INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 14 Control in Semiconductor Manufacturing

CHAPTER 2

# Siemens Polysilicon Reactor Simulator for Gain Scheduling and an Unmeasured Constraint

Pages 23–51

Edited by John D. Hedengren

**How to cite this chapter.** "Siemens Polysilicon Reactor Simulator for Gain Scheduling and an Unmeasured Constraint," in *Control in Semiconductor Manufacturing*, John D. Hedengren, Ed., vol. 14 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 23–51.

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

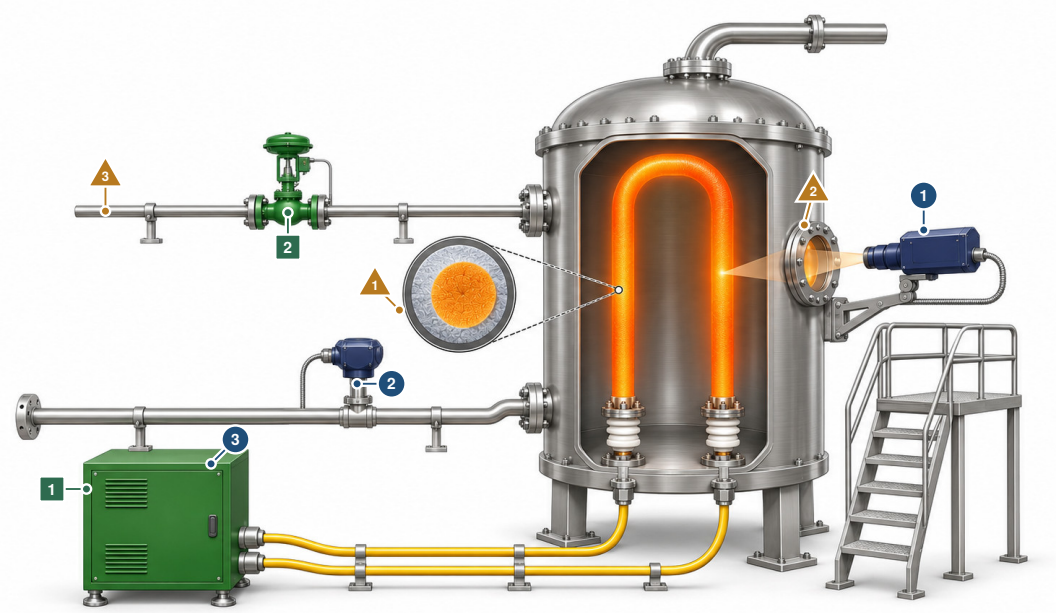
**The app for this chapter**
https://apmonitor.com/apps/polysilicon-cvd/

**The volume online**: contents, other chapters and citation records
https://apmonitor.com/apps/books/v14-semiconductor/#polysilicon-cvd

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

CHAPTER 2

# Siemens Polysilicon Reactor Simulator for Gain Scheduling and an Unmeasured Constraint




**The controller moves** 1 rod power supply, current regulated, set by the temperature loop
2 trichlorosilane feed valve, set by the composition loop

**The controller reads** 1 pyrometer on the rod surface, through a sight glass 2 trichlorosilane in the gas
3 the supply's voltage and current, from which the hidden center temperature is computed

**What disturbs it** 1 the rods growing thicker, which shrinks the loop gain a hundredfold
2 a silicon film on the sight glass, which makes the pyrometer read low 3 boron in the feed, which no instrument sees

*Notice the inset: the rod is hottest at its center, which no instrument reaches, while the pyrometer sees only the surface. Illustration, not to scale.*

### IN PLAIN TERMS A Heater That Is Also the Product

**The main idea.** Pure silicon for chips and solar cells grows on thin silicon rods in a sealed jar. A current keeps the rods glowing hot, and a gas leaves new silicon on them, so they grow many times thicker.

**The challenge.** As the rods thicken, a change in current moves their temperature about a hundred times less. The automatic control set up at the start slowly stops working. Worse, the hottest place is the middle of each rod, where no instrument can look. It can melt while the surface reading looks perfect.

**Why it matters.** The hidden middle, not the measured surface, decides how large the rods can grow.

### THREE WAYS TO READ THIS CHAPTER

| | | |
|---|---|---|
| **Explore** | 11 min | The first page, the history and fact boxes in Section 2.2, the try-it and fact boxes in Section 2.6, 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 2.6, Section 2.7, and Section 2.8 with its open challenge |

**The control question** The silicon rods of a Siemens reactor are the heating element and the product, and they grow from 8 to 150 mm in one batch. How does a temperature loop stay useful while its plant gain falls a hundredfold, and how is the batch ended at a center temperature that no instrument measures?

**Learning objectives**

- Compute the resistance, current, power and loss slope of a rod pair at a given diameter from the model equations.
- Explain why current regulation is self-regulating and voltage regulation runs away, from the sign of $\partial P/\partial \bar{T}$ against the loss slope.
- Build a gain schedule from a closed-form plant gain, and explain why a schedule quantity with the wrong magnitude and the right shape works.
- Predict the parabolic center-to-surface rise, and protect the unmeasured center with an override on an inferred estimate.
- Judge whether a predictive controller's gain in batch time comes from prediction or from constraint margin.

**Prerequisites** Sections 1.6.2 (velocity-form PID), 1.6.4 (scenarios and data) and 1.4.2 (what a reactor measures, and a pyrometer that is wrong). Section 1.6.3 is needed only for the predictive controller.

**The app** https://apmonitor.com/apps/polysilicon-cvd/

**See also** Volume 7, Chapter 2 for a heat-generation slope that decides stability. Volume 5, Chapter 11 for a soft sensor of a core temperature nothing reaches, and Volume 5, Chapter 10 for overrides on inferred constraints. Volume 4, Chapter 2 for a protective override on a limit.

## 2.1 Introduction

A silicon rod standing in a Siemens reactor is three things at once. It is the heating element, because the current that holds it near 1100 °C passes down one leg of a hairpin, across a bridge at the top and up the other. It is the deposition surface, because trichlorosilane decomposes on it and leaves silicon behind. And it is the product, because at the end of the batch the reactor is opened and the rods are broken up and sold. The third fact is the one that decides the control problem. A plant whose equipment is fixed and whose operating point drifts can be retuned once. A plant that is its own product changes while it runs, and here it changes by a factor that is not small.

The numbers are set by geometry alone. Resistance is resistivity times length divided by cross-sectional area, and the cross-section grows as the square of the diameter, so growing an 8 mm seed to 150 mm multiplies the area by 352. Measured on the simulator at a surface held at 1100 °C, the rod pair resistance falls from 50.9 Ω to 0.0649 Ω, the current rises from 21.8 A to 2646 A, the terminal voltage falls from 1111 V to 172 V, and the power rises from 24.2 kW to 454 kW. The steady-state gain from current to surface temperature falls from 10.46 °C/A to 0.0828 °C/A, a factor of 126. A proportional-integral controller commissioned in the first hour of the batch is still connected in the last one, and by then it is very nearly open loop.

The shape of the problem is general. Any plant that accumulates or consumes the material it is built around presents the same picture, including a batch crystallizer whose surface area grows, an exchanger whose fouling resistance grows, a fixed bed whose catalyst deactivates, a furnace tube that cokes and a cell whose electrodes age. Gain scheduling is the standard answer and it has a substantial analysis literature, covering when a family of linearizations may be interpolated, what the resulting controller guarantees, and the condition that the scheduling variable must move slowly compared with the closed-loop dynamics [1–3]. What a survey cannot supply is a plant where the schedule is not fitted. Most industrial schedules use a proxy variable, a lookup table and a long commissioning period. Here the scheduling variable is the diameter, it is the cause of the gain change rather than a correlate of it, and the gain itself is a closed-form function of it.

Two further properties come out of the same rod. Silicon is a semiconductor, so its resistivity falls steeply with temperature, and the electrical power therefore falls with temperature at constant current

and rises with temperature at constant voltage. Whether the operating point is stable is decided by which terminal of the same power supply is held constant, and no amount of tuning reaches that question. And heat is generated throughout the rod while it leaves only from the surface, so the radial profile is parabolic and the hottest point is the axis. The pyrometer reads the surface. Silicon melts at the center. The constraint that ends the batch has no instrument on it and never has had one, which is why the rod interior appears in the reactor-modeling literature as a computed field rather than as a measurement [4, 5].

However, a student handed a first-order transfer function and a tuning rule never meets a plant that will not hold still, and a block diagram that draws the power supply as a source of power hides the sign that decides stability. This chapter works with an interactive Siemens reactor in which the loop gain moves by two orders of magnitude during one batch, the regulation mode of the supply decides whether an operating point exists at all, and the binding constraint is available only through a model. Section 2.2 sets out the process and its few instruments. Section 2.3 gives the nine-state deposition and energy balance with its instrument chain and supply model, in the detail needed to rebuild it, and states where each assumption fails. Section 2.4 derives the three properties a controller has to be designed around: the sign that decides stability, a gain that is a closed-form function of the diameter, and a constraint reached only through a model. Section 2.5 describes the scheduled PI loop, the pyrometer, the override on the inferred center and a shrinking-horizon predictive controller of the surface setpoint. Section 2.6 runs the scenario sequence, from a seed rod held by hand to a graded challenge carrying four simultaneous faults. The predictive controller is a negative result there. Under the shipped override it lengthens the batch, and the hour it saves when the override band is narrowed is margin that the override alone recovers. Section 2.7 names the instruments, the supply and the control system behind the loops. Sections 2.8 to 2.10 state the limitations, discuss the energy trade and conclude.

## 2.2 Polysilicon Production and Its Instrumentation

Silicon for semiconductors and for most photovoltaic wafers begins as metallurgical-grade material at roughly 98 % purity, which is several orders of magnitude short of what a device needs. The industrial route to the remaining purity is chemical rather than metallurgical. The silicon is converted to a chlorosilane, usually trichlorosilane, the chlorosilane is distilled until its impurities are measured in parts per billion, and the purified gas is decomposed back to solid silicon on heated silicon rods [6, 7]. Both stages purify. The distillation does most of the work, and the growth surface rejects impurities again by segregation as the deposit forms. Neither stage is closed-loop within a batch, which is why the industry runs on incoming-material specification and statistical process control for purity rather than on feedback.

> ✦ DID YOU KNOW? **Most of the World's Polysilicon**
>
> Rod deposition from trichlorosilane, in its modern recycling form, makes about 80 percent of the world's polysilicon, and its deposition step typically consumes around 90 kW h for each kilogram, more than half of the energy used in polysilicon production [8]. A complete batch of this chapter's simulator spends 87.0 kW h/kg.

The deposition stage carries the name of the company that commercialized it and the arrangement has proved remarkably durable. Reviews of feedstock technology published decades after its introduction still identify rod deposition from trichlorosilane as the dominant industrial process, with the fluidized-bed silane route as the principal alternative [6, 9, 10]. What has changed is not the geometry but the energy accounting and the modeling. The reactor is among the largest single electricity consumers in the wafer supply chain, and polysilicon production is a substantial term in photovoltaic life-cycle energy analyses [11]. That pressure produced a body of work on where the energy goes, on the radiative exchange between rods and a cooled bell jar wall [12], on the deposition

radius at which the energy per kilogram stops improving [13], and on comparative thermal analysis against the fluidized-bed alternative [5].

> ⧖ HISTORY **Rods of the Same Silicon (1956)**
>
> A patent filed in Germany in June 1956, and granted in the United States in 1961, gave the deposition reactor the form this chapter models [14]. Silicon had been grown from a chlorosilane and hydrogen on a heated strip of tantalum, which then had to be removed from the crust at the risk of new impurities. The patent grew it instead on self-supporting rods of silicon, joined at their upper ends so that the current passed up one rod and down the other. Its rods were 0.5 m long and 3 mm across, held near 1150 °C. The temperature was judged by the glowing color of the rods, and the current was raised by hand with a rheostat, continuously and gradually, because the resistance of the rods fell as they thickened [14]. That falling resistance is the gain change this chapter schedules.

The instrumentation history is short because there is very little of it. The surface temperature is read optically through a sight glass, and everything else the operator can see is electrical. The rod interior has never had an instrument, and it is treated in the literature as a computed quantity. The analysis of temperature homogeneity of polysilicon rods is carried out numerically for exactly that reason [4], and the same is true of the deposition chemistry, where the split between surface kinetics and gas-phase transport of trichlorosilane is established from transport models and growth-rate measurements rather than from anything measured inside the reactor during a run [15]. A reactor with one critical measurement and one critical unmeasured constraint is the natural home of inferential control, which was formalized as the use of secondary measurements to estimate a primary variable that cannot be measured directly [16] and which is now a standard industrial tool under the name soft sensing [17]. The structure that enforces the resulting limit is a selector, and the systematic design of constraint switching with selectors is itself an active subject [18].

## 2.3 Reactor Model

Reactor R-401 is a single rod pair inside a water-cooled bell jar. Two seed rods of diameter $d$ stand on electrodes in the baseplate and are joined by a bridge at the top, giving an electrical path of length $L$ through solid silicon. Power supply PS-401 drives that path and regulates either its current or its terminal voltage. Trichlorosilane and hydrogen enter at the baseplate, and the exhaust carries away the silicon tetrachloride byproduct. Pyrometer TE-401 reads the rod surface through a sight glass and supplies the only temperature measurement in the machine. Figure 2.1 shows the arrangement, the instrument tags and every symbol used in the equations below.

Table 2.1 classifies what the loops interact with. The clock is in hours, the rod is in millimeters, temperatures are in degrees Celsius, the electrical quantities are in amperes, volts and kilowatts, and the deposit is in kilograms with its boron content in parts per billion atomic. Nine states are integrated. The run stops on its own when the diameter reaches the harvest value or when the computed center temperature reaches the melting point of silicon.

### 2.3.1 Resistance of a Semiconducting Rod

The resistivity of intrinsic silicon above 1000 °C is dominated by thermally generated carriers, so it falls as the rod heats, and it is represented by a saturating exponential with a high-temperature floor, as shown in Equations (2.1) and (2.2).

$$\rho(T) = \rho_{\infty} + A\,\mathrm{e}^{B/T} \tag{2.1}$$

$$R(\bar{T}, d) = \frac{\rho(\bar{T})\,L}{\pi d^2/4} \tag{2.2}$$

where $\rho$ is in $\Omega\,\mathrm{m}$, $T$ is absolute temperature in kelvin, $B$ is an activation temperature in kelvin, $L$ is the electrical path length of the rod pair in meters and $d$ is the rod diameter. Both sides of (2.2)

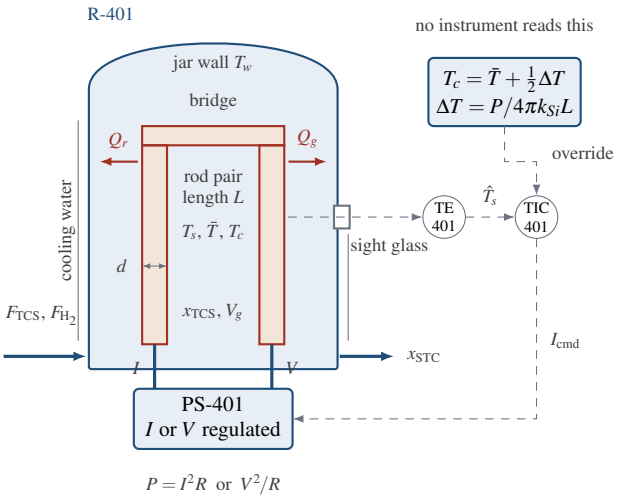



**Figure 2.1** Siemens reactor R-401 with one rod pair inside a water-cooled bell jar. Supply PS-401 drives current $I$ at terminal voltage $V$ through an electrical path of length $L$ and diameter $d$, dissipating $P$ throughout the rod volume. Heat leaves by radiation $Q_r$ to the jar wall at $T_w$ and by convection $Q_g$ to the hydrogen carrier. Pyrometer TE-401 reads the surface $\hat{T}_s$ through a sight glass and TIC-401 manipulates the current command $I_{\text{cmd}}$. The mean rod temperature $\bar{T}$ is the state that is integrated, and the surface $T_s$ and center $T_c$ sit half the parabolic rise $\Delta T$ either side of it. The center temperature is computed and never measured.

**Table 2.1** Classification of Variables. Nominal Values Are Those of a Seed Rod at the Design Surface Temperature.

| Symbol | Role | Nominal | Range |
|---|---|---|---|
| $T_s$ (TE-401) | CV, measured | 1100 °C | 950–1200 °C SP |
| $T_c$ | CV, inferred | 1117 °C | limit 1150–1414 °C |
| $x_{TCS}$ | CV, measured | 0.120 | 0.04–0.30 SP |
| $I$ (PS-401) | MV | 21.8 A | 0–3000 A |
| $F_{TCS}$ | MV | 2.50 kmol/h | 0.2–12 kmol/h |
| supply mode | DV, operator | current | current or voltage |
| $f_{loss}$ | DV, unmeasured | 1× | 0.8–1.3× |
| $v_g$ gas velocity | DV, unmeasured | 1× | 0.5–2× |
| $C_B$ feed boron | DV, unmeasured | 0.60 ppba | 0.1–4 ppba |
| $f_v$ vaporizer | DV, unmeasured | 1× | 0.85–1.15× |
| $\sigma_T$ pyrometer noise | DV, measured | 2 °C | 0–20 °C |
| $b_T$ pyrometer offset | DV, unmeasured | 0 °C | −40–40 °C |
| $d$ | state | 8 mm | 8–150 mm |
| $\bar{T}$ | state | 1108.4 °C | 20–2200 °C |
| $I_a, V_a$ | states | 21.8 A, 1111 V | supply ratings |
| $x_{TCS}, x_{STC}$ | states | 0.120, 0.020 | 0–0.9 |
| $m_{Si}$ | state | 0 kg | 0–214 kg |
| $B_{acc}$ | state | 0 | $\geq 0$ |
| $E$ | state | 0 kW h | 0–18 700 kW h |

carry units of ohms. The resistance is evaluated at the mean rod temperature $\bar{T}$ and not at the surface, because the whole cross-section conducts. Two separate effects shrink it over a batch. At a fixed 1100 °C the area term alone takes $R$ from 53.65 Ω to 0.1526 Ω, a factor of 352 that is exactly the area ratio. Holding the *surface* at 1100 °C instead raises $\bar{T}$ from 1108.4 to 1258.0 °C as the rod

thickens, and (2.1) roughly halves again over that span, so the measured resistance falls from 50.9 Ω to 0.0649 Ω, a factor of 785.

### 2.3.2 Electrical Power and the Energy Balance

The supply regulates one terminal quantity and the rod decides the other, as shown in Equation (2.3).

$$P = \begin{cases} I_a^2 R(\bar{T}, d) & \text{current regulation} \\ V_a^2 / R(\bar{T}, d) & \text{voltage regulation} \end{cases} \tag{2.3}$$

where $I_a$ and $V_a$ are the delivered current and voltage, each of which follows its command through a first-order supply lag $\tau_{ps}$ as shown in Equation (2.4).

$$\tau_{ps} \frac{\mathrm{d}I_a}{\mathrm{d}t} = I_{\mathrm{cmd}} - I_a, \qquad \tau_{ps} \frac{\mathrm{d}V_a}{\mathrm{d}t} = V_{\mathrm{cmd}} - V_a \tag{2.4}$$

Under voltage regulation the power is also limited to $I_{\max}^2 R$, which is the supply leaving voltage regulation for its current rating, and the dissipation in either mode is limited to $1.5 P_{\max}$ before it is used, which prevents a numerical excursion from injecting unbounded power during a runaway. The rod is one lumped thermal capacity, and its energy balance is Equation (2.5).

$$\rho_{Si} \frac{\pi d^2 L}{4} c_{p,Si} \frac{\mathrm{d}\bar{T}}{\mathrm{d}t} = 3600 \left( P - Q_{\mathrm{loss}} - q_{\mathrm{dep}} \right) \tag{2.5}$$

where $\rho_{Si}$ is the density of silicon, $c_{p,Si}$ its heat capacity, the product of the three leading factors is the rod heat capacity in J/K, and the factor 3600 converts a watt to a joule per hour so that the derivative is per hour. The rod heat capacity runs from 487 J/K at the seed to 171 kJ/K at harvest. Losses leave the surface by radiation to the jar wall and by convection to the hydrogen carrier, as shown in Equation (2.6).

$$Q_{\mathrm{loss}} = f_{\mathrm{loss}} A \left[ \varepsilon \sigma \left( T_s^4 - T_w^4 \right) + h_g \left( T_s - T_g \right) \right], \quad A = \pi d L \tag{2.6}$$

where $A$ is the lateral area of the rod pair in $\mathrm{m}^2$, $\varepsilon$ is emissivity, $\sigma$ is the Stefan-Boltzmann constant, $h_g$ is the convection coefficient to hydrogen, and $f_{\mathrm{loss}}$ is the disturbance multiplier on total loss. Temperatures in the radiative term are absolute. Differentiating (2.6) gives the loss slope used throughout the analysis, as shown in Equation (2.7).

$$S \equiv \frac{\partial Q_{\mathrm{loss}}}{\partial T_s} = f_{\mathrm{loss}} A \left( 4 \varepsilon \sigma T_s^3 + h_g \right) \tag{2.7}$$

The slope $S$ is 61.6 W/K at the seed rod and 1154 W/K at harvest, growing linearly with $d$ because the area does. The rod thermal time constant $\rho_{Si} \pi d^2 L c_{p,Si} / (4S)$ therefore also grows linearly, from 7.9 s to 148 s.

### 2.3.3 The Parabolic Radial Profile

Generation is volumetric and loss is superficial, so the rod cannot be isothermal across its section. Uniform generation $q'''$ in a long cylinder of constant conductivity gives a parabolic profile, and the difference between the axis and the surface is Equation (2.8).

$$\Delta T = T_c - T_s = \frac{q''' d^2}{16 k_{Si}}, \qquad q''' = \frac{P}{\pi d^2 L / 4} \tag{2.8}$$

where $k_{Si}$ is the thermal conductivity of hot silicon. Substituting $q'''$ collapses (2.8) to $\Delta T = P / (4 \pi k_{Si} L)$, so the rise depends on the total dissipation and the rod length alone and not explicitly on the diameter at all. That is worth stating carefully, because the $d^2$ in (2.8) invites the wrong

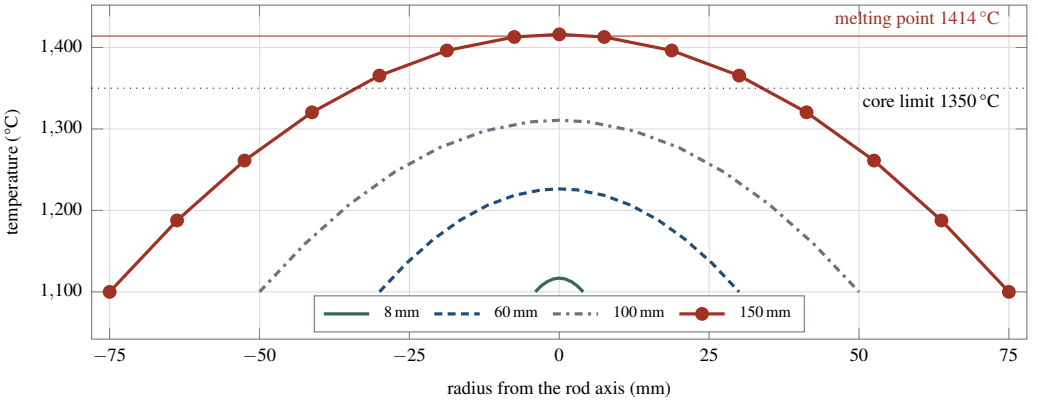



**Figure 2.2** Radial temperature profile of the rod at four diameters with the surface held at 1100 °C: 8 mm solid, 60 mm dashed, 100 mm dash-dotted and 150 mm solid with dots. The thin solid line is the melting point and the thin dotted line the core limit. The pyrometer reads the two endpoints of every curve and they are identical. The axis temperature rises from 1116.9 to 1416.1 °C across the same four cases, passes the default 1350 °C limit at 118.7 mm and reaches the melting point at 149.0 mm.

conclusion. At a held surface temperature the power is a surface loss and grows linearly with $d$, so $q'''$ falls as $1/d$ and the rise grows *linearly*. Measured on the simulator, $\Delta T$ is 16.9 °C at 8 mm and 316.1 °C at 150 mm, a ratio of 18.7 against a diameter ratio of 18.75. The area-weighted mean of a parabola over a circular section is exactly halfway between axis and surface, which gives the two algebraic outputs in Equation (2.9).

$$T_s = \bar{T} - \tfrac{1}{2}\Delta T, \qquad T_c = \bar{T} + \tfrac{1}{2}\Delta T \tag{2.9}$$

Only $\bar{T}$ is a state. The pyrometer sees $T_s$ and nothing sees $T_c$. Figure 2.2 draws the profile at four diameters against the melting point.

### 2.3.4 Deposition, Gas Phase and Purity

Deposition is a surface reaction in series with transport of trichlorosilane through the boundary layer, and the two resistances add, as shown in Equation (2.10).

$$\frac{1}{k_{ov}} = \frac{1}{k_s} + \frac{1}{k_m}, \quad k_s = k_{s0}\,\mathrm{e}^{-\frac{E_a}{R}\left(\frac{1}{T_s}-\frac{1}{T_{\mathrm{ref}}}\right)}, \quad k_m = k_{m0}\,v_g \tag{2.10}$$

where $k_s$ and $k_m$ are the surface-reaction and mass-transfer growth velocities in mm/h on the radius, $E_a/R$ is an activation temperature, $T_{\mathrm{ref}}$ is a reference surface temperature, and $v_g$ is the gas-velocity disturbance multiplier. The diameter and the deposited mass follow in Equation (2.11).

$$\frac{\mathrm{d}d}{\mathrm{d}t} = 2\,k_{ov}\,\frac{x_{\mathrm{TCS}}}{x_{\mathrm{ref}}}, \qquad \dot{m} = \rho_{Si} A\,\frac{1}{2}\,\frac{\mathrm{d}d}{\mathrm{d}t} \times 10^{-3} \tag{2.11}$$

where the factor $10^{-3}$ converts millimeters to meters so that $\dot{m}$ is in kg/h. At the design surface temperature and composition the diameter grows at 1.800 mm/h throughout the batch, because neither term in (2.10) depends on $d$. The reaction absorbs $\Delta H_{\mathrm{dep}}$ per kilogram, entering (2.5) as $q_{\mathrm{dep}} = \Delta H_{\mathrm{dep}}\dot{m}/3.6$ in watts.

The bell jar gas is one well-mixed volume holding $n_T = V_g/(10^3 v_m)$ kilomoles. Trichlorosilane is fed into a fixed hydrogen carrier, the rods take silicon out of it, and the overall stoichiometry sends four

trichlorosilane molecules to one silicon atom and three tetrachlorides, as shown in Equation (2.12).

$$n_T \frac{\mathrm{d}x_{\mathrm{TCS}}}{\mathrm{d}t} = f_v F_{\mathrm{TCS}} - F_{\mathrm{out}} x_{\mathrm{TCS}} - 4\,\frac{\dot{m}}{M_{Si}} \tag{2.12}$$

$$n_T \frac{\mathrm{d}x_{\mathrm{STC}}}{\mathrm{d}t} = 3\,\frac{\dot{m}}{M_{Si}} - F_{\mathrm{out}} x_{\mathrm{STC}} \tag{2.13}$$

where $F_{\mathrm{out}} = f_v F_{\mathrm{TCS}} + F_{\mathrm{H_2}}$ is the sweep leaving the jar, $f_v$ is the vaporizer calibration disturbance and $M_{Si} = 28.09\,\mathrm{kg/kmol}$. The gas holdup is 0.1786 kmol against a sweep of about 20.5 kmol/h, so the composition time constant is 31.4 s, four orders of magnitude shorter than the batch. Boron arrives dissolved in the feed and is incorporated less selectively as the surface gets hotter, as shown in Equation (2.14).

$$[B] = C_B\, k_{B0}\, \mathrm{e}^{k_{B1}(T_s - T_{\mathrm{ref}})/100}, \qquad \frac{\mathrm{d}B_{\mathrm{acc}}}{\mathrm{d}t} = \dot{m}\,[B] \tag{2.14}$$

where $[B]$ is the instantaneous boron content of the deposit in parts per billion atomic, $C_B$ is the boron in the feed and $k_{B1}$ is expressed per 100 °C. The reported purity is the running average $B_{\mathrm{acc}}/m_{Si}$ over the whole deposit, which is what a laboratory would measure on a broken piece of rod. Consumed energy accumulates as $\frac{\mathrm{d}E}{\mathrm{d}t} = P/1000$ in kW h/h.

### 2.3.5 Parameters, Initialization and Numerics

The constants are collected in Table 2.2. Twenty-two are adjustable in the application and ten are hidden, and the hidden ones are the material properties and reference conditions a reader would have to look up rather than read off the interface.

Every scenario opens at a genuine steady state rather than settling into one. The initialization solves the design point by fixed-point iteration, because the loss is a function of the surface, the resistance is a function of the mean, and the mean sits half the parabolic rise above the surface with a rise that depends on the power. Starting from $\bar{T} = T_{\mathrm{SP}}$, the iteration of Equation (2.15) is repeated eight times.

$$R \leftarrow R(\bar{T}, d), \quad I \leftarrow \sqrt{\frac{Q_{\mathrm{loss}} + q_{\mathrm{dep}}}{R}}, \quad \bar{T} \leftarrow T_{\mathrm{SP}} + \frac{I^2 R}{8\pi k_{Si} L} \tag{2.15}$$

The same expression supplies the manual-mode aid that sets the current for whatever diameter the rod has at that moment. Nine scenarios start the rod part way through a batch at 60, 70, 90, 100 or 140 mm so that the interesting part is reached without waiting for it. In automatic mode the composition loop starts at the trichlorosilane feed that replaces what the carrier sweeps out and what the rods deposit, so the gas opens steady as well.

The integrator is classical fourth-order Runge-Kutta on a fixed step of 0.0004 h (1.44 s). The controller executes every 0.0008 h (2.88 s) and the historian logs 20 samples per simulated hour, one every three minutes. The pseudo-random generator is seeded at 20260919, so a noisy scenario replays identically. Section 1.6.1 describes the engine that integrates, seeds and logs the run.

### 2.3.6 Assumptions

Each simplification below is deliberate and each one fails somewhere specific.

*One rod pair stands for a whole reactor.* Production machines carry 36 to 72 rods, and the rods near the middle see each other rather than the cooled wall, so they run hotter than the ones at the edge. Radiative exchange in (2.6) is rod to wall only, with no view factors between rods. The uniformity problem that omission removes is a large part of the real control engineering and it is the subject of the rod-homogeneity literature [4, 12].

*The rod is one lumped temperature plus an analytic profile.* Equation (2.8) is exact only for uniform volumetric generation and constant conductivity. Neither holds. The conductivity of silicon falls with temperature over this range, which steepens the real profile, and current crowds toward the hotter center because the resistivity there is lower, which concentrates the generation inward [19].

**Table 2.2** Model Parameters. Hidden Constants Never Appear on the Panel.

| Symbol | Meaning | Value | Unit | Visibility |
|---|---|---|---|---|
| $L$ | electrical path length | 5.2 | m | adjustable |
| $d_0$ | seed rod diameter | 8 | mm | adjustable |
| $d_{\text{end}}$ | harvest diameter | 150 | mm | adjustable |
| $k_{Si}$ | conductivity of hot Si | 22 | W/(m K) | adjustable |
| $B$ | resistivity activation temp. | 12120 | K | adjustable |
| $\varepsilon$ | rod emissivity | 0.70 | — | adjustable |
| $T_w$ | bell jar wall | 100 | °C | adjustable |
| $h_g$ | convection to hydrogen | 60 | W/(m$^2$ K) | adjustable |
| $T_g$ | bulk gas temperature | 350 | °C | adjustable |
| $k_{m0}$ | mass-transfer growth | 1.20 | mm/h | adjustable |
| $k_{s0}$ | surface-reaction growth | 3.60 | mm/h | adjustable |
| $E_a/R$ | deposition activation temp. | 12000 | K | adjustable |
| $F_{\text{H}_2}$ | hydrogen carrier flow | 18 | kmol/h | adjustable |
| $V_g$ | bell jar gas volume | 4.0 | m$^3$ | adjustable |
| $k_{B0}$ | boron incorporation | 0.35 | — | adjustable |
| $k_{B1}$ | selectivity loss with $T_s$ | 0.55 | per 100 °C | adjustable |
| $[B]_{\text{spec}}$ | electronic-grade limit | 0.30 | ppba | adjustable |
| $I_{\max}$ | supply current rating | 3000 | A | adjustable |
| $V_{\max}$ | supply voltage rating | 1800 | V | adjustable |
| $P_{\max}$ | supply power rating | 550 | kW | adjustable |
| $c_{kg}$ | value of the product | 22 | $/kg | adjustable |
| $c_{kWh}$ | electricity | 0.06 | $/kW h | adjustable |
| $\rho_\infty$ | resistivity floor | $2.0 \times 10^{-5}$ | Ω m | hidden |
| $A$ | resistivity prefactor | $7.32 \times 10^{-8}$ | Ω m | hidden |
| $\rho_{Si}$ | density of silicon | 2330 | kg/m$^3$ | hidden |
| $c_{p,Si}$ | heat capacity of silicon | 800 | J/(kg K) | hidden |
| $T_{\text{melt}}$ | melting point | 1414 | °C | hidden |
| $\Delta H_{\text{dep}}$ | heat of deposition | 120 | kJ/kg | hidden |
| $v_m$ | molar volume | 0.0224 | m$^3$/mol | hidden |
| $\tau_{ps}$ | supply response | 0.002 | h | hidden |
| $x_{\text{ref}}$ | reference composition | 0.12 | — | hidden |
| $T_{\text{ref}}$ | reference surface temp. | 1100 | °C | hidden |

Both effects steepen the profile and neither is modeled, so the computed $T_c$ should be read as an estimator that runs low rather than as a truth value.

*The gas phase is a single well-mixed volume.* There is no axial concentration profile up the rod and no reactor-scale flow field, so the deposition non-uniformity along the rod length that limits real rod quality is absent. The mass-transfer coefficient is a single number scaled by a velocity multiplier rather than a boundary-layer calculation.

*The chemistry is one overall stoichiometry.* Dichlorosilane, homogeneous nucleation to powder, and the hydrochlorination recycle that converts the tetrachloride byproduct back into feed are all absent, and the last of those is most of a real plant [6]. The byproduct term (2.13) therefore tracks an inventory that has nowhere to go.

*Purity is one segregation coefficient.* Equation (2.14) lumps species, surface condition and growth rate into two constants, and phosphorus, carbon and metallic contaminants behave differently and are not represented. The model can show that a hotter surface is less selective. It cannot be used to set a specification.

*Seed heating and the cold start are absent.* Equation (2.1) at room temperature gives a resistance of about $7 \times 10^{15}$ Ω, so a cold silicon rod is very nearly an insulator and a real reactor has to be struck at kilovolts or preheated before current regulation has anything to hold. Every scenario in the

application begins with the rod already hot.

*Disturbances are stylized.* A pyrometer offset that appears at once at 25 °C is not how a window fogs, but a step is a fault whose consequence can be read off a trend, while a slow drift would be indistinguishable from the plant. A feed impurity that triples between one sample and the next does not happen in a distillation train either, and it isolates the observation that matters, which is that no instrument in the reactor responds. Scenarios that begin at 60, 70, 90, 100 or 140 mm skip the first two days of a batch, which is defensible only because the diameter is the state that carries all of the history.

*The numerics are more repeatable than a control system.* A fixed step and a seeded generator reproduce one trajectory on every machine, while a distributed control system delivers jitter, quantization and a scan order that moves.

## 2.4 Three Properties of the Plant

### 2.4.1 One Term Decides Stability and Dominates the Gain

Differentiating (2.1) and (2.2) at fixed diameter gives $\partial R/\partial \bar{T} = -RB/\bar{T}^2$ in the exponential-dominated regime, and substituting into (2.3) gives the two slopes of Equation (2.16).

$$\left.\frac{\partial P}{\partial \bar{T}}\right|_I = -\frac{PB}{\bar{T}^2} < 0, \qquad \left.\frac{\partial P}{\partial \bar{T}}\right|_V = +\frac{PB}{\bar{T}^2} > 0 \tag{2.16}$$

A heated body holds an operating point when its dissipation rises with temperature more slowly than its losses do, so the criterion is $\partial P/\partial \bar{T} < S$ with $S$ from (2.7). The same comparison of a generation slope with a removal slope decides ignition in the exothermic reactor of Volume 7, Chapter 2. Under current regulation the left side is negative, the criterion is satisfied by a wide margin, and the rod is self-regulating before any controller is attached. Under voltage regulation the left side is positive and larger than $S$ at every diameter. Evaluated on the simulator at a surface of 1100 °C, the quantity $PB/\bar{T}^2$ is 154, 1068 and 2349 W/K at 8, 60 and 150 mm against loss slopes of 62, 462 and 1154 W/K, so the ratio is 2.50, 2.31 and 2.03. Scenario 5 runs that criterion out on a 60 mm rod sitting at its design point. With the terminal voltage held near 346 V the center reaches the 1414 °C melting point 6.50 min later. With the current regulated instead, starting from the same 525 A, the same rod holds 1100.0 °C for the two hours the case is run. This is a property of the plant and of the material, it is decided entirely by which terminal of one power supply is held constant, and it is invisible in a block diagram that draws the supply as a source of power.

**Pause.** Try Exercise 2.1 before reading on.

The same term sets the loop gain. Differentiating the steady-state balance $P = Q_{\text{loss}}$ with (2.16) and (2.9) substituted gives Equation (2.17).

$$\frac{\partial T_s}{\partial I} = \frac{2IR}{S + \dfrac{PB}{\bar{T}^2}\left(1 + \dfrac{S}{8\pi k_{Si} L}\right)} \tag{2.17}$$

where the bracket accounts for the fact that raising the power raises the mean temperature by more than it raises the surface. The self-regulation term $PB/\bar{T}^2$ is 2.50 times the loss slope at the seed rod and 2.03 times it at harvest, so it is the larger part of the denominator throughout. Dropping it gives the familiar $2IR/S$, which is the quantity the schedule is built from and which overstates the true gain by a factor of 3.449 at the seed and 3.593 at harvest. The application plots and exports both, labeled `sched_2IR_S` and `gain_dTdI` in the comma-separated record and drawn as two traces on the supply strip. The second is a central difference of the design-point calculation, and it agrees with an independent solution of the model's own equilibrium under a 0.1 % current perturbation to four decimal places at every diameter, giving 10.4645 °C/A at the seed, 0.4240 at 60 mm and 0.0828 at harvest. Equation (2.17) reproduces those to within 7 % over the whole range. Table 2.3 lists the electrical march of the batch and both versions of the gain.

**Table 2.3** The Plant at Five Diameters With the Surface Held at 1100 °C and $x_{\mathrm{TCS}} = 0.120$. Every Column Is a Steady State of the Simulator.

| $d$ (mm) | 8 | 30 | 60 | 100 | 150 |
|---|---|---|---|---|---|
| $R$ (Ω) | 50.9 | 3.155 | 0.659 | 0.189 | 0.0649 |
| $I$ (A) | 21.8 | 169.7 | 525.3 | 1265 | 2646 |
| $V$ (V) | 1111 | 535 | 346 | 239 | 172 |
| $P$ (kW) | 24.2 | 90.9 | 181.7 | 302.9 | 454.4 |
| $\bar{T}$ (°C) | 1108.4 | 1131.6 | 1163.2 | 1205.4 | 1258.1 |
| $T_c$ (°C) | 1116.9 | 1163.2 | 1226.4 | 1310.7 | 1416.1 |
| $S$ (W/K) | 61.6 | 230.9 | 461.7 | 769.6 | 1154.3 |
| $PB/\bar{T}^2$ (W/K) | 154 | 558 | 1068 | 1679 | 2349 |
| $2IR/S$ (°C/A) | 36.0959 | 4.6382 | 1.4986 | 0.6222 | 0.2975 |
| $\partial T_s/\partial I$ (°C/A) | 10.4645 | 1.3290 | 0.4240 | 0.1741 | 0.0828 |
| rod $\tau$ (s) | 7.9 | 29.7 | 59.4 | 98.9 | 148.4 |

### 2.4.2 A Schedule That Is Not Fitted

Every term on the right of (2.17) is a closed-form function of the diameter at a given surface temperature, and the diameter is the integral of a deposition rate the tool already computes. The application therefore schedules the proportional gain by the ratio of the plant gain at the seed to the plant gain now, as shown in Equation (2.18).

$$\gamma(d) = \mathrm{clamp}\left(\frac{g(d_0)}{g(d)}, 0.05, 400\right), \quad g(\delta) = \frac{2I(\delta)R(\bar{T},\delta)}{S(T_s,\delta)} \tag{2.18}$$

where $I(\delta)$ is the design-point current of (2.15) evaluated at diameter $\delta$ and the present effective setpoint, $R$ is evaluated at the mean rod temperature and $S$ at the surface temperature that the supply's meters infer (Section 2.5.3), and every term uses the declared parameters of Table 2.2 with the design heat loss and gas velocity, so that only the diameter varies between the numerator and the denominator. The multiplier is referenced to the seed so that the number the reader enters is the gain at the start of a batch. The self-regulation term of (2.17) is deliberately left out of $g(\delta)$, and the reason is the most useful thing this application has to say about gain scheduling. The schedule quantity is wrong in magnitude by a factor of about 3.5 and right in shape, because that factor moves only from 3.449 at the seed to 3.593 at harvest while the schedule quantity itself moves by 121. A schedule needs the shape. The one constant factor it gets wrong is absorbed once, by the proportional gain entered at the seed, and it never has to be revisited. The ratio of the simplified expression between seed and harvest is 121 against 126 for the measured derivative, a difference of 4 % over two orders of magnitude.

🕮 **Pause.** Try Exercise 2.2 before reading on.

Gain scheduling carries a standing warning that the scheduling variable must vary slowly against the closed-loop dynamics [1, 2]. Here that condition is met by a margin that is almost embarrassing. The rod time constant is between 7.9 and 148 s while the diameter takes 79 h to traverse its range, so the scheduling variable moves three orders of magnitude more slowly than the plant it schedules. It is also monotonic, unambiguous, and the physical cause of the gain change rather than a correlate of it [3].

### 2.4.3 A Constraint Reached Only Through a Model

Combining (2.8) and (2.9) gives a center temperature computed from the electrical power and two geometric constants. With the surface held at 1100 °C the center passes the default 1350 °C limit at 118.7 mm and reaches the 1414 °C melting point at 149.0 mm, one millimeter short of the harvest diameter. Dropping the surface setpoint to 1050 °C moves the 1350 °C crossing out to 162 mm, past harvest. The recipe for the last third of a batch is therefore a constraint problem rather than a tuning

problem, and the variable it is written on cannot be measured. Its estimator is the model in (2.8) applied to the supply's meters, with the mean temperature taken from the measured resistance $V/I$ through the declared resistivity law of (2.1) and the rise taken from the measured power $VI$ at the declared $k_{Si}$ and $L$, which is the classical inferential arrangement of a primary variable reconstructed from secondary measurements [16, 17]. The kiln of Volume 5, Chapter 11 estimates the core of a brick that no instrument reaches in the same way.

### 2.4.4 A Reaction on the Flat Part of an Arrhenius Curve

The two resistances in (2.10) are not comparable at the operating point. Evaluated at $x_{\mathrm{TCS}} = 0.120$ and $v_g = 1$, the surface reaction delivers $k_s = 1.812$, 2.588, 3.600 and 5.190 mm/h at surface temperatures of 1000, 1050, 1100 and 1160 °C, against a mass transfer of 1.20 mm/h that does not depend on temperature at all. Above about 1050 °C the transport step is the slower one and it sets the rate. Raising the surface from 1100 to 1160 °C therefore raises the diameter growth from 1.800 to 1.949 mm/h, a gain of 8.3 %, while the boron content of (2.14) rises from 0.2100 to 0.2921 ppba, a loss of 39 % against a specification of 0.30 ppba. The same sixty degrees adds 34 °C to the parabolic rise, so the center at 100 mm moves from 1311 to 1405 °C. Every one of those three consequences is available from the model before the application is opened, and none of them is visible in a rate equation read without asking which of two series terms is smaller.

## 2.5 Controller Design

### 2.5.1 The Implemented Algorithm

TIC-401 and the composition controller are both instances of the velocity-form PID of Section 1.6.2, with the derivative taken on the measurement. Equations (2.19) to (2.21) write it with the schedule multiplier of this plant.

$$e_k = a\left(T_{\mathrm{eff},k} - \hat{T}_{s,k}\right) \tag{2.19}$$

$$P_k = K_c\,\gamma_k\,e_k, \qquad D_k = -K_c\,\gamma_k\,\tau_D\,a\,\frac{\mathrm{d}\hat{T}_s}{\mathrm{d}t}$$

$$\Sigma_k = \Sigma_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{2.20}$$

$$\begin{aligned} c_k &= u_{\mathrm{bias}} + \Sigma_k + I_{k-1} \\ u_k &= \min(I_{\max},\ \max(0,\ c_k)) \\ I_k &= I_{k-1} + \left[\frac{K_c\,\gamma_k}{\tau_I}e_k + \frac{u_k - c_k}{T_t}\right]\Delta t_c \end{aligned} \tag{2.21}$$

where $a = +1$ for both loops because both process gains are positive and a rod reading below setpoint calls for more current, which is reverse action in the convention of Section 1.6.2, $\gamma_k$ is the schedule multiplier of (2.18) and equals one when scheduling is off, $T_{\mathrm{eff}}$ is the setpoint after any override, and $\Delta t_c$ is the controller period. The symbols map onto Chapter 1 as $\Sigma_k = S_k$, $c_k = v_k$ and $u_{\mathrm{bias}} = u_b$. The running sum is never clipped, and the back-calculation anti-windup, with $T_t = \tau_I$, acts on the integral alone, as Section 1.6.2 describes. It sees only the limits of (2.21). Under voltage regulation the current command is converted to a terminal voltage $V_{\mathrm{cmd}} = u_k R_{\mathrm{ref}}$, with $R_{\mathrm{ref}}$ the resistance at the design point for the present diameter and effective setpoint, and clamped at the 1800 V rating, and the supply caps the delivered power at $I_{\max}^2 R$ and $1.5 P_{\max}$. None of these is known to the controller, so while one binds the integral keeps winding, as Section 2.6.2 shows. Both loops run with $\tau_D = 0$, so the derivative term is present in the implementation and inactive in every scenario, and no rate limit is configured on either output. The temperature output is clamped to 0–3000 A and the feed output to 0.2–12 kmol/h. Transfer between Manual and Auto is bumpless in both directions, as in Section 1.6.2. The controller period is 2.88 s against a physics step of 1.44 s, and both are far shorter than the shortest rod time constant of 7.9 s.

The default tunings are $K_c = 0.0080\,\mathrm{A/°C}$ and $\tau_I = 9.0\,\mathrm{s}$ on temperature, and $K_c = 10$ kmol/h per mole fraction with $\tau_I = 31.3\,\mathrm{s}$ on composition. The composition reset time is set equal to the 31.4 s gas residence time of (2.12).

### 2.5.2 The Instrument Chain

The pyrometer is the only sensor in the loop. Its reading is formed from the true surface temperature by adding a Gaussian sample scaled to $\sigma_T$, adding the fixed offset $b_T$, and passing the sum through a first-order filter of time constant $\tau_p$, as shown in Equation (2.22).

$$\begin{aligned} \hat{T}_{s,k} &= \alpha\,\hat{T}_{s,k-1} + (1-\alpha)\left(T_{s,k} + \sigma_T n_k + b_T\right) \\ \alpha &= \mathrm{e}^{-\Delta t_c/\tau_p} \end{aligned} \tag{2.22}$$

where $n_k$ is a unit normal deviate drawn once per controller execution and $\tau_p$ defaults to 1.8 s. The exact zero-order-hold update is unconditionally stable at any sample period. The offset $b_T$ is the disturbance that matters, because an emissivity error or a silicon film on the sight glass is a bias and not noise, and a bias inside the loop is indistinguishable from a setpoint change. Section 1.4.2 shows how an emissivity error becomes such a bias in a pyrometer reading.

### 2.5.3 The Override on the Inferred Center

The center-temperature limit is enforced by retreating the surface setpoint inside a band $\Delta$ below the limit, as shown in Equation (2.23).

$$T_{\mathrm{eff}} = \begin{cases} T_{\mathrm{SP}} & \mu \geq \Delta \\ T_{\mathrm{SP}} + K_{ov}\,(\mu - \Delta) & \mu < \Delta \end{cases}, \quad \mu = T_{c,\max} - T_c \tag{2.23}$$

where $\mu$ is the margin to the limit in degrees, $\Delta$ defaults to 80 °C and the override gain $K_{ov}$ defaults to 1.0 °C/°C. The structure is a low selector written as a continuous retreat rather than as a switch, so the two objectives blend over the band instead of trading control of the actuator abruptly [18]. The compressor of Volume 4, Chapter 2 guards its surge limit with an override of another form, a loop whose output rests on its clamp until the limit approaches. The action here is proportional only, so the center settles a little short of the limit rather than exactly on it and the band is what decides by how much. A first-order override needs a band wide enough to arrest an approach that is itself accelerating, and Section 2.6.4 measures what 40, 80 and 120 °C each cost. The important property is the independence of the two paths. The override is computed from the supply's voltage and current and the diameter, the mean temperature from the resistance $V/I$ through the declared resistivity and the rise from the power $VI$ through (2.8) at the declared $k_{Si}$ and $L$, the pyrometer appears nowhere in that calculation, and a fault in the temperature measurement therefore shows up in the override as a constraint that is moving for no visible reason.

Every quantity the controllers compute is one the tool could know. The override, the schedule of (2.18), the open-loop current, the voltage reference, the starting trichlorosilane feed and the manual-mode aid of (2.15) are evaluated with the declared parameters of Table 2.2, the design heat loss and gas velocity, and the temperatures the supply's meters infer. None of them reads the plant's own mean rod temperature or the run's heat-loss factor, gas velocity and parameters, which no instrument on a real reactor reports. At the declared values the metered center equals the plant's own to within $1 \times 10^{-6}$ °C. A changed $k_{Si}$, resistivity activation temperature $B$ or rod length $L$ is therefore a model error that the override cannot see, and Section 2.6.5 measures how far each one moves the estimate.

### 2.5.4 A Shrinking-Horizon Predictive Controller

A model predictive controller (MPC) is provided as a third structure for the temperature loop, beside the open loop and the PI on the pyrometer. It moves one variable, the setpoint $u$ of the surface loop,

and leaves the PI with its schedule underneath and the override of (2.23) above, so the override retreats the planned setpoint exactly as it retreats the operator's. The problem it addresses is the one of Section 2.4.3 posed as a batch optimization, in which growth is the objective, the surface setpoint trajectory is the decision and the center temperature is a path constraint [20], with growth itself as the objective in place of a tracking error [21]. Studies of the Siemens reactor have tailored the rod temperature profile for throughput in the same spirit [22]. It departs from the linear model and the tracking objective of the option most apps share (Section 1.6.3), as the following paragraphs set out.

The controller's state is the diameter instrument and the center estimate of Equation (2.24), the inferential calculation of Section 2.5.3 written out.

$$\hat{T}_c = \frac{B}{\ln[(\rho_m - \rho_\infty)/A]} + \frac{VI}{8\pi k_{Si} L}, \qquad \rho_m = \frac{V}{I}\,\frac{\pi d^2}{4L} \tag{2.24}$$

where $V$ and $I$ are the supply's voltmeter and ammeter readings, $\rho_m$ is the measured resistivity, the first term inverts (2.1) and (2.2) for the mean rod temperature in kelvin, and every constant takes its declared value. The pyrometer does not enter it. Over the rod of scenario 7 from 100 to 150 mm the estimate departs from the true center by less than $1 \times 10^{-6}$ °C at the declared values, by $-21$ to $-30$ °C with $k_{Si} = 18$, by $+16$ to $+22$ °C with $k_{Si} = 26$, by $-134$ to $-139$ °C with $B$ raised 10 %, and by $-5$ to $-11$ °C with rods 6.0 m long.

The controller's model is not identified from step tests. It is the design point of (2.15) at the declared parameters and the design disturbances, which gives the center $f(d,u)$ and the current $I(d,u)$ at a diameter and a surface temperature, together with the diameter growth $g(u)$ of (2.11) at the composition setpoint. It assumes the surface loop holds its setpoint, which the scheduled PI does within minutes while the diameter moves less than 0.09 mm between decisions. Two output biases make it offset-free, $b_T = \hat{T}_c - f(d,\hat{T}_s)$ and $b_I = I - I(d,\hat{T}_s)$, each filtered by a first-order lag of 0.01 h. They absorb what the model leaves out, such as a heavier heat loss, a slower gas or a fogged sight glass, but they cannot absorb an error in $\hat{T}_c$ itself.

A decision is taken every 60 controller executions (0.048 h, counted in integers), and at once when $b_T$ has moved 3 °C or $b_I$ 2 % of the rating since the last one. The horizon is the time the previous plan predicts the rod needs to reach $d_{\text{end}}$, at most 120 h, cut into at most 16 graduated blocks. The first two last one decision period each, every later block is as long as the two before it together, and the last takes the remainder, so the horizon shrinks with the batch. Measured on scenario 11 it is 13 blocks and 28.55 h at 100.5 mm, seven blocks and 1.20 h at 148 mm and two blocks and 0.10 h at 149.95 mm. Each decision solves Equation (2.25) for the block setpoints $u_j$ and two slacks.

$$\begin{aligned}
\min_{u,s_T,s_I}\ & -\sum_{j=1}^{N_b} \frac{\tau_j\, g'(\bar{u}_j)}{g(T_{\text{SP}})}\,(u_j - \bar{u}_j) + S\sum_{j=1}^{N_b}(u_j - u_{j-1})^2 \\
& \quad + \rho\left(s_T^2 + s_I^2\right) + \sigma\,(s_T + s_I) \\
\text{s.t. } & 950\,^\circ\text{C} \le u_j \le T_{\text{SP}}, \qquad |u_j - \bar{u}_j| \le 40\,^\circ\text{C} \\
& f(d_j,u_j) + b_T \le T_{c,\max} - 5\,^\circ\text{C} + s_T \\
& \frac{I(d_j,u_j) + b_I - 0.98\,I_{\max}}{0.01\,I_{\max}} \le s_I, \qquad s_T,\ s_I \ge 0 \\
& d_j = d_{j-1} + \tau_j\, g(u_j)
\end{aligned} \tag{2.25}$$

where $\tau_j$ is the length of block $j$, $\bar{u}_j$ is the previous plan at that block, $u_0$ is the setpoint in force, $d_0$ is the measured diameter, and $d_j$ is capped at $d_{\text{end}}$ so that the center is judged only up to the harvest. The objective is in hours. The first term is the diameter the plan adds, expressed as hours of growth at the operator's setpoint $T_{\text{SP}}$, which is also the ceiling because a hotter surface is less pure. The move weight is $S = 10^{-4}$ h/°C$^2$, and the slacks, in degrees for the center and in percent of the rating for the current, are priced at $\rho = 10$ and $\sigma = 10$ hours per unit squared and per unit. A degree of setpoint is

**Table 2.4** Scenario Sequence and the Mechanism Each Teaches. The Starting Diameter Is 8 mm Unless Shown.

| # | Scenario | What it teaches |
|---|---|---|
| 1 | Seed rod by hand, MANUAL | arithmetic of a 51 Ω heating element |
| 2 | Loop closed on a small rod | closed-loop step on a fast plant |
| 3 | Same tuning at 140 mm | a plant gain 111 times smaller |
| 4 | Schedule on diameter, 140 mm | closed-form gain scheduling |
| 5 | Voltage regulation, 60 mm | thermal runaway to the melting point |
| 6 | Run to harvest, 100 mm | the unmeasured center constraint |
| 7 | Override on, 100 mm | constraint handling on an estimate |
| 8 | Surface at 1160 °C, 60 mm | rate against purity |
| 9 | Sight glass fogs, 90 mm | sensor bias inside the loop |
| 10 | Feed boron triples, 70 mm | a fault no instrument sees |
| 11 | MPC under the override, 100 mm | a plan on an estimated limit |
| 12 | Graded challenge, 100 mm | four faults at once |

worth 0.048 to 0.057 h of batch time over a 30 h batch, so the linear price is about two hundred times larger and the soft limits hold exactly whenever they can be held. The constraint rows are linearized along $\bar{u}$ by central differences of 0.5 °C and 0.5 mm, the condensed program of $N_b + 2$ variables and at most $4N_b + 2$ rows is solved by the dual active-set method [23], and the rows are relinearized along the new plan up to three times, stopping when no block moves more than 0.05 °C. Scenario 11 took two passes in 634 of its 642 decisions and three in the other eight, the nonlinear model’s center along every accepted plan stayed at or below the 1345 °C ceiling, and the largest slack was zero.

Only the first block is applied. Acceptance requires an optimal status from every pass, finite values of the applied move, the planned trajectory, its predicted center and current, its cost and both slacks, planned setpoints within their bounds, and the expected number of blocks. A rejected plan is marked in the historian and the request returns to the operator’s setpoint over a 0.25 h ramp, and the next accepted plan ramps back over the same interval. Five injected faults, one for each acceptance check, were each rejected with the reason stated and handed over with a largest request step of 0.36 °C per execution. The plan does not model the override above it, and its prediction of the hours to harvest and of the highest center is therefore its own outcome rather than that of the trimmed run. Measured in the browser on a shared and loaded host, the controller with its decisions takes between 6.2 and 6.7 % of one processor core at the fastest simulation speed and under 0.8 % at the default speed, with a mean decision of 0.15 to 0.41 ms and a largest of 1.9 ms.

## 2.6 Results

Table 2.4 lists the twelve scenarios, which are ordered as a laboratory sequence. A scenario fixes the starting mode, the initial diameter, the two tunings, the control options, the disturbance settings and any timed events.

Scenario 1 makes the point with arithmetic before any controller is involved. The manual aid returns 21.8 A for an 8 mm rod at 1100 °C, which is 1111 V across 50.9 Ω for 24.2 kW. An hour later the rod is 9.66 mm, the same 21.8 A holds only 1037.0 °C, and the calculation has to be done again. At harvest the same rod pair needs 2646 A at 172 V, so a supply for this machine has to deliver a kilovolt at tens of amps on the first day and a quarter of a kilovolt at more than two thousand on the fourth.

### 2.6.1 The Same Controller on a Plant That Grew

Scenarios 3 and 4 apply an identical 20 °C setpoint step to a 140 mm rod, differing only in whether (2.18) is active. With scheduling off, the surface reaches 1053.2 °C two and a half hours after a step to 1070 °C and never enters a ±2 °C band, leaving 16.8 °C of error outstanding. With scheduling on, the multiplier $\gamma$ is 50.9 at the controller scan that applies the step, on a rod that has

**Table 2.5** Four Complete Batches From an 8 mm Seed With Identical Tuning. IAE Is Accumulated Against the Effective Setpoint Over the Whole Run, Whose Length Is Given in the First Row.

| | PI only | PI + sched. | PI + ovr. | PI + both |
|---|---|---|---|---|
| Batch length (h) | 80.14 | 78.57 | 81.55 | 80.90 |
| Final diameter (mm) | 150.0 | 149.3 | 150.0 | 150.0 |
| Rod melted | no | **yes** | no | no |
| Peak $T_c$ (°C) | 1387 | 1414 | 1317 | 1323 |
| Final $T_s$ (°C) | 1083.8 | 1099.7 | 1043.7 | 1047.1 |
| IAE (°C h) | 760.4 | 43.1 | 625.6 | 38.8 |
| Max. deviation (°C) | 16.16 | 2.24 | 10.31 | 2.24 |
| Current travel (A) | 2432 | 2596 | 2001 | 2035 |
| Silicon (kg) | 213.5 | 211.5 | 213.5 | 213.5 |
| Energy (kW h/kg) | 87.5 | 88.4 | 86.6 | 87.0 |
| Boron (ppba) | 0.1972 | 0.2095 | 0.1826 | 0.1873 |

reached 140.8 mm, and the true surface settles inside the band in 331 s and overshoots by 0.40 °C, or 2.0 % of the step. The overshoot the scorecard reports from the logged pyrometer reading is 15.6 %, because the 2 °C pyrometer noise is a tenth of the step, and over six noise seeds it ranges from 11.7 to 22.7 % while the settling time ranges only from 330 to 359 s. Read at $t = 4$ h, the integrated absolute error is 65.0 °C h unscheduled against 5.9 °C h scheduled, a factor of 11.0, and the actuator travel rises from 206 A to 420 A. The loop was never unstable and was never mistuned. It became slow because the plant became sluggish underneath it.

> **TRY IT IN THE APP** **3 · The Same Tuning at 140 mm**
>
> Load it and press Run. At half an hour the setpoint steps from 1050 to 1070 °C, and two and a half hours later the unscheduled loop has the surface only at 1053.2 °C. At 4 h the scorecard's IAE reads 65.0 °C h. Load 4 · Scheduling on a Diameter the Tool Knows Exactly and run it to the same time: the step settles within minutes and IAE reads 5.9 °C h.

Running the whole batch separates the two further. Table 2.5 reports four complete runs from an 8 mm seed to harvest or failure, with the same PI tuning in every column and only the two option switches changed.

The first column is a loop that slowly gives up. The diameter is a ramp disturbance on the current requirement, a proportional-integral loop tracks a ramp with a standing offset inversely proportional to its gain, and the gain is falling, so the surface droops monotonically from 1096.2 at 20 mm to 1084.8 °C at 140 mm. The second column never lets the surface fall more than 2.24 °C below setpoint, a deficit it reaches in the first hour on the seed rod, holds it between 1099.0 and 1099.7 °C from 21 mm to the end of the batch, and cuts the accumulated error by a factor of 17.6 over runs of 80.14 and 78.57 h. Figure 2.3 draws the accumulated error and the peak center of the four batches side by side.

The second column also melts the rod. Holding 1100 °C accurately all the way to 149.3 mm is exactly what drives the unmeasured center to 1414 °C, while the drooping surface of the first column keeps the peak center at 1387 °C and survives. Improving the loop made the outcome worse, and the reason is that the regulated variable was not the one that ended the batch. The same situation is visible part way through scenario 6, where a pyrometer reading 1100.2 °C against a 1100 °C setpoint accompanies a computed center of 1393.4 °C, as shown in Figure 2.4. Run on to the end, scenario 6 takes the center past the 1350 °C limit at 118.9 mm and melts the rod at 149.2 mm, within a few tenths of a millimeter of the steady-state crossings of Section 2.4.3. Adding the override of Section 2.5.3 to the scheduled loop caps the peak center at 1323 °C, harvests the full 150 mm in 80.90 h and yields 213.5 kg at 87.0 kW h/kg. The override costs 2.34 h of batch time against the run that melted, which

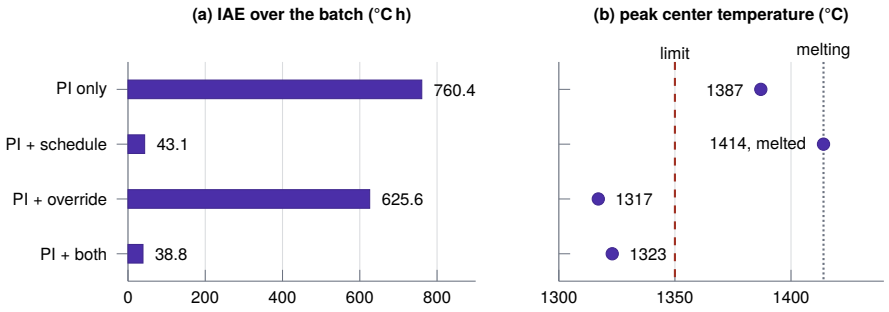



**Figure 2.3** The four complete batches of Table 2.5, with identical tuning. (a) The integrated absolute error over the whole batch: the schedule cuts it from 760.4 to 43.1 °C h, and from 625.6 to 38.8 °C h with the override in service. (b) The peak true center temperature against the 1350 °C limit (dashed) and the 1414 °C melting point (dotted). Notice that the run that tracks its setpoint best without the override is the one that melts: the regulated variable is not the one that ends the batch.

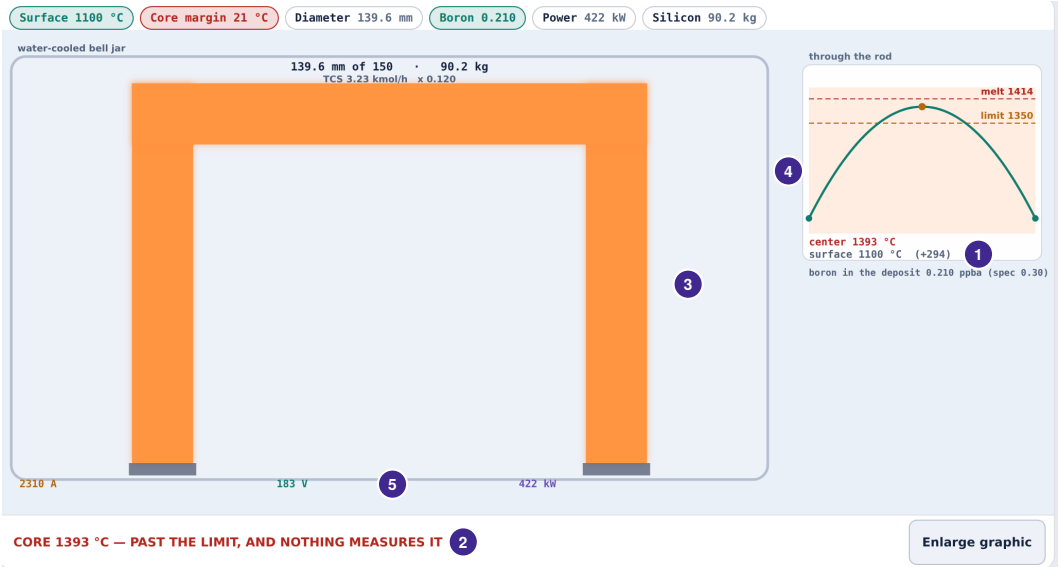



**Figure 2.4** Scenario 6 at $t = 22.0$ h, with the center-temperature override out of service, cropped to the plant view. The scheduled loop holds the pyrometer at 1100.2 °C against a setpoint of 1100 while the center stands at 1393.4 °C, past the 1350 °C limit and 21 °C from the melting point, on a rod that has reached 139.6 mm (3). The cut through the rod (4) draws the parabola of Figure 2.2 against the limit and the melting point, with the surface the pyrometer reports at its two ends (1) and a center-above-surface rise of 293.8 °C. The alarm line (2) states what the panel knows: the center is past the limit and nothing measures it. The supply's readings at the foot of the rods (5), 2310.0 A at 182.83 V for 422.35 kW, are what the override computes the center from.

buys the last 0.7 mm, a center that is still 91 °C from the melting point, and a deposit at 0.1873 ppba instead of 0.2095 ppba.

**DID YOU KNOW? The Loop That Tracks Worse Uses Less Energy**

From the simulator: in Table 2.5 the unscheduled loop with the override accumulates an error of 625.6 °C h, against 38.8 °C h with the schedule added. Yet it harvests the same 213.5 kg with the least energy of the four, 86.6 kW h/kg, and the least boron, 0.1826 ppba. Its surface ends the coolest, at 1043.7 °C.

### 2.6.2 Losing the Rod to the Choice of Terminal

Scenario 5 starts a 60 mm rod at its design point and runs it open loop, with the same steady-state calculation supplying the command under each regulation mode. Under current regulation the surface holds 1100.0 °C for the whole two hours, and the only thing that moves is the drift the growing rod causes, with the current rising from 525 to 579 A and the terminal voltage falling from 346 to 333 V as the deposit thickens. The power slope logged in the exported record (columns `dpdt_W_K` and `loss_slope_W_K`) is −1068 W/K against a loss slope of 462 W/K.

Under voltage regulation nothing else is changed and the rod is destroyed. The terminal voltage is held between 345 and 346 V, the surface leaves 1100 °C within four minutes, and the center reaches the 1414 °C melting point 6.50 min after the run starts. At that instant the surface is 1172.6 °C, the current has risen on its own from 525 to 1008 A and the power from 182 to 348 kW, all at a voltage that barely moved. That is the positive slope of (2.16) closing on itself, and the slope logged before anything happened was +1068 W/K against the same 462 W/K. The application stops the run at the melt. Continuing the integration past it, as a counterfactual model run, the delivered power reaches the 825 kW ceiling of $1.5P_{\max}$ at 6.91 min and the current reaches the 3000 A supply rating at 7.10 min, 0.6 of a minute after the center passed the melting point. The supply has left voltage regulation, the terminal voltage collapses from 345 to below 225 V by 7.4 min, and the surface peaks at 1678 °C and stands at 1618 °C, with the axis at 2047 °C, two hours in. The saturation flag is raised from the controller scan at 7.01 min onward, which is worth noticing, because the controller never asked for the rating. Its command stays between 528 and 595 A while the rod pulls the rest on its own, so the flag means the actuator is out of room rather than that the loop wanted more.

**TRY IT IN THE APP 5 · Current Against Voltage**

Load the scenario and press Run. The supply holds the voltage of a 60 mm rod at its design point, and no controller acts. Watch the center trace climb until the run stops itself, the rod melted 6.50 min after the start, while the current has risen on its own from 525 to 1008 A. Reset, set Power supply regulates to Current and run again: the surface holds 1100.0 °C.

Closing the temperature loop does not rescue it, and the way it fails is worth a sentence. The pyrometer sees the surface leave its setpoint and the controller raises the current command, which raises the voltage command, but the plant underneath is unstable in both directions and this scenario departs downward. The surface falls through 1085.1 °C at 3 min, 952.3 °C at 4.8 min and 521.2 °C at 9 min while the resistance of (2.1) climbs through its exponential, so $P = V^2/R$ falls the whole way. The voltage command reaches the 1800 V supply ceiling at 59.8 min, and because that clamp is applied after the controller, the integral keeps winding until the current command itself reaches 3000 A at 65.6 min. At 6 h the rod sits at 294.4 °C drawing 0.0 A. Feedback wrapped around an unstable plant whose actuator authority collapses with the state is not a fix, and no tuning changes the outcome.

### 2.6.3 A Bias Inside the Loop and a Fault With No Instrument

Scenario 9 deposits silicon on the sight glass at $t = 2$ h so the pyrometer reads 25 °C low. With the override off, the loop does precisely what it was asked to do. Two hours later the reading is 1097.6 °C, within the pyrometer noise of a setpoint it appears to be holding, while the true surface has risen from 1099.7 to 1124.3 °C and the true center from 1297 to 1343 °C. By $t = 20$ h the center stands at 1410 °C, four degrees below melting, and the boron average has risen from 0.2096 to 0.2378 ppba.

With the override on the rod is already being held back before the fault arrives, at 1089.1 °C rather than 1099.7, and two hours after the same event the true surface is only 6.8 °C higher, because the override sees the extra power and retreats the effective setpoint to 1071.3 °C. The center ends the run at 1321 °C instead of 1410 °C, and the deposit at 0.1947 ppba instead of 0.2378 ppba, for 72.1 kg against 76.9 kg. The estimator catches the sensor fault precisely because the faulty sensor is not one of its inputs. An inferred variable built from an independent set of measurements is a cross-check as well as a constraint handler, and on a plant with one critical instrument it is often the only cross-check available.

Scenario 10 is the case with no such luck. The upstream distillation lets three times the boron through at $t = 2$ h. The true surface temperature reads 1099.63 just before the event and 1099.38 °C two hours after it, which is the wander the pyrometer noise causes. The gas composition stays at 0.1200, the growth rate stays at 1.799 mm/h, and the current rises only as the diameter dictates. The boron in the deposit goes from 0.2096 to 0.629 ppba, twice the 0.30 ppba specification, and the average over the whole deposit reaches 0.5677 ppba by $t = 20$ h. No instrument in the reactor responds, the product is scrap, and the only way to establish it is to analyze a piece of the rod after the batch.

🕮 **Pause.** Try Exercise 2.3 before reading on.

### 2.6.4 Rate Against Purity, and the Graded Challenge

Scenario 8 runs a 60 mm rod with the surface setpoint at 1160 °C for throughput and the override in service. Over 20 h it produces 56.0 kg at 90.8 kW h/kg with a boron average of 0.2414 ppba. The same scenario at 1080 °C produces 51.2 kg at 86.9 kW h/kg with 0.1877 ppba. Eighty degrees of extra setpoint bought 9.4 % of mass, cost 4.5 % in energy per kilogram and 29 % more boron, and moved the center from 1270 to 1320 °C. The surface does not stay at 1160 °C, because the override trims the effective setpoint to 1144 °C within three minutes and takes the surface to 1109.7 °C by 20 h as the center approaches the limit, so the measured trade is smaller than the 8.3 % against 39 % of Section 2.4.4 only because the constraint intervened first.

The challenge starts a 100 mm rod with four things wrong at once. The loop is unscheduled, the setpoint is 1140 °C instead of 1100, the bell jar is losing 10 % more heat than the recipe assumed, and the sight glass fogs at $t = 6$ h. The target is to run the rod to the 150 mm harvest with the measured surface within 25 °C of its effective setpoint and the rod's true center below 1350 °C from the first hour on, and the deposit inside specification at harvest. Changing a model parameter or a disturbance, or firing an extra disturbance, voids an attempt. Run as configured, the rod melts at $t = 6.93$ h and 113.0 mm, with a peak deviation of 42.2 °C between the logged pyrometer reading and its setpoint. The suggested configuration turns scheduling and the override on and takes the setpoint back to 1100 °C. It reaches harvest at 29.7 h with a worst surface error of 5.4 °C after the first hour, an IAE of 39.7 °C h over the run, 118.95 kg deposited, and a boron average of 0.171 ppba against the 0.30 ppba specification. As loaded the center stands at 1396 °C, past the limit, but settings chosen before Run set the starting state, so the suggested run starts from the 1100 °C design point with the center at 1332 °C. The override pulls it down to 1293.5 °C about six minutes in, it stays below 1303 °C through the first six hours, and drifts up as the rod thickens, to 1342 °C at 24 h and a peak of 1349.0 °C just before harvest, 1.0 °C inside the limit. The completion rule judges the center on the logged record from the first hour on, where its peak is 1348.8 °C, 1.2 °C inside the limit, and the worst surface error is 19.6 °C inside its band. Over six pyrometer-noise seeds the judged center peak stays between 1348.6 and 1348.9 °C and every run passes. The width of the override band is what decides that margin, because the action of (2.23) is proportional and the offset grows with the diameter. Narrowing $\Delta$ from the default 80 °C to 40 reaches harvest 0.83 h sooner but lets the center peak at 1374.8 °C, 24.8 °C past the limit, and leaves it above the limit for 18.07 h, so the attempt fails. Widening $\Delta$ to 120 °C holds the peak to 1323.3 °C and costs 0.94 h of batch time. Because every run ends at the

**Table 2.6** The Predictive Controller Against the Override on the Rod of Scenario 7 (100 mm to 150 mm, Same Seed, Noise and Instruments). The Margin Is the Least Distance of the True Center Below the 1350 °C Limit, and the Band Δ and Gain $K_{ov}$ Are Those of (2.23).

| | Structure | Batch (h) | kW h/kg | Peak $T_c$ (°C) | Margin (°C) | Boron (ppba) | At $I_{max}$ (h) | Solves |
|---|---|---|---|---|---|---|---|---|
| *Standard recipe, 3000 A supply* | | | | | | | | |
| 1 | PI and override, $\Delta = 80$ (shipped) | 29.595 | 85.99 | 1323.0 | 27.0 | 0.172 | 0 | — |
| 2 | MPC under the override, $\Delta = 80$ (scenario 11) | 29.940 | 85.76 | 1310.3 | 39.7 | 0.167 | 0 | 642 |
| 2′ | MPC with the band narrowed to $\Delta = 5$ | 28.482 | 87.28 | 1345.6 | 4.4 | 0.193 | 0 | 609 |
| 3a | Override retuned, $\Delta = 50$, $K_{ov} = 1$ | 28.956 | 86.64 | 1342.0 | 8.0 | 0.183 | 0 | — |
| 3b | Override retuned, $\Delta = 20$, $K_{ov} = 4$ | 28.675 | 87.02 | 1340.9 | 9.1 | 0.189 | 0 | — |
| 4a | MPC, $\Delta = 80$, without the center row | 29.595 | 85.99 | 1323.0 | 27.0 | 0.172 | 0 | 628 |
| 4b | MPC, $\Delta = 5$, one-block horizon | 28.485 | 87.28 | 1345.5 | 4.5 | 0.193 | 0 | 611 |
| 4c | MPC margin 80 °C, $\Delta = 5$ | 31.035 | 85.15 | 1310.7 | 39.3 | 0.153 | 0 | 672 |
| 4d | Override alone, $\Delta = 10$, $K_{ov} = 10$ | 28.530 | 87.22 | 1344.2 | 5.8 | 0.192 | 0 | — |
| *Limited supply, 1700 A rating* | | | | | | | | |
| 1 | PI and override, $\Delta = 80$ (shipped) | 30.146 | 85.74 | 1310.3 | 39.7 | 0.165 | 12.28 | — |
| 2 | MPC under the override, $\Delta = 80$ (scenario 11) | 30.522 | 85.54 | 1310.3 | 39.7 | 0.160 | 0 | 656 |
| 2′ | MPC with the band narrowed to $\Delta = 5$ | 29.739 | 86.38 | 1345.3 | 4.7 | 0.174 | 0 | 643 |
| 3b | Override retuned, $\Delta = 20$, $K_{ov} = 4$ | 29.607 | 86.40 | 1333.8 | 16.2 | 0.175 | 17.02 | — |
| 4e | MPC, $\Delta = 80$, without the current row | 30.222 | 85.67 | 1310.3 | 39.7 | 0.164 | 9.79 | 638 |

same diameter, the deposit is the same 118.95 kg in all three and the band trades batch time against margin alone. The default and the wide band never put the center above the limit.

### 2.6.5 The Predictive Controller Under the Override

Scenario 11 runs the rod of scenario 7 (100 mm, 1100 °C, 2 °C pyrometer noise) under the controller of Section 2.5.4, with the schedule on and the override above it at its shipped band of 80 °C. Every row of Table 2.6 uses the same rod, seed, noise and instruments and runs to harvest. The comparison is a negative result. Under the shipped protection the MPC harvests in 29.940 h against 29.595 h for the override alone, 0.345 h slower, for 85.76 against 85.99 kW h/kg. Its first plan predicts 28.69 h to harvest with the center at 1345.0 °C, but the override does not know the plan is there and takes its own correction off the planned setpoint, so the center settles between 1296.7 and 1299.1 °C from 10 h on and the two margins add. At harvest the plan requests 1059.1 °C and the override applies 1031.4 °C. Without its center row the plan rides at the operator's setpoint and the run is the override's own to the last digit, 29.595 h.

Narrowing the band to 5 °C, so that the override is only a backstop behind the plan, gives the result the controller was built for. The MPC then harvests in 28.482 h with the center held at 1343.2 to 1345.6 °C from 10 h on, 1.11 h sooner than the override alone, for 1.5 % more energy per kilogram and 0.193 against 0.172 ppba of boron, as shown in Figure 2.5. The counterfactuals of Table 2.6 show that the hour is margin given up and not prediction. A one-block plan that looks no further than its next 0.048 h takes 28.485 h, because the center limit restricts only the present setpoint at the present diameter and the optimal profile rides that path constraint [20]. Held to a margin of 80 °C behind the same backstop, the plan takes 31.035 h. A proportional override with a high gain reproduces the constant margin without any model of the future. At $K_{ov} = 10$ and a 10 °C band it harvests in 28.530 h with a least margin of 5.8 °C against the plan's 4.4, and over least margins from 4 to 26 °C the plan and that override lie within 0.016 to 0.031 h of each other at equal margin. Of the 1.11 h, 0.55 h comes from narrowing the margin from 27 to 5 °C (the plan at a margin of 27 behind the backstop takes 29.036 h) and 0.56 h from holding the margin constant rather than letting it grow, which is what the high-gain override also does. The override retuned once by the plan's own rule, to the band at which its steady margin at harvest is the plan's 5 °C, recovers 0.64 and 0.92 h of the gain at gains of 1 and 4. The shipped and narrowed results stand without pyrometer noise (29.594, 29.939 and 28.479 h) and at half the integration step (identical at the printed digits).

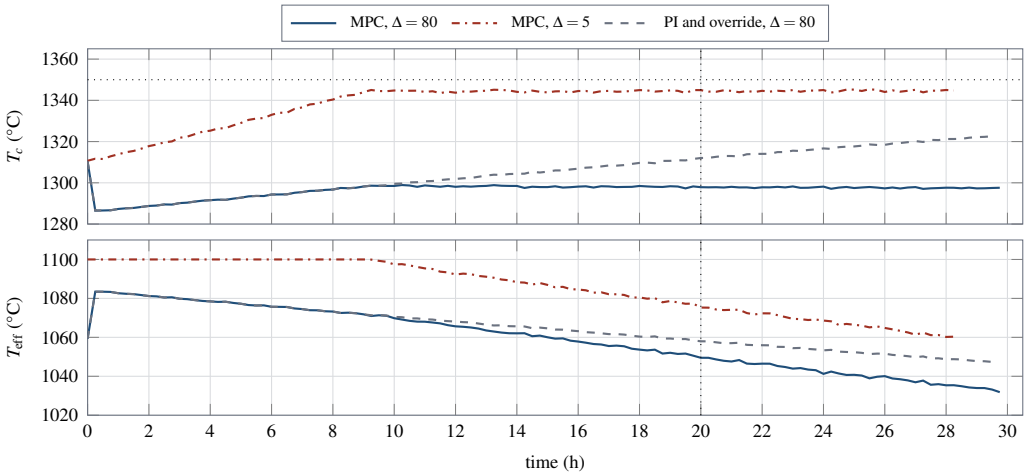



**Figure 2.5** True center temperature and effective surface setpoint of the rod of scenario 7 under scenario 11's MPC (solid), the same MPC with the override band at 5 °C (dash-dotted), and the override alone (dashed), sampled from the engine every 0.25 h. At $t = 20$ h (dotted) the MPC under the shipped band holds the center at 1297.9 °C on an effective setpoint of 1049.5 °C while its plan requests 1077.4 °C, the narrowed MPC holds 1345.0 °C on 1075.3 °C, and the override alone holds 1312.0 °C on 1058.0 °C.

**Table 2.7** Three Model Errors on the Rod of Scenario 7. Error Is the Metered Center Estimate Less the True Center From 100 to 150 mm. Each Run Gives the Hours Over the Limit and the Peak True Center.

| Error | Estimate (°C) | Override $\Delta = 80$ | MPC $\Delta = 5$ | MPC $\Delta = 80$ |
|---|---|---|---|---|
| $k_{Si} = 18$ | −21 to −30 | 9.57, 1363.8 | 29.21, 1375.8 | 0.00, 1356.9 |
| $B + 10\,\%$ | −134 to −139 | 17.28, 1411.5 | melts at 27.3 | 17.28, 1411.5 |
| $L = 6.0$ m | −5 to −11 | 0, 1326.1 | 15.33, 1353.8 | 0, 1310.4 |

A model error spends that margin. With the band at 5 °C the metered estimate of (2.24) is the only protection left, and Table 2.7 runs it against three parameter changes the controller is not told about. With $k_{Si} = 18$ the estimate reads 21 to 30 °C low and the narrowed MPC keeps the center past the limit for 29.21 h, against 9.57 h for the override alone. With $B$ raised 10 % it reads 134 to 139 °C low and the rod melts at 27.3 h, while the override alone peaks at 1411.5 °C, 2.5 °C short of melting. With rods 6.0 m long it reads 5 to 11 °C low and the center spends 15.33 h past the limit, against none. The shipped band prevents all three. Under it the MPC is never longer over the limit than the override alone (0.005 h in the conductivity case, all of it in the first integration step from the starting state, and none in the rod-length case), and in the resistivity case the override trims the plan to exactly its own run. The shipped band therefore removes both the danger and the gain, and it is the reason scenario 11 keeps it. Under the shipped protection the MPC is slower than the override alone, and a narrower override band gives the same batch time without the MPC.

With the supply rated at 1700 A, just under the 1749 A the shipped run draws two thirds of the way from 100 to 150 mm, the MPC under the shipped band takes 30.522 h and never reaches the rating, against 30.146 h with 12.28 h on the rating for the override alone. The current row is the difference. Without it the MPC takes 30.222 h with 9.79 h at the rating, so the 2 % headroom costs 0.30 h. On the challenge at its suggested settings the MPC takes 30.786 against 29.682 h.

The worst case of every limit was run under each shipped structure and under the MPC, as loaded and with the settings of scenario 11, across the existing scenarios, the challenge and eighteen extreme

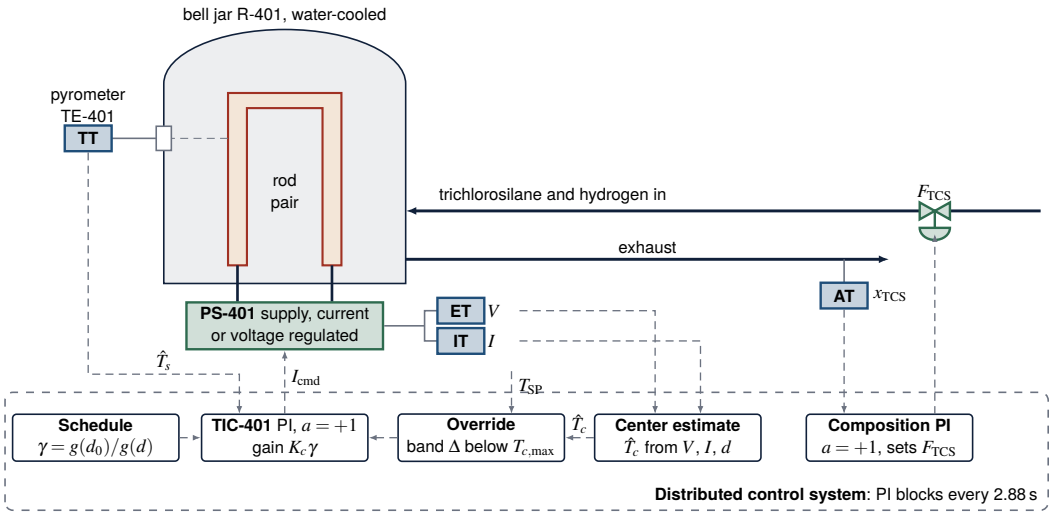



**Figure 2.6** The loops of reactor R-401 drawn with their hardware. Pyrometer TE-401 sights a rod through a sight glass, and TIC-401 sets the current command of supply PS-401 with its gain multiplied by the schedule $\gamma(d)$. The supply's own voltmeter and ammeter feed the center estimate, which drives the override that trims the surface setpoint. The composition loop reads the trichlorosilane fraction of the jar gas and sets the feed. Both PI loops are reverse acting ($a = +1$). Notice that the pyrometer enters only TIC-401: the limit that ends the batch is guarded by the electrical meters alone.

settings. Of the 60 MPC runs, 2 are less safe than the better shipped run, 11 are safer and 47 are the same. Neither flag is the plan's doing. On the challenge as loaded the center peaks at 1396.2 against 1395.2 °C at $t = 0.0004$ h, in the loaded state, while the MPC then spends 2.43 h over the limit against 6.93 h and a melt for the shipped structure. Voltage regulation on the 60 mm rod with the schedule of scenario 11 melts at 0.10 h, exactly as the scheduled PI with the same override does.

## 2.7 Technology in Practice

The loops of Section 2.5 run on a pyrometer, a power supply with its own meters, and a control system, drawn together in Figure 2.6.

### What Measures It

The only temperature instrument is pyrometer TE-401. The model forms its reading from the true surface temperature with 2 °C of noise, a fixed offset that a fault can set, and a first-order lag of 1.8 s. A wrong emissivity becomes the bias of Section 1.4.2.

> **TECHNOLOGY** **Sensor: A Radiation Pyrometer Through a Sight Glass**
>
> A radiation pyrometer concentrates the infrared light a hot surface emits onto a small sensing element, often a thermopile, through a lens or a mirror. The emitted power grows as the fourth power of the absolute temperature, a nonlinearity that confines good accuracy to a narrow range [24]. The instrument reads only the surface: a hotter region inside that does not show at the surface is invisible to it [24]. Its main weakness is inaccuracy. The emittance of the target has to be found once against a known temperature, and radiation reflected from other hot bodies taints the reading [24], and a bell jar holds many glowing rods. On a reactor it sights a single spot of one rod [25]. A film on the sight glass makes it read low and steady, the fault of Section 2.6.3.

The second set of measurements is electrical. The supply's voltmeter and ammeter give the rod resistance $V/I$ and the power $VI$, from which the override infers the mean and center temperatures,

and the deposition rate integrated over the batch gives the diameter. A design for a twelve-rod reactor uses the same two ideas, estimating the diameter from the resistance and the temperature law of silicon and using the resistance as a measurement in its own right [25]. Purity is found only by analyzing a piece of the rod after the batch.

### What Moves It

The final element of the temperature loop is supply PS-401. In the model it follows its current or voltage command through a first-order lag of 0.002 h and rests on ratings of 3000 A, 1800 V and 550 kW, with no rectifier dynamics. A rod reading below its setpoint calls for more current, so TIC-401 is reverse acting, $a = +1$ in the convention of Section 1.6.2.

> **TECHNOLOGY** **Final Element: The Rod Power Supply**
>
> A cold silicon rod barely conducts. The reactor of the original patent was heated up at high voltage from transformer taps that held the current to two amperes until the rods glowed red, and then switched to a winding rated for low voltage and high current, on which the current was raised as the rods grew [14]. A newer design rectifies the mains and drives the rods from a resonant transistor inverter above 50 kHz. At that frequency the skin effect moves most of the current toward the outside of the rod, which flattens the radial temperature profile [19]. A supply that regulates the wrong terminal shows on a trend as a current climbing on its own at a steady voltage, the runaway of Section 2.6.2.

The second manipulated variable, the trichlorosilane feed, enters through a vaporizer whose calibration error $f_v$ the model carries as a disturbance.

### Where the Controller Runs

Loops of this kind run in a distributed control system, or in programmable controllers acting as its local controllers, with cycle times from milliseconds to minutes [26]. The app executes both PI loops every 2.88 s against a physics step of 1.44 s, both far shorter than the shortest rod time constant of 7.9 s. In a control system the structure of this chapter becomes standard parts: a PID block whose gain is multiplied by the schedule, a calculation block that turns $V$, $I$ and $d$ into the center estimate, and the override as a trim on the setpoint ahead of the PID. The predictive controller of Section 2.5.4, which decides every 0.048 h, belongs to the supervisory level that sets the setpoints of the loops below it [26].

> **IN PRACTICE** **Operators Lower the Surface as the Rods Grow**
>
> A model study of a twelve-rod reactor found what this chapter finds: over a campaign the time constant of the temperature loop grows by an order of magnitude while its gain falls by more than two, enough to call for adaptive control [25]. Holding the surface constant needs ever more current, which the stresses in the rod and the melting of its core limit, so operators lower the surface temperature as the campaign goes on [25]. Because the pyrometers see single spots, the rod resistance itself can serve as the controlled variable, since it responds even to local temperature changes [25]. The override of Section 2.5.3 automates the operator's retreat.

### What Is Modernizing It

Three developments bear on this reactor. The first is adaptive control on electrical estimates: a 2015 design estimates the diameter and temperature from the rod resistance and tunes its loop with neuro-fuzzy methods [25]. The second is nonlinear model predictive control on a lumped model of the reactor, shown on a laboratory-scale Siemens reactor [27]. The shrinking-horizon controller of this chapter belongs to that family. The third changes the plant. A supply above 50 kHz moves the current outward by the skin effect and narrows the difference between center and surface [19], and polished or coated walls cut the radiant loss [8]. Either would let a rod grow further before its center reaches the limit of Section 2.4.3.

## 2.8 Limitations

The model is not a reactor design tool. A single rod pair with rod-to-wall radiation cannot predict the rod-to-rod uniformity that decides real yield, and the studies that do address it solve a radiation field over tens of rods [4, 12, 13]. The parabolic profile of (2.8) is the load-bearing approximation of the whole application, and it rests on uniform generation and constant conductivity, neither of which holds in hot silicon. The chemistry is a single overall stoichiometry with no dichlorosilane, no powder formation and no hydrochlorination recycle, so nothing in the application can speak to the conversion efficiency or the byproduct economics that dominate a real plant [6, 10]. The cold start is absent entirely, which removes the most awkward control problem the process has. The supply is two algebraic expressions and a first-order lag, with a current rating, a voltage rating and a power ceiling but no rectifier dynamics, no transformer tap and no minimum firing angle, so the transition from voltage regulation into the power ceiling and the current limit in the counterfactual run of Section 2.6.2 is instantaneous where a real supply would take a cycle or two. Disturbances arrive as commanded steps and never as the drifting behavior of a real plant. Repeatability is itself unrealistic, because a fixed integration step and a seeded generator reproduce one trajectory exactly while a distributed control system does no such thing. The three features that survived the simplification, the 785-fold change in resistance, the sign of (2.16) under the two regulation modes and the linear growth of the center-to-surface rise, survived because a controller has to be designed around each of them. Where the individual assumptions break is set out in Section 2.3.6. The predictive controller's model is the quasi-steady design point at the declared parameters. It leaves out the transients of the surface and composition loops, it knows nothing of a learner's conductivity, resistivity law or rod length, and its biases cannot absorb an error in the center estimate itself. It does not model the override above it, so its predicted batch time and center describe the plan rather than the trimmed run, and a plan that did model the override could at best reproduce the override's own run, which is what the plan without its center row does. Three linearization passes are not always enough. A decision taken 140 °C above its answer, at a setpoint of 1200 °C on a 150 mm rod, was accepted with 34 °C of slack and the next decision was exact, a case that no scenario reaches.

> ⚠ **OPEN CHALLENGE** **A Core That No Instrument Reaches**
>
> The temperature gradient inside the rods remains the main limit of rod deposition. Silicon conducts better as it heats, so the current crowds toward the hot center and heats it further, and the gradient stresses the rod and limits how large it can grow before its core melts [19]. The center is still estimated numerically, and the gradients inside the rods are not controlled [25]. Progress would be a center estimate that is checked against the plant during a batch, or a heating current that keeps the profile flat enough for the estimate to matter less.

## 2.9 Discussion

The energy number on the scorecard is not decoration. A complete batch under the scheduled loop with the override in service consumes 18 564 kW h for 213.5 kg of silicon, or 87.0 kW h/kg, and polysilicon production is a substantial term in the life-cycle energy of a photovoltaic module [11]. Both of the levers that improve it are constrained by something other than the loop. Running hotter raises the rate by 8.3 % for sixty degrees and costs 39 % more boron, and running to a larger diameter improves the energy per kilogram until the center constraint stops it, which is the trade the deposition-radius optimization literature studies directly [5, 13]. The application exposes the harvest diameter, the core limit and the thermal conductivity as adjustable parameters so that the trade can be computed rather than asserted. Moving $k_{Si}$ alone moves the diameter at which the batch has to stop, which is an unusually direct demonstration that a material property can be the binding constraint on a control problem. The override infers the center at the declared 22 W/(m K), so a different $k_{Si}$ is a model error it cannot see: at 18 its center reads up to 30 °C low and the true center spends 9.6 h up to 14 °C past the limit.

The model of this chapter also supports methods beyond the implemented loops. The predictive controller of Section 2.5.4 already treats the surface-temperature trajectory as the decision variable with the constraint of (2.8) enforced explicitly, and its result says where the remaining value lies. A constant margin is something a high-gain override already provides, so what a model could add is a margin sized to its own uncertainty, which needs an estimate that reports how wrong it may be. The estimator in (2.8) is the simplest member of a family, and moving horizon estimation on the same equations would reconstruct the loss multiplier and the pyrometer bias from the current and the composition, which is what would turn the cross-check of Section 2.6.3 into a diagnosis. A purity specification that is violated with every instrument nominal is the standing argument for soft sensing on the feed side rather than on the reactor side [17].

## 2.10 Conclusion

This chapter studied a Siemens polysilicon reactor as an interactive laboratory for plants that are non-stationary by construction. The simulator integrates nine states covering the rod diameter and mean temperature, two supply lags, the trichlorosilane and tetrachloride mole fractions, the deposited silicon, the incorporated boron and the consumed energy, with a semiconducting resistivity, a radiative and convective loss on the rod surface and a parabolic radial profile that places the hottest point on the axis. Growing an 8 mm seed to 150 mm takes the resistance from 50.9 to 0.0649 Ω, the current from 21.8 to 2646 A and the steady-state gain from 10.46 to 0.0828 °C/A, a factor of 126 with nothing disturbed. A gain schedule built from the same closed-form expressions cuts the batch integrated error from 760.4 to 43.1 °C h and holds the surface inside 2.24 °C over runs of 80.14 and 78.57 h, and it melts the rod at 149.3 mm because the constraint that ends a batch is at the center and not at the surface. Adding an override on that inferred center caps it at 1323 °C and harvests 213.5 kg at 87.0 kW h/kg. The same term $PB/\bar{T}^2$ that makes constant-voltage operation unstable by a factor between 2.03 and 2.50 against the loss slope is also the larger part of the loop-gain denominator, and holding the terminal voltage of a 60 mm rod near 346 V carries its center to the melting point in 6.50 min. A shrinking-horizon predictive controller of the surface setpoint takes 29.94 h against 29.60 h under the shipped override, and the 1.11 h it saves with the band narrowed to 5 °C is margin that a high-gain override recovers to within 0.03 h and that a 10 % error in the resistivity law spends to a melt.

A fuller model would add rod-to-rod view factors, a temperature-dependent conductivity, a cold-start sequence and a rectifier model for the supply. As it stands, the reactor gives practice in controlling a plant that is also the product. The Czochralski puller of Chapter 3, the next step of the wafer flow, also grows its product while it runs and also meets a limit that no instrument reports.

**KEY POINTS**

- The silicon rod is heater, deposition surface and product at once, so the plant changes while it runs.
- Growing from 8 to 150 mm takes the plant gain from 10.46 to 0.0828 °C/A, a factor of 126.
- Holding the current makes the rod self-regulating, while holding the voltage carries the center of a 60 mm rod to the melting point in 6.50 min.
- A schedule built from the closed-form gain cuts the batch IAE from 760.4 to 43.1 °C h but melts the rod, because the limit is at the unmeasured center.
- An override on the center inferred from the supply's meters holds it to 1323 °C and harvests 213.5 kg.

**Where else this applies.** An exact gain schedule handed over by the physics in Chapter 3, and a loop that holds its pyrometer reading while the product is wrong in Chapter 9.

## References

[1] J. S. Shamma and M. Athans, "Analysis of gain scheduled control for nonlinear plants," *IEEE Transactions on Automatic Control*, vol. 35, no. 8, pp. 898–907, 1990.

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

[3] D. J. Leith and W. E. Leithead, "Survey of gain-scheduling analysis and design," *International Journal of Control*, vol. 73, no. 11, pp. 1001–1025, 2000.

[4] G. del Coso, I. Tobías, C. Cañizo, and A. Luque, "Temperature homogeneity of polysilicon rods in a Siemens reactor," *Journal of Crystal Growth*, vol. 299, no. 1, pp. 165–170, 2007.

[5] A. Ramos, W. O. Filtvedt, D. Lindholm, P. A. Ramachandran, A. Rodríguez, and C. del Cañizo, "Deposition reactors for solar grade silicon: A comparative thermal analysis of a Siemens reactor and a fluidized bed reactor," *Journal of Crystal Growth*, vol. 431, pp. 1–9, 2015.

[6] B. Ceccaroli and O. Lohne, "Solar grade silicon feedstock," in *Handbook of Photovoltaic Science and Engineering*, 2nd ed., A. Luque and S. Hegedus, Eds. Chichester, UK: John Wiley & Sons, 2010, pp. 169–217.

[7] A. F. B. Braga, S. P. Moreira, P. R. Zampieri, J. M. G. Bacchin, and P. R. Mei, "New processes for the production of solar-grade polycrystalline silicon: A review," *Solar Energy Materials and Solar Cells*, vol. 92, no. 4, pp. 418–424, 2008.

[8] K. Shen, W. Jin, and J. Wang, "Numerical simulation of heat transfer process and heat loss analysis in Siemens CVD reduction furnaces," *Frontiers in Heat and Mass Transfer*, vol. 22, no. 5, pp. 1361–1379, 2024.

[9] G. Bye and B. Ceccaroli, "Solar grade silicon: Technology status and industrial trends," *Solar Energy Materials and Solar Cells*, vol. 130, pp. 634–646, 2014.

[10] W. C. Breneman and S. Julsrud, "From conventional polysilicon Siemens process to low-energy fluidized bed processes using silane," in *Solar Silicon Processes*, B. Ceccaroli, E. Ovrelid, and S. Pizzini, Eds. Boca Raton, FL: CRC Press, 2016, pp. 199–230.

[11] V. M. Fthenakis and H. C. Kim, "Photovoltaics: Life-cycle analyses," *Solar Energy*, vol. 85, no. 8, pp. 1609–1628, 2011.

[12] G. del Coso, C. del Cañizo, and A. Luque, "Radiative energy loss in a polysilicon CVD reactor," *Solar Energy Materials and Solar Cells*, vol. 95, no. 4, pp. 1042–1049, 2011.

[13] Z. Nie, P. A. Ramachandran, and Y. Hou, "Optimization of effective parameters on Siemens reactor to achieve potential maximum deposition radius: An energy consumption analysis and numerical simulation," *International Journal of Heat and Mass Transfer*, vol. 117, pp. 1083–1098, 2018.

[14] H. Schweickert, K. Reuschel, and H. Gutsche, "Production of high-purity semiconductor materials for electrical purposes," U.S. Patent 3,011,877, filed 11 June 1957 (priority Germany, 25 June 1956), 1961, granted 5 December 1961.

[15] H. Habuka, T. Nagoya, M. Mayusumi, M. Katayama, M. Shimada, and K. Okuyama, "Model on transport phenomena and epitaxial growth of silicon thin film in $SiHCl_3$–$H_2$ system under atmospheric pressure," *Journal of Crystal Growth*, vol. 169, no. 1, pp. 61–72, 1996.

[16] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. Steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

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

[18] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[19] K. A. Kozin, A. G. Goryunov, and F. Manenti, "Siemens-reactor's high-frequency power supply," *Chemical Engineering Transactions*, vol. 39, pp. 1651–1656, 2014.

[20] B. Srinivasan, S. Palanki, and D. Bonvin, "Dynamic optimization of batch processes," *Computers & Chemical Engineering*, vol. 27, no. 1, pp. 1–26, 2003.

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

[22] G. del Coso, C. del Cañizo, I. Tobías, and A. Luque, "Increase on Siemens reactor throughput by tailoring temperature profile of polysilicon rods," in *2007 Spanish Conference on Electron Devices*. Madrid, Spain: IEEE, 2007, pp. 25–28.

[23] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

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

[25] K. Kozin, A. Goryunov, F. Manenti, F. Rossi, and A. E. Stolpovskiy, "Development of an advanced control system for a chemical vapor deposition (CVD) reactor for polysilicon production," *Chemical Engineering Transactions*, vol. 43, pp. 1531–1536, 2015.

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

[27] L. Viganò, M. Vallerio, F. Manenti, N. M. N. Lima, L. Zuniga Linan, and G. Manenti, "Model predictive control of a CVD reactor for production of polysilicon rods," *Chemical Engineering Transactions*, vol. 21, pp. 523–528, 2010.

## 2.11 Exercises

**2.1** Concept check · 2 min

A 60 mm rod sits at its design point with the surface at 1100 °C and no controller acting. The supply is switched from current regulation to voltage regulation at the same terminal voltage. What does the model predict?

(a) The rod holds its temperature, because the operating point is an equilibrium in either mode

(b) The rod cools slowly, because a fixed voltage delivers less power than a fixed current

(c) The rod oscillates, because the supply lag and the rod time constant now interact

(d) The rod runs away, because the power then rises with temperature faster than the loss

**2.2** Concept check · 2 min

The schedule of Equation (2.18) is built from $2IR/S$, which overstates the measured plant gain $\partial T_s/\partial I$ by a factor of 3.449 at the seed and 3.593 at harvest. Why does the schedule still give the loop the same shape at every diameter?

(a) The factor is nearly constant, so the gain entered at the seed absorbs it once

(b) The factor is large, so the scheduled loop is detuned and robust at every diameter

(c) The integral action removes the factor, because a PI loop leaves no steady error

(d) The factor has no effect, because the multiplier is clamped between 0.05 and 400

**2.3** Concept check · 1 min

In scenario 10 the boron in the feed triples at $t = 2\,\mathrm{h}$. Which signal in the reactor responds?

(a) The pyrometer reading, because boron changes the emissivity of the rod surface

(b) None of them, because temperature, current, composition and growth do not depend on boron

(c) The rod current, because a doped deposit conducts better and lowers the rod resistance

(d) The gas composition, because boron displaces trichlorosilane in the feed

**2.4** Theory and calculation · 6 min

Use Equations (2.1) and (2.2) with $\rho_\infty = 2.0 \times 10^{-5}\,\Omega\,\mathrm{m}$, $A = 7.32 \times 10^{-8}\,\Omega\,\mathrm{m}$, $B = 12\,120\,\mathrm{K}$ and $L = 5.2\,\mathrm{m}$. With the surface held at 1100 °C, the mean rod temperature is 1108.43 °C on the 8 mm seed and 1258.03 °C at 150 mm.

(a) Find the resistivity and the resistance of the rod pair on the seed. *Checkpoint:* $4.925 \times 10^{-4}\,\Omega\,\mathrm{m}$ and $50.95\,\Omega$.

(b) The seed dissipates 24.23 kW. Find the current and the terminal voltage. *Checkpoint:* 21.81 A and 1111 V.

(c) Find the resistance at 150 mm, the ratio to the seed value, and how much of that ratio is area and how much is resistivity. *Checkpoint:* $0.0649\,\Omega$, a ratio of 785, of which 351.6 is area and 2.23 is resistivity.

Give each answer to within 0.5 %.

**2.5** Theory and calculation · 6 min

A 60 mm rod pair of length $L = 5.2\,\text{m}$ has its surface at 1100 °C, a mean temperature of 1163.2 °C and a dissipation of 181.7 kW. Take $\varepsilon = 0.70$, $\sigma = 5.670 \times 10^{-8}\,\text{W}/(\text{m}^2\,\text{K}^4)$, $h_g = 60\,\text{W}/(\text{m}^2\,\text{K})$, $f_{\text{loss}} = 1$, $B = 12\,120\,\text{K}$, $\rho_{Si} = 2330\,\text{kg}/\text{m}^3$ and $c_{p,Si} = 800\,\text{J}/(\text{kg}\,\text{K})$.

(a) Find the loss slope $S$ of Equation (2.7). *Checkpoint:* 461.7 W/K.

(b) Find $PB/\bar{T}^2$, its ratio to $S$, and say which regulation mode holds the operating point. *Checkpoint:* 1067 W/K, 2.31 times $S$. Current regulation.

(c) Find the rod thermal time constant $\rho_{Si}\pi d^2 L c_{p,Si}/(4S)$. *Checkpoint:* 59.4 s.

**2.6** Theory and calculation · 6 min

With the surface held at 1100 °C the dissipation is proportional to the diameter, 454.4 kW at 150 mm. Take $k_{Si} = 22\,\text{W}/(\text{m}\,\text{K})$ and $L = 5.2\,\text{m}$.

(a) Find the center-to-surface rise $\Delta T = P/(4\pi k_{Si} L)$ at 150 mm. *Checkpoint:* 316.1 °C.

(b) Find the diameters at which the center reaches the 1350 °C limit and the 1414 °C melting point. *Checkpoint:* 118.6 mm and 149.0 mm, each within 0.2 mm.

(c) The rod's true conductivity is 18 W/(m K), but the estimator of Equation (2.24) assumes 22. Find by how much the estimate of $T_c$ reads low per 100 kW of measured power, and the power range that the error of −21 to −30 °C in Table 2.7 implies. *Checkpoint:* 7.73 °C per 100 kW. About 270 to 390 kW.

**2.7** App challenge · 8 min

Load scenario 3, The same tuning at 140 mm, and change no setting. The setpoint steps from 1050 to 1070 °C at 0.5 h. Run with the default seed at the default speed of 0.5 h/s, pause within 0.05 h of $t = 4.0\,\text{h}$ and export the CSV file. Then do the same with scenario 4, Scheduling on a diameter the tool knows exactly.

(a) In the scenario 3 file, read the true surface, column `ts_C`, at $t = 3.0\,\text{h}$, and read the IAE tile of the scorecard. *Checkpoint:* 1053.2 °C, within 0.1 °C. IAE 65.0 °C h, within 1 °C h.

(b) In scenario 4, read the IAE and Overshoot tiles, and the schedule multiplier, column `sched_gain_x`, at $t = 1.0\,\text{h}$. *Checkpoint:* IAE 5.9 °C h, within 0.2 °C h. Overshoot 15.6 %. Multiplier 51.2, within 0.1.

(c) The true overshoot of scenario 4 is 0.40 °C, or 2.0 % of the step. Explain why the tile reads more.

**2.8** App challenge · 8 min

Load scenario 2, The loop closed, on a small rod. Before Run, set Gain scheduling on diameter to On. Run with the default seed at 10 h/s until the run stops itself, which takes about eight seconds.

(a) Read the status line, the Diameter tile and the IAE tile of the scorecard. *Checkpoint:* The rod melted, at 149.3 mm. IAE 43.1 °C h, within 0.2 °C h.

(b) Press Reset, set Core temperature override to On, and run again until the run stops. Read the Diameter, Silicon and Energy tiles, the IAE tile, and the largest value of column `t_core_C` in the exported file. *Checkpoint:* 150.0 mm, 213.5 kg, 87 kW h/kg, IAE 38.8 °C h, peak center 1322.7 °C within 0.5 °C.

(c) Compare both runs with the PI-only column of Table 2.5, and explain why the better loop of part (a) lost the rod.

**2.9** App challenge · 9 min

Load the graded challenge, Challenge · finish the batch without melting it. Before Run, set Gain scheduling on diameter to On, Core temperature override to On and Rod surface temperature to 1100 °C. Run with the default seed at 10 h/s until the run stops at harvest, and export the CSV file. Skip its comment lines, which start with #. Repeat with Override band below the core limit (in Advanced) at 40 °C and then at 120 °C, pressing Reset before each run.

(a) For each band, read the harvest time from the clock and, from the rows with `time_h` of 1 or more, the largest value of `t_core_C`. *Checkpoint:* Band 80: 29.68 h, 1348.8 °C. Band 40: 28.86 h, 1374.5 °C. Band 120: 30.62 h, 1323.1 °C. Times within 0.02 h, temperatures within 0.3 °C.

(b) Say which runs pass the challenge, and find the batch time each band trades for each degree of margin. *Checkpoint:* Bands 80 and 120 pass. About 0.032 h per degree from 40 to 80, and 0.036 h per degree from 80 to 120.

**2.10** Application area · 3 min

An engineer proposes raising the surface setpoint from 1100 to 1160 °C to shorten the batch. Use Equation (2.10) with $k_s = 3.60\,\mathrm{mm/h}$ at 1100 °C, 5.190 mm/h at 1160 °C and $k_m = 1.20\,\mathrm{mm/h}$, at $x_{\mathrm{TCS}} = x_{\mathrm{ref}}$. Use Equation (2.14) with $C_B = 0.60$ ppba, $k_{B0} = 0.35$ and $k_{B1} = 0.55$ per 100 °C.

(a) Find the diameter growth rate at each setpoint. *Checkpoint:* 1.800 and 1.949 mm/h, 8.3 % more.

(b) Find the boron content of the deposit at each, and give a recommendation. *Checkpoint:* 0.210 and 0.292 ppba, 39 % more, against a 0.30 ppba specification.

**2.11** Application area · 4 min

Table 2.3 lists 1111 V at 21.8 A on the seed and 172 V at 2646 A at harvest, where the rods draw 454.4 kW. A complete scheduled batch with the override consumes 18 564 kW h for 213.5 kg, and Table 2.2 prices electricity at $0.06 per kW h.

(a) Find the electricity cost of the batch and per kilogram. *Checkpoint:* $1114 per batch, $5.22 per kilogram.

(b) A supply that could deliver the seed voltage and the harvest current at once would be rated at their product. Find that rating, compare it with the largest power, and say what a supply designer does about it. *Checkpoint:* 2.94 MV A, 6.5 times the 454 kW.

**2.12** Review the agent · 5 min

An AI assistant reviews scenario 11 and writes: “With the override band narrowed to 5 °C, the MPC harvests in 28.48 h against 29.60 h for the shipped override alone. Its look-ahead saves 1.11 h on every batch. Narrow the band to 5 °C on the plant and run the MPC. The override is still there as a backstop, so the change adds no risk.” Find the flaw and say how you would check it in the app.