INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 14 Control in Semiconductor Manufacturing

CHAPTER 3

Czochralski Crystal Puller Simulator for Mid-Ranging a Fast Limited Actuator

Pages 52–81

Edited by John D. Hedengren

How to cite this chapter. "Czochralski Crystal Puller Simulator for Mid-Ranging a Fast Limited Actuator," in *Control in Semiconductor Manufacturing*, John D. Hedengren, Ed., vol. 14 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 52–81.

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

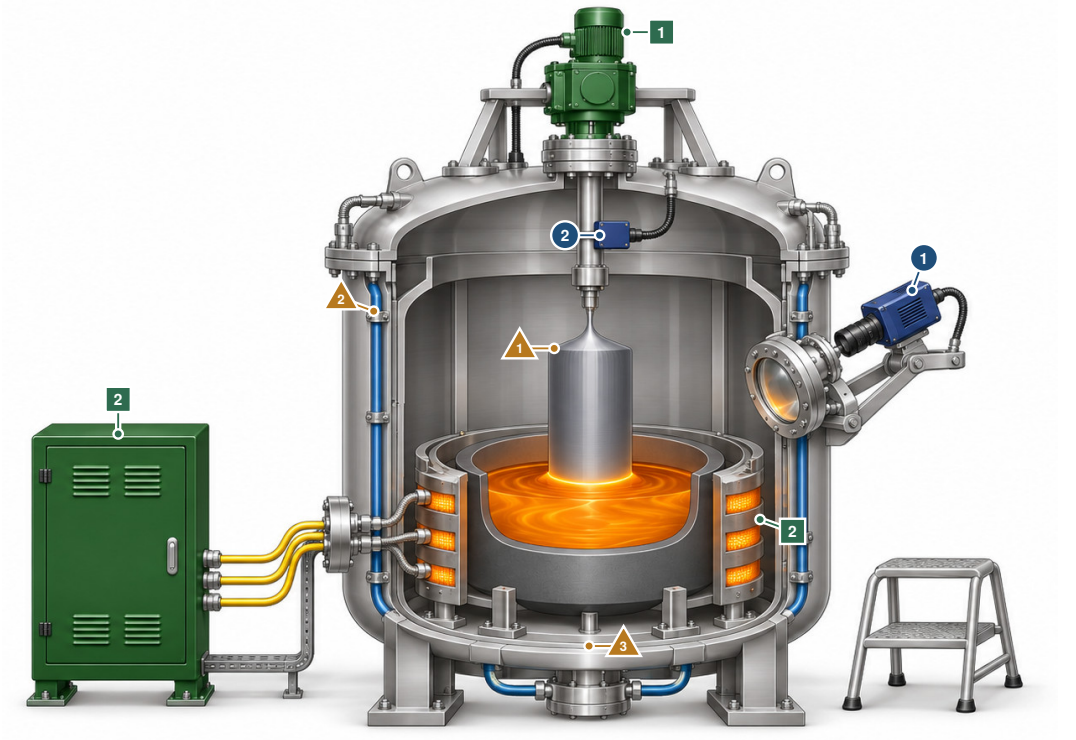
The app for this chapter
https://apmonitor.com/apps/cz-crystal-growth/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v14-semiconductor/#cz-crystal-growth

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

CHAPTER 3

Czochralski Crystal Puller Simulator for Mid-Ranging a Fast Limited Actuator



The controller moves 1 pull drive, set by the diameter loop
2 heater around the crucible and its power supply, set by the mid-ranging loop

The controller reads 1 camera on the bright ring, which gives the diameter
2 the pull rate the drive applies, read by the mid-ranging loop

What disturbs it 1 argon flow that changes how fast the crystal freezes
2 a hot zone that loses more heat than the recipe assumed 3 a crucible lift that drifts

Notice the bright ring at the foot of the crystal: the camera reads the diameter there, while the pull-rate ceiling that ends a good crystal has no instrument at all. Illustration, not to scale.

IN PLAIN TERMS **Pulling a Crystal at a Safe Speed**

The main idea. Computer chips are cut from one large crystal of silicon. It is drawn slowly upward out of a pot of melted silicon, and the pulling speed sets how thick it grows, through a short column of liquid that clings to its bottom.

The challenge. Pulling faster makes the crystal thinner, so an automatic control holds its size by changing the pulling speed. Each upset it corrects leaves the pull a little faster. Above a speed limit that no instrument shows, the crystal stops being one perfect crystal, and what grows after is scrap.

Why it matters. A second, slower control on the heater gives the pulling speed back its room and saves the whole crystal.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact and history boxes in Sections 3.1 and 3.2, the try-it and fact boxes in Section 3.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 3.6, Section 3.7, and Section 3.8 with its open challenge

The control question The pull rate reaches the crystal diameter only through the meniscus, two integrations away, and it has a ceiling that falls as the crystal fattens and that no instrument reports. How is the diameter held without spending that margin, and what does a slow heater loop closed on the pull rate buy?

Learning objectives

- Derive the double-integrator diameter plant and explain why derivative action is the whole of its phase margin.
- Derive the gain schedule on the pull rate that preserves the loop shape, and say what the fixed camera and drive lags leave out of it.
- Size the separation between a mid-ranging heater loop and the diameter loop from the static path from power to pull rate.
- Read the pull-rate margin and the prime crystal, not the diameter, to judge a run, and trade derivative noise gain against filter lag.
- Compare one predictive plan for both actuators with two retuned loops on the same batch.

Prerequisites Sections 1.6.2 (velocity form and loop action), 1.6.3 (the predictive option), 1.6.4 (scenarios, scorecard and CSV) and 1.4.8 (control within a batch). Chapter 2 schedules a gain along a batch.

The app https://apmonitor.com/apps/cz-crystal-growth/

See also Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging and valve-position control on power plants. Volume 7, Chapter 5 for derivative action on a double integrator. Volume 8, Chapter 4 and Volume 8, Chapter 6 for batches that follow a trajectory instead of a steady state.

3.1 Introduction

The only mechanical connection between a growing silicon crystal and the melt it comes from is a column of liquid a few millimeters tall, held up by surface tension. Everything the pull drive does to the diameter of the crystal it does through that column. Withdrawing the crystal faster than it solidifies lengthens the column, a longer column lowers the angle the solid makes with the free surface at the three-phase junction, and a lower angle narrows the crystal from that moment onward. Two integrations separate the actuator from the controlled variable, so a step in pull rate produces no immediate diameter change whatever and then an accelerating one, and a permanent pull-rate error is not a diameter offset but a cone. A loop built on that plant arrives at its crossover frequency with 180° of lag before the controller has contributed anything.

The second actuator is the heater, and it is the opposite in every respect. Heater power reaches the interface through a graphite element, a quartz crucible and a hundred kilograms of molten silicon, which in the model of this chapter is a dominant time constant of 84.8 min against a diameter loop whose crossover period is 36.3 min. The heater has authority without speed. The pull drive has speed without authority, and it has something worse than a small gain, which is a ceiling: above a pull rate set by how fast the latent heat of solidification can leave the interface up the crystal, the growth front breaks down and every kilogram grown afterward is polycrystalline. That ceiling falls as the crystal fattens and no instrument on the machine reports it.

✦ DID YOU KNOW? From 50 to 100 Kilograms

In 1985 production pullers melted charges of 50 kg of polysilicon, and the largest wafers in production were 150 mm across [1]. As of 2009, more than 95 percent of the 300 mm single-crystal silicon ingots were grown by the Czochralski process [2]. This chapter's puller grows a 300 mm crystal from a 100 kg charge.

A fast actuator with a hard limit paired with a slow actuator that has authority is a standard industrial arrangement, and the standard answer is not to choose between them. The fast loop is closed on the product and a second, slower loop is closed on the first loop's own output, so that the fast actuator is found near the middle of its range when the next disturbance arrives. The structure carries several names, valve position control and mid-ranging among them, and it is documented as a general design rather than a trick [3, 4]. Volume 4, Chapter 10 and Volume 4, Chapter 9 build the same structure on a power cycle and on a trigeneration plant. Readers will recognize the same shape in engine torque against transmission ratio, in a surge valve against a compressor recycle, and in every trim heater sitting behind a control valve. The Czochralski puller is an unusually honest instance of it, because the fast actuator's limit is a process constraint rather than a stop on the drive, and because the growth rate carries a two-sided product specification of its own [5, 6].

Czochralski dynamics and control have their own literature. The diameter response through the meniscus was written down and analyzed in the 1970s [7, 8], a state-space treatment with an explicit thermal model followed [9], and model-based designs with observers and scheduled parameters are current [10, 11]. However, a student handed the transfer function $-2V_p c_h/(s^2+\omega_n^2)$ and asked for a PID setting never meets the part of the problem that decides whether the run survives, which is that holding the diameter perfectly spends an actuator nobody is watching. This chapter sets out an interactive Czochralski puller in which the diameter loop always succeeds and the run is lost anyway, so that the remedy has to be a second loop rather than a better tuning. Section 3.2 traces how diameter control became automatic. Section 3.3 gives every equation and constant of the ten-state puller, with its meniscus, its thermal chain, its draining charge, its optical instrument and its actuators, at the level of detail a rebuild requires. Section 3.4 derives the double integrator, the exact gain schedule that the meniscus equations hand the diameter loop and the phase margin it holds, and the separation the heater path needs. Section 3.5 configures the two loops and the optional predictive controller. Section 3.6 runs the scenario sequence, in which the pull-rate margin, not the diameter, is the number that predicts the outcome. Section 3.7 names the camera, the drives and the controller behind the loops. Sections 3.8 and 3.9 state what the model leaves out and where model-based designs take the problem. The growth model, the mid-ranging structure and the predictive controller are all established methods.

3.2 Automatic Diameter Control

The method is named for a 1918 measurement of crystallization velocity in which a thin filament was drawn from molten metal and its length recorded against time [12]. What made it an industrial process was the demand for germanium and then silicon single crystals for solid-state devices, and by the early 1980s the pulling of silicon had its own detailed process literature covering the charge, the hot zone, the argon sweep and the defect structure of the product [13].

Diameter control was manual for the first three decades. An operator watched the bright ring of reflected light where the meniscus meets the crystal and moved the pull rate or the heater by hand, and the shape of the finished boule was a permanent record of how attentive the operator had been. Two independent routes to an automatic measurement appeared in the same volume of proceedings in 1972. One is optical, using a television camera aimed at the bright ring and an electronic threshold to convert its apparent width into a diameter signal [14]. The other weighs the crystal, or the crucible, and differentiates the load cell signal to recover the rate at which mass is being solidified [15]. The weighing route required its own theory before it was trustworthy, because the load cell reads the weight of the meniscus as well as the weight of the crystal and the two contributions move in opposite

directions during a diameter transient [16]. The application of this chapter takes the optical route and carries the camera filter and camera noise explicitly, because those two numbers set what derivative action costs.

Measurement was only half the problem. A controller needs a model, and the model that makes the diameter loop tractable is the dynamic meniscus, in which the growth angle is a function of the height of the melt column and the radius integrates the departure of that angle from its equilibrium value [7, 8]. That relation is what turns the plant into a double integrator and what forces a derivative time comparable to the closed-loop settling time. The two-actuator structure followed from furnace economics. As charge weights grew from a few kilograms to a hundred and more, a heater correction large enough to matter took longer than the diameter excursion it was meant to fix, so the fast correction moved permanently to the pull drive and the heater became the trim. Later work made the split explicit, treating the choice of manipulated variable and the scheduling of controller parameters along the batch as design questions in their own right [9, 11].

> ⧖ HISTORY **A Heater Loop That Waits and Sees (1972)**
>
> A paper received in October 1972 reported automatic diameter control of a silicon puller [17]. Wafers were moving from two to three inches, crystals free of dislocations had become the norm, and holding the diameter by hand had grown difficult. Optical methods already sighted the meniscus with a radiation pyrometer or a television camera. The paper weighed the crystal instead and moved the pull rate, with strong derivative action to keep that loop from oscillating. Control through the heater alone had not succeeded: its response was slower, its dead time about twice as long, and its gain too sensitive to temperature to measure [17]. So a second controller trimmed the heater power whenever the pull rate left a set band, acting slowly in a “wait and see” manner so that the two would not oscillate [17]. It is an early form of the mid-ranging structure of this chapter.

3.3 Puller Model

Puller CZ-300 grows one 300 mm silicon crystal from a 100 kg charge. A graphite heater surrounds a quartz crucible that holds the melt, the crucible is lifted as the charge drains so that the melt surface stays at a fixed height, and the seeded crystal is withdrawn through a meniscus of height h at pull rate V_p. Bright-ring camera DT-301 reports the diameter, diameter controller DIC-301 drives the pull rate, and mid-ranging controller VIC-301 reads the pull rate and drives the heater supply. Figure 3.1 shows the arrangement and every symbol used in the equations that follow. The application starts in body growth and the seed, neck, crown and shoulder are outside its scope.

Table 3.1 lists what the two loops interact with, with D read by DT-301, V_p driven by DIC-301 and P driven by VIC-301. The plant is written in the units of a puller control room, lengths in millimeters, rates in millimeters per minute, mass in kilograms, temperature in degrees Celsius, power in kilowatts and time in minutes. Ten states are integrated.

3.3.1 The Meniscus and the Radius

The melt column between the crystal and the melt surface lengthens whenever the crystal is withdrawn faster than it solidifies, and the crucible lift is supposed to hold the melt surface still while the charge drains. Both effects enter the same integrator, as shown in Equation (3.1).

$$\frac{\mathrm{d}h}{\mathrm{d}t} = V_p - V_g + \varepsilon_L \tag{3.1}$$

where h is the meniscus height in mm, V_p is the pull rate, V_g is the growth rate and ε_L is the crucible lift error, all in mm/min. Silicon grows at a fixed angle to the free surface of the meniscus, so the angle the solid makes with the vertical departs from zero only when the column departs from its equilibrium height, and the radius integrates that departure as shown in Equation (3.2) [7, 8].

$$\frac{\mathrm{d}r}{\mathrm{d}t} = V_p \tan\big[\alpha(h)\big], \qquad \alpha(h) = -c_h\,(h - h_0) \tag{3.2}$$

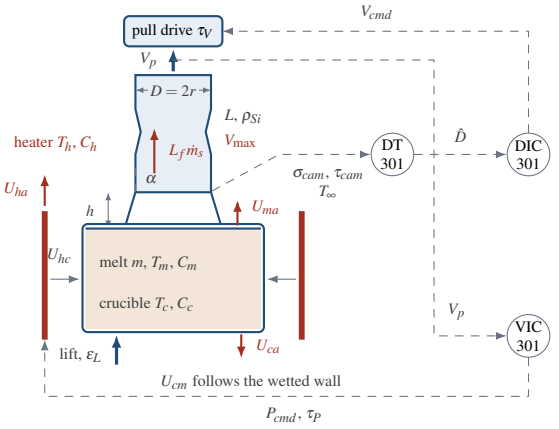


Figure 3.1 Czochralski puller CZ-300 in section, with the melt column that is the only path from pull rate to diameter. The crystal of radius r and length L is withdrawn at V_p through a meniscus of height h, and the growth angle α at the three-phase junction falls as h rises. Latent heat $L_f \dot{m}_s$ leaves the interface up the crystal and sets the pull rate $V_{\max}$ above which structure is lost. Heater power enters node T_h and reaches the melt T_m through the crucible node T_c. Bright-ring camera DT-301 supplies $\hat{D}$ with noise σ_{cam} and filter τ_{cam}, DIC-301 commands the pull drive, and VIC-301 reads the pull rate and commands the heater.

Table 3.1 What the Two Loops Interact With. The Nominal Column Is the Design Point That Opens Every Scenario but the Tail Scenario.

Symbol	Role	Nominal	Range	Unit
$D = 2r$	CV, measured	300.00	270–330 SP	mm
V_p	MV, and PV of VIC-301	1.000	0.05–2.50	mm/min
P	MV of VIC-301	118.26	30–220	kW
V_g	output, measured	1.000	0.72–1.32 spec	mm/min
$V_{\max}$	output, unmeasured	1.600	—	mm/min
η	DV, measured	1.000	0.700–1.300	—
λ	DV, measured	1.000	0.800–1.300	—
ε_L	DV, unmeasured	0	±0.100	mm/min
σ_{cam}	DV, unmeasured	0.35	0–3.00	mm
r	state	150.00	20–260	mm
h	state	8.00	0.2–40	mm
T_h	state	1695.96	—	°C
T_c	state	1566.73	—	°C
T_m	state	1425.00	—	°C
m	state	100.0	0.5–100	kg
L	state	0	—	mm
V_a	state	1.000	0.05–2.50	mm/min
P_a	state	118.26	30–220	kW
m_p	state	0	—	kg

where r is the crystal radius in mm, c_h is the growth angle change per unit meniscus height in rad/mm and h_0 is the height at which the angle is zero. The argument of the tangent is clipped to ± 0.5 rad, which is never reached in normal operation and keeps the right-hand side finite if a student drives the meniscus to a wall. Both sides of (3.2) carry mm/min.

Equations (3.1) and (3.2) are the whole reason the application exists. There is no term in (3.2) that contains V_p alone, so $\partial \dot{r} / \partial V_p = 0$ whenever $h = h_0$. The application's own test suite verifies that

algebraically by stepping the pull rate 20 % above design at the equilibrium meniscus and confirming that $|\dot{r}| < 10^{-12}$ mm/min, then displacing the meniscus by 1 mm and confirming that $\dot{r}$ becomes negative. Two integrations, no direct path, and the loop pays for it in Section 3.4.

3.3.2 Growth Rate and the Structure Limit

Latent heat released at the interface leaves by conduction up the crystal and radiation from its surface, and a thinner crystal removes it over a shorter path, so the rate at which the melt can freeze rises as the radius falls. A hotter melt carries superheat into the interface and slows it. The growth rate retains those two terms only, as shown in Equation (3.3).

$$V_g = \left[V_0 \frac{r_0}{r} - K_T \left(T_m - T_{m0} \right) \right] \eta \tag{3.3}$$

where V_0 is the growth rate at the design radius r_0 and the design melt temperature T_{m0}, K_T is the growth rate lost per degree of melt superheat in mm/min per °C, and η is a hot-zone multiplier representing how well the heat shield removes heat up the crystal. The radius in the denominator is floored at 10 mm. The same heat-removal argument bounds the pull rate, as shown in Equation (3.4).

$$V_{\max} = V_{\max,0} \frac{r_0}{r} \eta \tag{3.4}$$

where $V_{\max,0}$ is the pull rate at which structure is lost at the design radius. Above $V_{\max}$ the interface cannot shed the latent heat, the growth front breaks down and the crystal ceases to be a single crystal. The application latches a loss flag at that instant and stops accruing prime crystal.

Equation (3.4) deserves its own sentence, because the modeling choice behind it is where this application differs from a textbook saturation exercise. $V_{\max}$ is a *process* constraint and not an actuator stop. It is not the limit the controller output is clamped to, it falls as the crystal fattens, no instrument in the application measures it, and the controller can and will drive straight through it. The pull drive has a separate and larger mechanical limit of 2.50 mm/min, and that is what the diameter controller output is clamped to. Declaring $V_{\max}$ as the controller output limit would make the constraint uncrossable and would remove the only failure the application is built around.

3.3.3 Melt Inventory, Length and Product

The mass solidified per minute follows from the interface area and the growth rate, as shown in Equation (3.5).

$$\dot{m}_s = \rho_{Si} \pi r^2 \max(0, V_g), \qquad \frac{\mathrm{d}m}{\mathrm{d}t} = -\dot{m}_s, \qquad \frac{\mathrm{d}L}{\mathrm{d}t} = V_p \tag{3.5}$$

where ρ_{Si} is the density of solid silicon in kg per mm^3, m is the melt remaining in kg and L is the crystal length in mm. At the design point (3.5) returns $\dot{m}_s = 0.1647$ kg/min, so the 88 kg between m_0 and m_{end} takes 534 min at exactly 1.000 mm/min, and makes 534 mm of boule. Prime crystal is integrated rather than counted, so the figure is exact at any step size, as shown in Equation (3.6).

$$\frac{\mathrm{d}m_p}{\mathrm{d}t} = \begin{cases} \dot{m}_s, & |D - D_{\mathrm{SP}}| \le \delta \ \wedge \ V_{lo} \le V_g \le V_{hi} \\ 0, & \text{otherwise or after structure is lost} \end{cases} \tag{3.6}$$

where δ is the diameter tolerance and $[V_{lo}, V_{hi}]$ is the growth rate band inside which the crystal freezes in neither an excess of vacancies nor an excess of self-interstitials.

3.3.4 The Thermal Chain

Heater power reaches the interface through three lumped nodes, as shown in Equation (3.7).

$$\begin{aligned} C_h \frac{\mathrm{d}T_h}{\mathrm{d}t} &= 60\,[P_a - U_{hc}(T_h - T_c) - U_{ha}(T_h - T_\infty)\lambda] \\ C_c \frac{\mathrm{d}T_c}{\mathrm{d}t} &= 60\,[U_{hc}(T_h - T_c) - U_{cm}(T_c - T_m) \\ &\qquad - U_{ca}(T_c - T_\infty)\lambda] \\ C_m \frac{\mathrm{d}T_m}{\mathrm{d}t} &= 60\,[U_{cm}(T_c - T_m) - U_{ma}(T_m - T_\infty)\lambda - q_{lat}] \end{aligned} \tag{3.7}$$

where C_h, C_c and C_m are heat capacities in kJ/K, the U terms are conductances in kW/K, T_∞ is the hot-zone reference temperature, λ is an insulation multiplier, and the factor 60 converts kW to kJ/min because the clock is in minutes. The latent heat drawn from the melt is $q_{lat} = L_f \dot{m}_s / 60$ in kW. Heater power is floored at zero before it is used.

Three of the coefficients in (3.7) are functions of the charge remaining, and they are not functions of it in the same way, as shown in Equation (3.8).

$$C_m = c_p \max(2, m), \;\; U_{cm} = U_{cm,0}\,\phi^{0.10}, \;\; U_{ma} = \text{constant} \tag{3.8}$$

where $\phi = \mathrm{clamp}(m/m_0, 0.02, 1.2)$ is the melt fraction and c_p is the melt heat capacity in kJ/(kg K). The heat capacity falls in proportion to the mass. The crucible-to-melt coupling follows the wetted wall, which falls much more slowly. The loss from the melt surface does not fall at all, because the surface area is set by the crucible diameter rather than by how much melt is in it. Section 3.4 measures what that combination does to the plant over a run.

3.3.5 Actuators, Instrument and Limits

Both actuators are first-order lags on their commands, as shown in Equation (3.9).

$$\tau_V \frac{\mathrm{d}V_a}{\mathrm{d}t} = V_{cmd} - V_a, \qquad \tau_P \frac{\mathrm{d}P_a}{\mathrm{d}t} = P_{cmd} - P_a \tag{3.9}$$

where V_a is the pull rate actually applied and P_a is the power actually delivered. The pull rate used everywhere in the physics is $V_p = \max(0, V_a)$. The bright-ring camera adds Gaussian noise to the true diameter and filters the result, as shown in Equation (3.10).

$$\tau_{cam} \frac{\mathrm{d}\hat{D}}{\mathrm{d}t} = 2r + \sigma_{cam}\, n(t) - \hat{D} \tag{3.10}$$

where $\hat{D}$ is the measured diameter in mm, σ_{cam} is the camera noise standard deviation and $n(t)$ is a standard normal deviate redrawn at every controller execution from the seeded generator. Advancing that filter uses the exact zero-order-hold recursion $\hat{D}_k = a\hat{D}_{k-1} + (1-a)(2r + \sigma_{cam} n)$ with $a = \exp(-\Delta t_c / \tau_{cam})$, which cannot go unstable however coarse the scan becomes.

After every step the radius is clipped to 20–260 mm, the meniscus to 0.2–40 mm and the melt to at least 0.5 kg. Three latches then run in order. Structure is lost when V_p has stayed above $V_{\max}$ continuously for t_{lose}, the crystal detaches if $h > h_{\max}$, and the crystal touches down into the melt if $h < h_{\min}$. Any of the three ends the useful run. The dwell matters more than it looks. Dislocation nucleation is not instantaneous, and without it one noisy camera sample differentiated through a 45 min derivative time scraps a boule, which is an artifact of the sample rate rather than of the interface. At $t_{lose} = 0.5$ min scenario 9 run with a 10 min derivative filter crosses $V_{\max}$ once, for three physics steps, and survives, while the same scenario as loaded stays across it for the whole dwell and loses the crystal. The dwell is accumulated as a running sum of physics steps, and the comparison allows for a sum that falls short of t_{lose} by a rounding error, so ten consecutive steps above the limit trip the latch. Separately, the batch is declared finished when $m \le m_{end}$, at which point the mechanical derivatives are frozen so that the model is never asked to place a meniscus in an empty crucible.

Table 3.2 Constants of the Puller Model. A Hidden Constant Has No Entry in the Parameter Panel and Has to Be Recovered From an Experiment.

Symbol	Meaning	Value	Unit	Visibility
h_0	meniscus at zero angle	8.0	mm	adjustable
c_h	angle per unit h	0.050	rad/mm	adjustable
$h_{\min}$	floods below	3.0	mm	adjustable
$h_{\max}$	detaches above	15.0	mm	adjustable
V_0	design growth rate	1.00	mm/min	adjustable
K_T	growth per degree	0.025	(mm/min)/K	adjustable
$V_{\max,0}$	structure limit at r_0	1.60	mm/min	adjustable
C_h	heater capacity	300	kJ/K	adjustable
C_c	crucible capacity	900	kJ/K	adjustable
U_{hc}	heater to crucible	0.80	kW/K	adjustable
$U_{cm,0}$	crucible to melt, full	0.60	kW/K	adjustable
U_{ma}	melt surface loss	0.356	kW/K	adjustable
U_{ca}	crucible loss	0.050	kW/K	adjustable
U_{ha}	heater loss	0.030	kW/K	adjustable
m_0	melt charge	100	kg	adjustable
m_{end}	melt left at end	12	kg	adjustable
t_{lose}	dwell above $V_{\max}$	0.50	min	adjustable
τ_V	pull drive lag	0.30	min	adjustable
τ_P	heater supply lag	0.50	min	adjustable
δ	diameter tolerance	1.5	mm	adjustable
V_{lo}	defect band, lower	0.72	mm/min	adjustable
V_{hi}	defect band, upper	1.32	mm/min	adjustable
—	value of prime crystal	900	\$/kg	adjustable
r_0	design radius	150	mm	hidden
T_{m0}	design melt temp.	1425	°C	hidden
T_{mp}	silicon melting point	1414	°C	hidden
c_p	melt heat capacity	1.00	kJ/kg/K	hidden
T_∞	hot-zone reference	1200	°C	hidden
L_f	latent heat of fusion	1800	kJ/kg	hidden
ρ_{Si}	solid silicon density	2.33e−6	kg/mm^3	hidden

3.3.6 Parameters, Initialization and Numerics

Table 3.2 lists every constant. Seven are hidden from the parameter panel and are the quantities a student has to infer from experiments rather than read.

Every scenario but the tail scenario starts from a condition in which nothing is moving. The radius is set to r_0, the meniscus to h_0, the melt to m_0 and the melt temperature to T_{m0}, and the two remaining temperatures are solved backward up the chain so that all three derivatives of (3.7) vanish, as shown in Equation (3.11).

$$\begin{aligned} q_m &= U_{ma}(T_{m0} - T_\infty)\lambda + L_f \dot{m}_s/60 \\ T_c(0) &= T_{m0} + q_m/U_{cm} \\ q_c &= q_m + U_{ca}(T_c(0) - T_\infty)\lambda \\ T_h(0) &= T_c(0) + q_c/U_{hc} \\ P_a(0) &= q_c + U_{ha}(T_h(0) - T_\infty)\lambda \end{aligned} \tag{3.11}$$

At the design point (3.11) returns $T_c = 1566.73\,^\circ\mathrm{C}$, $T_h = 1695.96\,^\circ\mathrm{C}$ and $P_a = 118.26\,\mathrm{kW}$, and the same expression supplies the design-point aid available in manual mode. The application's test suite confirms that the largest state derivative at that point is below 10^{-9} in its own units and that the diameter is exactly 300 mm. One scenario overrides the initial state, starting the run at $m = 30\,\mathrm{kg}$ and $L = 1100\,\mathrm{mm}$.

States advance under fixed-step fourth-order Runge-Kutta with a step of 0.05 min (3 s). DIC-301 and VIC-301 both run on a 0.1 min (6 s) period, and the historian keeps two samples for every simulated minute. The noise sequence comes from a generator seeded at 20260918, so scenario 9

produces the same camera record on any machine. Section 1.6.1 describes the engine, the seed and the speed setting, which leaves the trajectory untouched.

3.3.7 Assumptions

The model is detailed where the two-actuator structure lives and thin everywhere else, and each simplification fails somewhere.

One-dimensional quasi-steady heat transfer through three nodes. Real pullers are dominated by three-dimensional buoyant and Marangoni convection in the melt, which is why production machines apply a magnetic field to suppress it, and that convection supplies a large part of the real diameter disturbance [13]. The assumption fails wherever the melt flow pattern itself changes, which is exactly where a real diameter excursion starts.

The meniscus is linearized about the constant growth angle. Equation (3.2) is the small-departure form of the Laplace-Young solution rather than the solution itself [8], so the flooding and detachment heights in Table 3.2 are placed from experience rather than derived, and the model is not trustworthy for the large angle excursions of a shoulder.

Growth rate is two terms. There is no axial temperature gradient field in the model, so the defect specification is applied to V_g directly rather than to the ratio of growth rate to gradient that actually governs which point defects freeze in [5, 6]. The substitution is only valid while the gradient is treated as fixed, which a change of hot zone would break.

Loss of structure is instantaneous. In practice the interface breaks down over some minutes and an experienced operator sometimes sees it coming in the ring image, so the latch in this model is pessimistic about warning and optimistic about recovery.

The crucible lift is perfect unless disturbed. A real lift profile is itself a tuned trajectory that interacts with the thermal field, and here it is a constant error term in (3.1) with no dynamics. A lift error that appears instantly at 0.03 mm/min and never changes is not what a real lift mechanism does, but it isolates an unmeasured ramp into the one integrator that reaches the diameter. In the same way a hot-zone multiplier that changes in one instant is a step because a step is the only upset whose response a trend reads back cleanly, while a real argon disturbance arrives with a shape of its own.

No oxygen, carbon or dopant transport. Resistivity control, striations and axial segregation are all real control problems on the same machine and none of them is present, so the economics in the application are a flat value per kilogram rather than the graded price a real boule carries.

Repetition is exact, which no control system manages. Freezing the integration step and the scan makes one trajectory the only trajectory, where a real puller control system carries scan jitter, a coarser measurement resolution and an execution order that shifts.

3.4 Dynamic Analysis

3.4.1 The Diameter Plant Is a Double Integrator

Linearizing (3.1) and (3.2) about a steady operating point at which $h = h_0$ and $V_p = V_g$, and holding the melt temperature fixed so that V_g does not respond, gives $\Delta\dot{h} = \Delta V_p$ and $\Delta\dot{r} = -V_p c_h \Delta h$. Combining them gives the plant in Equation (3.12).

$$\frac{D(s)}{V_p(s)} = \frac{-2V_p c_h}{s^2 + \omega_n^2}, \qquad \omega_n = V_p\sqrt{\frac{c_h}{r_0}} \tag{3.12}$$

where the ω_n^2 term comes from the feedback of radius into growth rate through (3.3), and the leading factor of two converts radius to diameter. At the design point $\omega_n = 0.018\,26\,\text{rad/min}$, which is a natural period of 344 min, and the steady-state gain is $2r_0/V_p = 300\,\text{mm}$ of diameter per mm/min of pull rate. The negative sign means that pulling faster makes the crystal thinner, so a crystal growing fat is answered by pulling faster, and the diameter loop is direct acting in the convention of Section 1.6.2.

Two facts about (3.12) decide the tuning. Its phase begins at $-180°$ and stays there, so proportional and integral action alone leave no margin at any gain and the loop rings until something saturates.

Table 3.3 Derivative Time Against Phase Margin and Measured Diameter Excursion. Margins Are Computed From (3.12) at $V_p = 1\,\text{mm/min}$ With the Shipped Filter $\tau_F = 8\,\text{min}$. The Excursion Is the Peak-to-Peak Diameter Over $120 \le t \le 500\,\text{min}$ After a 4 mm Setpoint Step at $t = 60\,\text{min}$ in Scenario 3. Without Derivative Action the Meniscus Floods at $t = 158.4\,\text{min}$, So That Column Spans $60 \le t \le 158.4\,\text{min}$.

τ_D (min)	0	15	30	45	90
Phase margin (°)	−6.2	25.4	30.1	29.2	24.4
Crossover period (min)	82.0	59.4	44.3	36.3	25.6
Peak-to-peak D (mm)	11.28	2.52	1.98	2.29	3.50
Peak D (mm)	309.8	306.5	306.3	306.2	306.3
Meniscus floods	158.4 min	—	—	—	—

Derivative action is therefore the entire stabilizing mechanism rather than a refinement, as it is for the ball and beam of Volume 7, Chapter 5, and the amount required is set by where the derivative zero has to sit relative to crossover. Table 3.3 evaluates the loop $L(s) = K_c[1 + 1/(\tau_I s) + \tau_D s/(\tau_F s + 1)]\,D(s)/V_p(s)$ at the default $K_c = 0.055\,\text{mm/min}$ per mm, $\tau_I = 120\,\text{min}$ and $\tau_F = 8\,\text{min}$, beside the diameter excursion measured in the application on a 4 mm setpoint step.

The first column of Table 3.3 is a negative phase margin and the run does not survive it. Without derivative action the diameter overshoots the new setpoint to 309.8 mm at $t = 101.9\,\text{min}$, swings back to 298.5 mm at $t = 141.9\,\text{min}$, and the growing oscillation in the meniscus floods it at $t = 158.4\,\text{min}$, so the crystal is lost through the meniscus while the pull rate never comes closer than 0.214 mm/min to the structure limit. With $\tau_D = 45\,\text{min}$ the same step is answered with an excursion of 2.29 mm and a phase margin of 29.2°. Forty-five minutes of derivative time sounds enormous until the crossover period of 36.3 min is set beside it. The rule a double integrator enforces is that the derivative time has to be comparable to the closed-loop response time rather than a tenth of it.

🕮 **Pause.** Try Exercise 3.1 before reading on.

Those margins also price the derivative filter. The same three gains at $\tau_F = 3\,\text{min}$ give 50.9° at a crossover period of 28.4 min, so opening the filter to the shipped 8 min costs 21.7° of margin, on a plant that began with none.

3.4.2 An Exact Gain Schedule

Both the gain and the natural frequency of (3.12) are proportional to the pull rate, which is unusual and useful. Substituting $s = (V_p/V_0)\,\sigma$ into (3.12) gives Equation (3.13).

$$\left.\frac{D}{V_p}\right|_{V_p}(s) = \frac{V_0}{V_p}\cdot\frac{-2V_0 c_h}{\sigma^2 + \omega_{n,0}^2} = \frac{V_0}{V_p}\cdot\left.\frac{D}{V_p}\right|_{V_0}(\sigma) \tag{3.13}$$

where $\omega_{n,0} = V_0\sqrt{c_h/r_0}$ is the natural frequency at the design pull rate. The plant at pull rate V_p is therefore the design-point plant with its frequency axis stretched by V_p/V_0 and its magnitude reduced by the same factor. A controller that stretches its own frequency axis by exactly the same amount and raises its gain by the reciprocal of the magnitude change reproduces the loop shape identically, as shown in Equation (3.14).

$$K_c \leftarrow K_c\,\frac{V_p}{V_0}, \qquad (\tau_D,\ \tau_I,\ \tau_F) \leftarrow (\tau_D,\ \tau_I,\ \tau_F)\,\frac{V_0}{V_p} \tag{3.14}$$

Equation (3.14) is a schedule the plant equations hand the engineer rather than one fitted during commissioning, and it is exact rather than approximate. The distinction matters, because the general difficulty with gain scheduling is that a family of linear designs interpolated on a measured variable carries no guarantee away from the design points [18]. Here the scheduling variable is the manipulated

variable itself, it is measured exactly, and it is algebraically the reason the gain moves instead of a stand-in for the reason. Chapter 2 schedules its loop on a measured cause in the same way, the rod diameter. The application implements Equation (3.14) on the applied pull rate V_a, with the ratio V_a/V_0 clamped to 0.2–4 so that a camera excursion cannot run the gain away.

🕮 **Pause.** Try Exercise 3.3 before reading on.

Figure 3.2 evaluates the phase margin of the diameter loop against pull rate three ways. Applying (3.14) holds the margin at 29.19° for every pull rate from 0.5 to 2.0 mm/min and not merely close to it, which is what exact loop-shape preservation means. A fixed tuning slides from 38.1° at 0.5 mm/min to 21.3° at 2.0 mm/min. The third curve is the reason the direction has to be derived rather than guessed. Scaling K_c by V_0/V_p instead, together with τ_D and with τ_I and τ_F left fixed, which is the intuitive choice because a faster plant seems to want a softer controller, gives 24.5° at 0.5 mm/min, worse than no schedule at all over the whole lower half of the drive range. The plant is faster at high pull rate and also weaker, and the gain has to follow the weakness.

A measured step confirms the shape. Settling the full charge for 250 min at mid-ranging targets of 1.00, 1.40 and 0.70 mm/min with the camera noise off and then stepping the diameter setpoint 3 mm gives a highest diameter over the following 120 min of 303.71 mm unscheduled and 303.70 mm scheduled at $V_p = 1.044\,\mathrm{mm/min}$, because the schedule is within 5 % of unity there. At $V_p = 1.343\,\mathrm{mm/min}$ the fixed tuning already sits 2.53 mm high before the step and peaks at 305.44 mm, against 1.45 mm high and a peak of 304.43 mm scheduled. At $V_p = 0.822\,\mathrm{mm/min}$ the scheduled loop is the slower of the two, reaching 301.81 mm in 120 min against 302.55 mm, and that is correct rather than a failure, because (3.13) stretches the frequency axis with the pull rate and what the schedule preserves is damping, not speed.

The flatness in Figure 3.2 is exact only because every time constant in that loop scales with the pull rate, and two of the real ones do not. The bright-ring camera filter of 0.50 min in (3.10) and the pull drive lag τ_V of 0.30 min in (3.9) both sit inside the loop and neither is a controller setting, so (3.14) cannot reach them. Adding the camera lag alone to the loop transfer function takes the unscheduled margin from 34.9 to 14.1° over the same 0.5 to 2.0 mm/min sweep and the scheduled margin from 26.7 to 19.6°. Adding the drive lag as well gives 33.1 to 9.7° unscheduled against 25.2 to 13.9° scheduled, so the honest margin of the complete loop at the design point is 21.4° rather than the 29.2° of the meniscus loop alone. Nothing is flat once the fixed lags are in, and the schedule's worth has to be restated because of it. What it buys is margin where the loop is weakest: at 2.0 mm/min it takes the complete loop from 9.7 to 13.9°, which is the difference between a loop that rings on every disturbance and one that does not. That is a more defensible claim than the flatness and it is the one that matters in operation, because the fast end is where a mid-ranged pull rate spends its time after a disturbance.

The boundary is worth stating in general terms. A gain schedule preserves loop shape only over the dynamics whose time constants it sets, so the residue is exactly what it cannot reach, and no amount of scheduling converts it into authority. Scenario 8 shows the second half of that plainly. At a mid-ranging target of 1.35 mm/min the diameter walks from 300.0 to 303.5 mm and leaves the 1.5 mm band at $t = 50.3$ min unscheduled and 52.3 min scheduled, with a worst true-diameter deviation of 3.46 mm unscheduled and 2.73 mm scheduled. The schedule narrows the excursion by a fifth and does not close it, because the loop has been given back its shape and not its authority.

3.4.3 The Heater Path and Why It Is the Outer Loop

The heater reaches the diameter only by changing the growth rate, which the diameter loop then has to match with the pull rate. Linearizing the thermal chain (3.7) together with the latent-heat feedback at the design point gives eigenvalues at -0.5342, -0.2161 and -0.0118/min, which are time constants of 1.87, 4.63 and 84.8 min. The slow mode dominates because the three capacities total 1300 kJ/K while the net conductance to the hot zone is only 0.3125 kW/K once the latent-heat feedback of 0.1235 kW/K is credited against the 0.436 kW/K of ambient loss. That lumped estimate is 69 min,

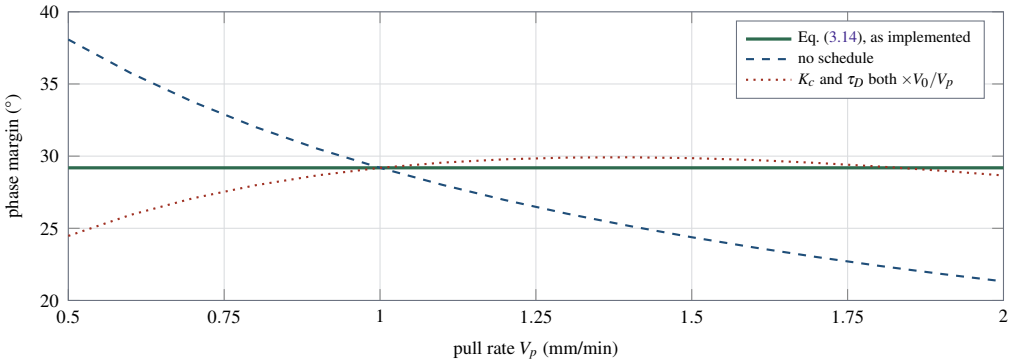


Figure 3.2 Phase margin of the diameter loop against pull rate for three scheduling choices, evaluated on Equation (3.12). All three curves start from the default tuning $K_c = 0.055\,\text{mm/min}$ per mm, $\tau_I = 120\,\text{min}$, $\tau_D = 45\,\text{min}$ and $\tau_F = 8\,\text{min}$ at $V_p = V_0$. The schedule of Equation (3.14) is the thick solid line, no schedule the dashed line and the reversed scaling the dotted line. Only Equation (3.14) holds the margin, because only it stretches the controller frequency axis and raises the gain together. Scaling the gain and the derivative time the other way, with τ_I and τ_F fixed, is worse than no schedule below the design pull rate.

short of the 84.8 min the full linearization returns but the right mechanism. On the linearized chain at the full charge a 5 kW heater step raises the melt by 14.09 °C, reaching 63 % of that change in 91.4 min and 95 % in 260.5 min, which is the slow mode with the two fast ones in series ahead of it. The application cannot show the whole of that response with nothing holding the diameter. Stepping the heater 5 kW in manual mode from the 118 kW at which scenario 1 opens warms the melt 9.39 °C above an otherwise identical unstepped run by $t = 88.3\,\text{min}$, when the slower growth has stretched the meniscus past $h_{\max}$ and the run stops. The linearized chain gives 8.72 °C at the same instant, and the difference is the latent heat no longer drawn by a crystal that has narrowed to 277 mm.

The static path from power to pull rate follows from (3.3). A change of 1 kW moves the melt by 2.819 °C, which moves the growth rate by $-K_T \times 2.819 = -0.0705\,\text{mm/min}$, and at steady state the diameter loop must set $V_p = V_g$. With the mid-ranging gain at 8 kW per mm/min the loop gain is $8 \times 0.0705 = 0.564$, so with $\tau_I = 120\,\text{min}$ the mid-ranging loop returns the pull rate with a time constant near $120/0.564 = 213\,\text{min}$. That is 5.9 times the diameter loop crossover period of 36.3 min, which is the separation the structure needs. The mid-ranging loop acts *through* the diameter loop, so if it were the faster of the two the loops would chase the same meniscus and the crystal would carry a slow waist. Cascade tuning states the rule as inner fast and outer slow. Mid-ranging is the same rule with the roles of speed and authority exchanged, and the outer loop is the one with the authority.

3.4.4 A Batch Has No Steady State

Nothing in the application is stationary, because the charge drains continuously and a completed run takes between 440 and 520 min depending on how fast the diameter loop is pulling. The batch crystallizer of Volume 8, Chapter 4 likewise follows a trajectory rather than a steady state. Using the melt node alone, the time constant is $C_m/[60(U_{cm} + U_{ma})]$ and the static gain from the crucible is $U_{cm}/(U_{cm} + U_{ma})$. Evaluating both from (3.8) gives 1.743 min and 0.628 at the full 100 kg charge against 0.293 min and 0.582 at 15 kg. The melt therefore answers 5.9 times faster at the tail than at the shoulder while its gain changes by only 7.2 %, which is an unusually clean separation of speed from gain and the reason a single mid-ranging tuning is wrong for part of every run.

The consequence for an operator is sharper than the arithmetic suggests. Scenario 1 opens in manual mode with the pull rate at the design 1.000 mm/min and the heater at 118 kW, within 0.26 kW of the design power, disturbs nothing, and leaves them there. Over the following four hours the

melt cools from 1425.0 to 1418.0 °C because the crucible-to-melt coupling is shrinking, the crystal therefore solidifies slightly faster than it is being pulled, the meniscus shortens from 8.00 to 3.39 mm, the diameter grows to 354.7 mm, and the crystal touches down into the melt at $t = 256.3$ min. Setting the heater to the 118.3 kW that the design-point aid computes only postpones the touch-down to $t = 312.0$ min. A perfectly executed recipe scraps the run.

3.5 Control Structure

3.5.1 The Shared PID Algorithm

Both controllers are the velocity-form PID of Section 1.6.2, with the derivative on the filtered measurement and back-calculation anti-windup on the integral alone. Equation (3.15) writes it in this chapter's symbols: γ is the action a of Equation (1.12), Σ_k the running sum S_k and u_k^* the output before limiting v_k.

$$
\begin{aligned}
e_k &= \gamma(\mathrm{SP}_k - \mathrm{PV}_k) \\
P_k &= K_c e_k, \qquad D_k = -K_c \tau_D \gamma \frac{\mathrm{PV}_k^f - \mathrm{PV}_{k-1}^f}{\Delta t_c} \\
\Sigma_k &= \Sigma_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\
u_k^* &= u_{bias} + \Sigma_k + I_{k-1} \\
u_k &= \min(u_{\max}, \max(u_{\min}, u_k^*)) \\
I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - u_k^*}{T_t} \right] \Delta t_c
\end{aligned}
\tag{3.15}
$$

where $\gamma = \pm 1$ is the action ($+1$ reverse acting and -1 direct acting), PV^f is the measurement after a first-order filter of constant τ_F advanced by the same zero-order-hold recursion as the camera, Δt_c is the controller period and T_t is the anti-windup tracking constant, which defaults to τ_I. With the schedule of Equation (3.14) on, T_t follows the scheduled τ_I. At $T_t = 120$ min and $\Delta t_c = 0.1$ min each scan would remove $1/1200$ of the clipped excess, but neither loop reaches its output limits in any run reported here, so the tracking time affects none of the results.

Transfer between manual and automatic is bumpless as Section 1.6.2 describes, and the derivative estimate is kept warm in manual so that it is correct the instant automatic is selected. Neither loop carries a rate limit.

3.5.2 The Two Loops as Configured

Diameter controller DIC-301 is direct acting, $\gamma = -1$, because opening the pull rate makes the crystal thinner, in the convention of Section 1.6.2. Its setpoint is D_{SP}, its measurement is $\hat{D}$ from (3.10), and its output is clamped to the pull drive limits of 0.05 to 2.50 mm/min. The defaults are $K_c = 0.055$ mm/min per mm, $\tau_I = 120$ min, $\tau_D = 45$ min and $\tau_F = 8$ min, with the camera filter $\tau_{cam} = 0.50$ min ahead of it.

Mid-ranging controller VIC-301 is also direct acting, $\gamma = -1$, because more heater power slows the growth and the diameter loop answers by pulling more slowly, so a pull rate above its target calls for more heat. Its setpoint is V_{SP}, the pull rate the heater loop aims to leave the drive at. Its measurement is not the diameter and not the controller output either, but the pull rate actually applied, V_p, which is the state of (3.9). Its output is clamped to 30 to 220 kW, and its defaults are $K_c = 8$ kW per mm/min and $\tau_I = 120$ min with no derivative action. The mid-ranging loop can be switched out, in which case the heater holds its last command and cannot be moved by hand. Manual mode hands both the pull rate and the heater to the operator.

Pause. Try Exercise 3.2 before reading on.

The optional schedule applies Equation (3.14) on every controller execution, multiplying K_c by V_a/V_0 and dividing τ_I, τ_D and τ_F by the same ratio, with V_a floored at 0.05 mm/min and the ratio clamped to 0.2–4. Figure 3.2 reports what it does to the margin and Section 3.6 what it does on the

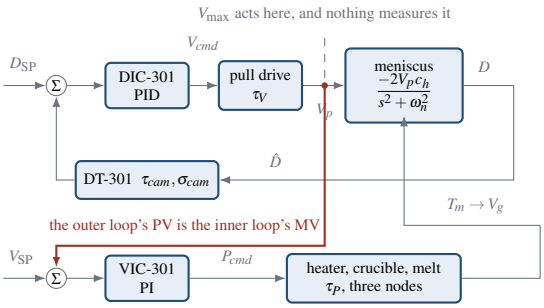


Figure 3.3 The mid-ranging structure as implemented. DIC-301 holds the diameter with the pull rate through a double integrator. VIC-301 never sees the diameter, takes the pull rate as its measurement, and walks the heater until the pull rate returns to V_{SP}. The unmeasured constraint V_{max} acts on the same signal VIC-301 regulates.

scenarios. Figure 3.3 shows the two loops and the single point at which they meet, which is the pull rate.

3.5.3 Model Predictive Control

An optional fourth structure, offered last in the Control list and loaded in charge only by scenario 11, replaces both loops in automatic mode with one model predictive controller (MPC) that moves the pull rate and the heater together. Every $T_s = 0.5$ min from $t = 0$ it solves one quadratic program over $N = 120$ steps (60 min) in move blocks of $\{1,1,1,1,2,2,3,4,6,9,15,75\}$ steps, then holds its first move until the following decision [19, 20], in the setup of Section 1.6.3. A 5 kW heater step moves the diameter only 0.88 mm in 30 min but 8.80 mm in 60 min.

The controller's model is not the puller of Section 3.3. It is the linearization of Equations (3.1) to (3.10) about the design point with the declared constants of Table 3.2 and the design disturbances ($\eta = \lambda = 1$, $\varepsilon_L = 0$), never the values of the run, so a constant a learner changes becomes model error hidden from the controller. Its nine states, written as departures from the design point with a tilde, are V_a, P_a, the three temperatures, h, r, the filtered diameter and one disturbance d, and its inputs are V_{cmd} and P_{cmd}. The lags are (3.9), the camera is (3.10) at its declared τ_{cam} without noise, the thermal chain is (3.7) frozen at the full charge ($C_m = c_p m_0$, $U_{cm} = U_{cm,0}$), and the remaining rows are Equation (3.16).

$$\begin{aligned} \frac{\mathrm{d}\tilde{h}}{\mathrm{d}t} &= \tilde{V}_a - \tilde{V}_g + d, \qquad \frac{\mathrm{d}\tilde{r}}{\mathrm{d}t} = -\bar{V} c_h \tilde{h}, \qquad \frac{\mathrm{d}d}{\mathrm{d}t} = 0 \\ \tilde{V}_g &= -\frac{V_0 r_0}{\bar{r}^2}\,\tilde{r} - K_T\,\tilde{T}_m \\ \tilde{q}_{lat} &= \frac{L_f \rho_{Si} \pi}{60}\left(2 r_0 V_0\,\tilde{r} + r_0^2\,\tilde{V}_g\right) \end{aligned} \tag{3.16}$$

where $\tilde{q}_{lat}$ enters the melt row of (3.7), and the radius path's operating point $\bar{V}$, $\bar{r}$ is reset at every decision to the measured V_a and $\hat{D}/2$, because the pull rate moves the gain and natural frequency of (3.12). The disturbance d carries, in growth-rate units, whatever the declared model omits, such as a lift error, a hot-zone change ($-V_g\,\Delta\eta$) or a draining melt's cooling ($K_T\,\Delta T_m$). An exact zero-order hold discretizes the model.

A central-difference Jacobian of the plant's derivative function at the design point agrees with every entry of the model to 9.8×10^{-11}, and its meniscus and radius pair, damped by the latent heat, sits at $-0.00149 \pm 0.01609\,i$ per minute. Stepped in manual mode against an unstepped run, noise off, from 30 to 120 min, the model's diameter is within 1.0 to 3.8 % of the plant's for pull-rate steps

of +0.01 and −0.05 mm/min and 3.2 to 6.9 % short for heater steps of ±2 and +10 kW, because the plant's melt drains.

Offset is removed by estimating d [21]. A steady-state Kalman filter at the 0.1 min controller period reads the two signals the shipped loops read, $\hat{D}$ at a variance of 0.012 mm^2 and V_a at 10^{-6}, with process noise of 10^{-5} on h, r and d only. On the model a step in d is 63 % estimated in 10 min, the slowest error mode is 84.7 min, and the filter runs under every structure. The estimate of d is the controller's only integral action. Each decision minimizes Equation (3.17) subject to Equation (3.18).

$$\begin{aligned} J = {} & \sum_{k=1}^{N} \left(2r_k - D_{\mathrm{SP}}\right)^2 + r_V \sum_{k=0}^{N-1} \left(V_{cmd,k} - V_{\mathrm{SP}}\right)^2 \\ & + \sum_{j=1}^{12} \left(s_V\, \Delta V_{cmd,j}^2 + s_P\, \Delta P_{cmd,j}^2\right) + \rho_D\, \varepsilon_D^2 + \rho_h\, \varepsilon_h^2 \end{aligned} \tag{3.17}$$

$$\begin{aligned} & 0.05 \le V_{cmd,j} \le \min\left(2.50,\ \hat{V}_{\max} - \delta_V\right) \\ & 30 \le P_{cmd,j} \le 220, \qquad \left|2r_k - D_{\mathrm{SP}}\right| \le \delta + \varepsilon_D \\ & h_{\min} + m_h - \varepsilon_h \le h_k \le h_{\max} - m_h + \varepsilon_h \end{aligned} \tag{3.18}$$

Here k counts prediction steps and j move blocks, the predicted true diameter $2r_k$ is weighted at 1 per mm^2, $r_V = 25$ per (mm/min)^2 prices the pull command's distance from V_{SP}, the target VIC-301 uses, and $s_V = 200$ per (mm/min)^2 and $s_P = 5$ per kW^2 price the change between blocks, the first from the commands applied now, with no terminal weight. The term in r_V is mid-ranging written as a price, and the heater has no reference of its own. Both input limits are hard and equal the loops' clamps, except that the pull rate is also held $\delta_V = 0.02\,\mathrm{mm/min}$ under $\hat{V}_{\max} = V_{\max,0}\, r_0/(\hat{D}/2)$, the controller's own estimate of (3.4) at the design hot zone, which the shipped loops are never given. The diameter band uses the tolerance δ at $\rho_D = 100$, and the meniscus band keeps the predicted meniscus $m_h = 1.5\,\mathrm{mm}$ inside the declared flooding and detachment heights at $\rho_h = 10^4$, both per mm^2. Each band has one nonnegative slack (ε_D, ε_h) for the whole horizon and is judged at every second step, so zero slack predicts a respected band without guaranteeing it at every step. The 24 block moves and two slacks are found by a dual active-set method, whose unconstrained first moves match independently formed normal equations to 1.2×10^{-11}.

A plan is rejected unless the solver reports it optimal with every move, prediction and cost finite and within bounds. On a rejection the commands hold for that scan, and DIC-301 and VIC-301, which track the applied commands, take over from the next scan until a valid plan returns. A rejection forced at 70.75 min of scenario 11 with the noise off moves neither command, and the loops then move at most 0.0076 mm/min and 0.050 kW per scan. The controller sees only $\hat{D}$, V_a, its commands, the setpoints, its settings and the declared constants, and a decision takes 0.40 to 0.47 ms on average in Node.js.

3.6 Results

Table 3.4 lists twelve scenarios in the order a laboratory session would take them. Each one pins the opening mode, the control architecture, the tuning constants, the four disturbance sliders and whatever timed events belong to it. Every figure reported below was read from the application's own scorecard using the physics, controller and logging periods of Section 3.3, and each run is carried to the end of the batch at $m = m_{end}$ or to the loss of structure, whichever arrives first, so an accumulated quantity such as IAE always carries its window with it. Pull-rate margin is the minimum of $V_{\max} - V_p$ over the run, evaluated at every 0.05 min physics step rather than at the two samples per minute the trend draws, because the quantity of interest is the closest the crystal actually came to losing its structure and not the closest a trend happens to show. The two differ by 0.018 to 0.025 mm/min on the runs of Table 3.5, always in the direction that flatters the loop. Every run starts from the scenario

Table 3.4 The Twelve Scenarios and What Each One Puts in Front of a Student

#	Scenario	What the student meets
1	The puller by hand	no steady state, meniscus flooding
2	Steady body growth	a plant that moves with nothing disturbed
3	4 mm setpoint step	derivative action on a double integrator
4	Argon upset, hot zone 1.04×	disturbance paid for in actuator position
5	Leaky hot zone plus argon upset	the structure limit the loop cannot see
6	Same run with mid-ranging	authority returned to the fast actuator
7	Crucible lift drifts –0.03 mm/min	an unmeasured ramp into the meniscus
8	Tail of the batch, $V_{SP} = 1.35\,\text{mm/min}$	a loop whose own output moves its plant
9	Camera noise 1.05 mm, $\tau_F = 4\,\text{min}$	derivative action on an optical measurement
10	Mid-range target 1.25 mm/min	an actuator that is also a specification
11	Two loops against one plan	both actuators in one predictive plan
12	Graded challenge, two upsets	the whole structure under a target

Table 3.5 Measured Outcome of Four Disturbance Sequences With the Diameter Loop Alone and With Mid-Ranging Added. Each Pair Is the Same Run With the Same Seed. Margin Is the Minimum of $V_{max} - V_p$ Over the Run at the 0.05 min Physics Step. Max $|\Delta D|$ Is on the Measured Diameter, Which Is What Both Loops and the Scorecard See, While the Prime Mass of (3.6) Is Scored on the True Diameter. The Last Pair Runs Scenario 9 With the Derivative Filter at 20 min.

Sc.	Heater	Margin (mm/min)	Max $\lvert\Delta D\rvert$ (mm)	Prime (kg)	Grown (kg)	Outcome
4	fixed	0.215	2.328	76.43	88.01	complete
4	mid-ranged	0.381	1.251	88.02	88.02	complete
5	fixed	0.221	2.349	73.46	88.00	complete
6	mid-ranged	0.358	1.375	88.00	88.00	complete
7	fixed	0.152	2.390	80.98	88.01	complete
7	mid-ranged	0.351	1.128	88.01	88.01	complete
9	fixed	−0.105	3.038	85.18	87.47	lost 483.0 min
9	mid-ranged	0.145	2.191	88.02	88.02	complete

as loaded, with any changed setting made at $t = 0$ before the run begins, which gives the same run as the setting carried through Reset.

3.6.1 The Loop Works and the Run Is Lost

Scenario 4 applies a hot-zone change of 1.04× at $t = 60\,\text{min}$, which makes the crystal solidify faster than it is being pulled. The meniscus shortens, the growth angle rises, the crystal fattens, and the diameter loop answers by pulling faster until growth and pull match again. The diameter is held 1.0 to 1.5 mm over 300 mm from $t = 130\,\text{min}$ to 390 min, inside its band, against 0.4 to 1.1 mm over in the same run without the upset, and everything looks resolved. It is not resolved. The pull rate is permanently higher, the margin to V_{max} is permanently smaller, and nothing in the interface reports that as a problem.

Table 3.5 runs four disturbance sequences with and without the mid-ranging loop. Every pair shares its disturbances, its tuning and its seed, and differs only in whether the heater is driven by the pull rate. Figure 3.4 draws the margin and the prime crystal of each pair.

Scenario 5 is the case the application is built around. The hot zone leaks 12 % more heat than the recipe assumed and an argon upset arrives on top of it at $t = 120\,\text{min}$. The diameter loop holds the diameter with a maximum deviation of 2.349 mm and never saturates, the pull rate ends the run at 1.378 mm/min against a structure limit of 1.651 mm/min, and the worst margin over the run is 0.221 mm/min, reached at $t = 447.3\,\text{min}$. Structure is never lost. The crystal survives, and only 73.46 kg of the 88.00 kg grown is prime, a yield of 83.5 %. Scenario 6 is the identical run with the

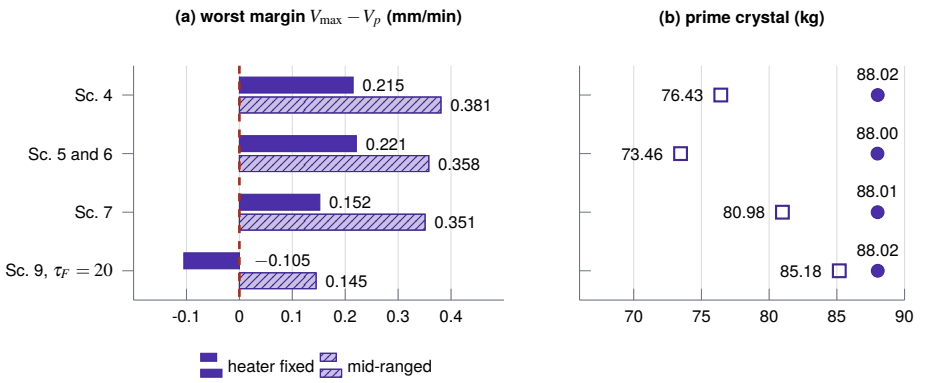


Figure 3.4 The four pairs of Table 3.5, each the same run with the heater fixed and with it mid-ranged. (a) The worst pull-rate margin over the run: solid bars for the fixed heater, hatched bars for the mid-ranged one, and the dashed line at zero, where the pull rate meets the structure limit. (b) The prime crystal: open squares for the fixed heater, filled circles for the mid-ranged one. Notice that mid-ranging widens the margin in every pair and keeps the whole of each crystal prime, while the fixed heater gives up prime crystal even in the runs that keep their structure.

heater driven by the pull rate. The margin never falls below 0.358 mm/min, the yield is 100 %, the maximum diameter deviation falls from 2.349 to 1.375 mm, and all 88.00 kg is prime, which at the flat valuation the application carries is \$79,200 against \$66,116. The diameter traces of scenarios 5 and 6 look alike, and their outcomes differ by 14.54 kg of product, so the actuator and not the controlled variable is the trace that tells them apart. The heater ends that run at 138.2 kW having started at 133.8 kW.

> **TRY IT IN THE APP** **5 · The Limit the Pull Rate Walks Into**
>
> Load the scenario and press Run. Watch the pull-rate strip: the pull rate climbs toward the red dashed structure limit while the diameter trace changes little. The closest approach, 0.221 mm/min, comes at 447.3 min, and only 73.46 kg of the 88.00 kg grown is prime. Reset, set Control to Diameter + mid-ranging and run again: the margin never falls below 0.358 mm/min, and all 88.00 kg is prime.

Scenario 7 makes the same point through a mechanism nobody is watching. A crucible lift that under-compensates by 0.03 mm/min is three percent of the pull rate, it enters the pure integrator of (3.1) directly, and no instrument in the application sees it. The diameter loop holds the diameter by offsetting the pull rate further and further, ending the run at 1.315 mm/min with a worst margin of 0.152 mm/min and 80.98 kg of 88.01 kg prime. With mid-ranging on, the heater absorbs the drift instead, moving from 118.26 to 121.59 kW over the run across a logged span of 117.56 to 122.09 kW, while the pull rate ends at 1.133 mm/min, the worst margin more than doubles to 0.351 mm/min, and the whole 88.01 kg is prime. A mid-ranged actuator that ends a batch somewhere it did not start is a diagnostic rather than a fault, because something upstream moved and reading the slow loop is often the only record of it.

The fourth pair is the one that loses a crystal. Scenario 9 degrades the camera to $\sigma_{cam} = 1.05$ mm, and with the derivative filter opened to 20 min so that the noise is not itself the failure, the diameter loop alone still walks the pull rate across V_{max} and holds it there past the dwell at $t = 483.0$ min, with 85.18 kg of 87.47 kg prime and everything after that instant polycrystalline. The same run mid-ranged finishes the batch at $t = 506.1$ min with a worst margin of 0.145 mm/min and 88.02 kg prime. Nothing about the diameter loop failed in the first of those runs. It held the diameter to within 3.038 mm using the only actuator it had, until that actuator ran out. The loss itself belongs to the shipped noise

Table 3.6 Noise Injected Into the Pull Rate by the Derivative Term, as the RMS Difference Between Paired Runs With and Without Camera Noise Over $60 \le t \le 240$ min in Scenario 6 With the Timed Upset Removed. 8 min Is the Default Filter, Which Every Scenario but Scenario 9 Ships With.

σ_{cam} (mm)	τ_F (min)	τ_D (min)	RMS ΔV_{cmd} (mm/min)	RMS ΔD (mm)
0.35	8	45	0.0429	0.0475
1.05	8	45	0.1288	0.1419
1.05	3	45	0.2886	0.1194
1.05	20	45	0.0650	0.1764
1.05	8	0	0.0275	0.3738

sequence. Over eight other seeds the diameter loop alone keeps its structure, with worst margins of -0.0004 to 0.174 mm/min and 83.9 to 85.8 kg prime, while the mid-ranged run keeps 88.0 kg prime and a worst margin of at least 0.186 mm/min on every seed, so the prime difference is robust and the loss is not.

3.6.2 Derivative Action Against an Optical Measurement

The price of the 45 min derivative time is paid on the camera. Table 3.6 measures it by running each case twice with identical seeds, once with the camera noise on and once with it off, and reporting the root-mean-square difference between the two trajectories over $60 \le t \le 240$ min. The differencing removes the drift of the batch and leaves only what the noise injected.

At the default filter of 8 min the camera noise carried into the pull-rate command rises from 0.0275 mm/min RMS without derivative action to 0.1288 mm/min with it, and to 0.2886 mm/min at $\tau_F = 3$ min. Recording the proportional and derivative terms of DIC-301 separately in the same paired runs isolates the derivative's own share, which is 5.4 times the proportional term's against the ratio $\tau_D/\tau_F = 5.6$ that the filtered derivative predicts, 13.7 times at $\tau_F = 3$ min against 15.0, and 2.2 times at 20 min against 2.25. The agreement is the point. Derivative noise gain on this loop is arithmetic, not a tuning surprise. The crystal itself barely registers any of it, because the meniscus integrates the command twice and the RMS diameter perturbation at the default filter is only 0.1419 mm, but the pull drive registers it and so does anybody reading the trend. The last row is not an option. With $\tau_D = 0$ the command noise falls to 0.0275 mm/min, but the loop has no phase margin and its own oscillation grows through the batch. The two runs of that pair drift apart as their oscillations grow, which is why its diameter difference is the largest in the table, and run to the end of the batch the noisy case leaves the diameter band at $t = 351.65$ min with a worst measured deviation of 3.99 mm.

Lengthening the filter is not free either. Going from 8 to 20 min halves the injected command noise to 0.0650 mm/min and adds phase lag exactly where (3.12) has none to spare, so the RMS diameter perturbation rises from 0.1419 to 0.1764 mm, and Section 3.4.1 prices the 3 to 8 min half of that move at 21.7° of phase margin. Scenario 9 ships with the schedule enabled, the camera degraded to $\sigma_{cam} = 1.05$ mm and the derivative filter shortened to 4 min, a noise gain of $\tau_D/\tau_F = 11.25$ that keeps the crystal on every one of 1001 noise seeds with the clean 0.35 mm camera. As configured the pull rate spikes to 1.79 mm/min, crosses the structure limit and stays there until the crystal is lost at $t = 113.3$ min, with 18.96 kg of prime crystal grown before that instant, and 1000 of 1001 seeds lose it. Raising the derivative filter to 12 min, the setting the scenario asks for, rescues the shipped run, which completes at $t = 506.1$ min with 88.01 kg prime but comes within 0.044 mm/min of the limit, and the rescue holds on 986 of 1001 seeds. That change costs 23.0° of phase margin on the loop of Table 3.3, from 44.3 to 21.4°. At the 8 min filter the application uses elsewhere the shipped seed also loses the crystal, at $t = 333.45$ min, but only 43 of 201 seeds do, and 20 min completes the shipped run with a worst margin of 0.145 mm/min. That trade is the permanent cost of stabilizing a double integrator with derivative action, and it is why a real puller spends its engineering on the camera rather than on the controller.

TRY IT IN THE APP 9 · A Noisy Bright Ring Against a Long Derivative Time

Run it as loaded. The pull rate spikes to 1.79 mm/min, crosses the structure limit and stays there, and the crystal is lost at 113.3 min with 18.96 kg of prime crystal. Reset, switch the panel to Advanced, set Derivative filter to 12 min and run again: the batch completes at 506.1 min with 88.01 kg prime, though the pull rate comes within 0.044 mm/min of the limit.

3.6.3 An Actuator That Is Also a Specification

The ratio of growth rate to the axial temperature gradient at the interface decides whether the crystal freezes in an excess of vacancies or an excess of self-interstitials, and only a narrow band of that ratio produces material free enough of both to make a device wafer [5]. The prediction was confirmed experimentally by varying the gradient and locating the defect boundaries against it [6]. The application applies the resulting specification directly to the growth rate through the band $[V_{lo}, V_{hi}] = [0.72, 1.32]\,\mathrm{mm/min}$ in (3.6).

Scenario 10 asks the mid-ranging loop to leave the pull rate at 1.25 mm/min instead of 1.00 mm/min for throughput. The diameter is held to a maximum deviation of 1.770 mm, the batch completes 65.9 min earlier at $t = 440.0\,\mathrm{min}$, and prime crystal falls from 88.01 to 74.51 kg. The growth rate leaves the band at $t = 385.2\,\mathrm{min}$, and 12.2 kg of the loss is grown above 1.32 mm/min from then on, while the other 1.2 kg is lost to a diameter excursion between 141 and 167 min. Throughput bought 65.9 min and sold 13.49 kg. The same run also drops its worst margin from 0.329 to 0.008 mm/min. With the diameter loop alone the heater stays at its design power and the target no longer acts, and the margin stays at or above 0.147 mm/min with 87.02 kg of 88.01 kg prime, so it is the throughput target and not the mid-ranging loop that spends the margin. An actuator that is simultaneously the manipulated variable of the primary loop and a controlled variable with a two-sided specification is an awkward object, and mid-ranging is what makes it tractable, because the diameter loop is free to move the pull rate in the short term while the heater sets where its average sits.

3.6.4 Model Predictive Control Against the Two Loops

Scenario 11 is the leaky hot zone of scenario 6 with the diameter setpoint stepped from 300 to 304 mm at 60 min and an argon upset at 240 min, loaded with the controller of Section 3.5.3 in charge. Table 3.7 runs it, scenario 3 and the ceiling case of scenarios 5 and 6 under the shipped loops, the same loops retuned once on scenario 11 and the controller as scenario 11 ships it. One rule chose the controller's weights from 300 candidates and the retuned loops by pattern search, the least true-diameter IAE on scenario 11 with no more travel on either actuator and no less worst margin than the shipped loops. For the loops it gives the schedule on, $K_c = 0.055$, $\tau_I = 20\,\mathrm{min}$ (its lower limit), $\tau_D = 30\,\mathrm{min}$ and $\tau_F = 8\,\mathrm{min}$ on DIC-301 and $K_c = 10\,\mathrm{kW}$ per mm/min, $\tau_I = 55\,\mathrm{min}$ on VIC-301. Nothing is retuned per case, each triple shares its seed and events, neither structure receives a preview, and each IAE has its own batch window. The loop comparison selects Diameter + mid-ranging explicitly in scenario 3, a case that ships with Diameter only, so "shipped loops" here names the gains, not the initial architecture of scenario 3.

Against the shipped loops the controller lowers the diameter IAE from 378.5 to 81.0 mm min on scenario 3, from 252.0 to 71.1 on the ceiling case and from 462.0 to 119.1 on scenario 11. Most of that gap is the shipped tuning. The retuned loops reach 149.6, 74.3 and 176.0 mm min, closing 77, 98 and 83 % of it, mainly through a 20 min integral time that keeps pace with the draining melt, and keep more margin than the controller on all three cases (0.338, 0.451 and 0.325 mm/min against 0.325, 0.403 and 0.313). No run in Table 3.7 comes within 0.08 mm/min of the structure limit, and removing the controller's ceiling bound leaves its three runs unchanged.

Table 3.7 Shipped Loops, Retuned Loops and Model Predictive Control on the Same Batches, Seed 20260918, Each Run to the End of Its Batch

Structure	IAE (mm min)	Margin (mm/min)	Pull (mm/min)	Heater (kW)	Prime (kg)	End (min)
Scenario 3, setpoint step						
shipped loops	378.5	0.386	1.46	12.9	84.63	494.9
retuned loops	149.6	0.338	1.28	15.2	85.01	504.6
predictive	81.0	0.325	1.65	11.8	85.80	512.2
Scenarios 5 and 6, ceiling case						
shipped loops	252.0	0.358	35.88	106.0	88.00	492.8
retuned loops	74.3	0.451	27.77	101.4	88.00	508.8
predictive	71.1	0.403	18.02	56.4	88.02	519.1
Scenario 11, setpoint step and argon upset						
shipped loops	462.0	0.309	35.60	105.9	84.44	484.4
retuned loops	176.0	0.325	27.53	104.0	85.02	497.3
predictive	119.1	0.313	18.32	58.2	85.64	507.8

IAE is the scorecard figure on the measured diameter, margin the minimum of $V_{max} - V_p$ at every 0.05 min physics step, and pull and heater the travel of the logged commands. The shipped ceiling row is scenario 6. The controller made 1016 to 1039 plans per run and none was rejected.

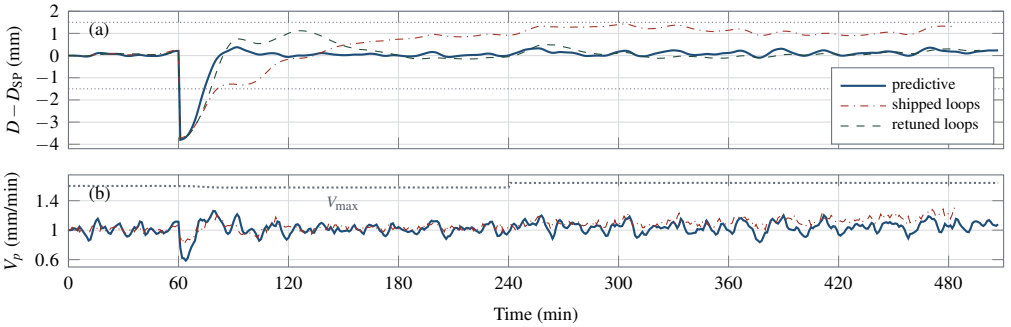


Figure 3.5 Scenario 11 under three structures to the end of each batch, with the setpoint step at 60 min and the argon upset at 240 min. The predictive run is the thick solid line, the shipped loops the dash-dotted line and the retuned loops the dashed line. (a) True diameter minus setpoint, with the tolerance as thin dotted lines. (b) Applied pull rate, with the structure limit of the predictive run as the thick gray dotted line.

✦ DID YOU KNOW? Better Tracking With Less Travel

From the simulator: on scenario 11 in Table 3.7 the predictive controller moves its actuators about half as far as the shipped loops, 18.32 mm/min of pull-rate travel against 35.60 and 58.2 kW of heater travel against 105.9, while it cuts the diameter IAE from 462.0 to 119.1 mm min. Better tracking need not cost more actuator movement.

The remainder, 176.0 against 119.1 mm min on scenario 11, belongs to the pull-rate plan, a model of the meniscus driven by the estimate of d. With the heater command frozen at its initial value inside the plan the IAE is still 130.5, while the margin falls to 0.136 mm/min and the prime crystal to 79.65 kg, which is what the heater buys, as under mid-ranging. With no process noise on d, and so no estimate, the IAE is 1700.4. The controller cuts the pull rate to 0.59 mm/min within four minutes of the step, and at 80 min the true diameter is 303.74, 303.08 and 302.41 mm under the controller, the retuned loops and the shipped loops, which reach 303.5 mm at 78.60, 81.85 and 111.35 min, as shown in Figure 3.5. From 250 to 300 min the shipped loops ride 1.14 to 1.40 mm high, because a 120 min integral time never catches the drain, against −0.02 to 0.33 under the controller and 0.05 to 0.49 under the retuned loops.

Table 3.8 Graded Challenge as Supplied and With the Suggested Configuration. Both Runs Carry the Same Disturbances, the Same Timed Events and the Same Seed, and Every Figure Is Read at the End of the Batch. Max $|\Delta D|$ Is on the Measured Diameter, While the Challenge Rule Judges the True Diameter.

	As supplied	Suggested
Architecture	DIC-301 only	DIC-301 + VIC-301
Schedule of Eq. (3.14)	off	on
K_c DIC-301 (mm/min per mm)	0.012	0.055
τ_I DIC-301 (min)	120	90
τ_D DIC-301 (min)	8	45
K_c VIC-301 (kW per mm/min)	2	8
τ_I VIC-301 (min)	600	120
Max $\|\Delta D\|$ (mm)	12.856	0.900
Final diameter (mm)	312.91	300.75
Worst margin (mm/min)	0.195	0.398
Prime crystal (kg)	12.80	88.02
Diameter IAE (mm min)	2313	189
Batch duration (min)	453.8	495.8
Value ($)	11,521	79,215

The controller does not win every measure. On the ceiling case the faster mid-ranging scenario 6 asks for (60 kW per mm/min, 10 min) keeps 0.454 mm/min of margin against 0.403. The margin of scenario 11 ties the shipped loops' only on the shipped noise sequence, and over ten other seeds it is 0.166 to 0.322 mm/min, below the shipped loops' 0.292 to 0.360 on all ten and the retuned loops' 0.208 to 0.308 on seven, while its IAE of 114.6 to 119.3 beats the retuned loops' 172.3 to 179.0 on all ten. Over nine cases (adding scenarios 4, 7, 8, 9, 10 and the challenge) the controller has the lower IAE on eight and the retuned loops the larger margin on eight, the challenge tying at 0.509 against 0.508. In the tail of scenario 8 it keeps 7.73 kg prime of 18 against 15.86 for the retuned loops, because its price draws the pull rate toward the 1.35 mm/min target above the defect band (all 18.00 kg with $r_V = 0$), as Section 3.6.3 found for mid-ranging.

Scenario 9 is the case the controller loses on IAE, 178.1 against 151.3 mm min for the retuned loops. As loaded, the shipped loops lose the crystal on all eleven noise seeds tested and the controller keeps it on all eleven, but on only six without its ceiling bound. Given $\hat{V}_{\max} - 0.02\,\mathrm{mm/min}$ as the output limit of DIC-301, the shipped loops also keep it on all eleven, so the crystal is saved by that estimate, which the loops can carry as a clamp, and not by the plan.

3.6.5 The Graded Challenge

The challenge scenario applies a hot-zone upset at $t = 90\,\mathrm{min}$ and a crucible lift drift at $t = 300\,\mathrm{min}$ on a hot zone already leaking 8 %, and asks for the diameter held within $\pm 1.5\,\mathrm{mm}$ of 300 mm for the whole batch with the pull rate still clear of the structure limit at the end. It opens with a soft diameter loop, a short derivative time, camera noise of 0.20 mm and the mid-ranging loop switched off, with a detuned mid-ranging setting waiting behind it. Table 3.8 reports both configurations.

The supplied tuning does not oscillate. It simply never catches up, and the crystal finishes 12.91 mm oversize with a maximum deviation of 12.856 mm and 12.80 kg of prime crystal out of 88.02 kg grown. Evaluating the loop of Table 3.3 at the supplied $K_c = 0.012$ and $\tau_D = 8\,\mathrm{min}$ gives a phase margin of 4.6° at a crossover period of 154.6 min, so the derivative time is a twentieth of the closed-loop response rather than comparable to it, and the soft proportional gain is the only thing keeping the loop out of oscillation. The suggested configuration meets the target with a maximum deviation of 0.900 mm, a worst margin of 0.398 mm/min and the full 88.02 kg prime. Running the same suggested tuning with the schedule switched off gives a maximum deviation of 0.999 mm

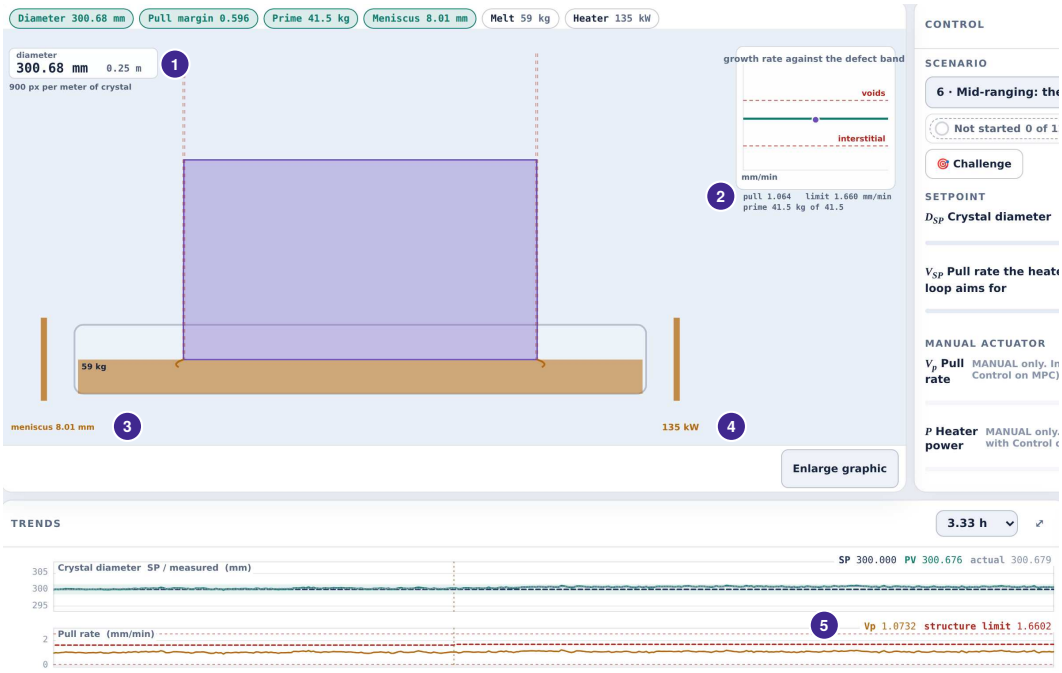


Figure 3.6 Scenario 6 running at $t = 4.00$ h, with the leaking hot zone and the argon upset already absorbed, cropped to the plant view and the first two trend strips. The diameter holds 300.68 mm against a 300.0 mm setpoint (1). The pull-rate command reads 1.073 mm/min on the second strip (5) against the structure limit drawn as a red dashed line at 1.660 mm/min, and the applied pull rate of 1.064 mm/min, printed with the limit under the growth-rate inset (2), leaves 0.596 mm/min of margin. The heater has been walked to 135.1 kW by the mid-ranging loop (4), the meniscus is at 8.01 mm between its flooding and detachment limits (3), and all 41.5 kg grown so far is prime (2).

at an essentially unchanged worst margin of 0.406 mm/min, so the schedule is worth 0.100 mm of diameter on this scenario and nothing at all in margin, which is what a loop-shape correction should be worth on a run that never approached the limit. The benefit is small but consistent, between 0.06 and 0.17 mm over nine noise seeds and positive on every one.

The challenge rule itself judges the true diameter at every logged sample and the margin at the end of the batch. Judged that way, the supplied settings end the batch at 453.8 min with a worst diameter error of 12.91 mm, more than eight times the band, and 0.195 mm/min of margin, so the attempt fails. The suggested settings end it at 495.8 min with a worst error of 0.75 mm, inside the band with 0.75 mm to spare, and 0.439 mm/min of margin.

Figure 3.6 shows scenario 6 four hours into the batch. The strip that carries the argument is the second one, where the pull rate is plotted against the moving structure limit, because the distance between those two traces is the quantity the mid-ranging loop exists to protect and the one the diameter trend above it cannot show.

3.7 Technology in Practice

The loops of Section 3.5 run on a camera, a pull drive, a heater supply and a controller, drawn together in Figure 3.7.

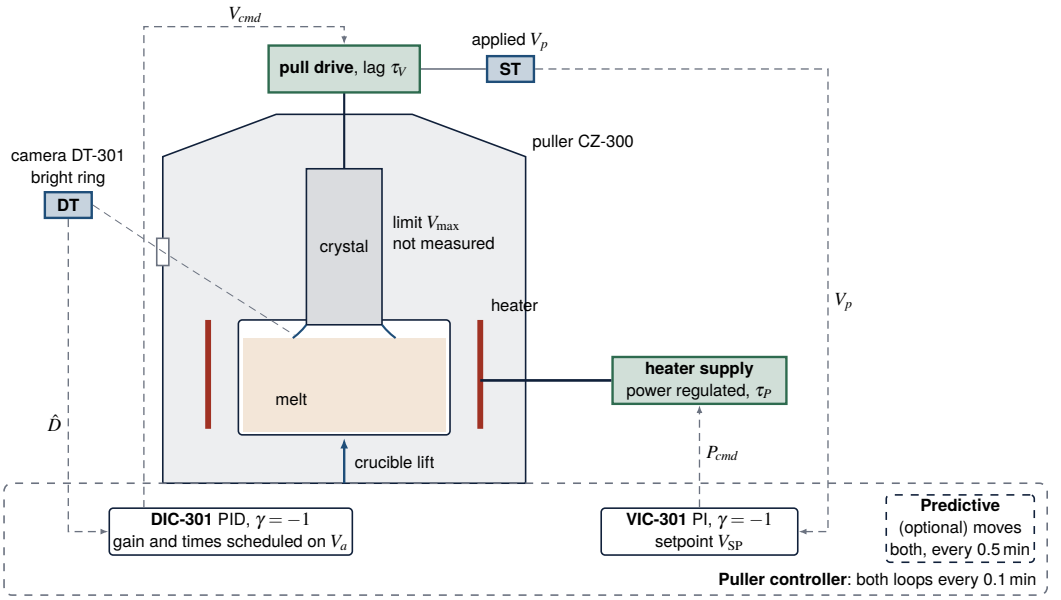


Figure 3.7 The loops of puller CZ-300 drawn with their hardware. Camera DT-301 sights the bright ring through a viewport, and DIC-301 sets the pull drive's command with its gain and times scheduled on the applied pull rate. VIC-301 reads that applied pull rate and sets the power of the heater supply. Both loops are direct acting ($\gamma = -1$). The optional predictive controller replaces both when it is selected. Notice that the structure limit has no instrument: VIC-301 protects it only by keeping the pull rate near V_{SP}.

What Measures It

The diameter loop reads camera DT-301, which the model reduces to noise of 0.35 mm by default and a filter of 0.50 min (Equation (3.10)).

> **TECHNOLOGY** **Sensor: A Camera on the Bright Ring**
>
> The bright ring is the reflection of the hot crucible wall in the curved meniscus where the crystal meets the melt. Optical pyrometers, photocells, line-scan cameras and two-dimensional cameras have all been used to find it. A camera in a viewport sees the ring at an angle, as an ellipse, and on a large crystal the crystal itself can hide the ring's widest point. Such instruments need frequent calibration by the operator [22]. A crystal that swings in a small orbit moves the ring, and the camera reports a diameter change that is not there. The loop then answers it with pull rate or heater and loses true control of the diameter [23]. A slower failure is a viewing window clouded by deposits, which weakens the signal a pyrometer receives through it [17].

The heater loop reads the pull rate the drive applies. The 1972 design weighed the crystal instead and controlled the ratio of weight to length, which is proportional to its cross-section [17], the weighing route of Section 3.2.

What Moves It

The pull drive follows its command through a lag of 0.30 min and is clamped at its mechanical limit of 2.50 mm/min, above the structure limit. A crystal growing fat is answered by pulling faster, so DIC-301 is direct acting, $\gamma = -1$ in the convention of Section 1.6.2.

TECHNOLOGY Final Element: The Pull Drive

The crystal hangs from a pull wire on a lifting mechanism at the top of the puller. The mechanism lowers a small seed into the melt and then withdraws it slowly, turning the crystal one way while the crucible turns the other [22]. A second drive raises the crucible. In the 1972 design one controller output set both drives through a distributor that kept the ratio of pull rate to crucible lift constant [17]. Because that loop needed strong derivative action, a step in its measurement became a jump in pull rate: a load cell with too large a dead band made the pull rate change abruptly and spoiled the resistivity along the crystal [17]. On a trend it is the pattern of Table 3.6, a busy pull rate under a calm diameter.

The heater supply follows its command through a lag of 0.50 min. The 1972 design already regulated power rather than temperature: an inner loop set the firing angle of thyristors from a measurement of the heater's power, which caught mains voltage swings before they warmed the heater and cut power fluctuations to a hundredth or less [17]. VIC-301 is direct acting as well, because a pull rate above its target calls for more heat.

Where the Controller Runs

Nested loops, in which one loop's setpoint is based on another loop's process variable, run with cycle times from milliseconds to minutes, and a programmable controller can serve as the primary controller of a smaller system [24]. Industrial pullers have used cascades of PID controllers [2]. The app executes both loops every 0.1 min against a physics step of 0.05 min, short beside the crossover period of 36.3 min. In a control system the structure is two standard blocks: a PID block whose gain and times the schedule multiplies, and a PI block whose process variable is the first block's actuator. The predictive controller of Section 3.5.3, which decides every 0.5 min, replaces both blocks with one plan, and the two blocks take over from the next scan whenever a plan is rejected.

IN PRACTICE What a Production Puller Controls

A study of an industrial 300 mm puller lists four controlled variables, the crystal radius, the growth rate, the heater temperature and the position of the melt surface, and three manipulated ones, the pull rate, the crucible lift rate and the heater power. Cascade control with PID controllers has been the industrial practice. The heater acts on the controlled variables with a long time constant, and the relation changes as the melt drains [2]. The diameter camera needs frequent calibration by the operator to keep the crystal within specification [22], and a camera bumped after calibration reads a wrong diameter [23]. The two loops of this chapter are the pull-rate and heater part of that cascade.

What Is Modernizing It

Three developments bear on this machine. The first is vision: two cameras on opposite sides of the crystal tell a swinging crystal, a moving melt level and a bumped camera from a real change of diameter [23]. The second is models trained on plant data. A gray-box model of a 300 mm puller, built toward predictive control, learns the hidden interface gradient from production records [2]. The third changes the hot zone. Magnetic fields that suppress melt convection and continuous feeding of liquid silicon were both listed as major developments in 1985, the second then far from a reliable production process [1]. A continuous feed would hold the melt inventory and remove much of the drift that makes this batch plant non-stationary.

3.8 Limitations

The model is not intended for hot-zone design. There is no radiation view factor, no heat shield geometry, no crystal or crucible rotation and no magnetic field, so the melt flow that dominates a real puller is absent and the thermal chain is three lumped nodes standing in for it [13]. The meniscus is the linearized form rather than the capillary solution, which restricts the model to body growth and places the flooding and detachment heights by hand. There is no axial gradient field, so the

defect criterion of (3.6) is applied to the growth rate alone and would have to be reformulated on V/G before a hot-zone change could be studied [5]. Oxygen, carbon and dopant transport are absent entirely, which removes resistivity control, striations and axial segregation, all of which are real control problems on the same machine.

Two limitations concern the control side specifically. The structure limit is a latch with a dwell rather than a progressive breakdown of the interface, so the model offers partial credit only in the sense that a short excursion is survivable, and it gives no early warning in the ring image at all. The dwell t_{lose} is one constant standing in for dislocation nucleation kinetics, and a real interface is more forgiving of a brief excursion at small diameter than of the same excursion at 300 mm, which the model does not represent. Exact repeatability is itself a fiction, since a frozen step and a frozen scan give every machine the same run and no furnace behaves that way. Three features came through every simplification, the double integrator through the meniscus, the unmeasured process constraint on the fast actuator and the non-stationary batch, and each one is present because it changes what the controller has to do. Section 3.3.7 says where each individual assumption gives out.

The predictive controller's model is linearized once at the full charge and the design hot zone, so the draining melt, a changed hot zone, a lift error and any edited constant reach it only through d. Its meniscus band rests on an estimate off by up to 0.65 mm in scenario 11 and 1.97 mm in scenario 9, the tangent of (3.2) and the latches are absent, and the weights were tuned once, on scenario 11.

> ⚠ **OPEN CHALLENGE** **The Gradient Nobody Measures**
>
> A predictive controller for a production puller needs a model of the whole batch, and most published control models were built for crystals under 200 mm. The temperature gradient in the crystal at the interface cannot be measured, and a first-principle model that held it constant fell short. A model that learned the gradient from crystal length and melt temperature in plant data cut the errors most in the late phase of the batch, where the gradient changes fastest [2]. The same gradient decides the defect band of Section 3.6.3. Progress would be a gradient estimate checked during the batch, so that the defect specification could be held as a ratio and not through the growth rate alone.

3.9 Discussion

The economics behind Table 3.5 are not subtle. A run occupies a furnace for eight hours, consumes a hundred kilograms of polysilicon and the electricity to melt and hold it, and the difference between the two scenario 9 rows in Table 3.5 is 2.84 kg of prime crystal and, with the shipped noise sequence, the single-crystal structure of everything grown after $t = 483.0$ min. Czochralski growth supplies essentially all of the single-crystal silicon used for semiconductors and a large share of that used for photovoltaics, where the energy consumed in crystal growth and wafering is one of the larger contributions to the energy payback time of a finished module [25]. A control structure that converts a lost run into a completed one is therefore an energy measure as well as a yield measure.

The equations given here also carry methods beyond the two loops. The ten-state system with an explicit constraint on a manipulated variable is a natural model predictive control problem, because $V_{\max}$ of (3.4) can enter as a state-dependent output constraint instead of being defended by a second loop, and the heater and the pull rate can be coordinated as one two-input problem rather than two, as Section 3.5.3 does in linear form. The unmeasured crucible lift error of scenario 7 is a textbook target for moving horizon estimation on the same equations, since it enters (3.1) where the diameter and the pull rate together carry enough information to separate it from a hot-zone change. Model-based designs of both kinds have been reported on real pullers [9–11], including the observer that reconstructs the radius from a weighing signal and the scheduling of controller parameters along the batch. Mid-ranging remains the useful comparison for all of them, because it is the answer available without a model and the one an engineer can commission in an afternoon [3, 4].

3.10 Conclusion

This chapter has used a Czochralski puller as an interactive laboratory for the case where the fast actuator is also the limited one. Ten states are integrated, the crystal radius and the meniscus height, a three-node thermal chain from heater to crucible to melt, the melt inventory, the crystal length, two actuator lags and the prime crystal grown so far, with the diameter reaching the pull rate only through the meniscus and therefore through two integrations. Without derivative action the loop has a phase margin of −6.2°, the diameter overshoots a 4 mm setpoint step by 5.8 mm, and the meniscus floods at 158.4 min, while 45 min of derivative time gives 29.2° of margin against a crossover period of 36.3 min. Adding a proportional-integral loop that drives the heater from the pull rate rather than from the diameter raises the worst pull-rate margin on the same disturbances from 0.221 to 0.358 mm/min and the prime crystal from 73.46 to 88.00 kg, and on a degraded bright-ring camera it is the difference between losing structure at 483.0 min and completing the batch with the shipped noise sequence. Both the gain and the natural frequency of the diameter plant are proportional to the pull rate, so scheduling K_c upward and τ_D, τ_I and τ_F downward in that variable holds the phase margin of the meniscus loop at 29.19° across the whole drive range, where a fixed tuning slides from 38.1 to 21.3° and scaling the gain and the derivative time the other way is worse than no schedule at all below the design pull rate. The camera filter and the drive lag are inside the loop and outside the schedule, and they are what leaves 25.2 to 13.9° in place of a flat curve on the complete loop. An optional predictive controller lowers the diameter IAE of scenario 11 from 462.0 to 119.1 mm min, and loops retuned by the same rule reach 176.0 with more margin. Further work would add a capillary meniscus, an axial gradient field so that the defect criterion becomes a ratio, melt convection as a disturbance with structure, and the shoulder as an open-loop trajectory problem. The Czochralski puller is an accessible browser platform for learning that an actuator limit is a control problem for a different loop from the one using the actuator.

KEY POINTS

- The pull rate reaches the diameter only through a short column of melt, two integrations away.
- Without derivative action the phase margin is −6.2° and the meniscus floods. A 45 min derivative time gives 29.2°.
- The diameter loop holds the diameter while the pull rate walks toward a structure limit that nothing measures.
- Driving the heater from the pull rate raises the worst margin from 0.221 to 0.358 mm/min and the prime crystal from 73.46 to 88.00 kg.
- A schedule on the pull rate holds the meniscus loop at 29.19°, but the camera and drive lags leave 25.2 to 13.9°.

Where else this applies. Mid-ranging a fast actuator with a slow one in Volume 4, Chapter 10 and Volume 4, Chapter 9, and a gain schedule along a batch in Chapter 2.

References

[1] J. H. Matlock, “Future application of Czochralski crystal pulling for silicon,” in *Proceedings of the Flat-Plate Solar Array Project Workshop*. Jet Propulsion Laboratory, 1985, NASA Technical Reports Server 19860010255.

[2] S. Kato, S. Kim, M. Kano, T. Fujiwara, and M. Mizuta, “Gray-box modeling of 300 mm diameter Czochralski single-crystal Si production process,” *Journal of Crystal Growth*, vol. 553, p. 125929, 2021.

[3] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York, NY: McGraw-Hill, 1996.

[4] B. J. Allison and A. J. Isaksson, “Design and performance of mid-ranging controllers,” *Journal of Process Control*, vol. 8, no. 5-6, pp. 469–474, 1998.

[5] V. V. Voronkov, “The mechanism of swirl defects formation in silicon,” *Journal of Crystal Growth*, vol. 59, no. 3, pp. 625–643, 1982.

[6] W. von Ammon, E. Dornberger, H. Oelkrug, and H. Weidner, "The dependence of bulk defects on the axial temperature gradient of silicon crystals during Czochralski growth," *Journal of Crystal Growth*, vol. 151, no. 3-4, pp. 273–277, 1995.

[7] D. T. J. Hurle, "Control of diameter in Czochralski and related crystal growth techniques," *Journal of Crystal Growth*, vol. 42, pp. 473–482, 1977.

[8] ——, "Analytical representation of the shape of the meniscus in Czochralski growth," *Journal of Crystal Growth*, vol. 63, no. 1, pp. 13–17, 1983.

[9] M. A. Gevelber and G. Stephanopoulos, "Dynamics and control of the Czochralski process," *Journal of Crystal Growth*, vol. 84, no. 4, pp. 647–668, 1987.

[10] J. Winkler, M. Neubert, and J. Rudolph, "Nonlinear model-based control of the Czochralski process I: Motivation, modeling and feedback controller design," *Journal of Crystal Growth*, vol. 312, no. 7, pp. 1005–1018, 2010.

[11] M. Neubert and J. Winkler, "Nonlinear model-based control of the Czochralski process III: Proper choice of manipulated variables and controller parameter scheduling," *Journal of Crystal Growth*, vol. 360, pp. 3–11, 2012.

[12] J. Czochralski, "Ein neues verfahren zur messung der kristallisationsgeschwindigkeit der metalle," *Zeitschrift für Physikalische Chemie*, vol. 92U, no. 1, pp. 219–221, 1918.

[13] W. Zulehner, "Czochralski growth of silicon," *Journal of Crystal Growth*, vol. 65, no. 1-3, pp. 189–213, 1983.

[14] K. J. Gärtner, K. F. Rittinghaus, A. Seeger, and W. Uelhoff, "An electronic device including a TV-system for controlling the crystal diameter during czochralski growth," *Journal of Crystal Growth*, vol. 13-14, pp. 619–623, 1972.

[15] W. Bardsley, G. W. Green, C. H. Holliday, and D. T. J. Hurle, "Automatic control of Czochralski crystal growth," *Journal of Crystal Growth*, vol. 16, no. 3, pp. 277–279, 1972.

[16] W. Bardsley, D. T. J. Hurle, and G. C. Joyce, "The weighing method of automatic Czochralski crystal growth," *Journal of Crystal Growth*, vol. 40, no. 1, pp. 13–20, 1977.

[17] T. Kimoto, A. Wachi, S. Sakurai, and M. Mikami, "Automatic diameter control of Czochralski single crystal rods," *Transactions of the Society of Instrument and Control Engineers*, vol. 9, no. 5, pp. 595–600, 1973, in Japanese, with an English abstract.

[18] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

[19] 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.

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

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

[22] R. H. Fuerhoff, "Non-distorting video camera for use with a system for controlling growth of a silicon crystal," U.S. Patent 5,656,078, filed 14 November 1995, 1997, granted 12 August 1997.

[23] A. LaBrie and M. Baba, "Two camera diameter control system with diameter tracking for silicon ingot growth," U.S. Patent 6,030,451, filed 12 January 1998, 2000, granted 29 February 2000.

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

[25] T. Wetzel and S. Borchers, "Update of energy payback time and greenhouse gas emission data for crystalline silicon photovoltaic modules," *Progress in Photovoltaics: Research and Applications*, vol. 23, no. 10, pp. 1429–1435, 2015.

3.11 Exercises

3.1 ○ Concept check · 2 min

Scenario 3 is run with the diameter derivative time set to zero. The setpoint steps from 300 to 304 mm at 60 min, the camera noise is off, and the run stops itself at 158.4 min. What ends it?

(a) The pull rate crosses the structure limit and stays above it for the whole dwell

(b) The pull-rate command reaches the drive's mechanical limit and the integral winds up

(c) The growing oscillation drives the meniscus below its flooding height

(d) Camera noise differentiated without a filter spikes the pull rate and scraps the crystal

3.2 Concept check · 1 min

Which signal does mid-ranging controller VIC-301 take as its measurement?

(a) The diameter error, filtered more slowly than DIC-301 reads it

(b) The melt temperature, which the heater moves most directly

(c) The pull-rate command that DIC-301 sends to the pull drive

(d) The pull rate the drive actually applies, V_p

3.3 Concept check · 2 min

During a batch the applied pull rate falls from the design 1.0 mm/min to 0.5 mm/min. Which change to the settings of DIC-301 keeps the loop shape on the meniscus plant of Equation (3.12)?

(a) K_c halved, and τ_I, τ_D and τ_F each doubled

(b) K_c doubled, and τ_I, τ_D and τ_F each halved

(c) K_c doubled and τ_D doubled, with τ_I and τ_F unchanged

(d) K_c halved, with τ_I, τ_D and τ_F all unchanged

3.4 Theory and calculation · 6 min

Use $c_h = 0.050\,\text{rad/mm}$, $r_0 = 150\,\text{mm}$, $V_0 = 1.00\,\text{mm/min}$ and the default DIC-301 settings $K_c = 0.055$ (mm/min per mm), $\tau_I = 120\,\text{min}$, $\tau_D = 45\,\text{min}$ and $\tau_F = 8\,\text{min}$. The mid-ranging loop holds the pull rate at the 1.35 mm/min target of scenario 8.

(a) From Equation (3.12), find the natural frequency, the natural period and the magnitude of the steady-state gain of the diameter plant at this pull rate, to three significant figures. *Checkpoint:* 0.0247 rad/min, 255 min and 222 mm per mm/min.

(b) Find the four settings the schedule of Equation (3.14) applies at this pull rate. *Checkpoint:* $K_c = 0.0743$, $\tau_I = 88.9\,\text{min}$, $\tau_D = 33.3\,\text{min}$, $\tau_F = 5.93\,\text{min}$.

(c) Name the action of DIC-301 in the convention of Section 1.6.2, with the value of γ, and say whether the pull rate changes it. *Checkpoint:* Direct acting, $\gamma = -1$, at every pull rate.

3.5 Theory and calculation · 7 min

Scenario 5 opens with a hot zone that leaks 12 % more heat than the recipe assumed ($\lambda = 1.12$) and the rest of the design point unchanged: $r = 150\,\text{mm}$, $V_g = 1.000\,\text{mm/min}$, $T_m = 1425\,^\circ\text{C}$, $m = 100\,\text{kg}$. Use the constants of Table 3.2.

(a) Find the solidification rate and the latent heat it draws from the melt, in kW. *Checkpoint:* 0.1647 kg/min and 4.94 kW.

(b) Work Equation (3.11) up the chain to the crucible temperature, the heater temperature and the heater power, each to one decimal. *Checkpoint:* $T_c = 1582.8\,^\circ\text{C}$, $T_h = 1727.9\,^\circ\text{C}$, $P_a = 133.8\,\text{kW}$.

(c) Compare the power with the design value at $\lambda = 1$ and say which term of the chain carries most of the difference. *Checkpoint:* 15.6 kW more, most of it the melt surface loss.

3.6 Theory and calculation · 6 min

At the design point a 1 kW change of heater power moves the steady melt temperature by 2.819 °C, and $K_T = 0.025$ (mm/min) per °C. The diameter loop crosses over with a period of 36.3 min.

(a) Find the static gain from heater power to the pull rate that the diameter loop settles at, and give the action of VIC-301 in the convention of Section 1.6.2. *Checkpoint:* −0.0705 mm/min per kW, direct acting.

(b) With the integral dominant, the mid-ranging loop returns the pull rate with a time constant near $\tau_I/(K_c|K_{PV}|)$, where K_{PV} is the gain of part (a). Find it and its ratio to the crossover period for

the shipped tuning ($K_c = 8$, $\tau_I = 120\,\text{min}$) and the retuned one ($K_c = 10$, $\tau_I = 55\,\text{min}$), with K_c in kW per mm/min. *Checkpoint:* 213 min and 5.9; 78 min and 2.1.

(c) Repeat for the challenge as supplied ($K_c = 2$, $\tau_I = 600\,\text{min}$) and compare with the length of a batch. *Checkpoint:* About 4260 min, some 71 hours, against a batch of about eight.

3.7 App challenge · 8 min

Load scenario 3, "A setpoint step, and why there is a derivative term" (default seed), and set Diameter derivative time to 0 before pressing Run. Run at the default 20 min/s until the run stops itself, then press Export CSV. The time column of the file is Time (h).

(a) Read the clock when the run stops, and the meniscus strip at that moment. What ended the run? Give the time to within 0.02 h. *Checkpoint:* 2.64 h; the meniscus has reached its flooding line.

(b) From the column `dia_mm` (the true diameter), find the highest value and the time of that row, to within 0.1 mm and 0.02 h. *Checkpoint:* 309.8 mm at 1.70 h.

(c) From the column `v_head_mm_min`, find the smallest pull-rate margin of the run to within 0.005 mm/min. Does it support (a)? *Checkpoint:* 0.214 mm/min: the structure limit was never reached.

3.8 App challenge · 8 min

Load scenario 5, "The limit the pull rate walks into" (default seed), and run it as loaded at 60 min/s until the run stops itself at the end of the batch. Read the Prime crystal tile and the clock, and export the CSV. Then press Reset, set Control to Diameter + mid-ranging, run to the end again, read the same two values and export a second CSV.

(a) Give the prime crystal and the end time of each run, to 0.1 kg and 0.02 h. *Checkpoint:* 73.5 kg at 7.75 h; 88.0 kg at 8.21 h.

(b) Find the smallest value of `v_head_mm_min` in each file, to within 0.003 mm/min. *Checkpoint:* 0.239 mm/min and 0.377 mm/min.

(c) Table 3.5 reports worst margins of 0.221 and 0.358 mm/min for the same pair. Explain the difference and say which number an engineer should trust. *Checkpoint:* The file logs two samples per minute, the table every physics step.

3.9 App challenge · 8 min

Load scenario 9, "A noisy bright ring against a long derivative time" (default seed), and run it as loaded at 20 min/s until it stops itself. Read the clock and the Prime crystal tile. Then switch the panel to Advanced, press Reset, set Derivative filter to 12.0 min, run at 60 min/s to the end of the batch, read the same two values and export the CSV.

(a) Give the stop time and prime crystal of the first run to 0.02 h and 0.1 kg, and the noise gain τ_D/τ_F of each run. *Checkpoint:* 1.89 h, 19.0 kg; 11.25 and 3.75.

(b) Give the end time and prime crystal of the second run, and the smallest `v_head_mm_min` in its file to within 0.005 mm/min. *Checkpoint:* 8.43 h, 88.0 kg, 0.086 mm/min.

(c) What did the longer filter cost, by the chapter's analysis? *Checkpoint:* 23.0° of phase margin, from 44.3 to 21.4°.

3.10 Application area · 3 min

A puller vendor proposes to replace the bright-ring camera with a load cell on the crystal, arguing that a weight is a cleaner signal than an image. From the chapter's history and model, what must a weighing measurement handle that the optical one does not, and what does the optical route cost the diameter loop of this chapter?

3.11 Application area · 4 min

(a) Prime crystal is valued at $900 per kg in the application. Using Table 3.5, what is one boule's difference in value between scenario 5 and scenario 6, and what else does the run cost that the valuation leaves out? *Checkpoint:* About $13,100 per boule, plus furnace time, polysilicon and energy.

(b) In scenario 7 with mid-ranging, the diameter stays in band all run and the pull rate is walked back toward its target, while the heater ends at 121.59 kW having started at 118.26 kW. What should a process engineer conclude? *Checkpoint:* Something upstream moved, here the crucible lift.

3.12 Review the agent · 5 min

An AI assistant reviews scenario 9 and writes: "The 8 min derivative filter costs 21.7° of phase margin against 3 min, and the loop has little margin to spare. Set Derivative filter to 3.0 min. The margin rises to about 51°, the loop becomes better damped and the pull rate stays further from the structure limit." Find the flaw, and say how you would check it in the application.