Pascault and Williams [12], and the viscosity a Castro–Macosko form [13] with two Arrhenius terms, as shown in Equation (6.5).

$$\begin{aligned}
\frac{T_g - T_{g0}}{T_{g\infty} - T_{g0}} &= \frac{\lambda\alpha}{1-(1-\lambda)\alpha} \\
\eta &= \eta_{01}\mathrm{e}^{E_{\eta1}/RT} \\
&\quad + \eta_{02}\mathrm{e}^{E_{\eta2}/RT}\Big(\frac{\alpha_{gel}}{\alpha_{gel}-\alpha}\Big)^{A_\eta}, \quad \alpha < \alpha_{gel}
\end{aligned} \tag{6.5}$$

The diffusion term is what makes this resin unlike a plain autocatalytic model, as shown in Figure 6.1b. Held at 180 °C from $\alpha = 0$ the resin reaches 0.80 in 32.42 min and 0.87 in 56.50 min, the diffusion rate equals the kinetic rate at $\alpha = 0.858$, where T_g is 205.0 °C, and at 0.90 the series rate is $2.90\times10^{-7}\,\mathrm{s^{-1}}$ against a kinetic $3.38\times10^{-5}\,\mathrm{s^{-1}}$. After 6 h the cure stands at 0.8992. A target of $\alpha \geq 0.95$ could never be met at that hold, so the application's cure target is $\alpha \geq 0.87$, which DiBenedetto's relation maps to $T_g \geq 208.7\,^\circ\mathrm{C}$, and the demold rule $T_g \geq 208\,^\circ\mathrm{C}$ is the same constraint seen through it.

🕮 **Pause.** Try Exercise 6.3 before reading on.

6.3.4 Pressure, Vacuum and the Porosity Index

The compressor moves the vessel toward its setpoint at no more than 0.5 bar/min from the pressure step's time t_P and vents at demold, and the bag's gas relaxes toward a leak-weighted share of the vessel's, as shown in Equation (6.6).

$$\begin{aligned} P_{k+1} &= P_k + \mathrm{clip}\big(P_0 + P_{cmd} - P_k,\ \pm r_P \Delta t\big) \\ p_{bag,k+1} &= \bar{p} + (p_{bag,k} - \bar{p})\,\mathrm{e}^{-\Delta t/\tau_{bag}}, \quad \bar{p} = \frac{\lambda_b P_{k+1}}{1+\lambda_b} \\ \frac{\mathrm{d}\Pi}{\mathrm{d}t} &= \begin{cases} r_v(1 - P_c/P_{min}) & \eta < \eta_{flow},\ P_c < P_{min} \\ -\Pi/\tau_v & \eta < \eta_{flow},\ P_c \geq P_{min} \\ 0 & \text{otherwise, or } \alpha \geq \alpha_{gel} \end{cases} \end{aligned} \tag{6.6}$$

where P_{cmd} is zero before t_P and at demold and P_{SP} between, $P_c = P - p_{bag}$ is the consolidation pressure, λ_b is 0.005 for a sound bag and 3 after a tear, and $\tau_{bag} = 60$ s. The porosity index Π of each node grows while the resin can flow and the consolidation is under 4 bar, collapses in 10 min once it is over, and freezes at gel. It is a teaching index that rewards pressure before gelation and nothing else, not a model of void growth.

6.3.5 Instruments

The controller reads a measurement packet and nothing else. Each part thermocouple reads its node through an exact first-order lag of 10 s with colored noise of 0.2 °C standard deviation and a 30 s correlation time, the air thermocouple likewise, and the packet also carries the coldest and hottest of the part thermocouples with their positions, the vessel pressure in bar gauge and the bag pressure. Every noise draw comes from one stream seeded by the case, one normal per instrument per step, so Reset replays a run exactly (Section 1.6.1). Truth columns in the comma-separated (CSV) export carry a `truth_` prefix. A grep-level test fails whenever the controller, the sequencer or the observer reads a truth field.

6.3.6 Parameters and Their Evidence

Table 6.2 lists every parameter. The 26 resin constants are those of the public material model [14], compared one by one with the report by the application's technical review and found identical. The report characterizes the resin sold as Hexcel 8552, whose name appears here only as the public name of that source. The laminate, the vessel, the parts and the limits are declared teaching values for a 3 m vessel and a typical carbon/epoxy prepreg. The standard recipe, 1 °C/min to 110 °C for 60 min and 1 °C/min to 180 °C for 120 min, cooled at 2 °C/min to 60 °C, follows the shape of the cycle the report quotes from the manufacturer, 30 to 60 min at 225 °F and 120 min at 350 °F, rounded to 110 and 180 °C. Its ramps of 1 °C/min are slower than the quoted 3 to 5 °F/min, a conservative choice of the simulated shop.

6.3.7 Integration, Initialization and Ledgers

The thin nodes of the skin, 0.3 mm apart, are far too stiff for an explicit method at a useful step, so the application integrates with its own operator split, one of the apps' own steps that Section 1.6.1 mentions. Each physics step of $\Delta t = 5$ s takes half a step of cure at every node, advancing α and its adiabatic heating together by fourth-order Runge-Kutta, then the linear heat network over the whole step by the L-stable second-order TR-BDF2 method, whose two stages share one matrix solved exactly by a tridiagonal sweep per part and a Schur complement on the air node, then the second half step of cure (a Strang split). The pressure and the bag move first, the porosity index last. Five ledgers integrate the heater, fan, coil, room and reaction energies with the same quadrature as the temperatures, so the stored heat closes against them to round-off. The controller runs every 10 s, and a log row is written at the same interval. Every case starts with the whole load, the air and the wall at

Table 6.2 Model Parameters and Their Evidence. S: Read in the Material Model Report [14], T: Teaching Value, D: Derived. A Marks Values Adjustable in the Parameters Panel

Symbol	Meaning	Value	Unit	Ev.
Resin and porosity				
A_1, E_1	first kinetic term	153 900.5, 64 929.5	s^{-1}, J/mol	S, A(E_1)
n_1, m_1, α_{01}	exponents, offset	2.347, 1.413, 0.1594	—	S
A_2, E_2	second kinetic term	3.963×10^{11}, 133 168.3	s^{-1}, J/mol	S
n_2, m_2	exponents	1.029, 5.586	—	S
K_d, B, a_d	diffusion term	4.0, 0.21, 4.8×10^{-4}	s^{-1}, –, K^{-1}	S
b	diffusion offset, T_g 120–195 °C	0.021–0.031	—	S
H_r	heat of reaction	600	J/g resin	S, A
T_{g0}, $T_{g\infty}$, λ	DiBenedetto	−7, 250, 0.78	°C, °C, –	S, A($T_{g\infty}$)
η_{01}, $E_{\eta 1}$	first viscosity term	7.5×10^{-11}, 81 908	Pa s, J/mol	S
η_{02}, $E_{\eta 2}$	second viscosity term	0.0481, 13 228	Pa s, J/mol	S
α_{gel}, A_η	gel point, exponent	0.545, 2.466	—	S, A(α_{gel})
q	adiabatic rise, $w_r H_r / c$	185.5	K	D
η_{flow}, P_{min}	porosity flow, pressure	1000, 4	Pa s, bar	T
r_v, τ_v	void growth, collapse	0.10, 10	%/min, min	T
Vessel, laminate, parts, instruments and limits				
C_g, C_w	air with heaters and fan, wall	0.25, 3.0	MJ/K	T, A(C_w)
A_w, UA_{amb}, T_{amb}	wall area, loss, room	25, 150, 20	m^2, W/K, °C	T
Q_{max}, $Q_{fan,0}$	heater rating, fan heat at 1 atm	300, 3	kW	T, A(Q_{max})
UA_c, $Q_{c,max}$, T_{cw}	coil, rated duty, water	6, 300, 15	kW/K, kW, °C	T
h_0, R_{bag}	coefficient at 1 atm, bag	14, 0.004	W/(m^2 K), m^2 K/W	T, A
h_c	part to tool contact	400	W/(m^2 K)	T
ρ, c, k_z	laminate density, heat, conduction	1580, 1100, 0.6	SI	T, A(c, k_z)
w_r	resin mass fraction	0.34	—	T, A
L, A_p	skin, spar, frame, shadowed	3, 25, 10, 6; 2.0, 0.6, 0.8, 0.8	mm; m^2	T
C_t	tools, same order	29.2, 231.8, 58.8, 54.0	kJ/K	T
A_t, s_p	tool areas; shadow	2.0, 1.8, 1.2, 1.0; 0.45	m^2; –	T
r_P, τ_{bag}, λ_b	compressor, bag lag, leak ratio	0.5, 60, 0.005	bar/min, s, –	T
τ_{TC}, σ_{TC}	thermocouple lag, noise	10, 0.2	s, °C	T
limits	peak, cure, voids, band	195, 0.87, 1, ±5	°C, –, %, °C	T

20 °C and the resin at its lot's initial cure. The state has 179 entries, and a whole challenge cycle of 438 min costs 228 ms of computing headless and about 22 s at the top speed of 1200 times real time in a browser.

6.3.8 Assumptions

Each part is one-dimensional, so edges, ply drops and corners, which run hotter in a thick part, are absent, and the tool is one lumped mass in perfect contact through h_c. The air is one well-mixed node and each part's heat transfer is a constant times a pressure power and a shadow factor, while a real vessel has a flow field whose survey finds its slow positions and whose radiation matters [17]. The laminate's properties are constant, although its specific heat rises with temperature. Resin flow, bleed, compaction, volatiles and residual stress are outside the model, so a dwell earns nothing here that it earns in a real cure. The porosity index is a teaching law. The 195 °C exotherm limit, the cure target, the 1 % porosity threshold and the ±5 °C hold band are teaching values, and the real ones come from the resin's specification and the part's process specification. A simulated pass is not a qualification of any part or process.

Table 6.3 Spar Peak Temperature (°C, Teaching Truth) on Air Control and the Standard Recipe, Against Thickness and Ramp

Ramp	10 mm	20 mm	25 mm	30 mm	35 mm	40 mm
1 °C/min	180.1	181.9	188.5	199.2	213.6	231.1
2 °C/min	180.0	193.8	210.9	233.3	260.4	279.3
3 °C/min	180.0	195.7	215.4	241.6	269.8	288.5

✦ DID YOU KNOW? Windward and Leeward in a Full-Size Autoclave

One full-size industrial autoclave has a cylindrical cavity 3.7 m across and 9.6 m long, with air or nitrogen driven by a fan. Fifteen thermocouples on a large steel mold read hottest on the windward side near the door and coldest on the leeward side [19]. The shadow factor of this chapter stands for that leeward lag.

6.4 What Sets the Lags and the Limit

The recipe asks every part for the same temperature, and the load cannot deliver it at once. On the standard recipe run on the air's clock the skin tracks the air within 1 °C (a peak of 181.4 °C against the air's 180.8 °C), while the spar's tool face reaches the hold band 37.5 min after the air, because both of its half-hour time constants lie in series. A recipe on the air's clock either cuts the spar's hold short or pads every hold for it. The standard recipe pads, the least-cured point still reaches $\alpha = 0.887$, and the padding is cycle time.

The heat release scales with thickness and ramp, as shown in Table 6.3. On air control and the standard recipe a 10 mm spar never exceeds 180.1 °C at any ramp, a 25 mm spar peaks at 188.5 °C at 1 °C/min and 215.4 °C at 3 °C/min, and a 40 mm spar reaches 231.1 °C even at 1 °C/min. The mechanism is a ratio of rates. When the air ramps faster than the core can follow, the core arrives late at the temperature where the cure accelerates, with most of its resin unreacted, and its heat must leave through half the laminate, L^2/a of half an hour or more, while the second kinetic term of Equation (6.4) doubles its rate every 8.8 K near 177 °C. Lowering the air then removes heat only from the faces.

> **DID YOU KNOW? Thickness Beats Ramp**
>
> From the simulator: on air control and the standard recipe, a 30 mm spar ramped at the slowest rate, 1 °C/min, peaks at 199.2 °C, hotter than a 20 mm spar ramped three times as fast, which peaks at 195.7 °C (Table 6.3).

Thermocouples lie on the wrong side of this mechanism. At 101.8 min of case 4 the spar's node at $z = 0.2L$ is at 195.8 °C while its thermocouples read 187.4, 182.6 and 148.8 °C, as drawn in Figure 6.1a, and the hottest of them passes 195 °C 1.5 min after the core does.

Pause. Try Exercise 6.1 before reading on.

6.5 The Control Layers as Implemented

Every 10 s the controller runs the cure observer, the recipe sequencer, the supervisor that writes the air setpoint and the air loop, in that order, on the measurement packet. The application offers four structures, Air only, Lead/lag, Lead/lag with the feedforward (FF) and the MPC of Section 6.5.6.

6.5.1 The Air Loop

The air loop is the velocity-form PI of Section 1.6.2, written for this loop in Equation (6.7).

$$\begin{aligned} e_k &= T_{air,SP,k} - T_{g,m,k}, \quad & s_k &= s_{k-1} + K_c(e_k - e_{k-1}) \\ \hat{u}_k &= s_k + I_{k-1}, & u_k &= \mathrm{slew}\big(\mathrm{clip}(\hat{u}_k, \pm 1),\ 1/(180\,\mathrm{s})\big) \\ I_k &= I_{k-1} + \Big[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{\tau_I}\Big]\Delta t_c \end{aligned} \tag{6.7}$$

with $K_c = 10\,\%$ of output per °C, $\tau_I = 10$ min, no derivative and $\Delta t_c = 10$ s. In the symbols of Equations (1.15) to (1.17), s_k is S_k, $\hat{u}_k$ is v_k with $u_b = 0$, and $T_i = \tau_I$. The output moves from zero to full in no less than 3 min, so the air overshoots a large step. In case 1 the air reaches 145.7 °C at 5.8 min on a step to 110 °C and stays within 2 °C of it from 13.6 min.

6.5.2 The Recipe Sequencer

The sequencer advances a reference r through ramp, dwell, ramp, hold, cool and demold for one clock thermocouple, the air on Air only and the lagging thermocouple T_{lag} on the lead/lag structures. On the lag clock the ramp is guaranteed. The reference advances at the ramp rate only while it is less than 10 °C ahead of T_{lag}, and a dwell or hold counts its time only while the clock thermocouple is within 5 °C of it. The hold ends on its time or, with the observer rule, once the observer's lowest $\hat{\alpha}$ is at least 0.875 and its $\hat{T}_g$ at least 210 °C, a guard on an estimate in the sense of Section 1.4.2. In the cool-down the reference falls at 2 °C/min, and demold comes when the hottest thermocouple reads 60 °C. The pressure is commanded on at t_P. The framework's shared stage sequencer was not used, because it switches on/off stages against a demand and has no ramp target, hold clock or choice of clock thermocouple.

6.5.3 Lead/Lag With a Maximum Selector

On the lead/lag structures a minimum selector across the twelve part thermocouples gives the lagging temperature T_{lag} and a maximum selector gives the hottest, T_{max} (Table 6.1). Two PI loops write candidate air setpoints from them, the lag loop on T_{lag} and the max loop on T_{max}, and a low selector

takes the smaller, as shown in Equation (6.8).

$$\begin{aligned} e_L &= r - T_{lag}, \qquad e_M = T_{lim} - T_{max} \\ \ell &= \mathrm{clip}(K_L e_L + R_L,\ -20,\ \Delta_{lead}) \\ o_M &= \mathrm{clip}(K_M e_M + R_M,\ -60,\ 40) \\ T_{lim} &= \min\left(190\,^\circ\mathrm{C},\ T_{seg} + 5\,^\circ\mathrm{C}\right) \\ T_{air,SP} &= \mathrm{clip}\big(\min(r+\ell,\ T_{lim}+o_M,\ T_F), \\ &\qquad 20\,^\circ\mathrm{C},\ 225\,^\circ\mathrm{C}\big) \\ R_X &\leftarrow R_X + \left(\bar{R}_X - R_X\right)\min(1,\ \Delta t_c/\tau_{I,X}) \end{aligned} \tag{6.8}$$

where T_{seg} is the dwell temperature during the first ramp and the dwell and the hold temperature after, T_F is the feedforward candidate (absent on Lead/lag), and $\bar{R}_L = T_{air,SP} - r$ and $\bar{R}_M = T_{air,SP} - T_{lim}$ express the applied setpoint in each loop's own terms. This is the external-reset form of override control [5], the integrator tracking that the limiters of the turbofan in Volume 4, Chapter 3 also use under their selector. For the selected loop $\bar{R} - R$ equals its own proportional term, so it acts as a PI, and the loop passed over sits $K_X e_X$ above the applied setpoint and takes over without a bump. The app's switch named Maximum selector turns the max loop's candidate off, and the Lead/lag card reads "max selector" while that candidate holds the air. The text calls it the max loop. The defaults are $K_L = 1$ with $\tau_{I,L} = 20\,\mathrm{min}$, a lead of at most $\Delta_{lead} = 40\,^\circ\mathrm{C}$, and $K_M = 2$ with $\tau_{I,M} = 10\,\mathrm{min}$. In Manual, on the first Auto period after Manual and after a change of structure, each reset term is set so that its loop's output equals the setpoint last applied. Measured on case 3, every transfer from Manual or from Air only into Lead/lag keeps the first period's setpoint within 0.01 °C of the one last applied, which the technical review added after finding steps of up to 7.3 °C. The selector cannot wind up the loop it passes over. In case 4 the lead sat at its 40 °C clip for 49.2 min with no windup, and the largest setpoint change in one period at a change of selection was 0.75 °C against 4.32 °C elsewhere in the run.

6.5.4 The Cure Observer and the Feedforward

The observer integrates Equation (6.4) at each of the twelve part thermocouples on its measured temperature, by fourth-order Runge-Kutta over each control period at the mean of the period's two readings, from the initial cure the load sheet states. It reports the lowest $\hat{\alpha}$ and its $\hat{T}_g$ for the demold rule, and for each part a predicted self-heating rate and excess, as shown in Equation (6.9).

$$\begin{aligned} \hat{s}_p &= \frac{w_r H_r}{c}\,\dot{\alpha}\left(T_{p,1}, \hat{\alpha}_{p,1}\right), \qquad \widehat{\Delta T}_p = \frac{\hat{s}_p L_p^2}{2a} \\ T_F &= T_{lim} - \min\left(80\,\mathrm{K},\ K_{ff}\max_p \widehat{\Delta T}_p\right) \end{aligned} \tag{6.9}$$

where $T_{p,1}$ is the mid-thickness reading. The excess is the quasi-steady rise of a slab of thickness L_p heating itself at the present rate and cooled through one face. The feedforward candidate T_F enters the low selector only on the Lead/lag + FF structure, in Auto, with the observer on and during the ramps, the dwell and the hold, with $K_{ff} = 0.4$. A feedforward subtracted after the selector made the lag loop's tracking ratchet during the build, which is why it is a candidate instead. The bound $\widehat{\Delta T}$ is loose and one-sided. On case 4 on Lead/lag it reads 656.8 K at 106.0 min while the true rise of the spar's core over its hotter face peaks at 62.4 K, and in the suggested challenge run it reads 101.1 K against a true 9.7 K, hence the small gain. With the feedforward acting it held in all 1011 control periods of case 4 in which it exceeded 1 K. Without the feedforward, once the reaction has burned out, the heat it left in the core exceeds the collapsed rate's estimate by up to 36.6 K at 115.3 min, because a rate that has stopped no longer predicts the heat already released. For the shipped structures no control model is identified, so this copy of the kinetics, with the load sheet's thicknesses and the declared

laminate and resin, is their only internal model, and the MPC's own model is given in Section 6.5.6. It knows the declared resin exactly but sees no node between the thermocouples and believes the lot's stated initial cure. A heat of reaction or a laminate changed in the Model dialog changes the plant and leaves the observer on its declared values. On case 4 with H_r raised to 700 J/g the spar then peaks at 215.2 °C on Lead/lag + FF.

6.5.5 The Recipe Optimizer

The optimizer is a grid search run offline on the application's own engine. For one structure it runs every recipe with hold temperatures of 175, 180, 185 and 190 °C and hold times of 10, 20, 30, 45, 60, 90 and 120 min on the lagging clock, without a dwell, to demold on the case's seed, and keeps the shortest feasible recipe judged on the truth with margins of 0.005 on the cure and 2 °C on the exotherm. Its result is shipped as the Optimized recipe, 185 °C for 10 min with the feedforward and 180 °C for 30 min on Lead/lag alone, and the application's tests rerun the search and confirm the same answers. The optimizer is not a receding-horizon controller and acts on no running cycle.

6.5.6 A Shrinking-Horizon Predictive Controller

The fourth structure replaces the lag loop with an MPC, a receding-horizon solution of a constrained optimal control problem [20] (Section 1.6.3), that writes the candidate u for the air setpoint. The air loop of Equation (6.7) stays beneath it, and three shipped protections of Lead/lag + FF stay above it in a low select, $T_{air,SP} = \min(u,\ r+\Delta_{lead},\ T_F,\ T_{lim}+o_M)$, so the lead clamp, the feedforward candidate and the max loop of Equations (6.8) and (6.9) can each lower the air under the plan, in the external-reset form already described. Industrial MPC is layered the same way, above the regulatory loops and beneath the plant's limit and safety logic [21]. While the MPC is selected, the lag loop's reset follows the applied setpoint, so its gain and integral time no longer act.

The controller's model is the heat network of Equation (6.3) written again from the declared constants of Table 6.2, never from the run's parameters, on 3 nodes through the load sheet's thickness of the skin, frame and shadowed part and 5 for the spar and the fitting, whose cores the thermocouples miss. Each part keeps its bag film $U_{s,p}A_p$, its contact h_cA_p and a lumped tool, with h_p from Equation (6.2) at the measured vessel pressure rounded to 0.5 bar and rebuilt when that rounding moves. The closed air loop is one lag, $\tau_a\dot{\hat{T}}_g = u - \hat{T}_g$ with $\tau_a = 44.6$ s, the sum of the time constant and dead time of a first-order-plus-dead-time fit ($K = 1.003$, 9.2 s, 35.4 s, residual 2.8 % of the response) to a ±10 °C step test on the air setpoint at 150 °C. One integrating air offset d_p per part enters its bag film and its tool's exposure. The reaction enters as a source along the plan, as Equation (6.10) states.

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{A}\mathbf{x}_k + \mathbf{b}\,u_k + \Gamma\bar{\mathbf{s}}_k, \qquad \hat{\mathbf{y}}_k = \mathbf{C}\mathbf{x}_k \\ \bar{s}_{k,i} &= \frac{w_r H_r}{c}\,\frac{\tilde{\alpha}_i(t_{k+1}) - \tilde{\alpha}_i(t_k)}{\Delta t_m} \end{aligned} \tag{6.10}$$

Here $\mathbf{x}$ holds the air, every node, every tool and every d_p (23 states, 30 with the fifth part), $\mathbf{A}$, $\mathbf{b}$ and Γ are the exact discretization at the decision period $\Delta t_m = 300$ s, and $\hat{\mathbf{y}}$ holds the 12 thermocouple positions and the two interior nodes of the spar (14 outputs, 19 with the fitting's three and two). The cure $\tilde{\alpha}_i$ of every model node is integrated with the declared kinetics of Equation (6.4) along the plan's own predicted temperatures, by fourth-order Runge-Kutta in 60 s steps, from the observer's $\hat{\alpha}$ at the thermocouples, linear between them. Each decision simulates this nominal path, solves, integrates the kinetics again along the new plan and solves again, with the source relaxed half and half between passes, up to six passes or until the plan moves less than 0.5 °C. Run open loop on the setpoints and cure increments of whole cycles, the model stays within 0.55 to 1.06 °C root mean square of each thermocouple on the standard and optimized recipes and with the fifth part (2.6 °C at worst), 0.78 to 1.49 °C on case 4, and 20.5 °C on case 6's shadowed part, whose shadow of 0.1 the model takes as 0.45. An independent 10 °C air step at 80 °C, with a twin run subtracted, keeps every thermocouple within 0.42 °C of the model over 3 h.

A steady-state Kalman filter, its gain from the discrete Riccati equation at 7 bar gauge, estimates **x** from the air and the twelve thermocouples with a measurement deviation of 0.25 °C and random walks per decision of 1.0 °C for the air, 0.3 °C for the nodes and tools and 1.5 °C for the offsets. Between decisions the estimate is propagated every control period with the setpoint actually applied and the observer's own cure increments, and it is corrected at each decision. The offsets make the plan offset free [22], as the fuel bias of the boiler-turbine MPC in Volume 9, Chapter 4 does, and absorb a wrong film, a heavier tool, a shadow and the reaction heat the model misses, all as air that a part does not receive.

Every 30 control periods, on the clock rounded to the control period and first on the priming call at $t = 0$, the plan covers what is left of the recipe, $N \le 72$ steps of 5 min (6 h), in 12 blocks re-spread over N at each decision, one step each while $N \le 12$ and from 1 to 11 steps at $N = 72$. The horizon ends at the end of the hold on its time or, with the observer's end rule, where the plan's predicted cure at every thermocouple reaches $\alpha \ge 0.875$ and $T_g \ge 210\,^{\circ}\mathrm{C}$. With less than one step left the last setpoint is held. The recipe itself is the preview: the reference r_k is the sequencer's on the lagging clock, extended ahead with its ramps, dwell and hold, which the lag loop also follows, so there is no preview to switch off. The plan solves Equation (6.11).

$$\begin{aligned} \min_{u,\varsigma} \; & \sum_{k=1}^{N} \sum_{j \in \mathscr{L}_k} (\hat{y}_{j,k} - r_k)^2 + S \sum_{\ell} \Delta u_\ell^2 + \sum_i \left(\rho_2 \varsigma_i^2 + \rho_1 \varsigma_i\right) \\ \text{s.t. } & \hat{y}_{j,k} \le T_{lim,k} - 0.5 - \tfrac{1}{6}\Delta T_{rxn,j,k} + \varsigma_j \\ & \hat{y}_{o,k} \le 193 - \tfrac{1}{6}\Delta T_{rxn,o,k} + \varsigma_o, \qquad \varsigma \ge 0 \\ & 20 \le u_k \le \min\left(225,\; r_k + \Delta_{lead},\; T_{lim,k} + 40\right) \end{aligned} \tag{6.11}$$

Temperatures are in °C. The set $\mathscr{L}_k$ holds the coldest predicted thermocouple position, tracked both ways as the lag loop tracks it, and every other position predicted below r_k, with unit weight, and $S = 0.1$ prices the block moves Δu_ℓ. The first ceiling holds the twelve thermocouple positions j 0.5 °C under the selector's limit $T_{lim,k}$, the lowest it takes in step k, and the second the spar's and fitting's interior nodes o 2 °C under the 195 °C exotherm limit. $\Delta T_{rxn,i,k}$ is the plan's own predicted rise of output i at step k from the reaction heat still to come (the response to the source term alone from zero), so the plan keeps room for a heat of reaction up to a sixth above the declared 600 J/g, the 700 J/g at the end of the parameter's range, and ς_i is the slack of output i. Both ceilings are soft. Each output carries a single slack for the whole horizon, priced at $\rho_2 = 10^3$ per °C^2 and $\rho_1 = 3 \times 10^4$ per °C, and the linear price is large enough that a ceiling the plan can meet is met exactly [23]. The bounds on u are hard: the setpoint's range, the lead clamp at the reference at the block's start, and with the selector on its loop's output clip. The first move may differ from the applied setpoint by 15 °C and each later block boundary by 15 °C per step between block centers, a ramp rather than one cap per boundary. The plan has at most 12 move variables and 14 to 19 slacks. Per-step slack cost 10 to 50 times the solve (case 4 115 ms against 6 ms) and moved no limit consistently, and no plan of the first build's failing runs predicted a violation (Section 6.7.7), so it was not adopted.

A plan is accepted when the solver reports optimal, when the first move, the planned inputs and outputs, the objective and the slacks are all finite with the expected dimensions, the slacks non-negative, the first move equal to the first planned move and every planned setpoint inside 20 °C to 225 °C. Otherwise Lead/lag + FF takes the air, whose lag loop has tracked the applied setpoint, at no more than 0.5 °C a control period until it meets its own request, the historian is marked, and the next accepted plan resumes. Faults injected at 60, 135 and 150 min handed over at 0.18, 0.26 and 0.09 °C a period with the noise off and resumed with steps of 5.19, 3.14 and 0.05 °C. No fallback occurred in the 80 runs of the MPC over the comparison, the worst cases and the parameter extremes. The plan reads only the packet, the load sheet, $\hat{\alpha}$, the recipe and its own state. A replay of a changed plant's packets into a fresh controller reproduced all 36 decisions, and the truth test catches 18 of 18 doctored copies.

Table 6.4 Code Verification and Agreement With the Source's Tests

Check	Result
Code verification	
Energy ledgers over the challenge cycle, air and lead/lag	6.6×10^{-12}, 6.2×10^{-12} of the heater energy
Two disturbed runs (fifth part, aged lot, heater loss, leak)	3.5×10^{-12}, 3.1×10^{-12}
Reaction ledger against $\rho w_r H_r \sum V \Delta\alpha$	equal to 4 decimals (9.9578 MJ)
Case 4 peak, app scheme at 0.1 s, explicit RK4 method of lines at 0.05 s	within 0.071 and 0.072 °C of 243.87, $\lvert\Delta\alpha\rvert \leq 3.33 \times 10^{-4}$
11, 21, 41 nodes: shipped, suggested, case 4 peaks	188.48/188.49/188.51, 189.30/189.31/189.31, 243.87/244.47/244.43, verdicts kept
Step 5 to 2.5 s, peaks	shipped 0.057, case 4 0.023 °C
Each case loaded first, from the list and after Reset	identical, 10 of 10
Agreement with the source's tests [14]	
26 resin constants against the report	identical
α after 120 °C, 480 min	0.636 against 0.636
α after 150 °C, 360 min	0.786 against 0.788
α after 180 °C, about 200 min	0.894 against 0.898 and 0.904
Peak rate, 120/150/180 °C	5.25×10^{-5}, 2.15×10^{-4}, 7.28×10^{-4} s^{-1} at 105, 24.9, 6.7 min, about the source's
T_g at α 0.52, 0.70, 0.90	110.7, 158.9, 218.0 against 108.2, 165.9, 218.1 and 226.7 °C
Lowest viscosity on 1, 2, 4 °C/min ramps	5.6, 4.1, 3.2 Pa s against about 6, 4, 3

A rule written before any search, judged on case 9 alone (verdict and margins, then the shortest cycle), set the weights from a flat 81-setting grid (280.3 to 284.8 min). A decision costs 5.3 ms on average in Node on case 9 (17.0 ms at most). In Chrome, in the page's own frame loop at 1200 times real time under a load average of 9.3 from parallel work, the logged decision averaged 9.9 ms (41.7 ms at most), 2.5 % of a core, with 59.9 frames per second. Because a decision costs more than 2 ms, a frame guard skips decisions, never frames, while the page lags real time.

6.6 Code Verification and Agreement With the Source

Code verification, agreement with the published fits and educational use are kept apart in Table 6.4. The first block checks that the code integrates the equations as written. The second compares the kinetics with the tests from which they were fitted, so it measures the fit and the transcription, not an independent validation. The source has two isothermal tests at 180 °C, which end at 0.898 and 0.904, and over its eleven isothermal tests the glass-transition relation is within 13.0 °C of the measured values with a root mean square of 5.3 °C, which is the scatter of the fit itself. No part or autoclave was measured for the model, and its use in teaching has been tried only in review runs. No learning study has been made.

Case 4's peak also depends on the noise draw. Over six seeds it spans 243.8 to 246.0 °C, and without noise it is 244.58, 244.04, 243.37 and 243.27 °C at steps of 5, 2.5, 1 and 0.5 s, so every run lies far past the 240 °C the case's prediction offers.

6.7 Nine Cases and a Challenge, Measured

The cases appear in Table 6.5, and Section 1.6.4 describes how a case is loaded and judged. Each starts from the standard recipe on Air only, pressure from $t = 0$, the default tunings and the default seed, and departs only as its row says. Each value below comes from the headless engine, advanced step by step as in the browser, with each change named in the text made before Run. Table 6.6 gathers the outcomes from the application's scorecard, judged on the truth at demold.

6.7.1 The Lag, by Hand and on the Air's Clock

Case 1 holds the air by hand. With the setpoint stepped to 110 °C at 1 min, the coldest thermocouples of the skin, the frame, the shadowed part and the spar reach 105 °C 10.1, 29.8, 57.8 and 82.4 min after the air does. At 30 min the spar reads 95, 78 and 71 °C from the bag side to the tool face, so the invar tool, not the laminate, holds it back. In case 2 the air enters the hold band at 214.9 min, the hold

Table 6.5 The Cases

#	Case and settings	Teaching point
1	Manual, air setpoint by hand	each part's lag
2	Standard recipe on Air only	the spar lags into the hold
3	The same, to be run on Lead/lag	slowest ramps, hottest limits
4	35 mm spar, Lead/lag, 3 °C/min, no dwell	a runaway the selector cannot stop
5	Pressure on at 240 min	voids frozen at gel
6	Shadow 0.1, Lead/lag	a cycle as long as its worst part
7	Lead/lag, no dwell, 180 °C for 10 min	the observer and demold
8	Aged lot (α_0 0.12), pressure at 170 min	gelation comes sooner
9	Optimized recipe, MPC	one plan against three rules
10	Challenge: the standard recipe on Air only	minimum cycle under limits

Table 6.6 Measured Outcomes at Demold, Default Seed. Cycle (min), True Peak (°C), Least-Cured Point and Minutes Outside the Hold Band

#	Configuration	Cycle	Peak	α_{min}	Band
2	as shipped	438.3	188.5	0.887	78.4
3	Lead/lag	463.2	186.1	0.896	4.2
4	as shipped (Lead/lag)	392.5	243.9	0.895	0.25
	Lead/lag + FF	392.2	184.8	0.895	—
	Lead/lag, 0.5 °C/min	541.3	184.7	—	—
5	as shipped, voids 12.49 %	438.2	200.5	0.875	142.2
6	as shipped (Lead/lag)	766.8	185.0	0.895	8.7
	Air only	534.3	—	0.865	179.7
7	as shipped	280.7	192.6	0.858	0.1
	hold ends on the observer	304.0	—	0.881	—
8	as shipped, voids 1.19 %	438.3	192.2	0.886	100.0

— not part of the case's measured record.

begins counting at 220.0 min, and the spar's tool face enters the band 37.5 min after the air. Over the cycle the part thermocouples spend 78.4 min outside the band, the worst 22.6 °C below it, and demold comes at 438.3 min.

6.7.2 Lead/Lag: the Slowest Ramps, the Hottest Limits

Case 3 runs the same recipe on Lead/lag. With the Maximum selector switch off, so with no max loop, the lag loop leads the air by up to 40 °C for the spar and the skin follows the air, to 139.1 °C in the 110 °C dwell and 202.7 °C three minutes into the hold, with 124.2 min outside the band. With the switch on, the max loop first takes the air at 63.6 min as the skin nears its 115 °C limit and holds it for 134 min of the cycle in all, the skin stops at 114.9 °C in the dwell and 185.1 °C in the hold, and only 4.2 min are spent outside the band. While the max loop held the air, the hottest thermocouple stayed within 5.30 °C under its limit. The cycle lengthens to 463.2 min, 24.9 min more than on the air's clock, because every hold now waits for the spar. Lead/lag buys the spar its recipe, not time.

🕮 **Pause.** Try Exercise 6.2 before reading on.

6.7.3 The Exotherm a Selector Cannot Stop

Case 4 thickens the spar to 35 mm and ramps it at 3 °C/min with no dwell, as shown in Figure 6.2. On Lead/lag the max loop has held the air since well before the runaway, and the air setpoint falls from 172.9 °C at 100 min to its floor of 125 °C (the limit less the 60 °C clip of the max loop) by 104.5 min. The core still passes 195 °C at 101.8 min, while the hottest thermocouple reads 187.3 °C, and peaks at 243.9 °C at 109.1 min. Retuning does not help. A limit of 180 °C still gives 231.0 °C, a max loop at $K_M = 6$ and 5 min gives 243.9 °C, and without the max loop the core reaches 301.5 °C. The screenshot of Figure 6.4 is taken at the moment the core crosses the limit.

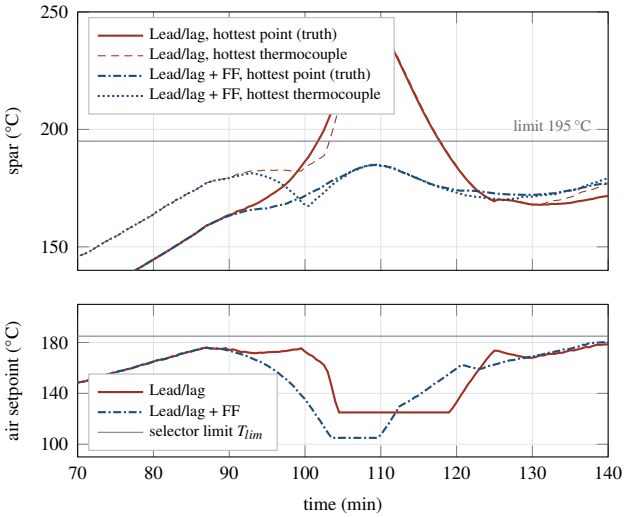


Figure 6.2 The exotherm of case 4, a 35 mm spar ramped at 3 °C/min. Top, the spar's hottest point (teaching truth) and its hottest thermocouple on Lead/lag (solid and dashed) and on Lead/lag + FF (dash-dotted and dotted), with the 195 °C limit as a thin line. Bottom, the applied air setpoint on Lead/lag (solid) and on Lead/lag + FF (dash-dotted) and the selector's limit (thin line). On Lead/lag the max loop cuts the air to its floor after the core is already running away.

> **TRY IT IN THE APP** **4 · The Exotherm: A Thick Spar at 3 °C/min**
>
> Load the scenario and press Run. On the Autoclave view, watch the spar's three thermocouples and the true peak badge. The core passes 195 °C at 101.8 min while the hottest thermocouple still reads 187.3 °C, and it peaks at 243.9 °C. Set Control structure to Lead/lag + FF, press Reset and run again: the spar now peaks at 184.8 °C.

The feedforward acts on the reaction instead of its temperature. Its candidate is first selected at 89.5 min, it reaches its 80 °C cap at 103.5 min, and at 100 min it has the air setpoint at 135.5 °C where the max loop had 172.9 °C. The core peaks at 184.8 °C in a cycle of 392.2 min against 392.5 min without it. The effect belongs to the observer's prediction and not to a static cap, because $K_{ff} = 0$ reproduces Lead/lag exactly (243.9 °C), and gains of 0.1, 0.2 and 0.4 give 239.2, 215.2 and 184.8 °C. A slower recipe also keeps the spar under the limit, though not at 1 °C/min (225.3 °C), and at 0.5 °C/min it peaks at 184.7 °C in a cycle 149 min longer, 541.3 min.

6.7.4 Pressure, Shadow, the Observer and an Aged Lot

Case 5 pressurizes the vessel at 240 min. The skin gels at 212.3 min and the frame at 238.0 min under vacuum alone, the voids formed while their resin flowed are frozen in, and the porosity index ends at 12.49 %. With the pressure on at 0, 165, 175 and 180 min the index ends at 0.00, 0.30, 0.81 and 1.30 %, so the step must come tens of minutes before the first gel. The late pressure also starves the parts of heat, the spar lags into the hold 87.7 min instead of 37.5 min, and when the pressure arrives its exotherm reaches 200.5 °C.

Case 6 pushes the shadowed part to a factor of 0.1, about 7.3 W/(m^2 K) at pressure. On Lead/lag it carries the lagging thermocouple for 94 % of the dwell and hold, the recipe waits for it, and the cycle takes 766.8 min against 463.2 min at 0.45. On Air only the cycle is 534.3 min, but the shadowed part reaches the band 126.6 min after the air and is left at $\alpha = 0.865$, under target. A lead/lag cycle is as long as its worst part, and on the air's clock the same part is under-cured.

Table 6.7 The Challenge, Default Seed. Cycle (min), True Peak (°C), Least-Cured Point, Minutes Outside the Band

Policy	Cycle	Peak	α_{min}	Band	Verdict
standard recipe, Air only (shipped)	438.3	188.5	0.887	78.4	fails
standard recipe, Lead/lag	463.2	186.1	0.896	4.2	fails
standard recipe, Lead/lag + FF	472.5	184.5	0.895	0.1	fails
standard hold, no dwell, Lead/lag	390.7	192.6	0.896	0.1	fails
185 °C, 10 min, Air only	275.5	203.4	0.808	9.8	fails
185 °C, 10 min, Lead/lag	285.8	195.3	0.878	0.2	fails
180 °C, 30 min, Lead/lag (its optimum)	301.2	192.6	0.879	0.1	passes
182 °C, 15 min, Lead/lag	288.0	193.8	0.872	—	passes
185 °C, 10 min, Lead/lag + FF (suggested)	289.8	189.3	0.878	0.5	passes

Case 7 holds 180 °C for 10 min on Lead/lag with no dwell. At demold (280.7 min) the least-cured point is at $\alpha = 0.858$ with T_g 205.0 °C, and the observer reports 0.858 and 205.1 °C. Ending the hold on the observer lengthens it from 21.7 to 45.0 min of clock time, because the hold counts only while the lagging part is in the band, and demolds at 304.0 min with $\alpha = 0.881$ and T_g 212.0 °C. Over both runs the observer stays within 0.014 of the true degree of cure at every thermocouple.

Case 8 starts an aged lot at $\alpha_0 = 0.12$ with the pressure step at 170 min, a time that suits a fresh lot. The skin gels at 200.8 min instead of 208.3 min, and the porosity index ends at 1.19 % against 0.50 % for a fresh lot. On six seeds it spans 1.178 to 1.188 %, all over the 1 % threshold. With the pressure at 150 min it ends at 0.16 %.

6.7.5 The Challenge

The challenge asks for demold within 310 min with no point of any part above 195 °C, every point at $\alpha \geq 0.87$, a porosity index of at most 1 % and no more than 10 min outside the hold band, on the load and resin as shipped. Table 6.7 lists the policies measured. The shipped standard recipe on the air's clock fails on cycle and band. The suggested configuration, Lead/lag + FF on the optimized 185 °C hold of 10 min, passes in 289.8 min with a peak of 189.3 °C and $\alpha = 0.878$ (T_g 211.1 °C), and on five other seeds of the noise in 289.7 to 291.2 min, all passing. The 310 min limit lies between the optimizer's best on Lead/lag alone with its margins, 301.2 min, and the standard recipe on any structure, 438 min or more.

The feedforward buys margin, not minutes. With the optimizer's margins it is what lets the hold run at 185 °C, since the same recipe on Lead/lag alone peaks at 195.3 °C and fails, and Lead/lag's own margined optimum takes 11.4 min longer. Without the margins Lead/lag alone keeps up, and 182 °C for 15 min passes in 288.0 min (287.7 to 289.8 min on six seeds, all passing) with its peak 1.2 °C under the limit and $\alpha = 0.872$, where the suggested run keeps 5.7 °C and 0.008 in hand. The optimum is set by the cure margin and not by the edge of the grid. Extending the grid to holds of 0 and 5 min and 2.5 °C steps keeps both optima, the shorter holds miss the cure margin at every temperature (the best, 187.5 °C for 5 min, takes 294.0 min), and a 1 min grid would save 2.3 min with 185 °C for 8 min. A 50 mm spar is infeasible within the limit. None of 38 settings, over three structures, feedforward gains of 0.4 to 2, four recipes and further dwells and ramps, passed, and the shortest that kept it under 195 °C took 440.8 min. Figure 6.3 draws the policies of Table 6.7 as bars against the two limits.

6.7.6 Case 9: One Plan Against Three Rules

Case 9 runs the challenge's suggested recipe, 185 °C for 10 min on the lagging clock with no dwell, under the MPC. Every row below shares the seed, instruments and noise, runs to demold and keeps case 9's MPC settings. The MPC demolds at 291.3 min with a peak of 189.3 °C and $\alpha = 0.879$ (T_g 212 °C) in 37 solves, against Lead/lag + FF's 289.8 min, 189.3 °C and 0.878, a tie within the spread of the noise. On six seeds it takes 290.0 to 291.3 min against 289.7 to 291.2 min, all passing. The MPC keeps the lagging thermocouple 1.47 °C behind the reference on the second ramp on average,

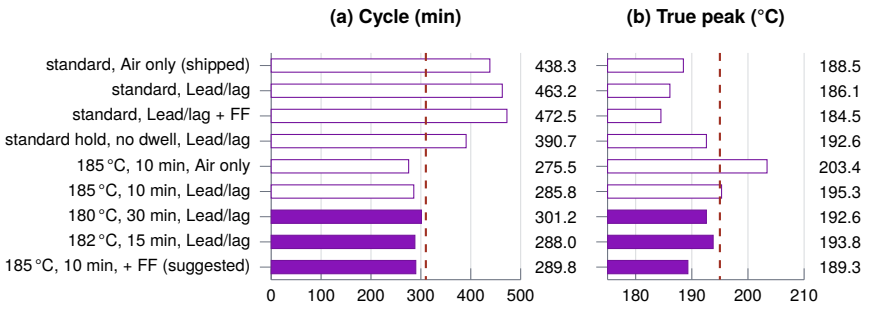


Figure 6.3 The policies of Table 6.7 on the challenge. (a) Cycle time with the 310 min limit dashed, (b) the true peak with the 195 °C limit dashed, dark bars for the policies that pass and white bars for those that fail. Notice that the three passing policies all sit just under 310 min, while the two fastest bars fail on the peak or the cure, and that the suggested run passes with the most room under the peak limit.

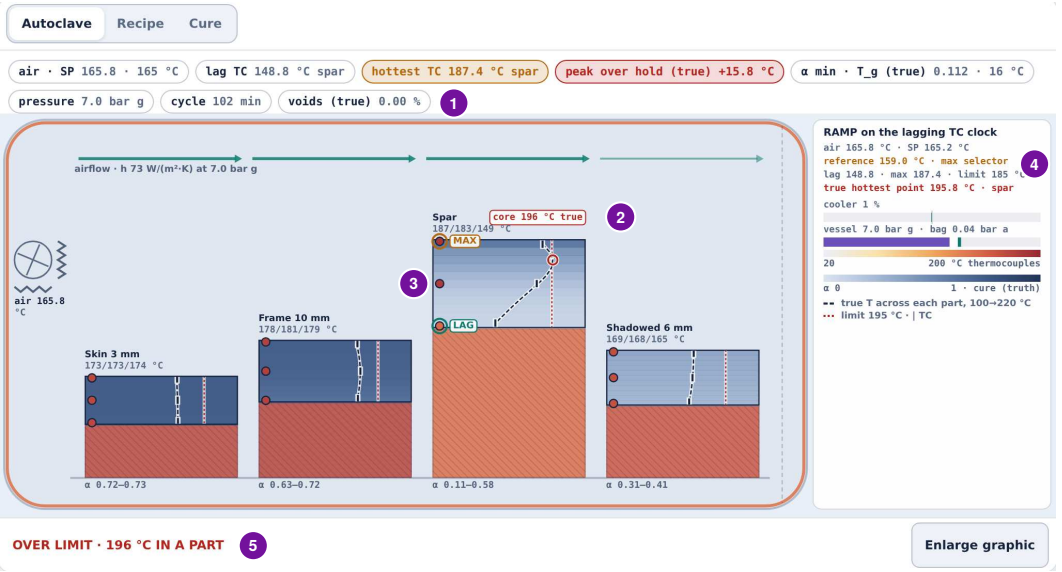


Figure 6.4 The core of a thick spar past the limit between its thermocouples, case 4 at $t = 101.8$ min on the Autoclave view, Lead/lag in Auto on the second ramp, cropped to the plant view. (1) The badges give the hottest thermocouple and the true peak over the hold. (2) The spar's dashed true profile crosses the red dotted 195 °C line between its thermocouple ticks, ringed as "core 196 °C true" against readings of 187, 183 and 149 °C, (3) with MAX on its bag side and LAG on its tool face, while the skin, frame and shadowed part read 165 to 181 °C. (4) The side panel gives the air at 165.8 °C against a setpoint of 165.2 °C chosen by the max loop (the panel reads "max selector"), the reference at 159.0 °C, the limit at 185 °C, the true hottest point at 195.8 °C and the vessel at 7.0 bar gauge with the bag at 0.04 bar absolute, and (5) the status strip reads OVER LIMIT.

against 3.96 °C, with 318 °C of setpoint travel against 362 °C. The feedforward candidate above it then holds the air for 62.5 min of the recipe (56.3 min under Lead/lag + FF), the last thermocouple enters the hold band at 179.6 min against 179.1 min, and both cool-downs take 101 min.

The three rows required beside the headline are in Table 6.8. Lead/lag + FF, retuned a single time under the rule that tuned the MPC and over all six of its settings ($K_L = 10$, $\tau_{I,L} = 40$ min,

Table 6.8 Cycle (min) and True Peak (°C) Under Lead/Lag + FF as Shipped, Retuned Once by the MPC's Rule, Fed the MPC's Prediction, and Under the MPC. Default Seed, Each Run to Demold

Case	FF	retuned	FF + pred.	ret. + pred.	MPC
9, optimized	289.8/189.3	281.3/190.0	291.5/189.2	281.5/190.0	291.3/189.3
standard recipe	472.5/184.5	452.7/185.4	474.7/184.4	468.5/185.0	456.8/184.5
4, 35 mm spar	392.2/184.8	367.5/240.0	401.7/184.0	384.7/194.6	414.8/184.0
50 mm, optimized	365.8/231.7	335.0/279.7	380.5/190.4	348.7/225.7	429.8/187.6
4, H_r 700 J/g	402.7/215.2	368.2/317.6	409.5/185.4	397.8/237.3	413.5/184.2
4, k_z 0.4	420.3/222.2	384.2/283.1	431.2/197.1	398.2/218.0	426.5/184.5
6, shadow 0.1	772.0/184.9	726.8/185.3	774.7/185.0	798.0/185.0	752.2/184.8

$K_M = 4$, $\tau_{I,M} = 2$ min, $K_{ff} = 0.1$, $\Delta_{lead} = 40$ °C) beats the MPC on case 9, 281.3 min at 190.0 °C, with 771 °C of travel, but the same tuning runs case 4's spar to 240.0 °C with 15.8 min above 195 °C. The strongest classical form of what the MPC knows feeds the max loop the MPC's own prediction, the hottest node over the next 30 min with the air held, from its estimator and its model with the kinetics. That row ties case 9 (291.5 min) and is faster than the MPC on the thick spars (case 4 in 401.7 against 414.8 min, the 50 mm spar in 380.5 against 429.8 min) at peaks up to 2.8 °C higher, but it loses case 4 with $k_z = 0.4$ (197.1 °C, 4.9 min above 195 °C) and the standard recipe (474.7 min).

The MPC's measurable gain is on the long recipes, as shown in Figure 6.5. On the standard recipe it demolds at 456.8 min against 472.5 min. The lagging thermocouple runs on average 0.6 °C behind the reference on the first ramp against 7.1 °C, and 2.3 °C against 7.3 °C on the second, so the hold begins at 225.9 min against 240.8 min, and the feedforward holds the air for 70.0 min under the MPC against 107.8 min. The retuned row is faster there (452.7 min) but leaves 13.9 min outside the band against 0.4 min. The same gain appears with the aged lot (460.8 against 475.2 min), the fifth part (455.8 against 473.2 min) and the shadow of case 6 (752.2 against 772.0 min, with no thermocouple over the selector's limit against 2.9 min). The aged lot reaches both structures through the same observer. With $\alpha_0 = 0.12$ and the load sheet at 0 both lose 7.0 min on the optimized recipe (MPC 295.5 against 288.5 min with the sheet right, Lead/lag + FF 293.3 against 286.3 min), and the fifth part leaves them at 289.5 and 289.7 min.

6.7.7 A Safety Defect in the First Build and Its Price

The first build kept only the max loop above the plan and set its ceilings without the reaction margin. Over 46 worst cases, every parameter of Table 6.2 that the application lets a reader change at both ends of its range plus the shadow of 0.1, a doubled spar tool, the aged lot and the fifth part, each on the optimized recipe and on case 4, the technical review found it worse than Lead/lag + FF on the selector's limit in 12. With $H_r = 700\,\mathrm{J/g}$ on the optimized recipe it peaked at 194.3 °C with a thermocouple more than 0.5 °C over the limit for 4.8 min (3.15 °C at most), where Lead/lag + FF kept 189.4 °C and none. The slack was not the cause. Every one of the 36 plans of that run had zero slack, so the plan believed its hottest thermocouple sat on its ceiling. At 125 min it predicted 178.4 °C five minutes ahead, the thermocouple read 184.7 °C at 130 min and 193.2 °C at 135 min, and the spar's estimated air offset grew from about 0.5 °C before 100 min to 7.7 °C at 130 min, because the filter read heat released in the core as air the part was not receiving. The same MPC told the true heat of reaction peaks at 190.4 °C with no minute over the limit, in 283.3 min, so the error was the model's heat of reaction. With $h_0 = 5$ on case 4 the first build also ran the air 114.5 °C above the reference, where the lead clamp allows 40 °C.

The fix was chosen by a rule written before any candidate ran. A row failed when the MPC spent more than 0.1 min longer than Lead/lag + FF either with a thermocouple 0.5 °C over the limit or with any point above 195 °C, and among the fewest failures the shortest case 9 won, then the fewest minutes above 195 °C. The rule first compared A, ceilings lowered by a sixth of the plan's predicted reaction rise, B, the shipped feedforward candidate above the MPC, and both together. Its result was AB, with 2 rows failing (case 4 with $h_0 = 5$ for 8.2 min at up to 3.9 °C, and $E_1 = 55\,\mathrm{kJ/mol}$

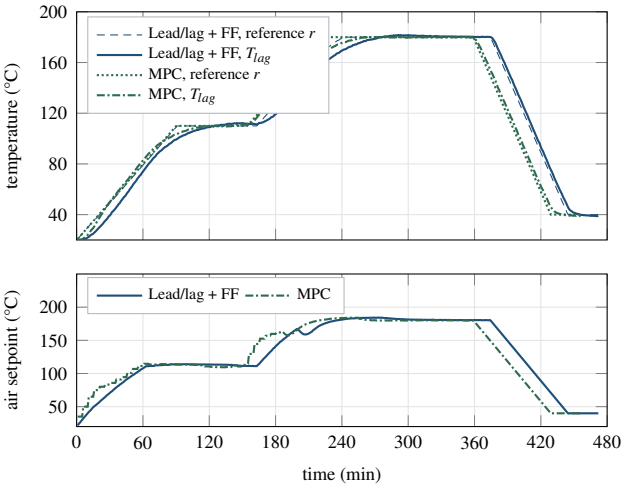


Figure 6.5 The standard recipe under Lead/lag + FF and under the MPC. Top, the recipe's reference on the lagging clock (dashed under Lead/lag + FF, dotted under the MPC) and the lagging thermocouple (solid and dash-dotted). Bottom, the applied air setpoint, the MPC's after the protections above it (solid under Lead/lag + FF, dash-dotted under the MPC). The MPC keeps T_{lag} closer to r on both ramps, so its hold begins 14.9 min sooner.

Table 6.9 Candidates for the Behavior Fix: Rows of 46 Worse Than Lead/Lag + FF, Case 9 and Case 4 Cycles (min), and Minutes Above 195 °C Summed Over the 46 Rows

Candidate	Rows	Case 9	Case 4	Σ min
first build	12	281.0	384.5	47.3
A, reaction margin	6	281.8	399.5	36.9
B, feedforward above	3	291.3	398.3	26.8
AB (original rule)	2	291.3	409.8	15.5
D, lead clamp above	7	281.0	389.0	43.0
AD	4	282.0	403.7	36.4
BD	0	291.3	402.0	16.1
ABD (kept)	0	291.3	414.8	8.8

by a margin) and case 9 at 291.3 min. The rule was then amended once, after these results, to add D, the lag loop's lead clamp kept above the MPC and modeled in its plan, a bound the shipped controller puts on its own command, with its combinations, as listed in Table 6.9. ABD was kept. Its price is case 9 from 281.0 to 291.3 min and case 4 from 384.5 to 414.8 min, against Lead/lag + FF's 392.2 min. The cheaper BD, the feedforward and the clamp without the margin, also fails no row and is 12.8 min faster on case 4, but spends 16.1 min above 195 °C summed over the rows where every structure goes over, against 8.8 min. The first build's late-pressure runs had also left porosity indices of 0.30 % (case 5) and 0.57 % (case 8) against none under Lead/lag + FF, from air run far ahead of the reference under vacuum alone. With the lead clamp above it both end at 0.00 %. Porosity is not in the plan, since no shipped protection acts on it.

6.7.8 Worst Cases Under Every Structure

Table 6.10 runs the worst cases the application offers under every structure, the first build and the MPC as kept, with the two limits that the challenge and the selector judge. After the fix the MPC never puts a thermocouple more than 0.5 °C over the selector's limit and no point above 195 °C in any of these rows, and it has the lowest or an equal peak wherever another structure passes 195 °C. Over the 46 parameter extremes it goes above 195 °C only where every structure does, with $k_z = 0.3$

Table 6.10 Worst Cases Under Every Structure: Cycle (min) / True Peak (°C) / Minutes Any Point Is Above 195 °C / Minutes a Thermocouple Is More Than 0.5 °C Over the Selector's Limit, the First Two Values Above the Last Two in Each Cell. Default Seed, Each Run to Demold. "First MPC" Is the Build Before the Fix of Section 6.7.7

Case	Air only	Lead/lag	Lead/lag + FF	FF retuned	first MPC	MPC
4, 35 mm spar, 3 °C/min	274.8/279.9 26.1/51.7	392.5/243.9 16.1/18.8	392.2/184.8 0/0	367.5/240.0 15.8/21.5	384.5/185.2 0/0	414.8/184.0 0/0
4 at 4 °C/min	261.0/279.3 26.0/55.4	391.3/244.3 16.1/18.9	391.3/184.8 0/0	366.0/239.6 15.8/21.2	373.7/184.7 0/0	405.2/183.9 0/0
50 mm spar, standard	452.7/265.5 39.5/55.1	583.7/209.5 15.8/53.7	558.8/183.0 0/0	536.3/208.9 15.1/43.3	513.0/186.4 0/0	547.3/182.9 0/0
50 mm spar, optimized	291.7/305.0 39.8/17.6	378.7/292.7 26.6/29.3	365.8/231.7 20.7/22.7	335.0/279.7 25.7/26.2	371.3/191.4 0/0	429.8/187.6 0/0
5, pressure at 240 min	438.2/200.5 10.0/19.3	633.0/184.7 0/0.3	641.8/184.5 0/0	626.7/185.1 0/19.8	618.7/184.8 0/0	634.7/184.4 0/0
8, aged, pressure at 170 min	438.3/192.2 0/14.8	574.2/184.8 0/0.1	589.7/184.4 0/0	570.0/185.2 0/3.8	557.5/184.7 0/0.8	582.7/184.4 0/0
6, shadow 0.1	534.3/188.5 0/8.5	766.8/185.0 0/6.2	772.0/184.9 0/2.9	726.8/185.3 0/5.0	723.8/184.8 0/0	752.2/184.8 0/0
optimized, H_r 700 J/g	275.5/219.8 15.8/17.2	307.0/200.2 5.2/8.7	289.5/189.4 0/0	282.7/192.8 0/3.2	284.8/194.3 0/4.8	288.5/189.4 0/0
4, H_r 700 J/g	274.8/363.7 33.2/52.6	394.8/311.7 17.8/20.6	402.7/215.2 11.8/15.7	368.2/317.6 17.4/25.7	389.0/192.9 0/9.4	413.5/184.2 0/0
4, k_z 0.4 W/(m K)	276.3/311.8 37.4/63.3	413.2/284.4 20.9/23.7	420.3/222.2 14.8/18.8	384.2/283.1 20.7/27.7	414.2/200.6 8.0/14.5	426.5/184.5 0/0
optimized, h_0 8	308.5/205.1 12.4/7.6	375.7/194.4 0/11.5	372.0/189.6 0/0	349.0/197.6 4.5/7.3	345.3/190.5 0/0	362.3/189.5 0/0
optimized, spar tool 2×	316.5/192.2 0/2.8	410.3/186.2 0/0	413.7/189.7 0/0	382.8/190.0 0/0	380.2/189.8 0/0	402.7/189.7 0/0
optimized, selector off	275.5/203.4 12.6/15.3	289.8/201.8 30.0/61.0	289.7/189.3 0/0	281.0/190.3 0/1.9	281.0/190.8 0/0	291.3/189.3 0/0
optimized, heater fails at 60 min	275.5/203.5 12.7/15.5	301.2/193.0 0/7.7	289.7/189.4 0/0	281.3/190.0 0/0	281.0/190.4 0/0	290.0/189.4 0/0

on the optimized recipe (197.5 °C for 4.8 min against Lead/lag + FF's 202.1 °C for 8.0 min) and with $E_1 = 55\,\mathrm{kJ/mol}$ on case 4 (210.8 °C for 4.1 min against 217.0 °C for 5.2 min), and with $w_r = 0.45$ on case 4 it peaks at 190.2 °C against 302.0 °C. Its heater saturates for at most 1.2 min (case 4 with a 100 kW heater), and every structure saturates 10.3 to 11.0 min with a 8 MJ/K wall. The vacuum leak leaves a porosity index of 4.30 % under the MPC and 4.23 to 5.41 % under the others, because the pressure is not the MPC's handle.

The MPC loses time wherever the protections above it bind. Case 4 takes 414.8 min against Lead/lag + FF's 392.2 min, and the 50 mm spar on the optimized recipe 429.8 min against 365.8 min, although there the feedforward lets the core reach 231.7 °C with 20.7 min above 195 °C and the MPC holds it at 187.6 °C. Two pre-existing defects were fixed in the same round. The log row at $t = 0$ carried the run's peak as $-\infty$ under every structure, because the referee's peaks rose only at the first step. They now start at the opening temperature, 20 °C, which moved one cell of the challenge's log and no other run. The observer and the feedforward read the run's H_r, w_r, c and k_z instead of the declared values, so a changed resin reached Lead/lag + FF. On case 4 its peak with $k_z = 0.4$ was 196.6 °C before the fix and is 222.2 °C now, and with $H_r = 700\,\mathrm{J/g}$ 202.6 °C and 215.2 °C, while the MPC, whose model was already declared, peaks at 184.5 °C and 184.2 °C.

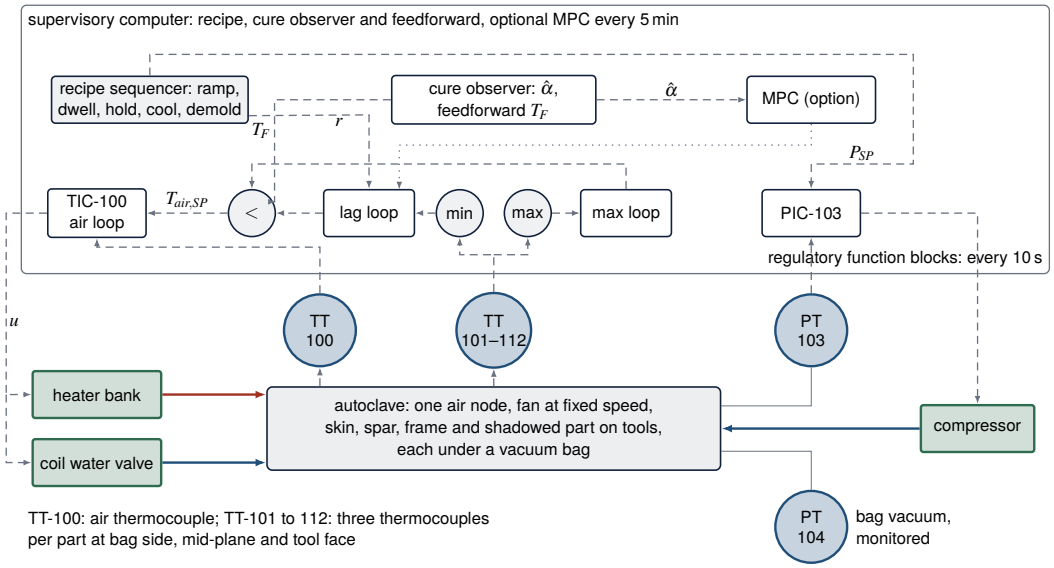


Figure 6.6 The chapter's loops on plant hardware. TIC-100 drives the heater and the coil's water valve on one split-range signal u. Minimum and maximum selectors feed the lag and max loops, and a low selector ($<$) passes the smallest candidate, with the observer's feedforward T_F, as the air setpoint. The sequencer sets the reference r and the pressure step of PIC-103. The optional MPC (dotted) replaces the lag loop. Notice that no loop reads the core of the spar.

6.8 Technology in Practice

What Measures It

The controlled variables of this chapter are thermocouple readings: the coldest keeps the recipe's clock and the hottest meets the limit. A process specification for cure panels states that all its temperatures are part temperatures and asks for thermocouples that monitor and record representative panels [24]. In one survey of a full-size autoclave, fifteen thermocouples on a large mold were recorded at intervals of 1 min, and one of them failed and showed an abnormal fluctuation, so its point was left out of the analysis [19]. The model's twelve part thermocouples, each with a 10 s lag, sit at the bag side, the mid-plane and the tool face, the places a real layup can reach. None sits where a thick core runs hottest.

> **TECHNOLOGY** **Sensor: Thermocouples Laid Up With the Part**
>
> A thermocouple is a junction of two dissimilar metals that makes its own voltage, and the instrument compensates for the second junction where the wires end. It is less accurate than a resistance thermometer but rugged, with a wide span [25]. On a cured panel the junctions sit at the laminate's mid-plane near an edge that is trimmed off after cure, or between the part and its caul plate [24]. The most common failure is an open circuit, or burnout, after which electrical noise can read as a wildly varying temperature, so the input is set to fail high or low, whichever is safer [25]. Behind the maximum selector of this chapter a reading failed high would cut the air, and behind the minimum selector one failed low would stall the recipe. The model's 10 s lag and noise stand for these sensors.

What Moves It

Two kinds of final element act on the load. One split-range signal drives the heater bank and the valve of the cooling-water coil, and a compressor raises the vessel pressure, while the vacuum bag holds the laminate against its tool. A process specification applies the pressure at a stated rate in pounds per

square inch per minute and keeps a minimum vacuum under the bag through the rest of the cure [24]. The model's compressor moves the vessel at no more than 0.5 bar/min, and a reader can halve the heater or tear the bag to see what each final element buys.

> **TECHNOLOGY** **Final Element: Heater and Cooler on One Signal**
>
> An autoclave heats its gas, commonly air or nitrogen, and a fan drives it over the tools and parts for forced convection [19]. When one controller output drives two final elements, they are split-ranged. In exclusive sequencing both are shut at half output, one opens as the output rises above it and the other as the output falls below it, so the two never act together [25]. The air loop of this chapter sequences its heater bank and its cooling-water coil in that way. The gas is a design choice too. In one simulation of a large framed mold, helium in place of air cut the mold's temperature difference by 68.27 % and raised its heating rate by 32.76 % [19].

Where the Controller Runs

Programmable logic controllers (PLCs) hold user programs for logic, timing, counting and PID control, and they serve as local controllers within a supervisory scheme [26]. Supervisory and lower-level loops run with cycle times from milliseconds to minutes [26]. The simulator's control layers run every 10 s, and its predictive controller decides every 5 min. Read as plant hardware, the air loop, the two selectors and the lag and max loops are regulatory function blocks, the recipe sequencer is the sequential layer that switches from ramp to dwell, hold, cool and demold, and the cure observer, the feedforward and the predictive controller belong to a supervisory computer. Figure 6.6 draws these layers on the vessel's hardware.

> **IN PRACTICE** **Running the Cure on the Lagging Thermocouple**
>
> A public process specification for curing the qualification panels of one aerospace prepreg shows how practice writes such a recipe. All temperatures are part temperatures, and the ramps and holds are based on the lagging thermocouple: a hold at 225 °F starts when that thermocouple reaches 215 °F. Before the cure the bag is leak checked, losing no more than 1 inch of mercury over 2 min, the vacuum is held, and the pressure is applied at a stated rate before the heating starts [24]. The chapter's sequencer, with its hold band and its pressure step, follows the same pattern.

What Is Modernizing It

Two developments aim at what the thermocouples miss. Embedded optical fibers have taken real-time spectra of a curing resin and composite, and Bragg gratings written into the same fibers respond to strain [27]. An ultrasonic scanner placed under the tool plate has been designed to watch defects form during cure, first in an oven as a step toward the autoclave, and to validate physics-based models that predict porosity and fiber waviness [28].

6.9 Limitations

The resin is one published fit, and its agreement with the tests it was fitted to is consistency, not validation. Everything else about the load and the vessel is a teaching value, and each part is one-dimensional on a lumped tool in one well-mixed gas, so corners, ply drops, gaps and the flow field are absent. The porosity index rewards pressure before gel and knows no resin flow, so the dwell earns nothing and the optimizer drops it. The observer shares the plant's kinetics exactly and sees no node between thermocouples, and its slab bound is loose while the reaction runs and fails after burn-out. The supervisor's gains were tuned on case 4 and the challenge, the optimizer searched one load on a 28-cell grid, and its optimum is optimal only for that load and seed. The MPC's model has 3 or 5 nodes per part where the plant has 11, takes the declared film, shadow, tool and resin whatever the reader sets, omits the air's overshoot on a large step, the heater's rating, the porosity, the gel point and a pressure step ahead, and linearizes the reaction along its own plan with a margin sized for one

parameter, the heat of reaction, up to a sixth above its declared value. Its tuning rule was judged on case 9 alone, and the gain it shows on the long recipes was not tuned for them.

OPEN CHALLENGE **Measuring the Cure Instead of the Temperature**

Indirect cure sensors such as fluorescence, dielectric loss and ultrasound are often hard to correlate with the chemistry that sets the material's properties. In one fiber-optic Raman record, a temperature program then used in commercial autoclaves ran about two hours past full cure [29]. Fibers stripped to touch the resin are fragile and unlikely to survive the heat and pressure of processing [27]. Progress would be a rugged reading of the cure inside the part, so that a hold could end on the state of the resin, which this chapter's observer can only estimate from thermocouples.

6.10 Conclusion

A mixed autoclave load is controlled through its slowest and its hottest thermocouple, and this simulator logs both beside the true temperature and cure through every part. On the air's clock the spar's tool face enters the hold band 37.5 min after the air and the standard cycle takes 438.3 min. Lead/lag gives the spar its recipe at 463.2 min, with the leading skin held at its limit by the max loop through a low selector. A 35 mm spar ramped at 3 °C/min passes 195 °C while its hottest thermocouple reads 187.3 °C and peaks at 243.9 °C under that max loop, while the observer's feedforward keeps it at 184.8 °C at no cost in cycle time. On the challenge an optimized hold of 10 min at 185 °C demolds in 289.8 min, but Lead/lag alone passes in 288.0 min with less margin, so the feedforward buys margin rather than time. The model is one-dimensional, its resin one fit and its limits teaching values. A fuller study would add a two-dimensional part with its corners, a flow field with radiation, resin flow and consolidation that give a dwell its purpose, an observer whose kinetics differ from the plant's, and a predictive controller whose margin covers the laminate and the film as well as the resin. With the shipped protections above it, the MPC ties the feedforward on the optimized recipe and shortens the standard recipe by 15.7 min, while its first build, planning to the limit on its own model, was less safe in 12 of 46 worst cases. The cases show that a cure recipe has to be run on the part, and that the hottest point of a thick part is one no thermocouple reads.

KEY POINTS

- A mixed load is cured through its slowest and its hottest thermocouple, not through the air.
- On the air's clock the spar enters the hold band 37.5 min after the air, and Lead/lag gives it its recipe at 463.2 min.
- A limit loop on thermocouples acts after the measurement moves, so a 35 mm spar peaks at 243.9 °C under the max loop.
- The observer's feedforward keeps that spar at 184.8 °C, and on the challenge it buys margin rather than time.
- With the shipped protections above it, the MPC ties the feedforward on the optimized recipe and shortens the standard recipe by 15.7 min.

Where else this applies. The freeze-dryer of Chapter 5 has the same shape: its hottest vials limit the shelf while the slowest decide the endpoint.

References

[1] A. C. Loos and G. S. Springer, "Curing of epoxy matrix composites," *Journal of Composite Materials*, vol. 17, no. 2, pp. 135–169, 1983.

[2] T. A. Bogetti and J. W. Gillespie, "Two-dimensional cure simulation of thick thermosetting composites," *Journal of Composite Materials*, vol. 25, no. 3, pp. 239–273, 1991.

[3] J. S. Kim and D. G. Lee, "Development of an autoclave cure cycle with cooling and reheating steps for thick thermoset composite laminates," *Journal of Composite Materials*, vol. 31, no. 22, pp. 2264–2282, 1997.

[4] J. H. Oh and D. G. Lee, "Cure cycle for thick glass/epoxy composite laminates," *Journal of Composite Materials*, vol. 36, no. 1, pp. 19–45, 2002.

[5] A. H. Glattfelder and W. Schaufelberger, *Control Systems with Input and Output Constraints*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2003.

[6] P. R. Ciriscioli and G. S. Springer, *Smart Autoclave Cure of Composites*. Lancaster, PA: Technomic Pub. Co., 1990.

[7] V. K. Pillai, A. N. Beris, and P. S. Dhurjati, "Implementation of model-based optimal temperature profiles for autoclave curing of composites using a knowledge-based system," *Industrial & Engineering Chemistry Research*, vol. 33, no. 10, pp. 2443–2452, 1994.

[8] M. M. Thomas, J. L. Kardos, and B. Joseph, "Shrinking horizon model predictive control applied to autoclave curing of composite laminate materials," in *Proceedings of 1994 American Control Conference*, vol. 1. IEEE, 1994, pp. 505–509.

[9] P. Dufour, Y. Toure, D. J. Michaud, and P. S. Dhurjati, "Optimal trajectory determination and tracking of an autoclave curing process: A model based approach," in *1999 European Control Conference (ECC)*. IEEE, 1999, pp. 2048–2053.

[10] N. G. Pantelelis, "Towards the dynamic optimisation for the cure control of thermoset-matrix composite materials," *Composites Science and Technology*, vol. 65, no. 7-8, pp. 1254–1263, 2005.

[11] M. R. Kamal and S. Sourour, "Kinetics and thermal characterization of thermoset cure," *Polymer Engineering & Science*, vol. 13, no. 1, pp. 59–64, 1973.

[12] J. P. Pascault and R. J. J. Williams, "Glass transition temperature versus conversion relationships for thermosetting polymers," *Journal of Polymer Science Part B: Polymer Physics*, vol. 28, no. 1, pp. 85–95, 1990.

[13] J. M. Castro and C. W. Macosko, "Studies of mold filling and curing in the reaction injection molding process," *AIChE Journal*, vol. 28, no. 2, pp. 250–260, 1982.

[14] A. Shahkarami, "Material characterization for processing: Hexcel 8552, material model development, final project wrap-up," Version 1.0, reviewed by D. Van Ee and approved by A. Poursartip. Public document of the NCAMP 8552 program, National Institute for Aviation Research, Wichita State University, Nov. 2009. [Online]. Available: https://www.wichita.edu/industry_and_defense/NIAR/Research/hexcel-8552/Additional-Documents-3.pdf

[15] N. Rai and R. Pitchumani, "Optimal cure cycles for the fabrication of thermosetting-matrix composites," *Polymer Composites*, vol. 18, no. 4, pp. 566–581, 1997.

[16] G. M. Maistros and I. K. Partridge, "Monitoring autoclave cure in commercial carbon fibre/epoxy composites," *Composites Part B: Engineering*, vol. 29, no. 3, pp. 245–250, 1998.

[17] J. N. E. Kluge, T. S. Lundström, L. G. Westerberg, and T. Nyman, "Modelling heat transfer inside an autoclave: Effect of radiation," *Journal of Reinforced Plastics and Composites*, vol. 35, no. 14, pp. 1126–1142, 2016.

[18] J. W. Hou, "Optimal cure cycle design for autoclave processing of thick composites laminates: A feasibility study," Old Dominion University, for NASA Langley Research Center, Norfolk, VA, Progress Report, NASA Contractor Report NASA-CR-181261, Nov. 1985, NASA Technical Reports Server 19870018301.

[19] G. Zhang, B. Zhang, L. Luo, T. Lin, and X. Xue, "Influence of mold and heat transfer fluid materials on the temperature distribution of large framed molds in autoclave process," *Materials*, vol. 14, no. 15, p. 4311, 2021.

[20] 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.

[21] 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.

[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] P. O. M. Scokaert and J. B. Rawlings, "Feasibility issues in linear model predictive control," *AIChE Journal*, vol. 45, no. 8, pp. 1649–1659, 1999.

[24] National Center for Advanced Materials Performance, "NCAMP process specification: Fabrication of NMS 380 qualification, equivalency, and acceptance test panels, Toray TC380 prepregs," Wichita State University, National Institute for Aviation Research, Wichita, KS, Process Specification NPS 83800, Revision B, Oct. 2024, distribution Statement A, approved for public release.

[25] 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.

[26] 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.

[27] K. H. Wood, T. L. Brown, M. C. Wu, and C. B. Gause, "Fiber optic sensors for cure/health monitoring of composite materials," NASA Langley Research Center, Hampton, VA, Tech. Rep., 2004, NASA Technical Reports Server 20040086070.

[28] T. B. Hudson, F. L. Palmieri, T. B. Abbott, J. P. Seebo, and E. R. Burke, "Design of an automated ultrasonic scanning system for in-situ composite cure monitoring and defect detection," in *SAMPE 2019 Conference*, Charlotte, NC, May 2019, NASA Technical Reports Server 20200002665, report NF1676L-31834.

[29] J. B. Cooper, K. L. Wise, and B. J. Jensen, "Methods used for monitoring cure reactions in real-time in an autoclave: Summary of research, NASA langley research center training grant NGT-1-52124," Old Dominion University, Norfolk, VA, Tech. Rep. ODURF-162661, Oct. 2000, NASA Technical Reports Server 20000108909.

6.11 Exercises

6.1 Concept check · 2 min

In case 4 a 35 mm spar ramped at 3 °C/min on Lead/lag peaks at 243.9 °C, although the max loop has already cut the air setpoint toward its floor. Why does the limit loop fail?

(a) Its gain is too low, and a gain of 6 with an integral time of 5 min would hold the core under the limit

(b) The low selector passes over the max loop while the lag loop still asks for more heat on the ramp

(c) The thermocouple noise hides the rise of the hottest reading until the core has passed the limit

(d) The core heats itself before any thermocouple moves, and a lower air cools only the faces of the spar

6.2 Concept check · 1 min

The standard recipe takes 438.3 min on Air only and 463.2 min on Lead/lag. What makes the Lead/lag cycle longer?

(a) Every dwell and hold counts its time only while the lagging spar thermocouple is in the band

(b) The max loop holds the air below the recipe for the whole cycle, so every ramp runs slower

(c) The lead clamp of 40 °C slows the ramp of the thin skin, which then reaches the hold late

(d) The lag loop's integral winds up while the max loop is selected and overshoots each hold

6.3 Concept check · 2 min

The application's cure target is $\alpha \geq 0.87$, which maps to $T_g \geq 208.7\,^{\circ}\mathrm{C}$, and not the 0.95 often quoted for a fully cured part. Why?

(a) The DiBenedetto relation is not defined above $\alpha = 0.9$, so a higher target has no glass transition

(b) At a 180 °C hold the diffusion term stalls the cure near 0.90, so 0.95 is never reached

(c) The porosity index freezes at gel, and this resin gels at $\alpha = 0.87$

(d) The observer's estimate drifts above 0.87, so a higher target could not be checked during the cure

6.4 Theory and calculation · 6 min

The spar is a 25 mm laminate with $\rho = 1580\,\mathrm{kg/m^3}$, $c = 1100\,\mathrm{J/(kg\,K)}$ and $k_z = 0.6\,\mathrm{W/(m\,K)}$, on an invar tool of $C_t = 231.8\,\mathrm{kJ/K}$ with $A_t = 1.8\,\mathrm{m^2}$. The heat transfer coefficient is $h_0 = 14\,\mathrm{W/(m^2\,K)}$ at one atmosphere, $P_0 = 1.013\,25\,\mathrm{bar}$, the bag adds $R_{bag} = 0.004\,\mathrm{m^2\,K/W}$, and the spar is not shadowed.

(a) Use Equation (6.2) to find h_p and U_s at 7 bar gauge, and the factor by which pressure raises h_p. *Checkpoint:* $h_p = 73.2\,\mathrm{W/(m^2\,K)}$, factor 5.23, $U_s = 56.6\,\mathrm{W/(m^2\,K)}$.

(b) Find the three time scales of the spar: the bag-side film $\rho cL/U_s$, the tool $C_t/(h_p A_t)$ and conduction L^2/a with $a = k_z/\rho c$. *Checkpoint:* 12.8 min, 29.3 min and 30.2 min.

(c) Repeat the film and conduction scales for a 35 mm spar and say which grows faster with thickness. *Checkpoint:* 17.9 min and 59.1 min. Conduction grows with L^2.

6.5 Theory and calculation · 7 min

The laminate has a resin mass fraction $w_r = 0.34$, a heat of reaction $H_r = 600\,\mathrm{J/g}$ of resin and a specific heat $c = 1100\,\mathrm{J/(kg\,K)}$, with $a = 3.45 \times 10^{-7}\,\mathrm{m^2/s}$.

(a) Find the adiabatic rise $w_r H_r/c$ of a laminate that cures fully, and the rise with $H_r = 700\,\mathrm{J/g}$, the top of the parameter's range. State the second as a fraction above the first. *Checkpoint:* 185.5 K and 216.4 K, one sixth higher.

(b) In case 4 on Lead/lag the observer's bound $\widehat{\Delta T}$ of Equation (6.9) for the 35 mm spar reads 656.8 K at 106.0 min. Find the self-heating rate $\hat{s}$ and the cure rate $\dot{\alpha}$ it implies. *Checkpoint:* $\hat{s} = 0.370\,\mathrm{K/s}$, $\dot{\alpha} = 2.0 \times 10^{-3}\,\mathrm{s^{-1}}$.

(c) With $K_{ff} = 0.4$ and $T_{lim} = 185\,^\circ\mathrm{C}$, find the feedforward candidate T_F at that moment, and compare the bound with the true rise of 62.4 K. *Checkpoint:* $T_F = 105\,^\circ\mathrm{C}$ at the 80 K cap. The bound is about ten times the true rise.

6.6 Theory and calculation · 5 min

The first kinetic term of the resin has an activation energy of 64 929.5 J/mol, E_1, and the second one of 133 168.3 J/mol, E_2. Its glass transition follows Equation (6.5) with $T_{g0} = -7\,^\circ\mathrm{C}$, $T_{g\infty} = 250\,^\circ\mathrm{C}$ and $\lambda = 0.78$.

(a) Near 450 K an Arrhenius factor doubles over about $\Delta T = \ln 2\, RT^2/E$. Find ΔT for each term. *Checkpoint:* 18.0 K and 8.8 K.

(b) Find T_g at $\alpha = 0.87$ and at 0.875, and relate them to the demold rule and to the observer's rule for ending the hold. *Checkpoint:* 208.7 °C and 210.2 °C.

6.7 App challenge · 9 min

Load scenario 3, Lead/lag: ramp for the slowest part, and set Control structure to Lead/lag. Run at the default seed and a speed of 1200, which plays 20 min of the cycle each second, so a cycle takes about 23 s. The run pauses itself at DEMOLD and writes a note with the cycle time, the peak, the least-cured point and the minutes outside the hold band.

(a) Run to DEMOLD and read the note. *Checkpoint:* 463 min, 186.1 °C, α 0.896, 4 min.

(b) Set the Maximum selector switch Off, press Reset and run to DEMOLD again. *Checkpoint:* 439 min, 202.2 °C, α 0.906, 124 min.

(c) Compare both runs with the standard recipe on Air only, which takes 438.3 min with 78.4 min outside the band, and say which part sets the ramp and which sets the limit. *Checkpoint:* The spar sets the ramp, the skin the limit.

Accept the cycle and the band within 1 min and the peak within 0.2 °C.

6.8 App challenge · 9 min

Load scenario 4, The exotherm: a thick spar at 3 °C/min, which opens on Lead/lag. Use the default seed and a speed of 1200, and run each case to DEMOLD, where the run pauses itself.

(a) Run as loaded and export the CSV file. Find the first row in which `truth_peak_c` exceeds 195 °C, its time `time_s` and the hottest thermocouple `max_c` in that row, and read the peak in the DEMOLD note. *Checkpoint:* 6110 s (101.8 min), 187.4 °C, peak 243.9 °C.

(b) Set Control structure to Lead/lag + FF, press Reset and run again. Read the peak and the cycle. *Checkpoint:* 184.8 °C in 392 min.

(c) Set the Exotherm feedforward gain to 0.1, press Reset and run again. *Checkpoint:* 239.2 °C.

Accept a temperature within 0.2 °C, a time within 10 s and a cycle within 1 min. Say what the three runs show about acting on the measurement and acting on the reaction.

6.9 App challenge · 8 min

Load the challenge, Cure the load in minimum time. Use the default seed and a speed of 1200, and run each case to DEMOLD, where the run pauses with the verdict.

(a) Run the shipped settings and read the verdict and the cycle. *Checkpoint:* Fails on cycle time and hold band, 438 min with 78 min outside the band.

(b) Press Reveal suggested tuning, then Reset, and run again. *Checkpoint:* Passes in 290 min, peak 189.3 °C, α 0.878.

(c) Set Control structure to Lead/lag, which leaves the Recipe at Optimized and writes Lead/lag's own optimum, then Reset and run again. Read the hold temperature and time the recipe now shows, the verdict and the cycle. *Checkpoint:* 180 °C for 30 min, passes in 301 min, peak 192.6 °C.

Accept a cycle within 1 min and a peak within 0.2 °C. Explain why Lead/lag + FF can run the hotter, shorter hold.

6.10 Application area · 3 min

A shop proposes to switch the autoclave pressure on later in the cycle to save compressor time. Use the measured results of case 5 and case 8, both on the standard recipe and Air only, with the 1 % porosity threshold.

(a) How late may the pressure come on for a fresh lot, and why must it come well before the first part gels? *Checkpoint:* About 175 min (0.81 %), tens of minutes before the first gel.

(b) What does an aged lot ($\alpha_0 = 0.12$) do to that answer? *Checkpoint:* At 170 min the index is 1.19 %. At 150 min it is 0.16 %.

6.11 Application area · 3 min

The suggested configuration of the challenge demolds in 289.8 min against 438.3 min for the standard recipe on the air's clock.

(a) Find the minutes saved per cycle, the fraction saved, and the cycles that fit in 24 h of vessel time with no loading time. *Checkpoint:* 148.5 min, 33.9 %, 3.29 against 4.97 cycles.

(b) Name two features of a real load that the model leaves out and that a thermocouple survey must check before the shorter recipe is adopted. *Checkpoint:* For example corners and ply drops, and the flow field.

6.12 Review the agent · 5 min

An AI assistant reviews the case 9 results and writes: "Lead/lag + FF retuned once ($K_L = 10$, $\tau_{I,L} = 40$ min, $K_M = 4$, $\tau_{I,M} = 2$ min, $K_{ff} = 0.1$, lead 40 °C) demolds the optimized recipe in 281.3 min at 190.0 °C, which is 10 min faster than the MPC and 8.5 min faster than the shipped tuning, with every limit met. Ship this tuning as the default for every load." Find the flaw and say how you would check it.