INTERACTIVE CONTROL ENGINEERING
Principles and Applications

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

CHAPTER 7

Injection Molding Simulator for Bumpless Switchover and Shot-to-Shot Mass Correction

Pages 170–192

Edited by John D. Hedengren

How to cite this chapter. "Injection Molding Simulator for Bumpless Switchover and Shot-to-Shot Mass Correction," 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. 7, pp. 170–192.

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

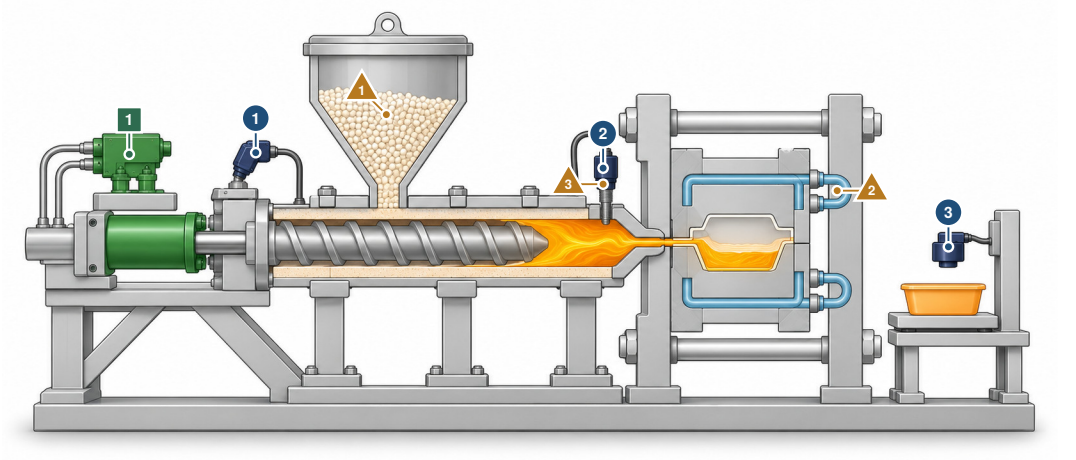
The app for this chapter
https://apmonitor.com/apps/injection-molding/

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

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

CHAPTER 7

Injection Molding Simulator for Bumpless Switchover and Shot-to-Shot Mass Correction



The controller moves 1 servo valve that drives the ram, for both fill and pack

The controller reads 1 screw position and speed 2 melt pressure ahead of the screw 3 weight of each finished part

What disturbs it 1 new resin lot or recycled feed 2 less coolant flow 3 drifting pressure reading

One valve fills the cavity at a set speed and then packs it at a set pressure, and only the scale at the right sees what each part weighs. Illustration, not to scale.

IN PLAIN TERMS One Valve, Two Jobs and a Scale

The main idea. A molding machine pushes melted plastic into a cold steel mold, first at a set speed, then at a set pressure while it shrinks. One valve does both jobs, so the controller hands it from speed to pressure in every shot without a jolt.

The challenge. A part's weight is fixed when the narrow entrance to the mold freezes, and no sensor sees that moment. A thicker batch of plastic can make every part light while speed and pressure stay on target.

Why it matters. Weighing each part and nudging the pressure for the next shot keeps parts right, if the nudges may grow big enough.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history and the first fact in Section 7.2, the second fact and the two try-its in Section 7.5, and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 7.5), Technology in Practice (Section 7.6), Limitations (Section 7.7) and the open challenge

The control question Two loops take turns on one valve inside every shot, and neither of them measures the part. How is the valve handed from the speed loop to the pressure loop without a bump, and how much authority does a correction from the scale need to hold part weight?

Learning objectives

- Explain how tracking the applied valve makes the switchover bumpless.
- Relate the switchover position to a short shot or to flash.
- Derive the loop gain, the stability bound and the variance factor of the shot-to-shot correction.
- Predict part length from ejection temperature, and explain why weight alone does not certify a part.
- Choose the gain and the limits of the correction for a drifting campaign.

Prerequisites Sections 1.4.2 (phases and the handover between controllers), 1.4.5 (run-to-run control, stability and delay), 1.6.2 (velocity form and tracking) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/injection-molding/

See also Volume 2, Chapter 2 for velocity-form PID with bumpless transfer between manual and automatic, Volume 2, Chapter 3 for a loop held at an actuator bound by back-calculation, and Volume 2, Chapter 12 for a drive whose mechanical resonance caps the speed loop.

7.1 Introduction

The loops met in a first control course keep their actuator for as long as the plant runs, and a loop that holds its setpoint is taken to be doing its job. Cyclic manufacturing breaks both habits. The actuator passes from one controller to another at events inside every cycle, and the quality of the product is decided by states that none of the loops measures. Transferring the actuator without a bump then matters as much as tuning either controller [1, 2], and so does correcting the recipe between cycles from a measurement of the product [3].

Injection molding is one of the most versatile processes for mass-producing plastic parts in net shape [4], and it breaks both habits in a cycle of a few seconds. A reciprocating screw pushes melt into a closed, cooled mold under speed control until a switchover point, after which a pressure loop holds the melt while the part shrinks and the gate freezes [5, 6]. Both loops drive the same hydraulic valve. The part weighs what the cavity holds when the gate freezes, and that depends on the pressure reaching the cavity through a narrowing, cooling gate and on the temperature of the melt at that moment [7, 8]. A more viscous resin lot, a drifting pressure transducer or a warmer mold therefore change the part while both loops sit on their setpoints. Recycled resin brings its own variation into the hopper, and the recycled share has room to grow, since about 9 % of the 6300 Mt of plastic waste generated by 2015 had been recycled [9].

The simulator of this chapter tests the belief that a process whose loops all hold their setpoints is making good product. In it the fill speed and the pack pressure can stay on their setpoints while the part comes out 1.4 % light, and the parts can weigh the right amount while their length walks out of tolerance. The broader class has two halves, and both are general. Any system that changes controllers at an event, such as a start-up handing over to production or a normal loop yielding to an override, needs the idle controller to follow the applied actuator so that the transfer produces no bump [1, 2], the handover at a phase change of Section 1.4.2. Any process that repeats a cycle, such as a batch reactor, a semiconductor process step or an additive build, can correct its recipe between cycles from a measurement of the product, and run-to-run, batch-to-batch and iterative learning control are one family of methods for doing so [3] (Sections 1.4.5 and 1.4.6). The additive build of Chapter 10 learns a correction along each layer from the layers before it.

This chapter works with an injection molding simulator in which the switchover point, the handover between the two loops, the hold time and a bounded shot-to-shot correction can each be changed and their consequence read off a scale. Velocity-to-pressure switchover, bumpless transfer by tracking and run-to-run correction are all established practice [4, 10]. The model of one shot has seven states.

It conserves mass exactly, lets a frozen skin close the gate on its own clock, and keeps the true states apart from what the encoder, the two pressure sensors and the scale report. Every equation and parameter is given, so that the model can be rebuilt. Eight scenarios and a thirty-part campaign separate what the inner loops control from what decides the part. Section 7.2 traces the control of injection molding. Section 7.3 sets out the model of one shot, its instruments and its disturbances. Section 7.4 describes the cycle sequencer, the two loops that share one valve and the shot-to-shot correction with its stability bound. Section 7.5 reports the runs of every scenario and of the campaign, Section 7.6 describes the transducers, valves and controllers of real machines, Section 7.7 states the limitations and Section 7.8 sums up. The exercises in Section 7.9 close the chapter.

7.2 Control of the Molding Cycle

Quantitative modeling of the cycle began with the cavity. A 1972 model wrote the continuity, momentum and energy equations for the polymer in the cavity through filling, packing and cooling, with a non-Newtonian melt whose density and viscosity depend on temperature, and it needed the pressure history at the cavity entrance as an input [7]. A finite-element and finite-difference simulation of the filling stage followed in 1980 [11]. For control the cavity had to be joined to the machine that feeds it. A physically based model of the whole cycle divided it into the hydraulic system, the ram and screw, the barrel and the delivery system. It had to account for the elasticity of the melt and its non-Newtonian flow, it included the growing solid skin in the sprue, runner and gate, and it agreed with experiments on high density polyethylene [8]. The simulator of this chapter keeps the same subsystems in lumped form.

By 1987 a review of injection molding process control could sort the controlled variables into those regulated in every phase, those triggered during a specific phase, and cycle-to-cycle control, in which data from earlier cycles are used to predict trends and take corrective action. The same review set the conventional proportional-integral and PID loops beside self-tuning, optimal and statistical process control [10]. The phase-dependent category contains the two loops of the simulator, and the cycle-to-cycle category contains its correction. Cycle-to-cycle and within-cycle adaptive control of the nozzle pressure during packing and holding followed at the end of the 1990s, with the nozzle pressure preferred to the cavity pressure because one transducer then serves different molds, and with dynamics found to be nonlinear and time varying [12]. Learning control was applied to the electrohydraulic injection unit [13], adaptive control to the filling velocity [14], and robust iterative learning control to the molding process [15].

> ⧖ HISTORY **1986 to 1994: Correcting the Next Shot**
>
> A doctoral thesis of 1996 on cavity pressure control reviewed how learning from one shot for the next had entered molding. A study published in 1986 and 1987 joined a hydraulic servo valve to a microprocessor to reduce the variation of the peak cavity pressure, acting from cycle to cycle rather than within the cycle. A 1991 proposal applied learning control to the cavity or nozzle pressure during filling, so that the error at each instant of one cycle corrected the same instant of the next. It worked in simulation but was not shown on a machine. A 1992 auto-adaptive controller concluded that resins of different viscosity could make parts of constant weight. A 1994 study moved the cavity pressure during cooling by the coolant over successive cycles, because the mold's large thermal mass rules out acting within one [16].

The switchover itself has been compared method by method. Seven triggers, from screw position and injection time through machine, nozzle, runner and cavity pressures to a cavity temperature at the end of flow, were set to give similar part weights and then subjected to perturbations of six standard deviations of the standard process. Conventional triggers such as screw position gave lower short-term variation, while the others rejected long-term variation better [5]. A 2005 review organized the field into process setup, machine control, process control and quality control, and it found that online quality control without human intervention had yet to be realized, primarily for

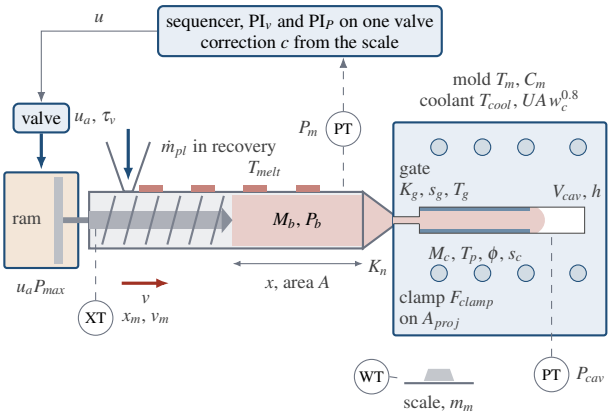


Figure 7.1 One shot of the simulated injection molding machine. The servo valve opening u_a follows its command u through a lag τ_v and sets the ram force, expressed as the melt pressure $u_a P_{max}$, which drives the screw at speed v against the melt pressure P_b of the mass M_b ahead of it, at position x from the front. The melt passes the nozzle K_n and the gate, of open fraction s_g and temperature T_g, into the cavity of volume V_{cav} and wall h, where it fills a fraction ϕ under a frozen skin (dark lines) with mass M_c and temperature T_p. The mold at T_m is cooled through $UA\,w_c^{0.8}$ by coolant at T_{cool}, and the clamp F_{clamp} over A_{proj} sets the flash limit. The encoder XT, the melt pressure transducer PT, the cavity sensor at 80 % of the flow path and the scale WT are all the controller reads. During recovery the rotating screw adds melt at $\dot{m}_{pl}$.

lack of transducers that measure quality in real time and of robust models linking the process to quality [4]. Weighing each part and correcting the next shot answers that gap directly. It belongs to the run-to-run and batch-to-batch family that surveys connect with iterative learning and repetitive control [3], whose exponentially weighted moving average (EWMA) form has a well-characterized stability and sensitivity [17]. The process control, monitoring and optimization of injection molding are now collected in a monograph [6].

Table 7.1 Variables of One Shot at the Nominal Recipe

Symbol	Role	Nominal	Range
v	CV in fill, screw speed	40 mm/s	SP 10–100
P_b	CV in pack, melt pressure	60 MPa	SP 20–120
m	CV per part, part mass	24.10 g	±1 %
u	MV, valve command (both loops)	54, 30 %	0–90 %
c	MV per part, pack correction	0 MPa	$\pm\Delta_{total}$
x	state, screw position	48 to 14.5 mm	0–48
u_a	state, applied valve opening	as u	0–1
M_b	state, melt ahead of the screw	35.6 g	—
M_c, M_cT_p	states, cavity melt, enthalpy	0 to 24.1 g	—
T_g	state, gate temperature	230 °C	—
T_m	state, mold temperature	68 °C	—
$\bar{\lambda}$	DV, resin lot viscosity factor	1	0.6–2.0
σ_r	DV, regrind variation (1σ)	0	0–25 %
w_c	DV, coolant flow	100 %	5–100 %
b, $\dot{b}$	DV, transducer bias, drift	0, 0	±30, 0–5
L	output, length of a 100 mm feature	99.54 mm	±0.12

✦ DID YOU KNOW? Where the Energy of Plastics Goes

An estimated 272 TBtu of energy was used in 2010 to make plastics and rubber products in the United States. Of the savings that the best technologies already available could bring, injection molding holds the largest share, 25.6 TBtu a year or 30 percent [18].

7.3 Model of One Shot

The simulated machine has a 30 mm screw and fills one cavity of 24 cm^3 with a 2 mm wall and 100 cm^2 of projected area, held shut by an 800 kN clamp. The resin is a generic amorphous thermoplastic close to ABS, melted at 230 °C, and the coolant enters the mold at 60 °C. The nominal recipe fills at 40 mm/s, switches to a 60 MPa pack pressure at a screw position of 14.5 mm, holds it for 3.5 s and cools for 4.0 s more, and it makes a 24.107 g part every 10.86 s. The machine, the mold and the instruments are shown in Figure 7.1.

Table 7.1 classifies the variables, with the bias in MPa and the drift in MPa per minute. Seven states are integrated: the screw position x, the applied valve opening u_a, the melt mass ahead of the screw M_b, the melt mass in the cavity M_c and its enthalpy content M_cT_p, the gate temperature T_g and the mold temperature T_m. The screw speed, the pressures, the fill fraction and the flow are algebraic in these states.

7.3.1 Injection Unit

The servo valve sets the force on the ram, expressed as the melt pressure u_aP_{max} it could hold with the screw at rest, and the valve opening follows its command u through a first-order lag. The screw and ram are pushed with tens of kilonewtons (60 MPa on the 30 mm screw is 42 kN), so their inertia is neglected and the screw speed follows the linear force-speed line of the drive, as shown in Equation (7.1).

$$\tau_v \dot{u}_a = u - u_a, \quad v = v_{max}\left(u_a - \frac{P_b}{P_{max}}\right), \quad \dot{x} = -v \tag{7.1}$$

where x is measured from the front stop, a positive speed moves the screw forward, and v is set to zero when the screw rests on its stop. The drive runs at full speed with no load and stands still at its force limit, and the same line pushes the screw backward whenever the melt pressure exceeds the drive force, which is how the screw recovers against the back pressure. All of the drive lag sits in the valve.

The melt ahead of the screw is compressible. Its mass is the state, so mass is conserved exactly, and its pressure follows from a linear pressure-volume-temperature (PVT) law, Equation (7.2).

$$P_b = K\left[1 - \frac{V_d + Ax}{M_b\, v_0(T_{melt})}\right], \quad \dot{M}_b = \dot{m}_{pl}\,\delta_{rec} - \dot{m} \tag{7.2}$$

where $A = \pi D^2/4$ is the screw area, V_d the melt held in the nozzle and screw tip, K the bulk modulus, $\dot{m}$ the mass flow into the mold and δ_{rec} one during recovery and zero otherwise. The specific volume at pressure P is $v_0(T)(1 - P/K)$, and the zero-pressure value v_0 is made of two straight lines, melt and glass, joined smoothly at the glass transition T_{gl}, Equation (7.3).

$$\begin{aligned} v_0(T) = v_{gl} &+ a_S(T - T_{gl}) \\ &+ (a_M - a_S)\, w \ln\left[1 + e^{(T - T_{gl})/w}\right] \end{aligned} \tag{7.3}$$

with $w = 3\,\mathrm{K}$. The melt at 230 °C has $v_0 = 1.040\,\mathrm{cm^3/g}$ and a density of $965\,\mathrm{kg/m^3}$ at the 6 MPa back pressure, so the $36.9\,\mathrm{cm^3}$ ahead of the screw at the shot size holds 35.6 g, and the solid at 23 °C has $v_0 = 0.950\,\mathrm{cm^3/g}$.

7.3.2 Flow Path and Frozen Skin

Melt passes from the barrel through the hot nozzle, the gate and the filled length of the cavity. Each section is treated as a power-law slit of index n, and the pressure drops of the three add, Equation (7.4).

$$P_b - P_c = \lambda\left[K_n a_n + \frac{K_g a_g}{s_g^{2n+1}} + \phi_f\,\frac{K_c a_c}{s_c^{2n+1}}\right]|Q|^{n-1}Q \tag{7.4}$$

where Q is the volumetric flow, P_c the compression pressure of a full cavity, λ the lot factor of the resin in this shot, $\phi_f = \min(1, \phi)$ the filled share of the flow length, and s_g and s_c the open fractions of the gate and of the filled cavity. The viscosity shift is Arrhenius about 230 °C, $a(T) = \exp[(E/R)(1/(T + 273.15) - 1/503.15)]$, taken at the melt temperature for the nozzle and at a core temperature $(T_{melt} + T)/2$ for the gate and the cavity, so a cooling section also grows stiffer.

Melt touching the steel freezes into a skin that grows inward. If the skin temperature runs linearly from the mold temperature to the no-flow temperature around a core near the melt temperature, a section of lumped temperature T keeps the open fraction of Equation (7.5).

$$s(T) = \mathrm{clamp}\left(1 - \frac{T_{melt} - T}{T_{melt} - T_s},\, 0,\, 1\right) \tag{7.5}$$

where $T_s = (T_m + T_{nf})/2$ is the mean skin temperature, and $s = 0$ once $T \le T_s$. The skin grows slowly while hot melt keeps arriving and quickly once it stops. No flow is computed while the nozzle is shut, while s_g or s_c is below 10^{-3}, or once the gate has been declared frozen. The mass flow is taken at the density of the melt it leaves, $\dot{m} = QM_b/V_b$ for $Q \ge 0$ with $V_b = V_d + Ax$, and $\dot{m} = QM_c/(\phi_f V_{cav})$ for backflow. The skin is what lets an early switchover leave a part short, because without it a slowing front would still reach the end of the cavity.

7.3.3 Cavity, Gate and Mold

The fill fraction is the zero-pressure volume of the cavity melt over the cavity volume, and past one the melt is compressed, Equation (7.6).

$$\phi = \frac{M_c\, v_0(T_p)}{V_{cav}}, \quad P_c = \begin{cases} K(1 - 1/\phi) & \phi > 1 \\ 0 & \phi \le 1 \end{cases} \tag{7.6}$$

The pressure falls linearly from the gate to the flow front. With ΔP_f the drop across the filled cavity, the last term of Equation (7.4), the mean cavity pressure is $\bar{P} = P_c + \Delta P_f/2$, and a sensor at 80 % of

the flow path reads $P_s = P_c + \Delta P_f(1 - 0.8/\phi_f)$ once the front has passed it and zero before. The three thermal states follow Equations (7.7) to (7.9).

$$\dot{M}_c = \dot{m}, \qquad \frac{\mathrm{d}(M_c T_p)}{\mathrm{d}t} = \dot{m} T_{in} - \frac{M_c(T_p - T_m)}{\tau_p} \tag{7.7}$$

$$\dot{T}_g = -\frac{T_g - T_m}{\tau_g} + \frac{|Q|}{V_{gh}}(T_{in} - T_g) \tag{7.8}$$

$$C_m \dot{T}_m = \frac{c_p M_c (T_p - T_m)}{\tau_p} - UA\, w_c^{0.8}(T_m - T_{cool}) - k_r(T_m - T_r) \tag{7.9}$$

where T_{in} is T_{melt} for inflow and T_p for backflow, and $\tau_p = h^2/(\pi^2\alpha) = 4.05\,\mathrm{s}$ is the slowest conduction mode of a slab of thickness h. The gate cools toward the mold with time constant τ_g and is renewed by the melt passing it, V_{gh} being the volume whose passage replaces it. The coolant flow fraction w_c, bounded to 0.02 to 2, scales the heat transfer as a turbulent film does, and $k_r = 3\,\mathrm{W/K}$ is a loss to the room at $T_r = 23\,^\circ\mathrm{C}$. The mold time constant $C_m/(UA + k_r)$ is 47.6 s at full coolant flow and 240 s at a tenth of it. The gate is reset to the melt temperature when the next shot begins, because the frozen gate leaves with the part.

7.3.4 Part Mass, Length and Defects

The gate is declared frozen once it would pass less than $0.01\,\mathrm{cm^3/s}$ under a 50 MPa drop, that is once $\lambda K_g a_g s_g^{-(2n+1)}$ exceeds $50 \times 10^6/(10^{-8})^n = 3.2 \times 10^{10}$ in the units of K_g. The test runs between the switchover and the ejection, and the latch holds until the next shot. Nothing enters or leaves the cavity afterward, so the mass weighed at ejection is the mass at gate freeze. In the nominal recipe the gate freezes 3.42 s after injection starts, 2.56 s after the switchover. The length of a nominally 100 mm feature is an illustrative proxy, Equation (7.10).

$$\begin{aligned} L &= 100\,\mathrm{mm}\left[1 - f s_V - k_{ej}(T_{ej} - T_{ej,0})/100\right] \\ s_V &= 1 - m\, v_0(T_r)/V_{cav} \end{aligned} \tag{7.10}$$

where m is the part mass, s_V its volume shrinkage from the cavity to the room, f the share of that shrinkage showing in the length while the mold restrains the part, and T_{ej} the part temperature at ejection, so that a part ejected warm finishes shrinking free of the mold. The nominal part shrinks 4.5 % in volume and measures 99.545 mm against $99.54 \pm 0.12\,\mathrm{mm}$.

A part is short if its fill fraction never reaches 0.999, light or heavy if its weighed mass is outside ± 1 % of the target, and labeled SIZE if its length is out of tolerance. It is flashed if the mean cavity pressure times the projected area ever exceeds the clamp force, $\bar{P} > F_{clamp}/A_{proj} = 80\,\mathrm{MPa}$, a flag that adds no mass. A shot whose valve command sits at its limit for more than 20 ms counts, like a flash, as a pressure excursion. The energy of a part is the drive work divided by an efficiency of 0.6, plus 720 kJ/kg for the melt plasticized and a 1200 W base load over the cycle, 32.8 kJ in the nominal recipe.

7.3.5 Instruments and Disturbances

The controller never reads a true state. The encoder rounds the screw position to 1 µm, and the measured speed v_m is its first difference over the 2 ms control period passed through an exact first-order filter of 8 ms. The transducer reading P_m follows the melt pressure through a 5 ms first-order lag and adds a bias b, a drift accumulated at the rate $\dot{b}$ from the start of the run, and white noise of 0.2 MPa. The cavity sensor adds 0.1 MPa of noise to P_s, and the scale adds 0.01 g to each part (1σ). The lot factor of every shot after the first is $\lambda_k = \bar{\lambda}\exp(\sigma_r z_k)$, log-normal and uncorrelated from shot to shot, where $\bar{\lambda}$ is the lot setting and σ_r the regrind variation. The deviates z_k and the scale errors are

Table 7.2 Machine and Material Parameters. A Marks a Value Adjustable in the Model Dialog and H a Hidden One

Symbol	Meaning	Value	Class
P_{max}	drive force as melt pressure	200 MPa	A
P_{lim}	injection pressure limit	180 MPa	A
v_{max}	screw speed at full drive, no load	400 mm/s	A
τ_v	valve and drive lag	20 ms	A
D	screw diameter	30 mm	A
V_d	melt in nozzle and screw tip	3 cm^3	H
x_{shot}	shot size, screw position	48 mm	A
$\dot{m}_{pl}$	plasticizing rate	8 g/s	A
P_{back}	back pressure	6 MPa	A
t_{open}	mold open, eject and close	2.5 s	A
T_{melt}	melt temperature	230 °C	A
T_{nf}	no-flow temperature	150 °C	A
T_{gl}, v_{gl}	glass transition, its volume	105 °C, 0.970 cm^3/g	H
a_M, a_S	melt, solid expansion	0.56, 0.24 mm^3/(g K)	H
K	melt bulk modulus	1.6 GPa	A
n	power-law index	0.35	A
E/R	viscosity activation temperature	6500 K	A
c_p	specific heat	2000 J/(kg K)	H
α	thermal diffusivity	0.10 mm^2/s	A

Table 7.3 Mold, Instrument, Quality and Energy Parameters

Symbol	Meaning	Value	Class
V_{cav}, h	cavity volume, wall thickness	24 cm^3, 2.0 mm	A
τ_g	gate thermal time constant	3.0 s	A
V_{gh}	melt volume that renews the gate	2 cm^3	H
K_n, K_g, K_c	consistencies, Pa $(s/m^3)^n$	$(3.4, 8, 15)10^8$	H
F_{clamp}, A_{proj}	clamp force, projected area	800 kN, 100 cm^2	A
C_m	heat capacity of the mold inserts	3.0 kJ/K	A
UA	coolant heat transfer at full flow	60 W/K	A
T_{cool}	coolant temperature	60 °C	A
$T_{m,0}$	mold temperature at the start	67.55 °C	A
—	transducer noise, lag	0.2 MPa, 5 ms	A
—	cavity sensor noise	0.1 MPa	A
—	encoder step, speed filter	1 µm, 8 ms	A
—	scale noise	0.01 g	A
—	mass tolerance	±1 %	A
L_{spec}	length specification	99.54 ± 0.12 mm	A
f	volume shrinkage seen in length	0.10	A
$T_{ej,0}$	ejection temperature reference	91.3 °C	A
k_{ej}	extra shrinkage of a warm part	0.015 %/K	A
—	drive efficiency	0.6	H
—	plasticizing energy	720 kJ/kg	A
—	base load	1200 W	A

drawn once from the scenario seed, so every configuration of a scenario molds the same material and reads the scale with the same errors.

The parameters are listed in Tables 7.2 and 7.3. Fixed in the code are the viscosity reference of 230 °C, the room at 23 °C, the 3 K join width, the sensor position, the room loss k_r and the gate-freeze threshold. Every scenario starts with the screw at the shot size, the melt ahead of it at the back

Table 7.4 Recipe and Controller Defaults, With Slider Ranges

Symbol	Meaning	Default	Range
v_{SP}	fill speed setpoint	40 mm/s	10–100
x_{sw}	switchover position	14.5 mm	5–40
P_{sw}, t_{sw}	switchover on P_{cav}, on time	15 MPa, 0.84 s	1–60, 0.2–3
t_{hold}, t_{cool}	hold, cooling time	3.5 s, 4.0 s	0.5–10, 1–20
P_{pack}	pack pressure setpoint	60 MPa	20–120
$K_{c,v}$, $\tau_{I,v}$	speed loop	1.0 %/(mm/s), 20 ms	0.05–4, 5–500
$K_{c,p}$, $\tau_{I,p}$	pressure loop	0.5 %/MPa, 50 ms	0.02–3, 5–1000
u_{fill}, u_{pack}	manual openings	40, 30 %	0–100
m_{SP}	part mass target	24.10 g	23–25.5
G	correction gain	40 MPa/g	0–200
Δ_{shot}, Δ_{total}	step, total limit	4, 15 MPa	0.5–20, 1–50
n_{start}	parts before correcting	3	0–10

pressure, an empty cavity, the gate at the melt temperature and the mold at $T_{m,0}$, which is the periodic operating temperature of the nominal recipe unless the scenario sets another.

7.3.6 Assumptions

The machine feeds one cavity straight from a hot nozzle, with no runner and a check ring that never leaks, so the shot size repeats exactly. Multi-cavity molds and runners add a balance problem that this model does not have.

The screw and ram have no inertia and the hydraulics no resonance. A machine with a soft hydraulic column would put a resonance inside the speed loop, as the coupling of the servo axis in Volume 2, Chapter 12 caps its speed gain.

Temperatures are lumped, one each for the part, the gate and the mold, and the skin of Equation (7.5) is a geometric construction on those temperatures rather than a solution of the conduction problem, which the cavity models of the literature solve [7]. The melt front has no temperature of its own, so a slow fill thickens the skin only through the mean part temperature. Real short shots also come from a front that stalls in a thin rib far from the gate.

The PVT law is linear in pressure and piecewise linear in temperature, and the viscosity is a power law with an Arrhenius shift. Near the glass transition a Tait equation and a WLF shift describe the melt better, while at the pressures and temperatures of this recipe the simple forms carry the trends.

Flash, the short-shot criterion and the length tolerance are illustrative thresholds, not measurements of a real mold, and the length proxy ignores anisotropic shrinkage and warpage. The scale weighs every part without delay, whereas many cells weigh a sample or report several shots late.

7.4 Sequence and Control as Implemented

The machine controller executes every 2 ms, and the seven states are integrated with a classical fourth-order Runge-Kutta step of 1 ms. Halving the step changes the part mass by 0.000 65 % and the peak mean cavity pressure by 0.08 %, inside the limits of 0.01 and 0.5 % that one of the application's own checks enforces. The historian logs at 50 Hz, and the playback speeds of 0.25 to 20 simulated seconds per wall second change only the display, as for every app of the framework (Section 1.6.1). The structure is drawn in Figure 7.2, and the recipe and controller defaults are listed in Table 7.4.

7.4.1 The Cycle Sequencer

A shot starts with the mold closed and the screw at the shot size. The fill ends at the switchover, triggered by the measured screw position reaching x_{sw}, by the cavity sensor reaching P_{sw}, or by the fill lasting t_{sw}, with a 3 s fill-time limit behind all three. The pack lasts the hold time. Recovery then rotates the screw, adding melt at 8 g/s, until the measured position is back at the shot size, while the drive holds $u = P_{back}/P_{max} = 3\,\%$. The cooling timer runs from the end of the hold, so recovery overlaps cooling. The mold then opens for t_{open}, the part drops onto the scale, the nozzle is shut, and

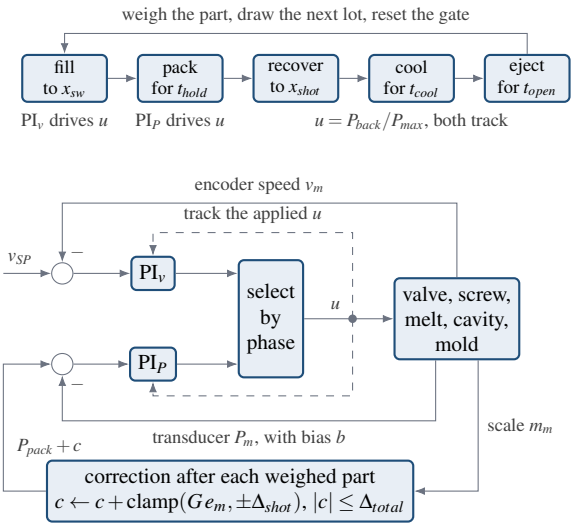


Figure 7.2 Sequencer and control structure. The fill ends at the switchover and the pack after the hold time, recovery refills the screw to the shot size, and the cooling timer runs from the end of the hold. The speed loop PI$_v$ drives the valve during the fill and the pressure loop PI$_P$ during the pack, while the loop not in charge tracks the applied valve command (dashed). Outside fill and pack the valve holds the back pressure. After each weighed part the correction moves the pack pressure setpoint by a bounded step computed from the mass error $e_m = m_{SP} - m_m$.

the next shot starts with a new lot factor and the gate at the melt temperature. The nominal cycle is 0.86 s of fill, 3.5 s of hold, 4.0 s of cooling and 2.5 s open, 10.86 s in all. In the terms of Section 1.4.2 the fill ends on a measured event, while the pack, the cooling and the mold opening end on a clock.

7.4.2 Two Loops on One Valve

Both loops use the shared PID of the framework, written in the velocity (incremental) form with derivative action on the measurement, Equation (7.11), which is Equations (1.14) to (1.17) of Section 1.6.2 with the action $a = +1$ and no output bias.

$$
\begin{aligned}
&e_k = r_k - y_k, \qquad P_k = K_c e_k, \qquad D_k = -K_c \tau_D \dot{y}_{f,k} \\
&S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\
&c_k = S_k + I_{k-1}, \qquad u_k = \text{clamp}\,(c_k,\ 0,\ u_{max}) \\
&I_k = I_{k-1} + \Delta t \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{\tau_t} \right]
\end{aligned}
\tag{7.11}
$$

where r is the setpoint, y the measurement, y_f its optionally filtered value, c_k the output the algorithm demands (written v_k in Chapter 1) and $\tau_t = \tau_I$ the tracking time. The running sum S_k is never clamped, the clamp acts on the emitted output only and the back-calculation on the integral alone, as Section 1.6.2 explains. The integral used at period k was accumulated through period $k-1$, and both integrals start from the 3 % back-pressure opening. Neither loop uses derivative action ($\tau_D = 0$) or a rate limit, the sample time is $\Delta t = 2$ ms, and the output is limited to $u_{max} = 100 \min(1, P_{lim}/P_{max}) = 90$ %, the machine's 180 MPa injection pressure limit. Both loops are reverse acting in the ISA convention of Section 1.6.2, because opening the valve raises both the screw speed and the melt pressure.

The speed loop acts on the filtered encoder speed, and the pressure loop acts on the transducer reading toward the setpoint $P_{pack} + c$. During the fill the speed loop drives the valve and the pressure

loop runs in its tracking mode, on its own setpoint and measurement, with the integral chosen at every period so that its output equals the applied command, $I_k = u_k - S_k$. During the pack the roles reverse, with the idle speed loop tracking at zero error, and outside the two phases both loops track the 3 % back-pressure opening. At the switchover the pressure loop therefore starts from exactly the valve the speed loop left, and its first output differs from it only by its own proportional increment over one period, $-K_{c,p}(P_{m,k} - P_{m,k-1})$. In manual mode the operator sets one opening for the fill and one for the pack, each held below the 90 % limit, the sequencer runs as before, and both loops track, so a return to automatic is bumpless as well [1], by the same tracking as the manual and automatic transfer of the cruise control loop in Volume 2, Chapter 2. In automatic mode the two manual openings follow the last fill and pack commands, so a switch to manual mid-run holds the valve where the loops left it.

Pause. Try Exercise 7.2 before reading on.

7.4.3 Shot-to-Shot Correction

After each part is weighed, in automatic mode only and once more than n_{start} parts have been made, the pack pressure correction moves by a bounded step, Equation (7.12).

$$\begin{aligned} \Delta c_k &= \mathrm{clamp}\big(G(m_{SP} - m_{m,k}),\, \pm\Delta_{shot}\big) \\ c_{k+1} &= \mathrm{clamp}\big(c_k + \Delta c_k,\, \pm\Delta_{total}\big) \end{aligned} \tag{7.12}$$

where $m_{m,k}$ is the weighed mass of part k, and the setpoint in use is $\mathrm{clamp}(P_{pack} + c, 0, P_{lim})$. The correction is an integral action in the shot index with one shot of delay, and it reads only the scale. Near 60 MPa the part gains $\partial m/\partial P = 0.0137\,\mathrm{g/MPa}$ of pack pressure (23.974 g at 50 MPa, 24.247 g at 70 MPa), so about 18 MPa moves it by 1 %. Inside its limits, with $e_k = m_{m,k} - m_{SP}$ the weighed deviation, d_k the disturbance including the scale error, and z_k the mass change the correction has made, the loop is the recursion of Equation (7.13).

$$e_k = d_k + z_k, \qquad z_{k+1} = z_k - k\,e_k, \qquad k = G\,\frac{\partial m}{\partial P} \tag{7.13}$$

The error from a persistent disturbance shrinks by the factor $1 - k$ per part, so the correction is monotone for $k < 1$, alternates for $1 < k < 2$ and diverges beyond. The correction is the integral form of an EWMA controller whose weight divided by its model gain is G, and $0 < k < 2$ is the stability condition established for that controller [17]. In the notation of Section 1.4.5 the loop gain k is the product $\lambda\xi$ of Equation (1.8), and the scale reports each part before the next shot, the case without metrology delay. The default gain gives $k = 0.55$, and the bound $k = 2$ lies at 146 MPa/g, inside the 200 MPa/g reach of the gain slider. For a disturbance drawn afresh each shot with variance σ^2, z_k depends only on earlier draws, and the stationary variance of the error is $2\sigma^2/(2-k)$. An integral correction therefore widens the spread of an uncorrelated disturbance at any gain, by 8 % in standard deviation at $k = 0.27$ and by 17 % at $k = 0.55$.

Pause. Try Exercise 7.3 before reading on.

The scorecard counts the parts in tolerance after five start-up parts, the rejects by cause, the pressure excursions, the cycle time, the energy per good part, the length of the last part, the spread of the weighed masses and the correction in use. The challenge is met when at least 29 of parts 6 to 35 are within ± 1 % of 24.10 g with no flash and no shot at the pressure limit. It is judged on the case as given, against the fixed 24.10 g specification, so changing the mass target, a model parameter or a disturbance setting voids the attempt.

7.5 Results

The application ships nine scenarios, each with a briefing and a completion rule (Section 1.6.4): 1, Nominal recipe (one cycle and its handover); 2, A thicker resin lot (a change the speed loop hides); 3, Recycled feed (noise that a correction amplifies); 4, Early switchover (a front that freezes); 5, Late

Table 7.5 Results Measured on the Engine

Case and condition	Result
1, switchover	96.0 % full at 0.860 s, valve 54.1 %
1, handover step	≤0.18 % in 4 shots, ≤0.44 % in 20
1, peaks	melt 87.4 MPa, mean cavity 47.5 MPa
1, part	24.107 g, 99.545 mm, 10.86 s, 32.8 kJ
1, hold 1 / 1.5 / ≥2.5 s	24.063 / 24.093 / 24.107 g
2, new lot, automatic	valve 84.7 %, 147.9 MPa, 23.755–23.769 g
2, manual, valve 40 %	short, 65.0 % at the time limit, 15.6 g
2, correction to 15 / 30 MPa	near 23.96 g / within 0.03 g
3, off / 40 / 120 MPa/g	weighed spread 0.065 / 0.073 / 0.080 g
4, switch at 30 mm	at 50.5 %, short at 98.1 %, 23.412 g
4, 20 mm / 80 MPa at 30 mm	24.038 g / 24.383 g heavy
5, switch at 11.5 mm	103.6 %, 94.6 MPa, flash, 24.560 g
5, cavity pressure 15 MPa	100.4 %, 56.7 MPa, 24.150 g
5, pressure loop 0.1, 300 ms	77.1 MPa, 24.574 g heavy
6, coolant 10 %, part 10 / 20	99.418 / 99.330 mm, 24.097–24.107 g
6, cooling 6.5 s from part 10	99.490–99.556 mm, 13.36 s, 35.7 kJ
7, bias 22 MPa	reads 60.0, true 38.0 MPa, 23.821 g
7, correction to 15 / 25 MPa	0.087 g light / within 0.03 g
8, cold mold, part 1 / 7	99.976 / 99.647 mm, 24.125 / 24.113 g
8, cooling time 3 s	four oversize parts instead of six

switchover (a screw that rams a full cavity); 6, Cooling restriction (right mass, wrong length); 7, Pressure sensor bias (an outer loop on the product); 8, Cold mold (start-up parts); and the challenge, a 30-shot production campaign that tests the authority of a correction. Cases 1 to 3 and 6 to 8 are experiments, cases 4 and 5 demonstrations. Their outcomes are listed in Table 7.5. All numbers here come from runs of the application's engine at each scenario's own seed, with any setting named for a case changed before the run starts. Part masses are true masses unless called weighed, and a value quoted for a shot that repeats is taken from the third shot.

7.5.1 One Cycle and Its Handover

From the second shot of the nominal recipe on, the switchover comes 0.860 s into the shot with the cavity 96.0 % full and the valve at 54.1 %. The melt ahead of the screw peaks at 87.4 MPa and the mean cavity pressure at 47.5 MPa, well under the 80 MPa flash limit. The first output of the pressure loop differs from the last output of the speed loop by at most 0.18 % in the first four shots and 0.44 % in twenty, one proportional increment on the change in the transducer reading over 2 ms. A pressure loop resuming from its own last output, the 29.8 % that held 60 MPa at the end of the previous hold, would instead move the valve by 24.3 % in one period. After the handover the melt pressure settles within ±2 MPa of its setpoint in 0.27 s. A gate-freeze study run on the simulator traces the curve a molder draws to set the hold time. Holding the pressure for 1 s makes a 24.063 g part and 1.5 s a 24.093 g part, while 2.5 s or more makes the 24.107 g part, ejected at 91.4 °C and 99.545 mm long, every 10.86 s for 32.8 kJ.

> ✦ DID YOU KNOW? **Holding Longer Adds Nothing**
>
> From the simulator: in the nominal recipe a hold of 1 s makes a 24.063 g part and a hold of 1.5 s a 24.093 g part, but every hold of 2.5 s or more makes the same 24.107 g part (Table 7.5). Once the gate has frozen, pressure on the screw no longer reaches the part.

7.5.2 The Resin Changes

In case 2 a lot 75 % more viscous reaches the hopper before the fourth shot. To hold 40 mm/s the speed loop opens the valve from 54.1 to 84.7 %, close to its 90 % limit, and the melt ahead of the screw peaks at 147.9 MPa instead of 87.4 MPa. The fill takes 0.874 to 0.876 s instead of 0.860 s and

looks the same. The part does not. It weighs 23.755 to 23.769 g, 1.37 to 1.43 % below the target, because the thicker melt loses more pressure on its way to the cavity during the pack. The cavity sensor peaks at 20.5 to 22.0 MPa instead of 41.4 MPa, and the gate freezes at 3.09 s instead of 3.42 s. The fill at the switchover also falls from 96.0 to 93.5 or 93.7 %, because the melt ahead of the screw is compressed further at the same screw position. With the valve fixed at 40 % in manual mode from the start, the same melt has filled only 65.0 % of the cavity when the 3 s fill-time limit forces the switchover, and every part of the new lot is short, 15.6 g on the fourth shot. Switched to manual after the first fill instead, the valve holds the 53.9 % the speed loop last applied, and the melt stops at 93.9 % of the cavity, a 22.3 g short shot. The speed loop hides the lot from the fill but not from the scale.

Pause. Try Exercise 7.1 before reading on.

The correction at its defaults raises the pack setpoint 4 MPa per part, brings the parts back inside the tolerance from the sixth, and stops at its 15 MPa limit with the parts near 23.96 g. With a 30 MPa limit it settles near 25.4 MPa, and from the tenth part the weighed masses stay within 0.03 g of the target.

TRY IT IN THE APP 2 · A Thicker Resin Lot

Load the scenario and press Run. Watch the injection valve trend as the new lot reaches the fourth shot: the speed loop opens the valve from 54.1 to 84.7 %, and the fill takes 0.874 to 0.876 s, hardly longer than before. Then watch the finished parts: from the fourth shot every part comes out about 1.4 % light.

Recycled feed is the opposite disturbance. In case 3 the regrind variation of 15 % spreads the lot factor from 0.73 to 1.40 over 24 parts and the peak melt pressure from 71.4 to 119.7 MPa, and the weighed masses spread by 0.065 g (1σ), all inside the tolerance. The correction at its default gain widens the spread to 0.073 g, and at 120 MPa/g with a 10 MPa step to 0.080 g, an ordering that holds at seven other seeds as well. Over 195 parts with the step and total limits widened to 20 and 50 MPa, the spread without correction is 0.0655 g, and it grows by 2, 6, 15 and 43 % at gains of 10, 20, 40 and 80 MPa/g, against 4, 8, 17 and 49 % from the recursion of Equation (7.13). A correction that reacts to a disturbance already gone can only add variance, and the more gain the more it adds.

7.5.3 Where to Switch Over

Moving the switchover to 30 mm hands the valve to the pressure loop at 50.5 % fill, 0.464 s into the shot, with a pack setpoint of 60 MPa below the 87.4 MPa the speed loop needs to finish the fill at speed. Under pressure control the front slows as the filled length grows, the skin gains on it, and the cavity reaches only 98.1 %. The part is short at 23.412 g, and the mean cavity pressure never passes 21.7 MPa. At 20 mm the switchover comes at 79.6 % and the part is good at 24.038 g. An 80 MPa pack at 30 mm fills the cavity instead but packs the part 1.2 % heavy (24.383 g), because the setpoint that finishes the fill also packs the part.

Moving it to 11.5 mm keeps the speed loop pushing at 40 mm/s, or 28 cm^3/s, into a cavity that is already full. The cavity is 103.6 % full at the switchover, the melt is compressed within milliseconds, and the mean cavity pressure peaks at 94.6 MPa, above the 80 MPa the clamp can hold. Every part flashes and is 1.91 % heavy (24.560 g). Switching over when the cavity sensor reads 15 MPa ends the fill at 100.4 % with a 56.7 MPa peak and a part in tolerance (24.150 g), which is the melt itself reporting that the cavity is full. A sluggish pressure loop does similar damage at the nominal 14.5 mm. With 0.1 % per MPa and 300 ms the melt pressure overshoots to 100.8 MPa after the handover, the mean cavity pressure peaks at 77.1 MPa, and the part is 1.97 % heavy (24.574 g). The switchover is correct only as a pair, a moment and a loop that takes the valve cleanly at that moment.

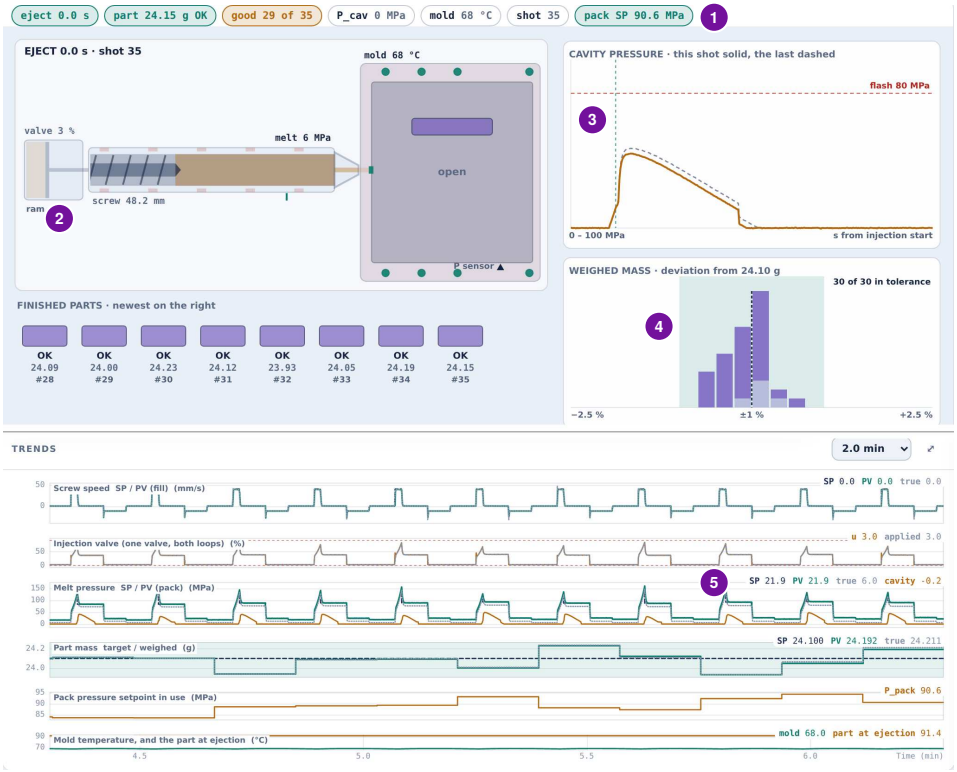


Figure 7.3 The injection molding application at the end of the challenge campaign, run with the suggested correction of Table 7.6 (40 MPa per g, 5 MPa per part, 40 MPa in total). The figure is a composite of one screen: the badges, the plant view and its panels above the trends. (1) Shot 35 has just been ejected at a weighed 24.15 g, and the badge counts 29 good parts of 35 because the first six, from the cold mold, are oversize. (2) The screw has recovered to 48.2 mm, and the valve holds 3 % against the back pressure. (3) The scope draws this shot's cavity pressure over the last one, both under the 80 MPa flash line. (4) The histogram counts 30 of 30 campaign parts in tolerance, and parts 28 to 35 are all OK. (5) The trends over the last 2.0 min show eleven shots of screw speed, valve and melt pressure, with the drifting transducer reading 21.9 MPa against a true 6.0 MPa, the weighed mass stepping about the 24.10 g target, the pack setpoint moved by the correction to the 90.6 MPa in force for the last shot, and the mold at 68 °C. The header and the scorecard of the full screen, cut here, show $t = 6.30$ min and report the challenge target met, 30 of 30 in tolerance, no short, flashed or light part, no pressure excursion and a 10.87 s cycle.

> **TRY IT IN THE APP** **5 · Late Switchover**
>
> Load the scenario, press Run and watch the cavity pressure trace against the flash line and the label of each finished part: every part flashes and is 1.91 % heavy. Set Switch to pack on to Cavity P and watch the next part. The fill now ends at 100.4 %, and the part, at 24.150 g, is in tolerance.

7.5.4 Mass and Length Part Ways

In case 6 the coolant flow falls to 10 % at 12 s. The mold warms by about 1.1 K per shot, from 69.0 °C at the second part to 78.0 °C at the tenth, and the part mass stays between 24.097 and 24.107 g. The tenth part leaves the mold at 99.7 °C, goes on shrinking free of it, and measures 99.418 mm, the first out of tolerance. By the twentieth part the mold is at 84.6 °C and the part 99.330 mm long. Raising

Table 7.6 Challenge Campaign, Parts 6 to 35. Gain G in MPa/g, Limits and Correction in MPa, Worst Weighed Deviation From the Target

G, Δ_{shot}, Δ_{total}	In tolerance	Worst (%)	Peak c
fixed recipe	9 of 30	2.29	0
40, 4, 15 (defaults)	26 of 30	1.37	15.0
40, 5, 20	29 of 30	1.09	20.0
40, 5, 25	30 of 30	0.81	25.0
40, 5, 40 (suggested)	30 of 30	0.69	34.2
20, 5, 40	30 of 30	0.79	32.3
80, 5, 40	30 of 30	0.74	36.3
120, 5, 40	30 of 30	0.79	34.7

the cooling time to 6.5 s after the tenth part ejects the parts at 90.4 to 94.8 °C and brings them back to 99.490 to 99.556 mm, for a 13.36 s cycle and 35.7 kJ per part instead of 10.86 s and 32.8 kJ.

Case 8 starts from a mold at 30 °C, which reaches 37.0 °C by the first ejection and 60.7 °C by the seventh. The masses are right, 24.125 g on the first part and 24.113 g on the seventh, within 0.08 % of the warm-mold part, because a colder mold both densifies the melt and freezes the gate sooner. The length is not. The first part, ejected at 63.1 °C, did most of its shrinking held in the mold and measures 99.976 mm, the first six parts are oversize, and the seventh, at 99.647 mm, is the first in tolerance. A 3 s cooling time ejects the parts warmer and cuts the oversize run to four. The correction stays within 1.7 MPa of zero in case 8 and within 1 MPa in case 6, because its only input, the mass, was never wrong. A scheme that monitors part weight alone passes all of these parts.

7.5.5 A Biased Transducer

In case 7 the melt pressure transducer starts to read 22 MPa high at 26 s. In the middle of the fifth shot's hold it reads 60.0 MPa while the melt is at 38.0 MPa, and the cavity sensor, a separate instrument, peaks at 29.5 MPa instead of 41.4 MPa. The parts fall from 24.107 to 23.821 g, 1.16 % below the target, while the pressure loop holds its reading perfectly. The correction never reads the transducer, but it reads the scale. At its defaults it raises the setpoint 4 MPa per part, restores the tolerance from the fifth part, and stops at its 15 MPa limit with the parts 0.087 g light. With a 25 MPa limit it settles at 21.7 MPa and holds the weighed parts within 0.03 g of the target from the ninth. An outer loop on an independent measurement corrects the biased sensor of an inner loop, within the authority it is given, and the size of the correction becomes the diagnosis. Past the stability bound the same case shows the step limit at work. At 200 MPa/g with a 20 MPa step and a 50 MPa total the correction alternates between 10.6 and 30.6 MPa from part to part and the weighed mass between 23.95 and 24.22 g, while the default 4 MPa step confines the swing to one step, 19.2 to 23.2 MPa at first, and the weighed parts to within 0.05 g of the target.

7.5.6 The Production Campaign

The challenge starts from a 30 °C mold with 8 % of regrind variation. From the twelfth shot the lot setting rises in five steps from 1.1 to 1.5, and the transducer drifts 2.5 MPa per minute, 15.75 MPa by the end of the 378 s campaign. Both disturbances make the parts lighter, and neither is visible to the speed and pressure loops. The outcomes are listed in Table 7.6 and drawn in Figure 7.4.

The fixed recipe keeps 9 of the 30 campaign parts in tolerance, 20 short of the target. The first light part is shot 15, the worst weighed part is 2.29 % light, and the last weighs 23.761 g. The correction at its defaults keeps 26, because it pins at its 15 MPa limit and the last parts drift out. With 5 MPa per part and a 40 MPa total it keeps all 30, the worst 0.69 % from the target, with no flash and no shot at the pressure limit. The same settings carried through a restart give the same result. The correction peaks at 34.2 MPa after the 33rd part and is 30.6 MPa while the last part is molded, the 90.6 MPa setpoint shown in Figure 7.3. That part weighs 24.15 g, and the run stops as it is weighed, before the

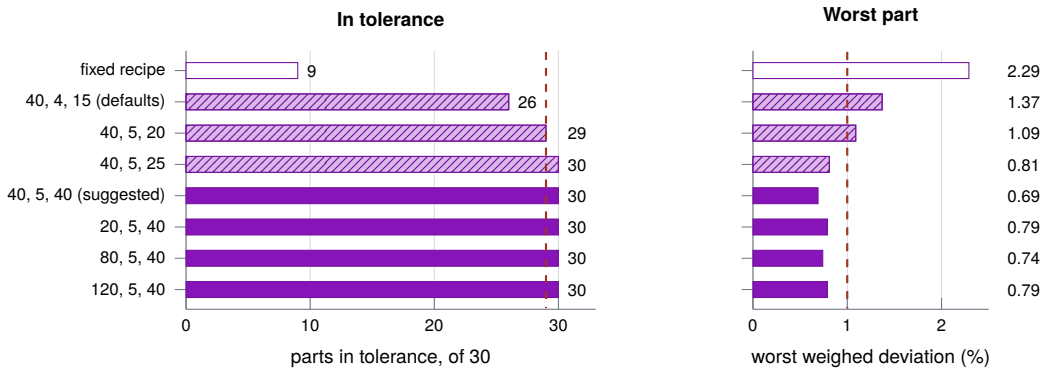


Figure 7.4 The challenge campaign of Table 7.6, labeled by gain G in MPa/g, step limit and total limit in MPa. The dashed lines are the target of 29 parts and the 1 % tolerance. Every setting allowed 40 MPa in total (solid bars) meets the target, while the default 15 MPa total (hatched) keeps 26 parts.

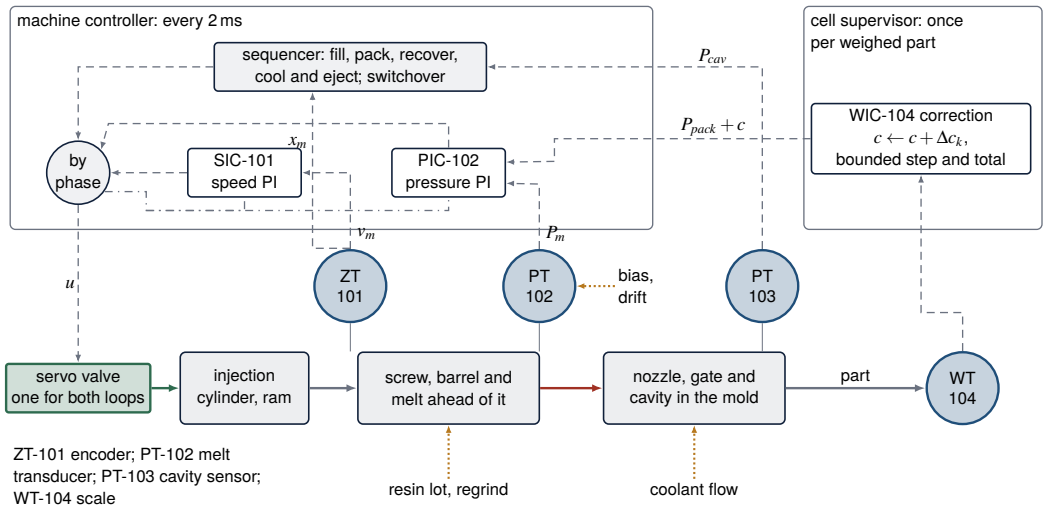


Figure 7.5 The chapter's loops on plant hardware. One servo valve drives the ram, and the sequencer hands its command u by phase to the speed loop SIC-101 on the encoder speed or to the pressure loop PIC-102 on the melt transducer, while the idle loop tracks the applied command (dash-dot). The switchover fires on the screw position, the cavity pressure or the time. After each part the scale WT-104 feeds the bounded correction that moves the pack setpoint. Tags are the book's.

correction acts on it. The mean cavity pressure peaks at 56.3 MPa and the valve at 84.4 % against its 90 % limit. A 20 MPa total already meets the target with 29 of 30 at the shipped seed, and at a 40 MPa total the worst part varies only from 0.69 to 0.79 % over gains of 20 to 120 MPa/g, smallest at the default, where the correction follows the drift without chasing much of the shot-to-shot variation. The cold mold still makes the first six parts oversize whatever the correction does, which is why the target scores mass and leaves the start-up parts out. Over five other seeds the fixed recipe keeps 9 to 11 parts, the default correction 26 to 29, the 20 MPa total 28 to 30 and the suggested settings 30, the only setting that meets the target at all six.

7.6 Technology in Practice

What Measures It

The loops of this chapter read three instruments inside the shot and one after it: an encoder on the screw, a transducer in the melt ahead of the screw, a pressure sensor in the cavity and a scale. A doctoral study of cavity pressure control weighed the candidates for the controlled variable. The hydraulic pressure is too far from the cavity to see the disturbances in between, the screw speed drops to about zero in packing, and the melt flow rate has no practical transducer, so it chose the cavity pressure, while noting that it is highly nonlinear and time varying [16]. In the model the encoder rounds to 1 µm, the cavity sensor adds 0.1 MPa of noise and the scale 0.01 g, and none of the four sees the gate freeze that decides the part.

> **TECHNOLOGY** **Sensor: Pressure Transducers From Pump to Cavity**
>
> Most electronic pressure sensors bond a strain gauge to a diaphragm. Pressure bows the diaphragm, the gauge's resistance changes, and a bridge circuit turns the change into a voltage. A metal gauge pressed past its elastic limit loses accuracy and can give the false impression that it still works [19]. One instrumented research machine carried transducers on the pump, ahead of the servo valve, on the injection cylinder, in the nozzle and in the cavity near the gate, each calibrated on a dead weight tester before installation [16]. A sensor in the cavity has to be built into the mold, which a review counts among the destructive ways of monitoring [20]. The model's melt transducer, with its 5 ms lag, 0.2 MPa of noise and an optional bias, stands for the nozzle sensor.

What Moves It

On a hydraulic machine the servo valve is one element of a circuit with a pump, relief valves and solenoid valves that divert the oil in the direction each phase of the cycle needs [16]. In this chapter the sequencer does the diverting in software: the valve command goes to the speed loop, the pressure loop or the back-pressure opening by phase, and the valve itself never changes. Figure 7.5 draws the loops on that hardware.

> **TECHNOLOGY** **Final Element: A Servo Valve on the Injection Cylinder**
>
> One research machine was refitted with a hydraulic servo valve for its injection cylinder, and with a second one for packing, because one valve did not give enough room to maneuver. The openings were proportional to the supplied current, and only 0.5 to 99.5 percent of the range was used, to keep clear of the dead zone. A servo valve is a complex, high-order, nonlinear device, but below about 100 cycles per second a first-order lag describes it adequately [16]. Its spool is machined to tight tolerances, so even small concentrations of particles in the oil can drastically shorten its life [19]. The chapter's valve is that first-order lag, 20 ms, shared by both loops.

Where the Controller Runs

The research machine of the thesis was sequenced by timers, limit switches and relays, and a microcomputer added data acquisition, machine sequencing and the control of each phase. It sampled fast signals such as pressures and slow ones such as barrel temperatures on separate boards, and the models of ram speed and nozzle pressure it reviewed were identified at a sampling period of 10 ms [16]. The app's machine controller runs every 2 ms. Control loops in industry run with cycle times from milliseconds to minutes, and in discrete part manufacturing an intermediate supervisor can oversee each cell with its machine controller [21]. Read as plant hardware, the speed and pressure loops and the phase sequencer belong to the machine controller, while the once-per-part correction from the scale fits the cell's supervisory layer.

IN PRACTICE **Where Practice Puts the Switchover**

Traditionally the phases of a shot were switched open loop, by limit switches, relays and timers, which leaves a large variability in part weight and dimensions. Switching on ram position is unreliable when a cushion of melt is left ahead of the screw, because the position at the end of filling then varies from shot to shot, and most process engineers avoid running without a cushion, because it risks short shots. One working rule keeps the ram under position or speed control until the cavity is about 90 to 98 percent full and then switches. The rate of rise of the cavity pressure marks the end of filling well [16]. The nominal recipe of this chapter switches at 96.0 % fill.

What Is Modernizing It

The first development is the drive. A national energy study of plastics counts the switch from hydraulic to all-electric molding machines as a 74 % saving of the machine's energy, the largest percentage saving of the processes it studied [18]. The second is cooling, which takes 50 to 80 % of the cycle. Conformal channels that follow the shape of the cavity, some made by additive manufacturing, are reported to shorten the cooling phase by 15 to 50 % [20]. The third is the control of quality. The same review collects norm-optimal iterative learning control of the cavity pressure, model predictive control combined with a neural network, quality prediction from in-mold sensors by support vector regression, and monitoring that reads the machine's own hydraulic pressure and screw position or measures the cavity pressure by ultrasound instead of a sensor built into the mold [20].

7.7 Limitations

One cavity is fed straight from a hot nozzle, without the runners, multi-cavity balance and check-ring leakage that change the shot delivered from one cycle to the next.

The temperatures are lumped and the skin is a geometric construction, so the model reproduces the timing of gate freeze but not the thickness distribution that a conduction solution gives [7, 8].

The fill uses one speed setpoint. Velocity profiles over the stroke, profiled hold pressures and a switchover point corrected from shot to shot are not modeled, and the correction acts on the pack pressure alone.

The drive has no inertia or hydraulic resonance, dynamics that controllers for real electrohydraulic machines must address [13].

The length is a proxy with an illustrative tolerance, without warpage or anisotropic shrinkage. The scale weighs every part at once, while a sampled or delayed weight would add dead time to the correction and move its stability bound, as Equation (1.9) shows for the integral form. The vial filler of Chapter 9 meets that delay: its checkweigher samples one vial per head every few cycles and reports it seconds later.

OPEN CHALLENGE **Closing the Quality Loop**

A molding machine nests three loops: machine control of speed, pressure and temperature, process control of the pressure and temperature in the mold, and a set point loop fed back from part quality. The machine loop is the most developed and the quality loop the least, its development having started only in the decade before the review. Learning, predictive and data-driven controllers on these loops have given good results, yet they are not widely used, because they lack standardization [20]. Progress would be a common framework that closes the quality loop for any machine, mold and resin, as this chapter's scale closes it for one part and one measurement.

7.8 Conclusion

Injection molding hands one valve between two loops inside every shot and fixes the part at a moment neither loop observes, and this chapter made both properties measurable on the simulator. Tracking the applied valve keeps the handover to one proportional increment, at most 0.44 % over twenty shots. A resin lot 75 % more viscous leaves the fill looking the same and the part about 1.4 % light, and

in manual mode it short-shots at 65 %. A switchover at 50.5 % fill makes a short part, and one at 103.6 % flashes at 94.6 MPa. A warmer or colder mold moves the length out of tolerance while the mass stays on target.

The correction from the scale removes what persists, within its authority, and widens the spread of what does not, by the 17 % its default loop gain of 0.55 predicts against 15 % measured over 195 parts. In the campaign the fixed recipe keeps 9 of 30 parts, the default correction 26, and a correction allowed 40 MPa all 30.

The model is lumped, with one cavity, a linear PVT law, a length proxy and a scale without delay. A fuller model would add runners and cavities, velocity and hold profiles, a delayed weight and a corrected switchover point. Within those limits the application shows how a product that no loop measures is still held to specification.

KEY POINTS

- Two loops take turns on one valve in every shot, and the part is fixed at a moment neither of them observes.
- Tracking the applied valve keeps the handover to one proportional increment, at most 0.44 % over twenty shots.
- A resin lot 75 % more viscous leaves the fill looking the same and the part about 1.4 % light.
- The correction from the scale removes what persists, within its authority, and widens the spread of what does not.
- In the campaign the fixed recipe keeps 9 of 30 parts, the default correction 26, and a correction allowed 40 MPa all 30.

Where else this applies. The converter of Volume 16, Chapter 4 corrects its heat-loss factor between heats with the same exponentially weighted update, and the lithography line of Volume 14, Chapter 6 corrects each lot's recipe while its measurements wait in a queue.

References

[1] R. Hanus, M. Kinnaert, and J.-L. Henrotte, "Conditioning technique, a general anti-windup and bumpless transfer method," *Automatica*, vol. 23, no. 6, pp. 729–739, 1987.

[2] S. F. Graebe and A. L. B. Ahlén, "Dynamic transfer among alternative controllers and its relation to antiwindup controller design," *IEEE Transactions on Control Systems Technology*, vol. 4, no. 1, pp. 92–99, 1996.

[3] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009.

[4] Z. Chen and L.-S. Turng, "A review of current developments in process and quality control for injection molding," *Advances in Polymer Technology*, vol. 24, no. 3, pp. 165–182, 2005.

[5] D. O. Kazmer, S. Velusamy, S. Westerdale, S. Johnston, and R. X. Gao, "A comparison of seven filling to packing switchover methods for injection molding," *Polymer Engineering & Science*, vol. 50, no. 10, pp. 2031–2043, 2010.

[6] Y. Yang, X. Chen, N. Lu, and F. Gao, *Injection Molding Process Control, Monitoring, and Optimization*. Carl Hanser Verlag, 2016.

[7] M. R. Kamal and S. Kenig, "The injection molding of thermoplastics part I: Theoretical model," *Polymer Engineering & Science*, vol. 12, no. 4, pp. 294–301, 1972.

[8] M. Rafizadeh, W. I. Patterson, and M. R. Kamal, "Physically-Based Model of Thermoplastics Injection Molding for Control Applications," *International Polymer Processing*, vol. 11, no. 4, pp. 352–362, 1996.

[9] R. Geyer, J. R. Jambeck, and K. L. Law, "Production, use, and fate of all plastics ever made," *Science Advances*, vol. 3, no. 7, p. e1700782, 2017.

[10] A. R. Agrawal, I. O. Pandelidis, and M. Pecht, "Injection-molding process control—A review," *Polymer Engineering & Science*, vol. 27, no. 18, pp. 1345–1357, 1987.

[11] C. A. Hieber and S. F. Shen, "A finite-element/finite-difference simulation of the injection-molding filling process," *Journal of Non-Newtonian Fluid Mechanics*, vol. 7, no. 1, pp. 1–32, 1980.

[12] Y. Yang and F. Gao, "Cycle-to-cycle and within-cycle adaptive control of nozzle pressure during packing-holding for thermoplastic injection molding," *Polymer Engineering & Science*, vol. 39, no. 10, pp. 2042–2063, 1999.

[13] H. Havlicsek and A. Alleyne, "Nonlinear control of an electrohydraulic injection molding machine via iterative adaptive learning," *IEEE/ASME Transactions on Mechatronics*, vol. 4, no. 3, pp. 312–323, 1999.

[14] Y. Yang and F. Gao, "Adaptive control of the filling velocity of thermoplastics injection molding," *Control Engineering Practice*, vol. 8, no. 11, pp. 1285–1296, 2000.

[15] F. Gao, Y. Yang, and C. Shao, "Robust iterative learning control with applications to injection molding process," *Chemical Engineering Science*, vol. 56, no. 24, pp. 7025–7034, 2001.

[16] M. Rafizadeh, "Physically-based dynamic model for the control of cavity pressure in thermoplastics injection molding," Ph.D. dissertation, McGill University, Department of Chemical Engineering, Montreal, 1996.

[17] A. Ingolfsson and E. Sachs, "Stability and Sensitivity of an EWMA Controller," *Journal of Quality Technology*, vol. 25, no. 4, pp. 271–287, 1993.

[18] U.S. Department of Energy, Advanced Manufacturing Office, "Bandwidth study on energy use and potential energy savings opportunities in U.S. plastics and rubber manufacturing," U.S. Department of Energy, Washington, DC, Report DOE/EE-1573, Sep. 2017, oSTI 1512369.

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

[20] O. Ogorodnyk and K. Martinsen, "Monitoring and control for thermoplastics injection molding: A review," *Procedia CIRP*, vol. 67, pp. 380–385, 2018.

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

7.9 Exercises

7.1 Concept check · 2 min

In scenario 2 a resin lot 75 % more viscous reaches the hopper before the fourth shot. In automatic mode the speed loop still fills at 40 mm/s, and the fill takes 0.874 to 0.876 s instead of 0.860 s. Every part of the new lot is nevertheless about 1.4 % light. Why?

(a) The speed loop runs into its valve limit, so less melt enters the cavity before the switchover

(b) The scale reads the thicker resin low, because a more viscous melt is less dense at room temperature

(c) The thicker melt loses more pressure on its way to the cavity during the pack, and the gate freezes sooner

(d) The pressure loop cannot reach its setpoint through the thicker melt, so the sequencer ends the hold early

7.2 Concept check · 1 min

At the switchover of the nominal recipe the pressure loop takes the valve from the speed loop. By how much does its first output differ from the last output of the speed loop?

(a) By $K_{c,p}$ times the gap between the pack setpoint and the melt pressure reading

(b) By the 24.3 % between its output at the end of the last hold and the valve now

(c) By nothing at all, because the clamp holds the emitted output at the previous command

(d) By its own proportional increment over one period, $-K_{c,p}(P_{m,k} - P_{m,k-1})$

7.3 Concept check · 2 min

In scenario 3 the lot factor of each shot is drawn afresh, with no correlation from shot to shot. What does the shot-to-shot correction at its default gain, loop gain $k = 0.55$, do to the spread of the weighed masses?

(a) It widens it by about 17 % in standard deviation, because each step answers a lot already gone
(b) It narrows it by the factor $1-k$ per part, as it does for a disturbance that persists
(c) It leaves it unchanged, because the correction can move only the mean of the masses
(d) It narrows it at this gain, and it would widen it only for a loop gain above one

7.4 Theory and calculation · 6 min

In the nominal recipe a pack pressure of 50 MPa makes a 23.974 g part and one of 70 MPa a 24.247 g part. Use the recursion of Equation (7.13).

(a) Compute the sensitivity $\partial m/\partial P$ from the two parts and the loop gain k at the default gain $G = 40\,\text{MPa/g}$. *Checkpoint:* 0.013 65 g/MPa and $k = 0.546$.

(b) Find the gain G at which a persistent step is removed in one part, and the gain at which the correction stops converging. *Checkpoint:* 73 MPa/g and 146 MPa/g.

(c) For an uncorrelated disturbance, find the factor by which the correction widens the standard deviation at G = 20, 40 and 80 MPa/g, and compare with the increases of 6, 15 and 43 % measured over 195 parts in case 3. *Checkpoint:* 1.076, 1.173 and 1.484.

7.5 Theory and calculation · 6 min

In case 2 the parts of the new lot weigh 23.755 to 23.769 g without correction, 23.762 g on average. The correction runs at its defaults ($G = 40\,\text{MPa/g}$, $\Delta_{shot} = 4\,\text{MPa}$, $\Delta_{total} = 15\,\text{MPa}$, $n_{start} = 3$), and the fourth part is the first of the new lot. Take $\partial m/\partial P = 0.013\,65\,\text{g/MPa}$ and the lower tolerance $0.99 \times 24.10 = 23.859\,\text{g}$.

(a) Find the correction that would bring the new lot back to 24.10 g. *Checkpoint:* 24.8 MPa.

(b) Using Equation (7.12), list the correction in force for parts 5 to 8 and find the first part back in tolerance. *Checkpoint:* 4, 8, 12 and 15 MPa. The sixth part.

(c) Predict the mass at which the parts settle, and compare with the 25.4 MPa at which the correction settles when its total limit is 30 MPa. *Checkpoint:* 23.967 g.

7.6 Theory and calculation · 5 min

Use Equations (7.3) and (7.10) with the parameters of Tables 7.2 and 7.3. The length specification is $99.54 \pm 0.12\,\text{mm}$.

(a) Compute v_0 at 23 °C, the volume shrinkage of a 24.107 g part and its length when ejected at $T_{ej,0} = 91.3\,°\text{C}$. *Checkpoint:* $0.9503\,\text{cm}^3/\text{g}$, $s_V = 0.0454$, 99.546 mm.

(b) Find the highest and the lowest ejection temperatures at which that part stays in tolerance. *Checkpoint:* 99.7 °C and 83.7 °C.

(c) In case 8 the first part weighs 24.125 g and leaves the mold at 63.1 °C. Compute its length. *Checkpoint:* 99.976 mm.

7.7 App challenge · 9 min

Load scenario 4, Early switchover, which switches over at 30 mm. Use the default seed, make every change before pressing Run, and run each case to 35 s at 5× (three parts, 7 s of wall time). Read the label of the third part in the parts row and its true mass, column `mass_true_g` of the exported file at 35 s.

(a) Run the scenario as shipped, and read also the fill fraction in the short-shot note. *Checkpoint:* SHORT, 23.412 g, 98.1 %.

(b) Press the reload button, set Switchover position to 20 mm and run again. *Checkpoint:* OK, 24.038 g.

(c) Load scenario 5, Late switchover (11.5 mm), run it as shipped, and read also the largest `p_cav_mean_MPa` between 21 and 30 s. *Checkpoint:* FLASH, 24.560 g, 93.8 MPa.

(d) Press the reload button, set Switch to pack on to Cavity P and run again. *Checkpoint:* OK, 24.150 g, 56.6 MPa.

Accept masses within 0.003 g and pressures within 1 MPa.

7.8 App challenge · 8 min

Load scenario 2, A thicker resin lot. The new lot arrives at 26 s, and the fourth part, the first of the new lot, is weighed at about 41 s. Use the default seed.

(a) Run the scenario as shipped to 45 s. Read the label and the weighed mass of the fourth part. *Checkpoint:* LIGHT, 23.769 g.

(b) Press the reload button, switch to Manual before pressing Run, which holds the valve at its 40 % fill opening, and run to 45 s. Read the label and the mass of the fourth part and the fill fraction in the short-shot note. *Checkpoint:* SHORT, 15.62 g, 65.4 %.

(c) Press the reload button again, set Shot-to-shot correction to On and Largest total correction to 30 MPa before pressing Run, and run to 220 s at 20×. Read the first part back in tolerance, the Pack correction tile from the thirteenth part on, and the largest deviation of a weighed mass from 24.10 g from the tenth part on. *Checkpoint:* The sixth part, 25.1 to 25.7 MPa, 0.03 g.

7.9 App challenge · 8 min

Load the challenge, a 30-shot production campaign from a cold mold. The run stops itself when the 35th part is weighed, at 377.9 s, about 19 s of wall time at 20×. Use the default seed and make each change after pressing Reset and before pressing Run. For each run read the tiles Challenge target, Mass in ±1 % after 5 parts and Pack correction.

(a) Run the campaign as shipped, with the correction off. *Checkpoint:* Not met, 9 of 30, 0.0 MPa.

(b) Set Shot-to-shot correction to On at its defaults (40 MPa/g, 4 MPa per part, 15 MPa in total) and run again. *Checkpoint:* Not met, 26 of 30, 15.0 MPa.

(c) Press Reveal suggested tuning (40 MPa/g, 5 MPa per part, 40 MPa in total) and run again. *Checkpoint:* Met, 30 of 30, 30.6 MPa.

Explain from the disturbances why run (b) fails and run (c) passes.

7.10 Application area · 4 min

In case 6 the coolant flow of the mold falls to 10 %. The parts keep their mass, and a cooling time raised from 4.0 to 6.5 s restores their length at a cycle of 13.36 s instead of 10.86 s and 35.7 kJ per part instead of 32.8 kJ.

(a) Find the parts per hour at each cycle and the cost of the longer cooling in output and in energy per part. *Checkpoint:* 331 and 269 parts per hour, 19 % fewer, 8.8 % more energy.

(b) Name two measurements that would reveal the fault before the length gauge does, and the lasting fix.

7.11 Application area · 3 min

A plant moves its scale to the end of a conveyor, so that the weight of a part reaches the correction two shots later than in the app. The integral form then has a delay of $D = 2$ parts in Equation (1.9), which is stable for $k = \lambda\xi$ below 0.618. With $\partial m/\partial P = 0.0137\,\mathrm{g/MPa}$, find the largest stable gain G and judge the default of 40 MPa/g.

7.12 Review the agent · 5 min

An AI assistant is asked to make the correction of case 7, Pressure sensor bias, as fast as possible and replies as follows. "The loop gain is $k = G\,\partial m/\partial P$ with $\partial m/\partial P = 0.0137\,\mathrm{g/MPa}$, and the correction is stable for $0 < k < 2$. The fastest correction is the largest stable gain, so set $G = 145\,\mathrm{MPa/g}$ ($k = 1.99$) with the step limit at 20 MPa and the total at 50 MPa. The parts will be back on target within one shot of the fault." Find the flaw and say how you would check it.