INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 14 Control in Semiconductor Manufacturing

CHAPTER 9

Rapid Thermal Processing Simulator for Ramp Feedforward and Pyrometer Emissivity Error

Pages 212–239

Edited by John D. Hedengren

How to cite this chapter. "Rapid Thermal Processing Simulator for Ramp Feedforward and Pyrometer Emissivity Error," in *Control in Semiconductor Manufacturing*, John D. Hedengren, Ed., vol. 14 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 212–239.

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

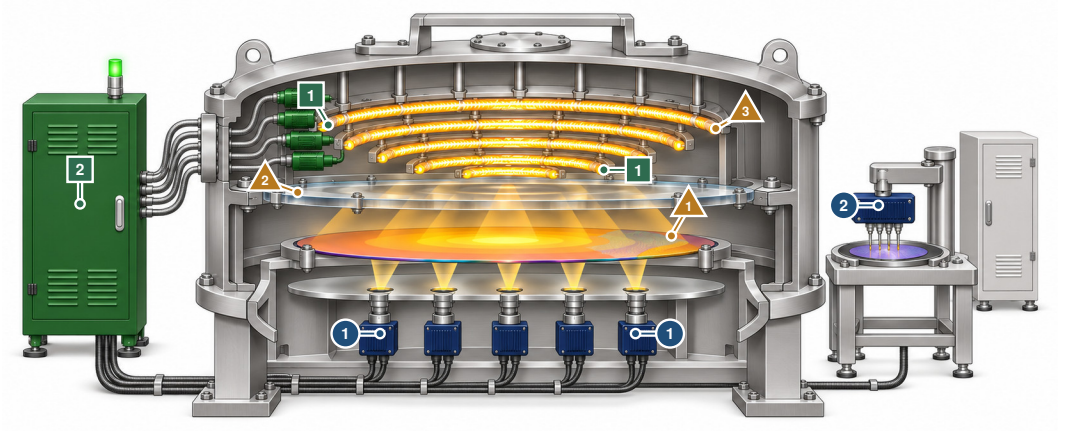
The app for this chapter
https://apmonitor.com/apps/rtp-wafer-anneal/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v14-semiconductor/#rtp-wafer-anneal

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

CHAPTER 9

Rapid Thermal Processing Simulator for Ramp Feedforward and Pyrometer Emissivity Error



The controller moves 1 the power of each of the five rings of lamps
2 the soak time, set from the power cabinet and corrected from wafer to wafer

The controller reads 1 five thermometers under the wafer that judge its heat by its glow
2 a probe that measures the finished wafer's sheet resistance

What disturbs it 1 a film on the wafer that makes it glow brighter than the thermometers assume
2 a quartz window that warms across a lot when its cooling water is turned down
3 lamps that age or fail, so a ring delivers less than it is told

Notice that the thermometers under the wafer read only its glow: a film like the patch at the right makes the wafer look hotter than it is, and only the probe station measures what the anneal did. Illustration, not to scale.

IN PLAIN TERMS Lamps That Bake a Wafer in Seconds

The main idea. A chip factory heats each thin silicon disk with rings of bright lamps, from warm to glowing hot and back in a few seconds. The heat wakes up the atoms planted in the silicon, and controllers steer the lamps from light sensors under the disk.

The challenge. The sensors judge heat by how brightly the disk glows. A thin coating can make a disk glow brighter without being hotter, so every reading looks perfect while the disk stays too cool. The lamp rings also light each other's parts of the disk.

Why it matters. Only a test of the finished disk shows whether the heat did its job.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 9.2, the fact boxes of Sections 9.4 and 9.6, the try-it boxes of Section 9.6, and the key points
Study	55 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 9.6, Section 9.7, Section 9.8 with its open challenge

The control question How can five lamp zones follow a spike anneal at 250 °C/s when the wafer's gain moves with the cube of its temperature, the zones light each other's pyrometers and the pyrometers may be converting radiance with the wrong emissivity? And what does a loop on its setpoint prove about the wafer it anneals?

Learning objectives

- Derive the radiative conductance and time constant of a wafer and explain why loops that ask for a heating rate need no gain schedule.
- Compute the power a ramp needs and the lag of a feedback loop without feedforward, and compare both with the inverse-model feedforward.
- Predict the temperature a loop holds when the emissivity is wrong, and name the measurements that reveal and correct the error.
- Relate a kelvin of peak temperature to the thermal budget, the sheet resistance and the soak correction between wafers.
- Compare a five-zone MPC with decoupled zone loops and a slip-aware ramp limit on a ramp too fast for the edge lamps.

Prerequisites Sections 1.6.2, 1.4.5, 1.4.2 and 1.4.3 of Chapter 1, and Section 1.6.3 for the MPC; Chapter 5 for multizone thermal control.

The app https://apmonitor.com/apps/rtp-wafer-anneal/

See also Volume 3, *Feedforward, Cascade, and Ratio Control*: Volume 3, Chapter 6 and Volume 3, Chapter 12 (a model feedforward that leaves feedback a trim). Volume 5, *Delays, Slow Measurements, and State Estimation*: Volume 5, Chapter 11 (a product temperature no instrument reaches). Volume 6, *Multivariable and Plantwide Control*: Volume 6, Chapter 2 (interacting loops and a static feedforward) and Volume 6, Chapter 3 (a sheet property measured after the step).



9.1 Introduction

An implanted wafer does not yet hold working transistors. The dopant atoms that will form each source and drain sit in silicon that the implant has damaged, and they become electrically active only when the wafer is heated, while the same heat lets them diffuse and deepens the junction. Rapid thermal processing (RTP), which heats one wafer at a time with lamps, offered an alternative to the furnace tube for this step [1, 2], and the anneal was later cut down to a spike, a ramp up and straight down, whose ramp-up rate is a process variable of its own [3]. The implant itself is the subject of Chapter 8. The chamber of this chapter takes a 300 mm wafer 775 µm thick from 600 to 1050 °C at 250 °C/s, holds it for 1.0 s and lets it cool. Near the top of that ramp the heat capacity alone asks for 31.5 kW, and the five lamp zones draw 106.8 kW of electricity to put 47.5 kW into the wafer.

Two properties of this plant are simple to state and hard to feel. The first is that its loss is radiative. The slope of the loss, the conductance that sets how quickly the wafer answers its lamps, rises from 8.2 W/K at 400 °C to 53.1 W/K at 1050 °C while the heat capacity rises only from 111.9 to 126.0 J/K, and the open-loop time constant drops from 13.7 to 2.4 s as a result. The second is that the temperature a loop regulates is a pyrometer's reading, computed from radiance with an emissivity that somebody supplied. The emissivity of silicon changes with temperature [4], and film stacks and patterns change it again [5]. In this model a nitride cap that the pyrometers take for bare silicon leaves them peaking at 1048.4 °C while the true center peaks at 1017.0 °C and the sheet resistance comes out 26.3 % high. The belief the simulator is built to dislodge is that a loop holding its setpoint proves the product was made. A second belief falls with it, that a gain moving with the cube of temperature must be scheduled, because loops that ask for a heating rate overshoot a 10 K step by 3.5 K at 600 °C and by 3.3 K at 1050 °C with no schedule at all.

None of this is peculiar to one chamber. RTP temperature control is a distributed thermal problem with a stringent uniformity, fast ramps and a thermal budget that shrinks as devices shrink, and the physics-based design route worked out for it applies equally to slower hot-wall batch furnaces [6].

Model-based control has been pursued since low-order models of the chamber's dynamics were identified [7], because it offers the greatest potential for tight control of the wafer temperature [8, 9]. Microelectronics manufacturing adds a product that is measured only after the step, so a slower correction between runs sits above the fast loops [10], and a run-by-run controller combining statistical process control with feedback was demonstrated on a silicon epitaxy reactor [11]. The chamber brings together five pieces that recur across this class of thermal processes, a gain that moves with temperature, a trajectory faster than feedback, heaters that light each other's targets, a non-contact sensor whose error depends on the material, and a quality read afterwards. A transfer function identified at one temperature carries none of them.

However, the published work on RTP control assumes a chamber in hand, and the standard course treatment linearizes a nonlinear plant about one operating point [12], so a student meets neither a loop that holds its reading while the wafer misses its budget nor the radiation that makes the ramp hard. This chapter works with a spike anneal chamber that runs in a browser, in which the true annulus temperatures are logged apart from the five pyrometers the controller reads, the product is an Arrhenius budget integrated along each wafer's own trajectory, and the lamps, the edge ring and the window each carry their own dynamics. The radiative model, the inverse-model feedforward, the decoupler, the correction between wafers and the model predictive controller all follow published practice. The model has 29 states with closed energy ledgers, a pyrometer chain that stray light and a wrong emissivity can fool, a supervisor that refuses a lot it cannot make and a five-zone model predictive controller measured against the loops it replaces, all written out in enough detail to rebuild the application. Nine cases and a graded lot pull apart a time constant and a gain, a ramp and its error, one zone and its neighbors, and a reading and the wafer.

Section 9.2 traces lamp heating, pyrometry and model-based control of the chamber. Section 9.3 sets out the chamber as modeled, from the annuli of the wafer to the thermal budget and the slip limit. Section 9.4 derives what radiation does to the loop: a time constant that falls with temperature, the power a ramp needs, zones that light each other and an error the loop cannot see. Section 9.5 describes the zone loops, the inverse-model feedforward, the supervisor, the emissivity modes, the soak correction and the model predictive controller. Section 9.6 reports what each case measures, the graded lot and the comparison of the predictive controller with the zone loops. Section 9.7 describes the real thermometers, lamps and controllers behind the loops, and Sections 9.8 and 9.9 give the limitations and the conclusions.

9.2 Lamps, Pyrometers and the Spike Anneal

Lamp heating entered the fab as an alternative to the furnace [1, 2], as the box below tells.

> ⧖ HISTORY **From the Furnace Tube to the Lamp (1988 to 1990)**
>
> Wafers were once heated in batches, in furnace tubes. In 1988 a review described rapid isothermal processing, heating one wafer with incoherent light, as a superior alternative to furnace processing. Its attractions were low cost, a minimum overall thermal budget, low power consumption and high throughput, and the review named more accurate temperature measurement as a development still needed [1]. Two years later a review written with its emphasis on temperature control summarized the options of 15 equipment makers, with the techniques of the day for measuring and controlling the wafer temperature and new optical, fiber-optic and photoacoustic ones [2]. In 2006 a dissertation still called the difficulty of measuring the wafer's temperature the biggest obstacle to the technology, which had kept its advantage over batch processing [13].

The rim was an early difficulty. A wafer lit uniformly loses more heat at its edge than in its middle, and a calculation of the temperature and stress fields in 1984 showed that the thermoelastic stress from that edge drop is much higher under uniform irradiation than under a strip heater. The same study resolved the stress onto the slip planes of the crystal and found the onset of damage on 2-inch wafers held for 15 s between 1030 and 1090 °C [14]. Heat transfer models with radiation and

convection were soon compared against the yield stress of silicon at the temperature and strain rate of the process, together with methods for improving the uniformity [15]. A lamp head divided into radial zones, with an outer zone that lights both the wafer's rim and the ring that supports it, is the arrangement the simulator carries.

Measurement was the second difficulty, because a pyrometer reads radiance and not temperature. The total hemispherical emissivity of lightly doped silicon 390 µm thick was measured to rise from 0.12 at 280 °C to a limiting value of 0.7 at 650 °C [4], and electromagnetic multilayer calculations coupled to a reactor model showed that film patterns change a wafer's radiative properties enough to matter for the temperature uniformity of shallow-junction anneals [5]. An emissivity assumed for one film stack can therefore be wrong for the next, and the error enters the loop through its only measurement.

Control followed the measurement into models [6–9, 16], as the box at the end of this section tells. Between runs, the run-by-run controller joined statistical process control to feedback and recovered an epitaxy process within 3 runs of a disturbance [11], and the practice of automatic control across microelectronics manufacturing was surveyed at length [10]. Spike annealing for ultra-shallow junctions was studied with the ramp-up rate as a variable of its own [3], and it is the recipe the simulator runs.

> ⧖ HISTORY **From Lamp Recipes to Model-Based Control (1993 to 2005)**
>
> In 1993 a dynamic model of a circularly symmetric chamber with several lamp zones, carrying radiation, conduction, convection and the heating of the chamber itself, was paired with a multivariable nonlinear model predictive controller. The study asked how the choice of the objective function changed the control of such a distributed system [16]. In 1997 an experimental chamber ran on a state-space model identified from data, a temperature trajectory shaped by convex optimization so that high-gain feedback could respect the lamps' limits, and linear quadratic Gaussian feedback [8]. By 2003 physics-based models, reduced to low order, ran as real-time embedded controllers on commercial chambers [6]. Yet in 2005 many tools still held one pyrometer at the wafer's center with one loop, the zone powers fixed as shares stored in the recipe [17].

9.3 The Chamber as Modeled

The simulated tool anneals a lot of 25 wafers, one every 60 s. Each slot loads a wafer at 25 °C with the lamps idle, preheats it open loop until all five pyrometers read, stabilizes it at 600 °C, ramps it at 250 °C/s to 1050 °C, soaks it for 1.0 s, lets it cool to 500 °C and unloads it. Figure 9.1 draws a half section of the chamber with the symbols of the equations below, and Table 9.1 classifies each variable as a controlled, manipulated, disturbance or state variable.

The 29 states are the eight annulus temperatures, the edge ring, the five delivered lamp powers, the window, eight thermal budgets and six energy ledgers of the wafer's flows and the lamps' electricity. Temperatures are in degrees Celsius on the panels and in kelvin inside every radiative and Arrhenius term.

9.3.1 The Wafer

The wafer is eight annuli of equal width $\Delta r = 18.75\,\text{mm}$ out to $R_w = 150\,\text{mm}$, each with one temperature. Annulus j balances the lamp power it absorbs against radiation and convection from both faces and conduction to its neighbors, as shown in Equation (9.1).

$$\begin{aligned} m_j c_p(T_j)\frac{\mathrm{d}T_j}{\mathrm{d}t} = &\sum_{k=1}^{5} G_{jk}P_k - \varepsilon(T_j)\,\sigma\,2A_j\,(T_j^4 - T_{ch}^4) \\ &- h\,2A_j\,(T_j - T_{gas}) + K_{j-1,j}(T_{j-1} - T_j) \\ &- K_{j,j+1}(T_j - T_{j+1}) + \delta_{j7}\,q_{rim} \end{aligned} \tag{9.1}$$

where $A_j = \pi(r_{j+1}^2 - r_j^2)$ with $r_j = j\Delta r$, $m_j = \rho\,t\,A_j$ with $\rho = 2329\,\text{kg/m}^3$ and $t = 775\,\mu\text{m}$, σ is the Stefan–Boltzmann constant, $T_{gas} = 25\,^\circ\text{C}$, $h = h_0 f_g^{0.6}$ with $h_0 = 8\,\text{W/(m}^2\,\text{K)}$ and f_g the gas flow as a

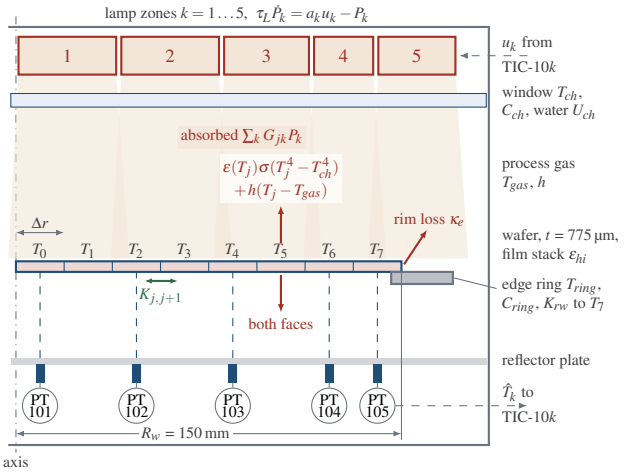


Figure 9.1 Half section of the lamp-heated chamber, drawn to scale in radius. Five lamp zones deliver P_k through a quartz window at T_{ch}, the eight annuli at T_0 to T_7 absorb $\sum_k G_{jk} P_k$, radiate and convect from both faces and conduct through $K_{j,j+1}$, and the rim loses heat through κ_e and exchanges it with the SiC edge ring at T_{ring}. Pyrometers PT-101 to PT-105 look up through the reflector plate at radii r_k and report $\hat{T}_k$ to the zone loops, which return the lamp commands u_k.

Table 9.1 Variables of the Chamber by Role, With Nominal Values of the Reference Configuration

Variable	Role	Nominal	Range
$\hat{T}_k, k = 1 \ldots 5$	CV, PT-101 to 105	recipe SP	350–1250 °C
T_{pre}, T_{peak}	recipe setpoints	600, 1050 °C	450–800, 950–1150
r, t_{soak}	recipe rate, soak	250 °C/s, 1.0 s	50–400, 0–4
R_s	quality, after the wafer	400 Ω/sq	±2 % band
ΔT	constraint, true	$< \Delta T_{slip}$	—
u_k	MVs, lamp commands	feedforward	0–100 % of rating
δ	MV, soak correction	0 s	−0.9 to 2.0 s
ε_{hi}	DV, film stack	0.70	0.63, 0.70, 0.77
a_3, a_5	DVs, lamp output	100 %	30–100
f_w, f_g	DVs, water, gas flow	100 %	10–150, 20–200
κ_e	DV, rim loss factor	1	0–3
σ_p	DV, pyrometer noise	0.3 K	0–3
T_j, $j = 0 \ldots 7$	states, annuli	25 °C at load	—
T_{ring}, T_{ch}	states, ring, window	281.5, 30 °C at reset	—
P_k	states, delivered power	1 % idle	—
θ_j	states, budgets	0 at load	—
$E_{abs}, \ldots$	states, six ledgers	0 at load	—

fraction of design, and the conductance between annuli is $K_{j,j+1} = k(\bar{T})\, 2\pi r_{j+1} t / \Delta r$ at the mean of the two temperatures. The terms at the rim, q_{rim}, follow in Equation (9.3). Equation (9.2) supplies the properties of silicon.

$$\begin{aligned} c_p(T) &= \left(a_0 + a_1\tau + a_2\tau^2 + a_3\tau^3 + a_4\tau^{-2}\right)/M_{Si} \\ k(T) &= \left(-4.55\times10^{-3} + 3.74\times10^{-5}\,T\right)^{-1} \\ \varepsilon(T) &= \varepsilon_{hi} - \frac{\varepsilon_{hi} - 0.30}{1 + \exp[(T - 430\,°\mathrm{C})/55\,\mathrm{K}]} \end{aligned} \tag{9.2}$$

The heat capacity is the NIST-JANAF Shomate fit for crystalline silicon in $\tau = T/1000\,\mathrm{K}$ with $a_0 = 22.81719$, $a_1 = 3.899510$, $a_2 = -0.082885$, $a_3 = 0.042111$ and $a_4 = -0.354063$ in J/(mol K) and $M_{Si} = 0.028\,085\,5\,\mathrm{kg/mol}$, which gives the whole wafer 90.8 J/K at 25 °C, 116.9 J/K at 600 °C

and 126.0 J/K at 1050 °C. The conductivity follows the phonon conductivity of silicon, 150 W/(m K) at 300 K and 24.8 W/(m K) at 1200 K, with T in kelvin. The total emissivity rises from 0.30 when cold, because lightly doped silicon is partly transparent in the infrared, to the film stack's high-temperature value ε_{hi} above about 600 °C, a shape taken from the measured rise of silicon's emissivity [4]. The three stacks are teaching values, 0.70 for bare silicon, 0.63 under a thermal oxide and 0.77 under a nitride cap, and each wafer draws its own value with a scatter of 0.1 %.

9.3.2 The Rim and the Edge Ring

The outer annulus carries two more terms. It loses heat through its edge face and through its weaker view of the reflector, lumped as an equivalent band $w_E = 3\,\text{mm}$ wide, and it exchanges radiation with a silicon carbide ring that supports it, as shown in Equation (9.3).

$$\begin{aligned} q_{rim} &= K_{rw}\,\sigma\,(T_{ring}^4 - T_7^4) \\ &\quad - \kappa_e\,\varepsilon(T_7)\,\sigma A_E\,(T_7^4 - T_{ch}^4) \\ C_{ring}\frac{\mathrm{d}T_{ring}}{\mathrm{d}t} &= \sum_k G_{rk}P_k - \varepsilon_r\sigma A_r(T_{ring}^4 - T_{ch}^4) \\ &\quad - K_{rw}\,\sigma\,(T_{ring}^4 - T_7^4) \end{aligned} \tag{9.3}$$

where $A_E = 2\pi R_w w_E = 2.83\times10^{-3}\,\text{m}^2$, κ_e is the rim loss factor, $K_{rw} = 2.8\times10^{-3}\,\text{m}^2$ is the exchange area between the ring and the rim, and the ring runs from 146 to 170 mm at a thickness of 0.7 mm with a density of 3210 kg/m^3 and a heat capacity of 1100 J/(kg K), so $C_{ring} = 58.9\,\text{J/K}$. It radiates from $A_r = 0.0439\,\text{m}^2$ with $\varepsilon_r = 0.85$ and absorbs 85 % of the lamp light that lands on it.

9.3.3 Lamps and the Gain Matrix

Each zone follows its command through a first-order filament lag, and each lights a blurred annulus of the wafer plane, as shown in Equation (9.4).

$$\begin{aligned} \tau_L\frac{\mathrm{d}P_k}{\mathrm{d}t} &= a_k u_k - P_k \\ G_{jk} &= \alpha_w\,\eta\,\frac{\int_{r_j}^{r_{j+1}}\phi_k(r)\,2\pi r\,dr}{\int_0^{R_p}\phi_k(r)\,2\pi r\,dr} \\ \phi_k(r) &= \tfrac{1}{2}\big[F_k(r) + F_k(-r)\big] \\ F_k(x) &= \operatorname{erf}\big[(x - b_{k-1})/s\big] - \operatorname{erf}\big[(x - b_k)/s\big] \end{aligned} \tag{9.4}$$

where the command u_k is in watts, a_k is the zone's output as a fraction of new, $\tau_L = 0.20\,\text{s}$, the footprints end at b_k = 40, 80, 115, 140 and 172 mm with $b_0 = 0$, the blur is $s = 19\,\text{mm}$, the mirrored term $F_k(-r)$ makes zone 1 a disc, $\eta = 0.85$ of a zone's power reaches the plane out to $R_p = 185\,\text{mm}$, and the wafer absorbs $\alpha_w = 0.70$ of what lands on it. The ring's row G_{rk} is the same integral from R_w to 170 mm with absorptivity 0.85. The ratings are 9, 22, 32, 30 and 52 kW, 145 kW in all, and an idle filament holds 1 % of each rating.

9.3.4 The Window and the Pyrometers

The quartz window is one lumped temperature. It absorbs a share of the lamp power and of the radiation leaving the wafer and the ring, and water cools it, which Equation (9.5) expresses.

$$\begin{aligned} C_{ch}\frac{\mathrm{d}T_{ch}}{\mathrm{d}t} &= \beta_L\sum_k P_k + \beta_W(q_{rad} + q_{edge} + q_{ring}) \\ &\quad - U_{ch}f_w(T_{ch} - T_{cool}) \end{aligned} \tag{9.5}$$

where $C_{ch} = 6\,\text{kJ/K}$, $\beta_L = 0.06$, $\beta_W = 0.10$, q_{rad}, q_{edge} and q_{ring} are the wafer's face radiation, its rim loss and the ring's radiation, $U_{ch} = 12\,\text{W/K}$, f_w is the water flow as a fraction of design with a

floor of 0.05, and $T_{cool} = 25\,^{\circ}\mathrm{C}$. Five pyrometers look at the back of the wafer at r_k = 9.4, 46.9, 84.4, 121.9 and 140.6 mm, the centers of annuli 0, 2, 4, 6 and 7, so annuli 1, 3 and 5 have no pyrometer of their own. Equation (9.6) gives what each one measures and how it converts the radiance.

$$\begin{aligned} \frac{R_k}{\sigma} &= \varepsilon(T)\,T^4 + f_s\,(T_{ch}^4 - T_{cal}^4), \qquad T_{bb,k} = (R_k/\sigma)^{1/4} \\ \hat{T}_k &= \left(T_{bb,k}^4/\varepsilon_{set}\right)^{1/4} + \sigma_p\, n_k \approx T\,(\varepsilon/\varepsilon_{set})^{1/4} \end{aligned} \tag{9.6}$$

where T is the true temperature at r_k, $f_s = 1$ scales an illustrative stray radiance from the window beyond its value at the 30 °C calibration, so the approximation holds while the window is at that temperature, $T_{bb,k}$ passes a first-order lag of 5 ms, n_k is unit colored noise with a 4 ms correlation time, and ε_{set} comes from the emissivity mode of Section 9.5. With the film-stack table the conversion solves $\varepsilon(\hat{T})\hat{T}^4 = T_{bb,k}^4$ with the table's own temperature dependence. A reading outside 350 to 1250 °C is flagged out of range and clamped.

9.3.5 Thermal Budget, Sheet Resistance and Slip

Dopant activation is represented by a thermal budget per annulus, counted in equivalent seconds at $T_n = 1050\,^{\circ}\mathrm{C}$, and a sheet resistance that falls as a power of it, as shown in Equation (9.7).

$$\frac{\mathrm{d}\theta_j}{\mathrm{d}t} = \exp\Big[-\frac{E_a}{k_B}\Big(\frac{1}{T_j} - \frac{1}{T_n}\Big)\Big], \qquad R_{s,j} = R_{s,T}\Big(\frac{\theta_T}{\theta_j}\Big)^m \tag{9.7}$$

where $E_a = 3.5\,\mathrm{eV}$, k_B is the Boltzmann constant, $R_{s,T} = 400\,\Omega/\mathrm{sq}$, $m = 0.30$, and $\theta_T = 1.3365\,\mathrm{s}$ is the budget of the nominal spike followed exactly by a uniform wafer, so that the nominal recipe makes the target by construction. The wafer's sheet resistance is the area mean of $R_{s,j}$, its range is taken at the eight annulus centers, and the four-point probe reports the mean with a repeatability of 0.3 %. A wafer is in specification when its true mean lies within 2 % of the target and it has not slipped. Slip is judged on the true temperatures at the eight annulus centers, as shown in Equation (9.8).

$$\begin{aligned} \Delta T &= \max_i T_i - \min_i T_i \;>\; \Delta T_{slip}(T_{max}) \\ \Delta T_{slip}(T) &= 12\,\mathrm{K}\cdot\exp\left[(1000\,^{\circ}\mathrm{C} - T)/110\,\mathrm{K}\right] \end{aligned} \tag{9.8}$$

held for 20 ms, where T_{max} is the hottest of the eight. The limit is 29.8 K at 900 °C, 12 K at 1000 °C and 7.6 K at 1050 °C. The budget, the power law and the slip limit are illustrative teaching values that rank thermal histories. They are not models of activation, diffusion or dislocation, and real slip limits come from comparing the thermal stress resolved on the slip planes with the yield stress of silicon at the temperature and strain rate of the process [14, 15].

9.3.6 Constants, Initial State and Numerics

Table 9.2 collects the constants. A visible constant can be moved with its slider in the Show Model panel (Section 1.6.4), a hidden one is not offered there, and a fixed one is written into the source. A lot opens with every annulus at 25 °C, the window at 30 °C, the lamps idle and the ring at its idle equilibrium of 281.5 °C, and each new wafer arrives at 25 °C with its budgets at zero. The held wafer of case 1 opens at the steady state of its lamp stations, solved by Newton's method on the nine wafer and ring balances with the window iterated to its own balance.

All 29 states are integrated together with a fourth-order Runge-Kutta scheme on a fixed 1 ms step, the controller runs on a 10 ms scan, and 20 rows a second reach the historian. The fastest state is the filament, with $\Delta t/\tau_L = 0.005$. Because the ledgers are integrated with the states, the wafer's stored enthalpy and its integrated flows agree over a whole wafer to 2.7×10^{-13} of the energy crossing its faces, and over that wafer it absorbs 213.4 kJ, radiates 145.3 kJ, convects 29.6 kJ and loses 2.9 kJ at the rim. The application's own checks set the limits for refinement. Sixteen annuli instead of eight move the first wafer of case 3, with the decoupler and the feedforward, by 0.16 % of sheet

Table 9.2 Constants of the Model and How the Interface Exposes Each One

Symbol	Quantity	Value	Access
R_w, t	wafer radius, thickness	150 mm, 775 µm	fixed
Δr, ρ	annulus width, density	18.75 mm, 2329 kg/m^3	fixed
ε_{hi}	bare, oxide, nitride	0.70, 0.63, 0.77	fixed
—	emissivity scatter, 1σ	0.1 %	visible
h_0	convection at design flow	8 W/(m^2 K)	visible
w_E, T_{gas}	rim band, gas	3 mm, 25 °C	fixed
τ_L	filament lag	0.20 s	visible
$P_{max,k}$	zone ratings	9, 22, 32, 30, 52 kW	visible
—	idle filament power	1 %	hidden
s, η, α_w	blur, reach, absorption	19 mm, 0.85, 0.70	fixed
C_{ring}, ε_r	edge ring	58.9 J/K, 0.85	fixed
K_{rw}	ring to rim exchange	2.8×10^{-3} m^2	fixed
C_{ch}	window heat capacity	6 kJ/K	hidden
β_L, β_W	window absorption	0.06, 0.10	hidden
U_{ch}	window cooling	12 W/K	visible
f_s	stray radiance	1.0	visible
—	pyrometer lag, range top	5 ms, 1250 °C	hidden
—	pyrometer range bottom	350 °C	visible
E_a, m	budget, power law	3.5 eV, 0.30	visible
$R_{s,T}$	target and band	400 Ω/sq, ±2 %	visible
—	probe repeatability	0.3 %	hidden
—	slip limit at 1000 °C	12 K	visible
—	e-folding, dwell	110 K, 20 ms	hidden

resistance, 0.02 K of center peak and 0.19 K at the edge pyrometer's radius, and halving the step moves it by 4×10^{-13} of sheet resistance and 6×10^{-5} K of center peak, the finer sampling of the peak. The engine, its seed and its speeds work as in Section 1.6.1. Here the seed is 20260927, so a lot replays identically, and any speed between a quarter of real time and twenty times real time leaves the trajectory unchanged.

9.3.7 Assumptions

Eight annuli do not resolve the rim's own dip. The lateral diffusion length of silicon at 1050 °C, $\sqrt{kt/(8\varepsilon\sigma T^3)}$, is about 4.8 mm, a quarter of an annulus, which is why slip and uniformity are judged at the annulus centers out to 141 mm rather than at the rim itself.

Radiation is gray and band averaged. One total emissivity serves both faces, both faces see the window's temperature, and the reflector plate adds no cavity enhancement. The conversion $T(\varepsilon/\varepsilon_{set})^{1/4}$ follows from that assumption, and it overstates the error of a real narrow-band instrument, as Section 9.4 quantifies.

The wafer absorbs a fixed 70 % of the lamp light at every temperature and for every stack. A cold, lightly doped wafer is partly transparent beyond 1.1 µm and absorbs less, which here would only slow the open-loop preheat. The footprints, the gain matrix, the ratings and the filament lag are teaching values, and the delivered power is proportional to the command.

The ring and the window are single lumped temperatures, and the stray radiance of the window is an illustrative term. The chamber is axisymmetric and the wafer does not rotate, so azimuthal nonuniformity, wafer bow, the gas flow pattern and pattern-dependent absorption are all absent, although patterns are known to matter [5].

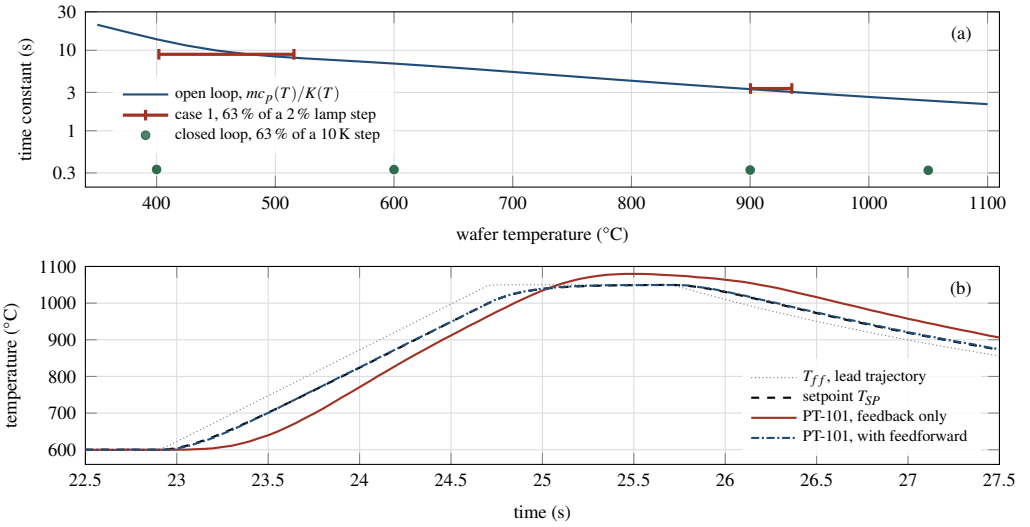


Figure 9.2 (a) Open-loop time constant of a uniform wafer from Equation (9.9) (solid line), with the 63 % times of the two lamp steps of case 1 drawn as bars across the temperatures each step spans and the 63 % times of a 10 K setpoint step under the zone loops as dots (pyrometer noise off). (b) First ramp of case 2 with one loop on the center pyrometer, as shipped (solid) and with the model feedforward (dash-dotted), against the setpoint (dashed) and the feedforward's lead trajectory (dotted).

9.4 What Radiation Does to the Loop

9.4.1 A Time Constant That Falls With Temperature

Linearizing Equation (9.1) for a uniform wafer yields the conductance $K(T)$, the slope of the loss with temperature, and with it the open-loop time constant of Equation (9.9).

$$\begin{aligned} K(T) &= \frac{\mathrm{d}}{\mathrm{d}T}\left[\varepsilon(T)\sigma 2A_w(T^4 - T_{ch}^4) + 2hA_w(T - T_{gas})\right] \\ \tau(T) &= m_w c_p(T)/K(T) \end{aligned} \tag{9.9}$$

where A_w and m_w are the area and mass of the whole wafer. With the chamber at 30 °C, K is 8.2, 17.1, 37.4 and 53.1 W/K at 400, 600, 900 and 1050 °C, and τ is 13.7, 6.8, 3.3 and 2.4 s, as shown in Figure 9.2. Above 800 °C the conductance is $4\varepsilon\sigma 2A_w T^3$ plus 1.13 W/K of convection. Below it the rising emissivity adds a slope of its own, so at 400 °C, where $\varepsilon = 0.447$, the fourth-power term alone gives 4.4 W/K and the emissivity's slope and the convection supply the rest. The heat capacity barely moves by comparison, from 111.9 to 126.0 J/K between 400 and 1050 °C, a factor of 1.13 against 6.5 for the conductance.

A plant gain that moves by 6.5 is the textbook argument for a gain schedule, the remedy Chapter 2 applies to a plant that moves across a whole run. The zone loops of Section 9.5 blunt it, because they ask for a heating rate and carry the heat capacity inside their own output, so their proportional action of 4 K/s per kelvin of error competes with a radiative rate $K/m_w c_p$ of only 0.15 to 0.42 s^{-1} between 600 and 1050 °C. Without pyrometer noise, the 63 % time of a 10 K setpoint step on the held wafer of case 1 in automatic is 0.33 s at every one of 400, 600, 900 and 1050 °C, as the dots in Figure 9.2(a) show.

🕮 **Pause.** Try Exercise 9.1 before reading on.

9.4.2 The Power a Ramp Needs

At 250 °C/s the heat capacity of 126.0 J/K asks for 31.5 kW near the peak, against the 17.15 kW that a uniform wafer at 1050 °C radiates into a 30 °C chamber and the 1.16 kW it convects. A feedback loop has to build that power from its own error, on a ramp that lasts 1.8 s. Less than half of the lamp power reaches the wafer. Zones 1 to 3 put 59.5 % of their power into the wafer, zone 4 54.4 % with 6.0 % into the ring, and zone 5 only 19.9 % into the wafer and 34.2 % into the ring. On the nominal ramp of the reference configuration the bank peaks at 106.8 kW electric for 47.5 kW absorbed, with the zones at 73 to 81 % of their ratings.

9.4.3 Zones That Light Each Other

Read at the five pyrometer spots, the gain matrix becomes a 5×5 matrix Φ of absorbed flux per watt, and it is far from diagonal. Section 1.4.5 treats the coupled zones of a heated plate in general, and Chapter 5 does so for a plate of a hundred zones. Per watt, the edge pyrometer receives 1.30 times more light from zone 4 than from its own zone 5, and the second pyrometer receives 1.19 times more from zone 1 than from zone 2. The ring, with a heat capacity of 58.9 J/K and nearly all of its light from zone 5, lags on a fast ramp, and the rim, radiating into it, lags with it. In case 3 with the decoupler and the feedforward, the ring is at 842 °C when the soak begins while the wafer averages 1001 °C, and still only at 914 °C when the soak ends.

9.4.4 An Error the Loop Cannot See

Equation (9.6) makes the pyrometer error multiplicative, as Equation (1.1) of Chapter 1 states in general. An emissivity 10 % higher than the pyrometer assumes makes it read $1.1^{1/4} = 1.024$ times the absolute temperature, so a loop holding the reading on 1050 °C holds a nitride-capped wafer at 1018.8 °C, 31.2 K lower. The same arithmetic runs the other way for an oxide stack at 0.63, whose wafer would be held near 1085.3 °C. A real pyrometer working in a narrow band at 0.95 µm follows Wien's law, $\Delta T/T = (\lambda T/c_2)\,\Delta\varepsilon/\varepsilon$ with $c_2 = 14\,388\,\mu\mathrm{m\,K}$, and makes an error of 0.87 %, or 11.6 K, about a third of the band-averaged model's, so the simulator overstates the effect on purpose. The window's own radiance adds a second, additive error once the chamber warms beyond its 30 °C calibration.

> **DID YOU KNOW? Two Degrees at a Thousand**
>
> In 2001 the roadmap of the chip industry asked the thermometers of rapid thermal tools for an accuracy better than 2 °C and a reproducibility better than 0.25 °C at 1000 °C [18]. In 2006 the best commercial systems near 1000 °C were estimated to be within 10 to 15 °C [13].

9.4.5 What One Kelvin Costs the Product

Near the peak the budget of Equation (9.7) changes by 2.3 % per kelvin and the sheet resistance by 0.70 % per kelvin, so the 2 % band on 400 Ω/sq is about 3 K of peak. Along the nominal setpoint, the 8 s at 600 °C contribute 1.1×10^{-6} equivalent seconds of the 1.3365 s budget, while 81.1 % of it accrues within 30 K of the peak, 61.1 % in the soak and 35.6 % in the cooling.

9.5 Implemented Controller

The controller reads only the five pyrometers and, on a calibration wafer, five thermocouples, besides the film stack that the lot's records declare. In each 1 ms step the order is the controller when its scan of $\Delta t_c = 10\,\mathrm{ms}$ is due, the integration, the slip test and wafer exchange, the outputs and the historian. The signal path of one zone is drawn in Figure 9.3.

9.5.1 Zone Loops That Ask for a Heating Rate

Each zone runs its own copy of the velocity-form PID of Section 1.6.2, TIC-101 to TIC-105, on its own pyrometer. Here it runs as a PI loop whose output is a heating rate v_k in kelvin per second, as

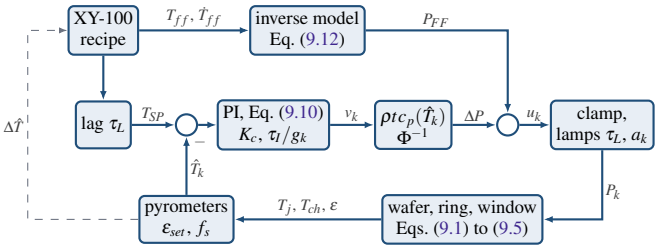


Figure 9.3 Signal path of one zone loop in the reference configuration. The supervisor's lead trajectory feeds the inverse model and, through the filament lag, the setpoint. The PI output is a heating rate that the areal heat capacity and the decoupler turn into lamp power, and the pyrometers' spread returns to the supervisor for the slip-aware ramp limit. Without the decoupler, Φ^{-1} becomes $1/\Phi_{kk}$.

shown in Equation (9.10).

$$\begin{aligned} e_n &= T_{SP,n} + \Delta T_k - \hat{T}_{k,n}, \qquad p_n = K_c e_n \\ S_n &= S_{n-1} + (p_n - p_{n-1}), \qquad w_n = S_n + I_{n-1} \\ v_n &= \mathrm{clamp}(w_n, v_{min}, v_{max}) \\ I_n &= I_{n-1} + \left[\frac{g_k K_c}{\tau_I} e_n + \frac{g_k (v_n - w_n)}{\tau_I} \right] \Delta t_c \end{aligned} \tag{9.10}$$

where ΔT_k is a zone offset from the recipe, zero by default, $K_c = 4\,\mathrm{K/s}$ per kelvin, $\tau_I = 1.5\,\mathrm{s}$, and $g_k = 1$ unless the schedule below is on. These are Equations (1.12) to (1.15) with p_n for P_k, w_n for v_k, v_n for u_k, no bias, no derivative and a tracking time equal to the integral time, so the sum S_n is never clamped and the back-calculation acts on the integral alone. There is no rate limit on the output, whose speed the filament lag already sets. A low reading raises its zone's heating, so each loop has action $a = +1$ and is reverse acting in the convention of Section 1.6.2. The app shows no action label.

The heating rate becomes lamp power through the controller's areal heat capacity and the gain matrix read at the spots, as shown in Equation (9.11).

$$\begin{aligned} \Delta f_k &= \rho\, t\, c_p(\hat{T}_k)\, v_k \\ \Delta P &= \Phi^{-1} \Delta f \quad \text{with the decoupler} \\ \Delta P_k &= \Delta f_k / \Phi_{kk} \quad \text{without it} \\ u_k &= \mathrm{clamp}\big(P_{FF,k} + \Delta P_k,\ P_{idle,k},\ P_{max,k}\big) \end{aligned} \tag{9.11}$$

where Δf_k is the absorbed flux the loop asks for at its spot and $P_{FF,k}$ is the feedforward below. With the decoupler all five zones together deliver exactly Δf at all five spots, the static decoupler of Section 1.4.5 written for absorbed flux. Without it, zone k alone answers loop k, and the loop's own limits v_{min} and v_{max} in Equation (9.10) are its lamp's idle floor and rating, less the feedforward, divided by $\rho t c_p(\hat{T}_k)/\Phi_{kk}$. With the decoupler the lamps are shared, the loops carry no limits of their own, and while any lamp is clamped each integral is driven back by $g_k(\bar{v}_k - v_k)\Delta t_c/\tau_I$ toward the output $\bar{v}_k$ that reproduces the lamps actually applied. The single-loop architecture of case 2 puts one loop on PT-101 and spreads its flux over the zones in the pattern that heats the whole wafer evenly. The idle floor gives way to zero in the cool phase. In MANUAL the operator sets the five commands, and every loop tracks by setting its integral so that its output reproduces the applied lamps, as it also does in the load, preheat and unload phases, so the return to automatic produces no bump, and a feedforward that appears or is switched is absorbed into the integrals in the same way.

The optional schedule multiplies each loop's integral gain by $g_k = [(\hat{T}_k + 273.15)/1323.15]^3$, the radiative conductance relative to its value at 1050 °C, so the integral time on the panel is the one at the peak.

9.5.2 The Inverse-Model Feedforward

The feedforward inverts a uniform-wafer version of Equation (9.1) along the recipe, as shown in Equation (9.12).

$$\begin{aligned} q_j &= m_j c_p(T_{ff})\,\dot{T}_{ff} + \varepsilon(T_{ff})\,\sigma\,2A_j\,(T_{ff}^4 - T_c^4) \\ &\quad + h_0\,2A_j\,(T_{ff} - T_{gas}) + \delta_{j7}\,\varepsilon\sigma A_E (T_{ff}^4 - T_c^4) \\ P_{FF} &= (G^\top W G)^{-1} G^\top W\, q, \qquad W = \mathrm{diag}(1/A_j) \end{aligned} \tag{9.12}$$

where the model assumes a cold chamber at $T_c = 30\,^\circ\mathrm{C}$ and the weighting makes the fit a least-squares match of absorbed flux over area. The filament lag is inverted by the reference itself. The supervisor keeps a lead temperature T_{ff} that moves at the recipe rate, computes the feedforward from it, and gives the loops the setpoint T_{SP} obtained by passing T_{ff} through the same lag, $\tau_L \dot{T}_{SP} = T_{ff} - T_{SP}$, which is what a wafer driven through real filaments by that feedforward can do. During cooling T_{ff} follows the model's free cooling of a uniform wafer with the lamps off. The model knows nothing of the edge ring, the warming window, the gas flow or a changed rim loss, and it uses the emissivity the pyrometers are set to. Its h_0, the lag τ_L of the reference and the exponent m of the soak correction below are the declared values of Table 9.2, so a learner who changes one of those constants changes a plant that the controllers do not know, and the mismatch reaches them only through the pyrometers and the probe.

9.5.3 Recipe Supervisor

XY-100 sequences each 60 s slot. It loads the wafer with the lamps idle for 6 s, preheats it open loop at 12 % of the bank in the uniform pattern until all five pyrometers read above 360 °C, and then hands over without a bump. The setpoint starts at the pyrometers' mean and continues at the heating rate the controller's model assigns to the preheat power, and each zone's setpoint starts from its own reading and joins the common one with a time constant of 0.8 s. The wafer is held at 600 °C for 8 s, ramped, soaked for 1.0 s from the moment T_{ff} reaches the peak, and cooled until T_{SP} reaches 500 °C. In case 3 with the decoupler and the feedforward, the phases change at 6.0, 10.3, 22.9, 24.7, 25.7 and 36.7 s. With the slip-aware limit on, the ramp rate is multiplied by a factor that follows, with a 20 ms lag, the target in Equation (9.13).

$$\rho_r = \mathrm{clamp}\Big(2\Big[1 - \frac{\max_k \hat{T}_k - \min_k \hat{T}_k}{\Delta T_{allow}(T_{peak})}\Big],\ 0.1,\ 1\Big) \tag{9.13}$$

where ΔT_{allow} is Equation (9.8) with the operator's 10 K at 1000 °C in place of the model's 12, evaluated at the peak, where the limit is tightest. The ramp therefore slows once the pyrometers' spread passes half of 6.3 K. A zone commanded to its rating for a total of 30 ms in the ramp or the soak marks the wafer lamp-limited, and, with the hold on, the supervisor stops the lot after that wafer rather than anneal wafers it cannot make.

9.5.4 Emissivity, Calibration and the Soak Correction

The emissivity the pyrometers use comes from one of three modes, a constant for every wafer, the film-stack table for the stack the lot's records say is loaded, or the last calibration wafer's estimate. A calibration wafer carries thermocouples at the five pyrometer radii, each lagging by 0.15 s with an offset of up to 0.5 K, and soaks for at least 3 s. Over the last 1.5 s of its soak each spot's estimate is the mean of $T_{bb,k}^4 / T_{tc,k}^4$, and the five are averaged. The estimate therefore absorbs the window's stray radiance as it was at calibration. The wafer-to-wafer correction adds δ to the soak after each wafer's probe reading, as shown in Equation (9.14).

$$\delta_n = \mathrm{clamp}\Big(\delta_{n-1} + \lambda\,\frac{\theta_T}{m}\ln\frac{\hat{R}_{s,n}}{R_{s,T}},\ -0.9\,\mathrm{s},\ 2.0\,\mathrm{s}\Big) \tag{9.14}$$

where $\lambda = 0.5$, because at the nominal spike one more second of soak changes $\ln R_s$ by $-m/\theta_T$. It is an integral controller between runs in the manner of run-by-run control [11], the incremental form of Equation (1.3) with the soak as the recipe and $\ln R_s$ as the result, which Chapter 6 applies to an etch trim. The soak rises when the sheet resistance is above target, so the correction is direct acting in the convention of Section 1.6.2.

9.5.5 Model Predictive Control of the Five Zones

The zone control may instead be a model predictive controller (MPC), the technology that carries most multivariable constraint control in the process industries [19], set up as Section 1.6.3 describes. It acts in AUTO through the stabilization, the ramp, the soak and the cool, and it replaces only the zone loops. The supervisor, its recipe, the lamp-limit flag and the hold are unchanged, the feedforward is not applied, and the five loops track the applied lamp commands so that they can take over without a bump. Its model is the zone loops' own, written at the five pyrometer spots with ten states, the spot temperatures $\tilde{T}_i$ and the delivered powers P_k, as shown in Equation (9.15).

$$\begin{aligned} c_a \frac{\mathrm{d}\tilde{T}_i}{\mathrm{d}t} &= \sum_{k=1}^{5} \Phi_{ik} P_k - L_i(T_i^\star) - \gamma_i(\tilde{T}_i - T_i^\star) + c_a d_i \\ \tau_L \frac{\mathrm{d}P_k}{\mathrm{d}t} &= u_k - P_k \\ L_i(T) &= (2+\chi_i)\,\varepsilon(T)\,\sigma\,(T^4 - T_c^4) + 2h_0(T - T_{gas}) \end{aligned} \tag{9.15}$$

where $T_i^\star(t)$ is the reference at spot i (the recipe setpoint plus the zone's offset), $c_a = \rho\, t\, c_p$ is the areal heat capacity at the mean of the five references at the moment of the decision, $\gamma_i = dL_i/dT$ at $T_i^\star$ at that moment, $\chi_i = 0$ except $\chi_5 = A_E/A_7 = 0.171$ for the rim's extra loss at the edge spot, $\varepsilon(T)$ is the emissivity the pyrometers are told, $T_c = 30\,^\circ\mathrm{C}$, h_0 and τ_L are the declared values, and d_i is an unmeasured disturbance in kelvin per second. The slope γ_i, the cubic gain of Section 9.4 per unit area, is 242 W/(m^2 K) at 600 °C and 752 at 1050 °C, and it and c_a are evaluated again at every decision, while the loss L_i itself enters along the reference over the horizon as a known time-varying term, so the plan knows that the top of the ramp costs more power than its start. The model has no radial conduction, no ring, no window and no aged lamps, and d_i absorbs them.

Every 50 ms, five scans, the controller discretizes Equation (9.15) for a zero-order hold on a 0.1 s grid and solves the quadratic program of Equation (9.16) over $N = 50$ steps, a 5 s horizon, with each command held over eight blocks, short near the present (1, 1, 2 and 2 steps) and long toward the end (4, 6, 10 and 24 steps), so the plan has 40 decision variables.

$$\begin{aligned} &\min \sum_{n=1}^{N} \Big[\sum_{i=1}^{5} e_{i,n}^2 + \sum_{i=1}^{4} (e_{i,n} - e_{i+1,n})^2 \Big] + w_u \sum_{\ell=1}^{8} \|\Delta u_\ell\|^2 \\ &\text{s.t. } \underline{u}_{k,\ell} \le u_{k,\ell} \le P_{max,k}, \qquad e_{i,n} = \tilde{T}_{i,n} - T_{i,n}^\star \end{aligned} \tag{9.16}$$

where both temperature weights are 1/K^2, the second sum is a weight on the difference between neighboring spots that makes the across-wafer spread count as much as tracking, $w_u = 0.04/\mathrm{kW}^2$ prices a 5 kW move per block like 1 K of error at one step, Δu_ℓ is the change of the five commands at the start of block ℓ (from the applied commands for the first block), and the floor $\underline{u}_{k,\ell}$ is the idle filament power for every block that starts before the preview reaches the cool phase and zero after it, as the loops clamp their own lamps. Every limit is hard and there is no soft constraint. The quadratic program is solved and applied in receding horizon as in Section 1.6.3 [20]. The preview is the supervisor's own reference over the horizon, replayed with the declared lag, the soak it will set, the ramp at the slip-aware factor of the moment and the free cooling after the soak. With the preview Off the reference is held at its present value, and the loops' feedforward receives the same recipe one lamp lag ahead.

Table 9.3 Controller Settings of the Reference Configuration, With the Ranges the Panel Allows

Setting	Value	Range
architecture	five zone loops	one loop on PT-101, MPC
MPC preview	on	off
K_c, τ_I at 1050 °C	4 K/s per K, 1.5 s	0–20, 0.1–20 s
feedforward, decoupler, schedule	on, on, off	on or off
T_{pre}, preheat power	600 °C, 12 %	450–800, 5–60
slip-aware limit, ΔT_{allow}	on, 10 K at 1000 °C	1–40 K
hold on a lamp limit	on	on or off
emissivity mode, constant	film-stack table, 0.70	constant, calibration
soak correction, λ	off, 0.5	0–1.5
scan, output, direction	10 ms, K/s, reverse	—

The disturbances come from a steady-state Kalman filter on each spot, run at every 10 ms scan. The lamp powers are propagated through the declared lag from the applied commands, and each pair $(\tilde{T}_i, d_i)$ is corrected by its pyrometer with the gain (0.0965, 0.317) of a Riccati solution for the spot model at 1050 °C, a pyrometer noise of 0.3 K and random walks of 0.02 K and 0.1 K/s per step on the temperature and the disturbance, so the estimation error decays with a time constant of 0.19 s. An input disturbance of this kind makes the controller offset-free at the five spots [21]. Acceptance of a plan requires an optimal status, a finite first move, plan, prediction and cost, and every planned command inside its block's limits. Otherwise the historian is marked, the five loops, which have been tracking the lamps, carry on from the applied commands, and the next accepted plan takes the lamps back over 0.3 s. With the noise off, the hand-over moves a lamp by at most 0.47 % of its rating, and a second rejection inside that return by at most 0.15 %.

Step tests check the model against the plant. On a held wafer, each zone stepped by 1 % of its rating for 20 s with the pyrometers held at 1050 °C gives a gain to its own spot within 6 % of the model's (−0.3 to +5.8 %) and a 63 % time of 2.60 to 2.69 s against 2.39 to 2.57 s. At 600 °C the model's gains are 15 to 26 % high and its 63 % times of 6.5 to 7.0 s are longer than the plant's 5.7 to 6.1 s, because conduction, which the model leaves out, spreads each zone's heat to the neighboring spots (zone 2 heats spot 1 by 20.0 K/kW against the model's 4.2). On the ramp of case 3 the model's 0.5 s prediction from the applied commands runs 2.9 K warm on average at the four inner spots (1.9 K with c_a taken along the path instead of at the decision), and in case 9 the aged edge lamps put the edge spot's prediction up to 21.6 K off until d_5 reaches about $-46\,\mathrm{K/s}$.

9.5.6 Three Models Kept Apart

Section 9.3 describes the plant that is integrated. What the controller carries inside is much smaller, namely the inverse model of Equation (9.12), the matrix Φ and its inverse, the areal heat capacity $\rho t c_p(\hat{T}_k)$, the free cooling of a uniform wafer and the soak sensitivity $-m/\theta_T$, and none of these contains the ring, the window or the true emissivity. The MPC's model of Equation (9.15) is assembled from the same pieces, and its estimator sees the plant only through the pyrometers and the applied commands. No identified control model exists in the application, and the open-loop steps of case 1 measure the plant the loops see, 63 % in 8.9 s near 400 °C and 3.4 s near 900 °C. Table 9.3 lists the shipped settings.

9.6 Measured Behavior of Each Case

Table 9.4 summarizes what each case changes and what it teaches. Every case loads the reference configuration of Table 9.3 and then applies only the changes in its row, and the outcomes come from the engine at its default seed with no interface attached, timed in simulated seconds. A variant is the case with one setting changed before the run starts, as each briefing asks after a Reset, which keeps the learner's settings (Section 1.6.4). A setting changed before Run and the same setting carried through a Reset give one run. Table 9.5 gives the first wafer of every single-wafer case and variant under the

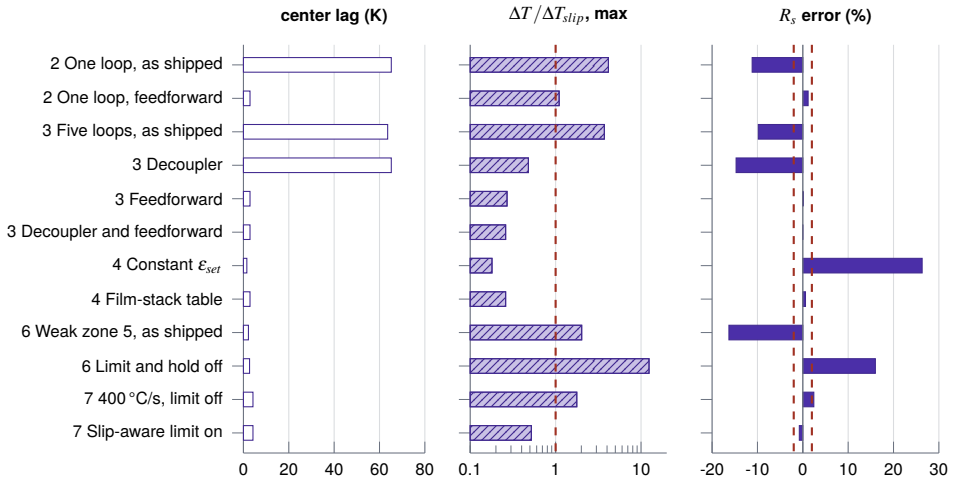


Figure 9.4 Three columns of Table 9.5 redrawn: the largest center lag on the ramp (open bars), the worst ratio of the true radial difference to the slip limit on a logarithmic axis (hatched, the slip limit of 1 dashed) and the sheet resistance error of the first wafer (solid, the 2 % band dashed). Notice that every row with the feedforward has a small lag and a sheet resistance inside the band, except case 4 with the constant emissivity.

Table 9.4 Scenario Sequence, From a Held Wafer to the Graded Lot

#	Scenario, changes from the reference	Teaching point
1	MANUAL, held wafer, +2 % at 400 and 900 °C	the cubic plant
2	one loop on PT-101, no feedforward[a]	a late ramp
3	five loops, no feedforward or decoupler[a]	loops that fight
4	nitride stack C, constant $\varepsilon_{set} = 0.70$	a trend that lies
5	stack C, calibration wafer first	an independent sensor
6	zone 5 at 35 % of its output	honestly infeasible
7	400 °C/s, zone 5 at 85 %[a]	rate against gradient
8	cooling water at 25 %, 25 wafers	a slow drift
9	case 7 under the MPC[a]	planning at the limits
10	constant ε, no feedforward or decoupler, stack C from wafer 10, zone 3 at 85 % from 16	the graded lot

[a] with the slip-aware ramp limit and the hold on a lamp limit off

zone loops, and Section 9.6.9 compares the MPC with them. Figure 9.4 draws three columns of that table as bars.

9.6.1 One Lamp Step at Two Temperatures

Case 1 holds one wafer in MANUAL with the stations at 1.7 to 1.8 %, 1.13 kW absorbed, where the center settles at 401.9 °C, the edge at 398.3 °C and the ring at 376.7 °C. At 5 s every station rises by 2 points, 1.28 kW more absorbed, and the center climbs 113.8 K by 45 s, 63 % of the way in 8.9 s. At 45 s the stations move to the feedforward's steady powers for 900 °C, 17.1 to 19.3 %, the wafer settles at 900.3 °C, and the same step at 75 s raises it only 34.9 K, 63 % of the way in 3.4 s. The conductance at 900 °C is 4.6 times that at 400 °C, so the same power buys 3.26 times less rise, 2.66 times faster, and the large step at 400 °C shortens its own time constant as it climbs, from 13.5 to 8.1 s between its two ends. In AUTO the loops hide most of this. On the same held wafer, a 10 K step of the hold setpoint overshoots on the center pyrometer by 3.5 K at 600 °C and 3.3 K at 1050 °C (3.1 and 2.9 K without pyrometer noise) and reaches 90 % in 0.45 s at both, and the schedule trims the overshoot at 600 °C to 2.2 K.

Table 9.5 First Wafer of Each Single-Wafer Case and Variant. Lags Are the Largest Setpoint Minus Pyrometer Reading on the Ramp, the Peak Is the True Center Peak, and ΔT Is the Largest True Radial Difference on the Ramp and Soak

#	Configuration	Ramp (s)	Center lag (K)	Edge lag (K)	Peak (°C)	ΔT (K)	$\Delta T/\Delta T_{slip}$ max	R_s error (%)	Range (%)
2	one loop, as shipped	1.80	65.2	77.6	1079.5	25.1	4.15	−11.17	17.80
2	one loop, feedforward	1.80	2.9	5.1	1049.1	8.4	1.10	+1.17	5.56
3	five loops, as shipped	1.80	63.6	41.4	1078.6	34.8	3.72	−9.81	7.87
3	decoupler	1.80	65.2	65.4	1079.5	2.9	0.48	−14.71	1.48
3	feedforward	1.80	2.9	2.6	1049.1	2.0	0.27	+0.06	1.41
3	decoupler and feedforward	1.80	2.9	2.9	1049.1	2.0	0.26	−0.05	1.40
4	constant ε_{set}, as shipped	1.80	1.5	1.5	1017.0	1.9	0.18	+26.31	1.32
4	film-stack table	1.80	2.9	3.0	1049.1	2.0	0.26	+0.59	1.37
6	weak zone 5, as shipped	9.61	2.2	12.4	1049.5	12.6	2.02	−16.29	1.57
6	limit and hold off	1.80	2.7	99.4	1049.3	104.9	12.41	+16.00	50.39
7	400 °C/s, limit off	1.13	4.2	15.6	1048.6	18.6	1.77	+2.43	1.48
7	slip-aware limit on	1.94	4.2	8.2	1048.6	6.3	0.52	−0.77	1.68

9.6.2 A Ramp Built From Its Own Error

Case 2 drives all five zones from one loop on PT-101 with no feedforward. On the ramp from 22.9 s the center reading falls up to 65.2 K behind its setpoint, because a proportional loop needs an error to ask for the ramp's power, and when the setpoint stops the integral that has finally caught up overshoots by 29.9 K, taking the true center to 1079.5 °C, as shown in Figure 9.2(b). The edge pyrometer lags by 77.6 K, the wafer slips at 24.7 s at 4.15 times the limit, and its sheet resistance is 355.3 Ω/sq, 11.2 % low, with a range of 17.8 % across the wafer. The feedforward supplies the ramp's power one lamp lag ahead, the center lag falls to 2.9 K with no overshoot and the sheet resistance is 404.7 Ω/sq. The edge still lags by 5.1 K and the wafer still slips, at 1.10 times the limit, because a single loop on the center cannot see the rim.

> **TRY IT IN THE APP** **2 · One Loop on the Center**
>
> Load the scenario and press Run. Watch the center pyrometer strip and the fast scope as the first ramp starts at about 23 s. The reading falls up to 65.2 K behind its setpoint and then overshoots by 29.9 K when the setpoint stops. At 60 s the Lot view gives the first wafer's sheet resistance, 355.3 Ω/sq, against a target of 400 Ω/sq.

9.6.3 Five Loops That Fight

Case 3 gives each zone its own loop without the feedforward or the decoupler. Each loop's correction also lands on its neighbors' spots, the middle loops get help while the center does not, and on the ramp the center lags 63.6 K while the edge lags only 41.4 K. The radial difference reaches 34.8 K and the wafer slips at 24.8 s, 3.72 times the limit. The decoupler spreads each correction through Φ^{-1} so that it moves only its own pyrometer. The wafer then lags by 65 K as one piece, with the radial difference at 2.1 K on the ramp and no slip, but it still overshoots by 29.9 K and comes out 14.7 % low. With the feedforward as well, the radial difference is 1.5 K on the ramp, the sheet resistance is 399.8 Ω/sq and its range 1.40 %. The feedforward alone does nearly as much, 1.6 K on the ramp and 0.06 % of error, because once the model supplies the power the loops have little left to fight over.

Pause. Try Exercise 9.3 before reading on.

9.6.4 A Perfect Trend on the Wrong Wafer

Case 4 loads nitride-capped wafers, emissivity 0.77, under pyrometers that assume 0.70, with everything else as the reference. The pyrometers read 2.4 % hot and the loops hold them on the recipe,

so the pyrometers peak at 1048.4 °C while the true center peaks at 1017.0 °C, 31.4 K lower. The first wafer's sheet resistance is 505.2 Ω/sq, 26.3 % high, and the second also 505.2 Ω/sq. Nothing on the loops' own trends shows it, as Figure 9.5 records during the first soak, and only the dashed true temperature, which no real tool has, and the probe after the wafer do. With the film-stack table the first wafer is 402.3 Ω/sq, 0.6 % high.

> ✦ DID YOU KNOW? **The Best Trend in the Table**
>
> From the simulator: in Table 9.5 the nitride-capped wafer read with the constant emissivity has the smallest lags of any row, 1.5 K at the center and at the edge, and the smallest radial difference, 1.9 K. Its sheet resistance is 26.31 % high, the largest error of any wafer in the table that did not slip.

Case 5 runs the same stack with the calibration mode, and the first wafer carries the thermocouples. Over the last 1.5 s of its 3 s soak the thermocouples read 31.8 K below the center pyrometer, which still assumes 0.70. The five estimates lie between 0.7705 and 0.7720 and average 0.7712, against the wafer's own 0.7706, and from the second wafer the pyrometers use them. Wafer 2 is 401.7 Ω/sq, 0.42 % high, and wafer 3 401.4 Ω/sq, 0.34 % high, where the wafers of case 4 were 26 % high. The estimate carries the window's stray light as it was at calibration, so it goes out of date as the chamber warms or the stack changes.

🕮 **Pause.** Try Exercise 9.2 before reading on.

9.6.5 An Edge Zone That Cannot Make the Recipe

Case 6 leaves zone 5 at 35 % of its output, 18.2 kW of its 52 kW rating, a third of which heats the ring. The zone sits at its rating for 1.36 s of the ramp and 0.48 s of the soak, so even the soak at 1050 °C needs more than it has. The slip-aware limit stretches the ramp from 1.80 to 9.61 s, the edge still falls 12.4 K behind, the wafer slips at 34.7 s at 2.02 times the limit, and the long slow ramp leaves its sheet resistance 16.3 % low. When it leaves at 60 s the supervisor holds the lot. With the limit and the hold off the edge pyrometer falls 99.4 K behind its setpoint at 250 °C/s, the radial difference reaches 104.9 K in the soak and the sheet resistance ranges over 50.4 % of its mean. Slower recipes do not rescue it. At 150, 100 and 50 °C/s with the limit off, zone 5 is at its rating through the whole soak and the wafer slips at 7.95, 4.51 and 1.95 times the limit. Some faults are not control problems, and this case is built to show one.

9.6.6 A Ramp Too Fast for the Rim

Case 7 ramps at 400 °C/s with zone 5 aged to 85 % and the slip-aware limit off. All five zones reach their ratings on the ramp, zone 5 for 0.55 s, the edge pyrometer lags by 15.6 K, the radial difference reaches 17.3 K on the ramp and 18.6 K in the soak, and the wafer slips at 24.2 s, 1.77 times the limit. With the limit on, the ramp slows while the spread is large and takes 1.94 s instead of 1.13, the edge lag stays at 8.2 K, the wafer reaches 0.52 of the limit and its sheet resistance is 396.9 Ω/sq, 0.8 % low, inside the band.

9.6.7 A Window That Warms Across the Lot

Case 8 runs 25 wafers with the cooling water at a quarter of its design flow. The window reaches 170.4 °C by the end of the lot, against 90.0 °C when the same lot runs at design flow with all 25 wafers in specification, and its radiance beyond the 30 °C calibration makes the pyrometers read high. The true center peak falls from 0.9 K below 1050 °C on the first wafer to 5.0 K below on the last, and the sheet resistance rises from −0.05 to +2.87 %. Wafer 20 is the first outside the band, wafer 22 falls back inside it at 1.94 %, and 20 of the 25 are in specification, against 18 to 20 at six other seeds of the random generator. With the wafer-to-wafer correction the soak grows to 1.10 s and all 25 are in specification, the worst 0.75 % off (Table 9.6). It repairs the product without knowing the cause, and the peak stays 5.0 K low.

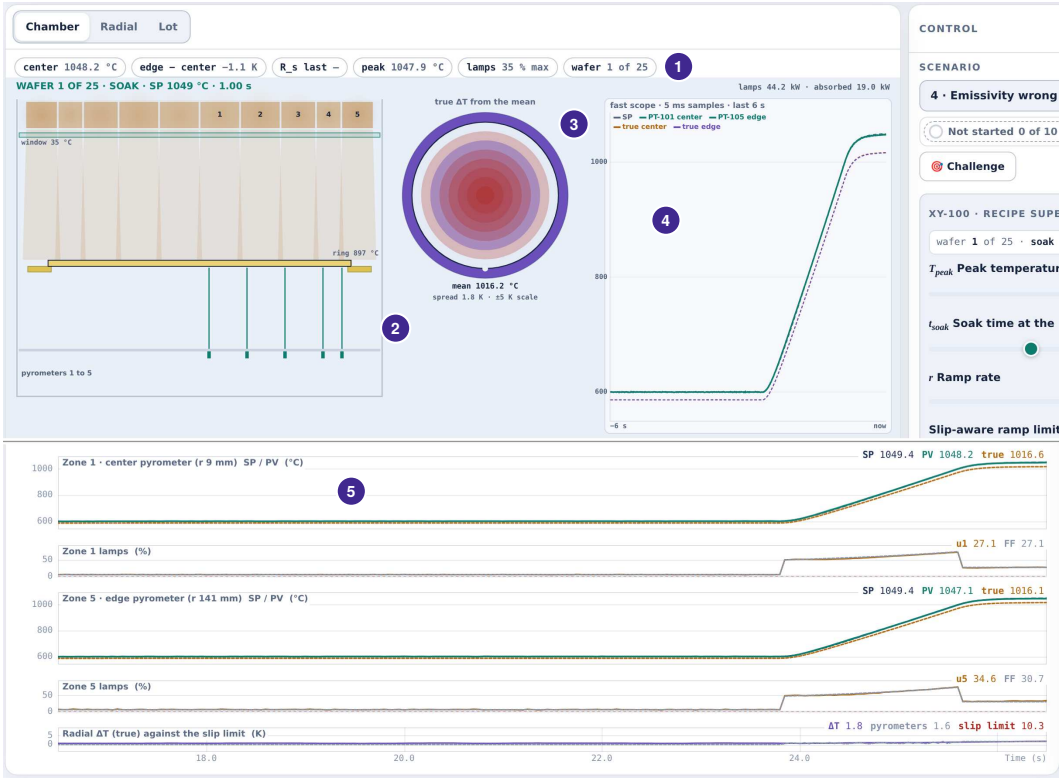


Figure 9.5 The chamber in case 4, emissivity wrong by 10 %, at $t = 26.50$ s, every control as the case sets it, in the Chamber view with a 10.0 s trend window, cropped to the plant view and the trends. Wafer 1 of 25, nitride capped and read by pyrometers that assume bare silicon, is in the last tenth of its 1.00 s soak at a setpoint of 1049.4 °C. The cutaway (2) shows the five lamp zones, the window at 35 °C, the wafer, the edge ring at 897 °C and all five pyrometer beams, with the bank delivering 44.2 kW of which the wafer absorbs 19.0 kW. The disc (3) paints each annulus by its true difference from the wafer mean of 1016.2 °C, a spread of 1.8 K in which the annulus from 94 to 113 mm, which no pyrometer sees, runs 1.3 K cold. The badges (1) give the center pyrometer at 1048.2 °C, the edge 1.1 K below it, the peak so far at 1047.9 °C and no sheet resistance yet. The fast scope's (4) last 6 s run through the end of the stabilization, the ramp from 23.8 s and the soak, with both pyrometers on the setpoint throughout and the true center and edge 13.5 K below them at 600 °C and about 31 K below them in the soak. The trends (5) show zone 1 at 1048.2 °C with its true temperature of 1016.6 °C dashed and its lamps at 27.1 % beside a feedforward of 27.1 %, zone 5 at 1047.1 °C with 1016.1 °C true and its lamps at 34.6 % beside 30.7 %, and the true radial difference at 1.8 K, 1.6 K between the pyrometers, under a slip limit of 10.3 K.

9.6.8 The Graded Lot

The challenge runs 1500 s, with the nitride stack from wafer 10 and a zone 3 lamp string failing to 85 % at wafer 16. It ships with five independent loops, no feedforward, no decoupler and the constant emissivity, and it asks for all 25 wafers within 2 % of 400 Ω/sq, none slipped, and every true center peak within 3 K of 1050 °C, with the case's model parameters and disturbances unchanged. As shipped no wafer is in specification. Wafers 1 to 9 overshoot to true center peaks of 1059.1 to 1060.6 °C, 9.1 to 10.3 % low, and all slip. Wafers 10 to 15 peak at 1026.8 to 1028.3 °C while their pyrometers still reach about 1060 °C, 13.5 to 15.3 % high, and four of them slip. Wafers 16 to 25 peak at 1022.2 to 1023.0 °C, 5.0 to 8.1 % high, and all slip, 23 slips in all (21 to 23 at the six other seeds).

Table 9.6 Lots of 25 Wafers. The Peak Is the True Center Peak Less 1050 °C, and the Energy Is the Lamps' Electricity Over 1500 s

Configuration	In spec	Slips	R_s error (%)	Peak (K)	Energy (kWh)
case 8, cooling water at 25 %					
as shipped	20	0	−0.22 to +2.87	−4.95 to −0.60	3.29
soak correction	25	0	−0.65 to +0.75	−4.96 to −0.58	3.30
graded lot					
as shipped	0	23	−10.27 to +15.34	−27.8 to +10.6	4.06
shipped, table on	0	25	−14.66 to −7.70	+3.7 to +10.5	4.12
suggested[a]	25	0	−0.60 to +1.23	−2.32 to −0.55	3.42
with the schedule	25	0	−1.26 to +1.33	−2.29 to −0.63	3.45
no decoupler	25	0	−0.32 to +1.38	−2.35 to −0.56	3.41
no feedforward	0	0	−14.89 to −12.65	+24.8 to +29.8	3.50
constant ε	9	0	−0.14 to +27.26	−34.6 to −0.8	3.39
and soak correction	19	0	−0.58 to +26.67	−33.7 to −0.8	3.61

[a] film-stack table, feedforward and decoupler, and the rows below it each change one element of it

With the film-stack table, the feedforward and the decoupler, all 25 are in specification between −0.60 and +1.23 %, the worst wafer reaches 0.54 of the slip limit, the true center peak stays 2.32 to 0.55 K below 1050 °C as the window warms to 93.3 °C, and the lamps use 3.42 instead of 4.06 kW h. The worst sheet resistance leaves a margin of 0.77 points to the band and the worst peak a margin of 0.68 K. The same lot passes at the six other seeds, its worst wafer 1.13 to 1.44 % off.

The variants in Table 9.6 say which element carries the lot. Once the feedforward is on, the lot passes without the decoupler (worst 1.38 %, 0.55 of the slip limit) and with the schedule (1.33 %), and the differences among these three worst wafers lie within their spread over the six other seeds. Without the feedforward no wafer is in specification, every one overshooting by 25 to 30 K with the decoupler and all 25 slipping without it, and without the table the lot is lost from wafer 10. The soak correction does not replace the table. With the constant emissivity it lengthens the soak to 2.53 s and brings wafers 16 to 25 back inside the band, but wafers 10 to 15 are already lost and every nitride wafer still peaks about 33 K low.

9.6.9 The Predictive Controller Against the Zone Loops

Table 9.7 compares the first wafer under ten structures on two recipes, the nominal 250 °C/s ramp of case 3 and the 400 °C/s ramp of case 9 with zone 5 aged to 85 %, which is case 7 under the MPC. Every row shares the seed, the noise sequence, the pyrometers and the supervisor, whose recipe drives both the loops' feedforward and the MPC's preview. The MPC keeps the settings shipped in case 9 on both recipes. Row 4 retunes the loops once with the SIMC rule applied to the decoupled channel, an integrator behind the 0.2 s filament lag with $\tau_c = 0.2$ s, which gives $K_c = 2.5$ per second and $\tau_I = 1.6$ s, with the cubic schedule and the same feedforward. Rows 5 to 7 remove one mechanism from the MPC each, and rows 9 and 10 turn the slip-aware ramp limit on.

On case 9 the MPC holds the wafer where every loop structure without the slip-aware limit slips. Its largest radial difference on the ramp is 10.89 K, 0.79 of the limit, against 17.27 K and 1.77 times the limit for the decoupled loops with the feedforward, and with less than half of their lamp travel (685 against 1614 kW). A plan that cannot give the edge more power gives the inner zones less, as shown in Figure 9.6. The wafer pays for it in budget, because its sheet resistance is 3.34 % high, outside the band, against the loops' 2.43 % on a wafer that slipped. The retuned loops slip further (2.28 times the limit), so retuning the loops does not recover the edge.

Table 9.7 First Wafer Under Each Structure on Case 3's Recipe and on Case 9. ΔT Is the Largest True Radial Difference on the Ramp, the Overshoot Is the True Center Peak Less T_{peak}, Saturation Is the Time Any Zone Is Commanded at Its Rating in the Ramp and Soak, the Ratio Is the Worst $\Delta T/\Delta T_{slip}$, and the Travel Is the Summed Change of the Five Commands

		case 3 recipe, 250 °C/s					case 9, 400 °C/s, zone 5 at 85 %					
#	Structure	ΔT (K)	Overshoot (K)	R_s error (%)	Ratio	Travel (kW)	ΔT (K)	Overshoot (K)	Saturation (s)	R_s error (%)	Ratio	Travel (kW)
1	five loops, feedforward	1.63	−0.91	+0.06	0.27	1240	15.25	−1.40	0.49	+2.36	1.60	1248
2	and the decoupler	1.50	−0.91	−0.05	0.26	1633	17.27	−1.44	0.55	+2.43	1.77	1614
3	MPC	2.11	+0.82	+1.11	0.52	568	10.89	+0.42	0.75	+3.34	0.79	685
4	loops retuned once	1.46	−0.71	−0.07	0.25	1113	20.44	−0.92	0.43	+3.41	2.28	1131
5	MPC without lamp rows	2.07	−0.88	+0.41	0.29	537	11.69	−0.91	0.70	+2.51	0.84	607
6	MPC, preview off	2.11	−0.96	+0.89	0.45	544	6.76	−1.12	0.75	+3.35	0.76	602
7	MPC without spread weight	2.33	+0.85	+0.98	0.78	490	11.12	+0.51	0.65	+3.12	1.05	628
8	loops and decoupler only	2.14	+29.46	−14.71	0.48	1522	22.79	+43.06	0.89	−16.59	4.30	1496
9	row 2, slip-aware limit	1.50	−0.91	−0.05	0.26	1633	6.29	−1.36	0.00	−0.77	0.52	2244
10	MPC, slip-aware limit	2.11	+0.82	+1.11	0.52	568	4.85	+0.76	0.40	+1.65	0.74	951

A ratio above 1 slipped the wafer. No lamp reaches its rating on case 3's recipe, where the slip-aware limit never acts. Each MPC row solved 528 programs on case 3's wafer and 511 on case 9's (516 in row 10), none was rejected, and a solve took 2.1 to 2.4 ms on average in Node on a shared host.

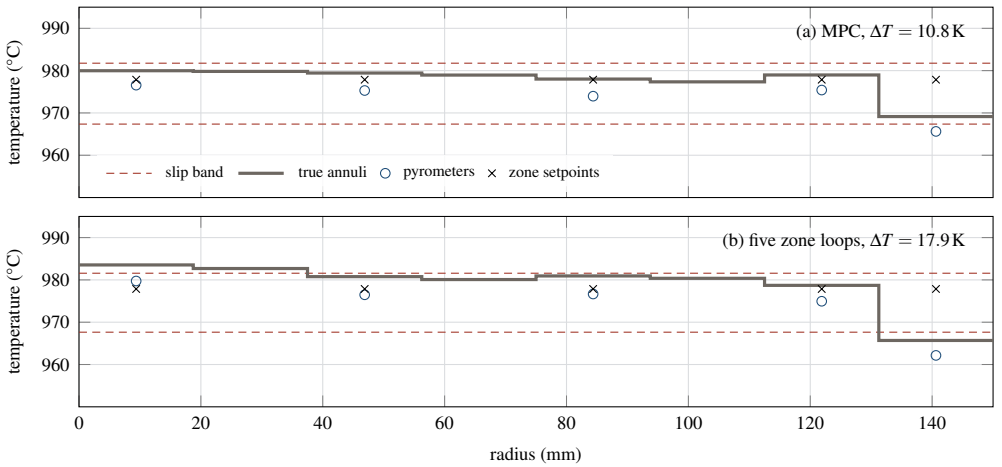


Figure 9.6 Radial profile of the first wafer of case 9 at $t = 24.2$ s, as the 400 °C/s lead trajectory reaches the peak and the setpoints stand at 977.9 °C, under (a) the MPC and (b) the five zone loops with the feedforward and the decoupler. Steps are the true annulus temperatures, circles the pyrometers, crosses the zone setpoints, and the dashed lines the slip band, as wide as the limit at the hottest annulus and centered on the profile. The MPC holds its edge lamps at their rating and the inner zones at 58 to 73 %, for 10.8 K against a limit of 14.4 K. The loops have cut every zone to 26 to 57 % with the center at 983.5 °C and the edge annulus at 965.7 °C, 17.9 K against 13.9 K.

TRY IT IN THE APP **9 · MPC Against the Zone Loops on a Fast Ramp**

Press Run and watch the Radial view during the ramp at about 23 s. Under the MPC the radial difference peaks at 10.89 K and the wafer does not slip, but its sheet resistance is 3.34 % high. Press Reset, set Zone control to Five zone loops and run again: the difference reaches 17.27 K and the wafer slips.

The counterfactual rows say where the smaller spread comes from. It survives the removal of each mechanism that the MPC adds to the loops, at 11.12 K without the weight on neighboring spots, 11.69 K without the lamp limits in the plan and 6.76 K with the preview off, so it comes from

Table 9.8 Graded Lot Under the Zone Loops and the MPC, at the Shipped Seed

Configuration	In spec	Worst R_s (%)	Worst peak (K)	Slips
as shipped	0	15.34	27.83	23
suggested loops	25	1.23	2.32	0
MPC, constant ε	9	28.45	30.94	0
MPC, film-stack table	19	2.67	1.14	0
retuned loops, table	15	3.16	1.88	0

planning the five lamps together on the model while they saturate. The spread weight decides only the verdict on slip. Graded from 0 through 0.25, 0.5 and 1 to 2/K^2, it moves the ratio from 1.05 (a slip) through 0.92, 0.83 and 0.79 to 0.77 while the spread stays between 10.6 and 11.1 K. Without the preview the plan, which then sees only the setpoint of the moment, lags the ramp and keeps the wafer more uniform at the same budget (+3.35 %). The shipped structure already receives the same recipe, through its feedforward, in row 2.

The supervisor's own answer is better for the product. With the slip-aware limit on, the decoupled loops stretch the ramp from 1.13 to 1.94 s, hold the difference to 6.29 K (0.52 of the limit) and make a wafer 0.77 % low, inside the band, which beats every MPC row. A slower ramp asks less of the edge lamps, and no plan of the same fast ramp can do that. With the same limit the MPC has the smallest spread of all, 4.85 K, but is still 1.65 % high.

On the nominal recipe the MPC loses. Its spread is 2.11 K against 1.50 K, its true center peak 0.82 K over the recipe instead of 0.91 K under, and its sheet resistance 1.11 % high against −0.05 %, with only its lamp travel lower (568 against 1633 kW). The budget is lost at the end of the soak. Of the 0.0508 s of equivalent budget that the MPC's wafer lacks (3.8 % of the loops'), 0.0319 is lost in the first 0.5 s of the cool and 0.0169 in the next 1.5 s, against 0.0016 in the soak. After the cut the setpoint follows the model's free cooling through the lamp lag, which a wafer with its lamps off cannot follow exactly, and the least-squares plan, which knows the lamps cannot go below off, spreads that error across the corner. It lowers the lamps from 46.1 to 25.8 and 12.8 kW over the last 0.2 s of the soak, where the loops hold 44 kW, and the true center then falls 5.2 K below its setpoint (the loops' 0.5 K). Allowing negative lamp power after the cut, a counterfactual that makes the setpoint reachable in the model and changes nothing else, brings the sheet resistance to +0.41 % on case 3 and +2.49 % on case 9, where the spread stays at 10.93 K. Removing the weight after the cut gives +0.31 %, while a move weight ten times lighter, a hundred times lighter or ten times heavier gives +1.14, +1.14 and +1.18 %, and one move block per step +1.12 %.

The emissivity error stays the lesson under every structure, because the MPC reads the same pyrometers. On case 4's nitride wafers it makes +27.18 % on the first two wafers against +26.31 and +26.29 % for the decoupled loops and +26.18 % for the retuned loops, with its true center 29.9 K below the peak instead of 33.0, and its trend looks just as perfect. On the graded lot, Table 9.8, the MPC with the film-stack table puts 19 wafers in specification and fails, because wafers 10 to 15 on the nitride stack end 2.3 to 2.7 % high, although no wafer slips and every true center peak is within 1.14 K. The suggested loops put all 25 in specification.

Every existing case was also run for 1501 s under each structure, to find the worst case of the hottest annulus, the slip ratio and the sheet resistance. No plan was rejected anywhere. The MPC never slips a wafer that the loops keep, and it holds every wafer they slip except the weak edge of case 6, which the supervisor holds after the first wafer under both. It is nonetheless less safe than the best shipped structure in four cases (4, 5, 8 and the suggested lot), where its hottest annulus is 0.5 to 3.3 K higher, late in the soak, and its worst slip ratio higher (0.54 against 0.19, 0.65 against 0.26, 0.81 against 0.27 and 0.67 against 0.54), although below 1 wherever the loops' is. In case 2 its slip ratio of 0.65 also exceeds the five zone loops' 0.48, with its hottest annulus 28 K lower. Its worst wafers are also further from the target, 5.16 % against 2.87 % in case 8 and 2.74 % against 1.43 % in case 5.

The application states this in the texts of case 9 and in its model description.

The cost fits the browser. On the page's own frame loop in case 9, timed twice on a shared host (load averages of 3 to 6 and of 12 to 13), a decision takes 2.43 to 2.57 ms on average and at most 11.3 to 13.5 ms, and at real time the solves use 4.9 to 5.1 % of a processor core and all the stepping 8.9 to 9.5 %, at 60 frames per second. At 20 times real time a frame guard holds the run to 16.6 to 17.0 times at about 42 frames per second, where the zone loops run the full 20 times on 4.6 to 5.2 % of a core, so a lot of 1500 s takes 75 s, and the results are identical at 1, 5 and 20 times real time.

9.7 Technology in Practice

The loops of Section 9.5 run on five pyrometers under the wafer, a lamp head of five zones, the tool's own controller and a probe after the chamber. Figure 9.7 draws one zone with that hardware.

What Measures It

Each zone reads one pyrometer, which the model of Section 9.3 reduces to a band-averaged reading behind a 5 ms lag with 0.3 K of colored noise over 350 to 1250 °C, converted with the emissivity it is told. Its radiation is gray and band averaged and its reflector adds no cavity enhancement, so its emissivity error is about three times that of a narrow-band instrument.

> **TECHNOLOGY Sensor: A Lightpipe Radiation Thermometer**
>
> A straight rod of quartz or sapphire looks at the back of the wafer and carries its glow through an optical fiber to a photodetector. Most such thermometers read a band from 0.9 to 1 µm, where silicon is opaque, so lamp light passing through the wafer stays out [13]. Against a blackbody one is calibrated with an uncertainty below 0.3 °C. In a tool the wafer's emissivity and its reflections off the chamber must be accounted for, or the uncertainty grows by orders of magnitude, and a calibration in the tool reached 2.1 °C or less [18]. The cold rod also cools the spot it reads: depressions of up to 25 °C were seen with its tip 2 mm from the wafer [22]. A wrong emissivity leaves no mark on the trend.

The model's pyrometers never cool their spots and err only through the emissivity they are given and the window's radiance, the faults of cases 4 and 8. Its probe reads every wafer at once with a repeatability of 0.3 %, where a fab measures later and on sampled wafers.

What Moves It

Each zone is modeled as a power proportional to its command behind a filament lag of 0.20 s, with ratings from 9 to 52 kW, 145 kW in all, and an idle filament holds 1 % of each rating.

> **TECHNOLOGY Final Element: Tungsten-Halogen Lamps on Thyristor Power**
>
> A tungsten-halogen lamp is a tungsten filament in a quartz tube filled with halogen gas. With the quartz above 250 °C, the halogen cycle keeps tungsten off the envelope, which prevents blackening and lengthens the lamp's life. The filament runs near 2900 K, and the quartz, warmed to about 400 °C, radiates heat of its own. A university chamber of 2005 carried 37 such lamps in three concentric zones, and its controller sent analog commands to the lamp power supplies [17]. Alternating current to a lamp can be throttled by phase control: the later in each half-cycle a thyristor is fired, the less average power the load receives, and large power circuits favor silicon-controlled rectifiers [23]. A zone that has lost lamps shows up as a command held at its rating.

The model's proportional power stands for the average such a controller delivers, and its aged zones stand for lost lamps.

Where the Controller Runs

In a university chamber of 2005 a programmable logic controller read the thermocouples and pyrometers, ran the PID loops and sent the lamp commands, while a workstation above it ran

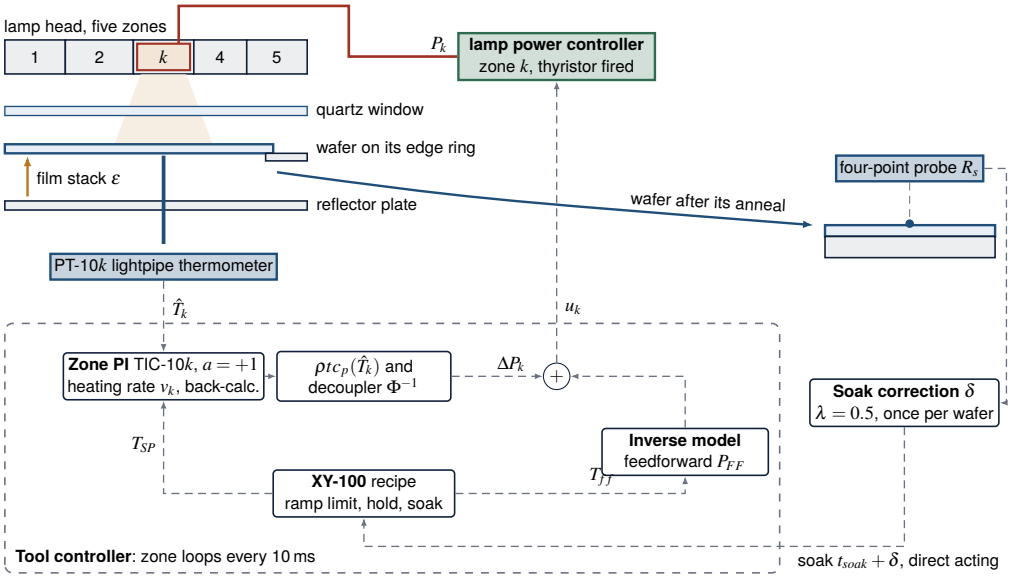


Figure 9.7 One of the five zone loops of the simulated chamber drawn with its hardware, and the soak correction beside it. The lightpipe thermometer PT-10k under the wafer feeds TIC-10k, a velocity-form PI loop with $a = +1$ every 10 ms whose output is a heating rate. The areal heat capacity and the decoupler turn it into lamp power, the inverse-model feedforward from the recipe of XY-100 is added, and the sum drives the power controller of lamp zone k. Once per wafer the probe's sheet resistance lengthens or shortens the soak, a direct-acting correction. Notice that no instrument of the inner loop sees the true wafer temperature, and that a film stack changes what the thermometer reads.

identification and supervisory tasks. Its multiloop PI tests ran at 1 and 10 Hz, and the faster rate gave less overshoot [17]. Control loops in operational technology run with cycle times from milliseconds to minutes [24]. The app integrates every 1 ms, runs its zone loops every 10 ms and the MPC every 50 ms, and corrects the soak once per wafer, every 60 s.

> **IN PRACTICE** **Calibrating the Thermometers Against a Thermocouple Wafer**
>
> Thermometers in a lamp tool are calibrated in place with a wafer that carries thermocouples. The difference between a thermometer's display and its thermocouple then corrects the later readings, and the calibration should use the same spacing between wafer and reflector as production [18]. A film on the measured side of a wafer changes its emissivity. Such wafers are screened out by hand, and a 2005 study asked for a systematic alarm on the tool instead [17].

In such a controller the structure is five PI blocks with back-calculation, a five-by-five matrix block for the decoupler, a feedforward block that reads the recipe, the sequence of XY-100 with its ramp limit and hold, and a run-to-run update on the soak.

What Is Modernizing It

The first development is multivariable, model-based control. By 2003 controllers built on physics-based models reduced to low order ran on commercial chambers [6], like the MPC of Section 9.6.9, while two years later many tools still kept one loop on a center pyrometer, the architecture of case 2 [17]. The second is shorter annealing. Flash lamp annealing, a non-equilibrium method on the sub-second time scale, has been used in microelectronics mostly after ion implantation, to activate dopants, to recrystallize amorphous layers and to anneal out defects [25]. Measuring the wafer during a flash is the open challenge of Section 9.8.

9.8 Limitations

The thermal model is a teaching model, with eight annuli that leave the rim's own dip unresolved, gray band-averaged radiation that makes the pyrometer error about three times a narrow-band instrument's at 0.95 µm, a fixed lamp absorptivity, a lumped ring and window, and an axisymmetric chamber with no rotation, bow, gas flow pattern or pattern-dependent absorption [5].

The product and the constraint are illustrative. The budget, the sheet resistance law and the slip limit rank thermal histories and are not models of activation, diffusion, junction depth or dislocation. Metrology is immediate and complete, where a fab measures later and on sampled wafers, so the soak correction of Equation (9.14) converges faster than a production loop would. Every size, rating and limit is a teaching value and not a tool's specification.

The MPC's model leaves out radial conduction, the edge ring, the window, the pyrometers' lag and aged lamps, so its gains are 15 to 26 % high at 600 °C and the disturbance estimate carries the rest. Its slope and heat capacity are taken at the reference rather than at the estimate, and the slope comes from the emissivity the pyrometers are told, so it inherits the emissivity error of case 4 unchanged. Its objective weighs the profile against the setpoint and not the thermal budget, which is why a plan that respects the lamps' floor loses budget at every cut, and an objective written on the budget itself would need the Arrhenius term that the plan does not carry.

> **OPEN CHALLENGE** **Measuring a Wafer That Is Hot for Milliseconds**
>
> Annealing has moved below the spike, to flash lamps and lasers that heat and cool a wafer within milliseconds. Real-time temperature control within a wafer and from wafer to wafer is one of the key challenges such tools must meet in production [26]. One prototype hid the flash behind a layer of flowing water under the lamps, read the wafer at 1.95 µm with a thermometer that compensates emissivity from two polarizations, and found its peak reading tracking the sheet resistance whatever the surface emissivity [27]. Progress would be such a reading on every product wafer, an independent sensor.

One limit is a property of the application, not of the model. The choice between a lot and one held wafer belongs to each case and is shown read-only, because a held wafer and a lot start from different chamber states, so a student who wants the held wafer loads case 1. The lot size applies at once, never below the wafers already annealed.

9.9 Conclusion

This chapter studied a spike anneal chamber, a browser plant whose loops read pyrometers while its product is computed from the true temperatures. The open-loop time constant falls from 13.7 to 2.4 s between 400 and 1050 °C, yet loops that ask for a heating rate answer a 10 K step in 0.33 s at all four temperatures tried, so the cubic schedule is offered and not needed. Feedback alone lags the 250 °C/s ramp by 65.2 K and overshoots by 29.9 K, and the inverse-model feedforward brings the lag to 2.9 K. Five independent loops pull the radial difference to 34.8 K, which the decoupler holds to 2.1 K. A 10 % emissivity error leaves the wafer 31.4 K below its pyrometers and 26.3 % high in sheet resistance until a table or a calibration wafer corrects it, and the graded lot improves from none of 25 wafers in specification to all 25. A five-zone MPC holds the 400 °C/s ramp of case 9 to 10.9 K without slip where the decoupled loops reach 17.3 K and slip, but it leaves the wafer 3.3 % high, and the supervisor's slower ramp makes a better wafer than either.

The model overstates the pyrometer error and uses illustrative product laws. A fuller model would add wafer rotation, spectral emissivity with a narrow-band pyrometer and a sampled metrology queue. The chamber is an accessible plant for learning that a loop on its setpoint has measured its sensor, not its product.

KEY POINTS

- A loop that holds its pyrometer on the recipe has measured its sensor, not the wafer it anneals.
- Feedback alone lags the 250 °C/s ramp by 65.2 K and overshoots by 29.9 K, and the inverse-model feedforward brings the lag to 2.9 K.
- A 10 % emissivity error leaves the wafer 31.4 K below its pyrometers and 26.3 % high in sheet resistance.
- A five-zone MPC holds the 400 °C/s ramp without slip where the decoupled loops slip, but the supervisor's slower ramp makes a better wafer than either.

Where else this applies. A pyrometer that sees only the surface in Chapter 2, coupled heating zones in Chapter 5, and a product temperature no instrument reaches in Volume 5, Chapter 11.

References

[1] R. Singh, "Rapid isothermal processing," *Journal of Applied Physics*, vol. 63, no. 8, pp. R59–R114, 1988.

[2] F. Roozeboom and N. Parekh, "Rapid thermal processing systems: A review with emphasis on temperature control," *Journal of Vacuum Science & Technology B: Microelectronics Processing and Phenomena*, vol. 8, no. 6, pp. 1249–1259, 1990.

[3] A. Agarwal, A. T. Fiory, H.-J. L. Gossmann, C. S. Rafferty, and P. Frisella, "Ultra-shallow junction formation by spike annealing in a lamp-based or hot-walled rapid thermal annealing system: effect of ramp-up rate," *Materials Science in Semiconductor Processing*, vol. 1, no. 3–4, pp. 237–241, 1998.

[4] P. J. Timans, "Emissivity of silicon at elevated temperatures," *Journal of Applied Physics*, vol. 74, no. 10, pp. 6353–6364, 1993.

[5] J. P. Hebb and K. F. Jensen, "The effect of multilayer patterns on temperature uniformity during rapid thermal processing," *Journal of The Electrochemical Society*, vol. 143, no. 3, pp. 1142–1151, 1996.

[6] A. Emami-Naeini, J. L. Ebert, D. de Roover, R. L. Kosut, M. Dettori, L. M. L. Porter, and S. Ghosal, "Modeling and control of distributed thermal systems," *IEEE Transactions on Control Systems Technology*, vol. 11, no. 5, pp. 668–683, 2003.

[7] C. D. Schaper, Y. M. Cho, and T. Kailath, "Low-order modeling and dynamic characterization of rapid thermal processing," *Applied Physics A Solids and Surfaces*, vol. 54, no. 4, pp. 317–326, 1992.

[8] Y. M. Cho and P. Gyugyi, "Control of rapid thermal processing: a system theoretic approach," *IEEE Transactions on Control Systems Technology*, vol. 5, no. 6, pp. 644–653, 1997.

[9] K. S. Balakrishnan and T. F. Edgar, "Model-based control in rapid thermal processing," *Thin Solid Films*, vol. 365, no. 2, pp. 322–333, 2000.

[10] T. F. Edgar, S. W. Butler, W. J. Campbell, C. Pfeiffer, C. Bode, S. B. Hwang, K. S. Balakrishnan, and J. Hahn, "Automatic control in microelectronics manufacturing: Practices, challenges, and possibilities," *Automatica*, vol. 36, no. 11, pp. 1567–1603, 2000.

[11] E. Sachs, A. Hu, and A. Ingolfsson, "Run by run process control: combining SPC and feedback control," *IEEE Transactions on Semiconductor Manufacturing*, vol. 8, no. 1, pp. 26–43, 1995.

[12] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[13] Y. Qu, "Silicon wafer surface temperature measurement using light-pipe radiation thermometers in rapid thermal processing systems," Ph.D. dissertation, The University of Texas at Austin, Austin, TX, 2006. [Online]. Available: https://hdl.handle.net/2152/2794

[14] G. G. Bentini, L. Correra, and C. Donolato, "Defects introduced in silicon wafers during rapid isothermal annealing: Thermoelastic and thermoplastic effects," *Journal of Applied Physics*, vol. 56, no. 10, pp. 2922–2929, 1984.

[15] H. A. Lord, "Thermal and stress analysis of semiconductor wafers in a rapid thermal processing oven," *IEEE Transactions on Semiconductor Manufacturing*, vol. 1, no. 3, pp. 105–114, 1988.

[16] T. Breedijk, T. F. Edgar, and I. Trachtenberg, "A model predictive controller for multivariable temperature control in rapid thermal processing," in *Proceedings of the 1993 American Control Conference*. IEEE, 1993, pp. 2980–2984.

[17] W. Cho, "Temperature control and modeling of rapid thermal processing," Ph.D. dissertation, The University of Texas at Austin, Austin, TX, 2005. [Online]. Available: https://hdl.handle.net/2152/1527

[18] C. W. Meyer, D. P. DeWitt, K. G. Kreider, F. J. Lovas, and B. K. Tsai, "ITS-90 calibration of radiation thermometers for RTP using wire/thin-film thermocouples on a wafer," *AIP Conference Proceedings*, vol. 550, pp. 254–258, 2001.

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

[20] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

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

[22] K. G. Kreider, D. H. Chen, D. P. DeWitt, W. A. Kimes, and B. K. Tsai, "Lightpipe proximity effects on Si wafer temperature in rapid thermal processing tools," in *Proc. 11th IEEE Int. Conf. Advanced Thermal Processing of Semiconductors*. IEEE, 2003, pp. 125–129.

[23] T. R. Kuphaldt, *Lessons In Electric Circuits, Volume III: Semiconductors*. Open textbook (CC BY), ibiblio.org, 2023, chapter 7, Thyristors.

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

[25] L. Rebohle, S. Prucnal, Y. Berencén, V. Begeza, and S. Zhou, "A snapshot review on flash lamp annealing of semiconductor materials," *MRS Advances*, vol. 7, no. 36, pp. 1301–1309, 2022.

[26] S. Chen, J. Hebb, A. Jain, S. Shetty, and Y. Wang, "Wafer temperature measurement and control during laser spike annealing," in *Proc. 15th Int. Conf. Advanced Thermal Processing of Semiconductors*. IEEE, 2007, pp. 239–244.

[27] Y. Yamada, T. Aoyama, H. Chino, K. Hiraka, J. Ishii, S. Kadoya, S. Kato, H. Kiyama, H. Kondo, T. Kuroiwa, K. Matsuo, T. Owada, T. Shimizu, and T. Yokomori, "In situ Si wafer surface temperature measurement during flash lamp annealing," *Japanese Journal of Applied Physics*, vol. 49, no. 4S, p. 04DA20, 2010.

9.10 Exercises

9.1 ○ Concept check · 2 min

The open-loop time constant of the wafer falls from 13.7 s at 400 °C to 2.4 s at 1050 °C. Yet a 10 K setpoint step on the held wafer of case 1 in automatic reaches 63 % in 0.33 s at all four temperatures tried. Why does the response not change with temperature?

(a) The pyrometer lag of 5 ms sets the response at every temperature.
(b) The loops ask for a heating rate, and their gain dwarfs the wafer's radiative rate.
(c) The cubic gain schedule is on in case 1 and cancels the change of the conductance.
(d) The heat capacity rises with temperature about as fast as the conductance does.

9.2 ○ Concept check · 1 min

In case 4 the pyrometers assume an emissivity of 0.70 on nitride-capped wafers whose emissivity is 0.77. On a real tool, which signal reveals the error?

(a) The center pyrometer's trend, which runs above its setpoint through the soak.
(b) The difference between the pyrometers, which grows on the ramp.
(c) The sheet resistance measured on the wafer after it leaves the chamber.
(d) The lamp commands, which reach their ratings during the ramp.

9.3 ○ Concept check · 2 min

In case 3 the static decoupler alone cuts the largest radial difference on the ramp from 34.8 K to 2.1 K, yet the first wafer's sheet resistance comes out 14.7 % low. What explains the low sheet resistance?

(a) The decoupler drives zone 5 to its rating, which leaves the rim cold through the soak.
(b) The decoupler inverts the gain matrix of 1050 °C only, so it is wrong on the ramp.
(c) The decoupler slows every loop, so the wafer never reaches the recipe peak.
(d) The wafer still lags the ramp by about 65 K as one piece and then overshoots.

9.4 Theory and calculation · 6 min

The pyrometers convert radiance with an assumed emissivity ε_{set}, and in the band-averaged model a loop that holds the reading on T_{SP} holds the wafer at $T = T_{SP}(\varepsilon_{set}/\varepsilon)^{1/4}$ in kelvin (Equation (1.1)). The setpoint is 1050 °C and $\varepsilon_{set} = 0.70$. Give temperatures to 0.1 K.

(a) Find the true wafer temperature for a nitride cap with $\varepsilon = 0.77$, and how far it lies below the setpoint. *Checkpoint:* 1018.8 °C, 31.2 K low

(b) Find the true wafer temperature for an oxide stack with $\varepsilon = 0.63$. *Checkpoint:* 1085.3 °C

(c) A narrow-band pyrometer at $\lambda = 0.95\,\mu\text{m}$ follows Wien's law, $\Delta T/T = (\lambda T/c_2)\,\Delta\varepsilon/\varepsilon$ with $c_2 = 14\,388\,\mu\text{m K}$. Find its relative and absolute error at 1050 °C for $\Delta\varepsilon/\varepsilon = 0.10$, and the ratio of that error to the one of part (a), to two decimals. *Checkpoint:* 0.87 %, 11.6 K, ratio 0.37

9.5 Theory and calculation · 6 min

The recipe ramps the wafer from 600 to 1050 °C at 250 °C/s. Near the peak the wafer's heat capacity is 126.0 J/K and its conductance K is 53.1 W/K. Each zone loop outputs a heating rate $v = K_c e$ plus its integral, with $K_c = 4\,\text{K/s}$ per kelvin.

(a) Find the duration of the ramp and the absorbed power the ramp alone asks for near the peak. *Checkpoint:* 1.8 s, 31.5 kW

(b) Ignore the integral action, the losses and the lamp lag. What error must a proportional loop carry to ask for 250 °C/s? Compare it with the center lag of 65.2 K that case 2 measures without feedforward, and say why the integral cannot remove the lag on this ramp ($\tau_I = 1.5\,\text{s}$). *Checkpoint:* 62.5 K

(c) Find the open-loop time constant $m_w c_p/K$ at 1050 °C, to two decimals, and compare it with the ramp duration. *Checkpoint:* 2.37 s

9.6 Theory and calculation · 8 min

The budget of Equation (9.7) has $E_a = 3.5\,\text{eV}$, $k_B = 8.617 \times 10^{-5}\,\text{eV/K}$, $m = 0.30$ and $\theta_T = 1.3365\,\text{s}$. The soak correction of Equation (9.14) uses $\lambda = 0.5$.

(a) At $T = 1323.15\,\text{K}$, find the relative change of the budget rate per kelvin, $E_a/(k_B T^2)$, the resulting change of R_s per kelvin of peak, both in percent to two decimals, and the peak error that matches the $\pm 2\,\%$ band, to 0.1 K. *Checkpoint:* 2.32 %, 0.70 %, 2.9 K

(b) The probe reads a wafer 2.87 % high. By how much does the correction lengthen the next soak, in seconds to three decimals? Is the correction direct or reverse acting in the convention of Section 1.6.2? *Checkpoint:* 0.063 s, direct acting

(c) One more second of soak changes $\ln R_s$ by $-m/\theta_T$, so the correction is a loop $e_{k+1} = e_k - K e_{k-d}$ on $e = \ln(R_s/R_{s,T})$, with d the wafers by which the probe result is late (Chapter 1's metrology delay, Equation (1.6)). Find K. Using Equation (1.7), is the correction stable if the result arrives two wafers late, and if it arrives three wafers late? *Checkpoint:* $K = 0.5$; stable at two, unstable at three

9.7 App challenge · 8 min

Load scenario 2, "One loop on the center", default seed, at speed 5. Run to $t = 61\,\text{s}$, the end of the first wafer.

(a) Read the first wafer's true mean sheet resistance on the R_s last badge. Export the CSV and find the largest value of `sp1_C` minus `pyro1_C` over the rows whose `phase` is `ramp`. Give the sheet resistance to 0.1 Ω per square and the lag within 0.5 K. *Checkpoint:* 355.3 Ω/sq, 64.9 K

(b) Press Reset, set Model feedforward On on the TY-100 card before pressing Run, run to 61 s again and read the same two values, the lag within 0.3 K. *Checkpoint:* 404.7 Ω/sq, 2.6 K

(c) The wafer of part (a) lagged its setpoint on the ramp, yet its sheet resistance is low. Explain why, from the trend of the center pyrometer at the top of the ramp. *Checkpoint:* an overshoot of about 30 K over-anneals the wafer

9.8 App challenge · 8 min

Load scenario 4, "Emissivity wrong by 10 %", default seed, at speed 5. Run to $t = 61$ s and export the CSV.

(a) For the rows with `wafer` equal to 1, find the largest `T_ann0_C` (the true center annulus) within 0.1 K and the largest `pyro1_C` within 0.3 K. Read the sheet resistance on the R_s last badge to 0.1 Ω per square. *Checkpoint:* 1016.9 °C, 1048.2 °C, 505.2 Ω/sq

(b) Apply $T = T_{read}(0.70/0.77)^{1/4}$ in kelvin to the largest pyrometer reading of part (a) and compare the result with the largest `T_ann0_C`. Then estimate the sheet resistance error from 0.70 % per kelvin of peak and a peak 31.2 K low, and say why the estimate falls short of the measured error. *Checkpoint:* 1017.1 °C; about 22 % against 26.3 %

(c) Press Reset, set Pyrometer emissivity to Film-stack table before pressing Run, run to 61 s and read the badge again. *Checkpoint:* 402.3 Ω/sq

9.9 App challenge · 10 min

Load the challenge, "A lot of 25 wafers", default seed, at speed 20 (the 1500 s lot takes about 75 s). The run stops by itself when the lot is complete.

(a) Run the lot as shipped. Read the In spec, Worst R_s error, Worst peak error and Slip flags tiles. *Checkpoint:* 0, 15.34 %, 27.8 K, 23

(b) Press Reset. Before Run set Pyrometer emissivity to Film-stack table, Model feedforward On and, in Advanced, Static decoupler On. Run the lot and read the same four tiles and the Lamp energy tile. *Checkpoint:* 25, 1.23 %, 2.3 K, 0, 3.4 kW h

(c) Which of the three changes carries wafers 10 to 15, and why can the other two not replace it? Use Table 9.6. *Checkpoint:* the film-stack table

9.10 Application area · 3 min

In case 6, zone 5 delivers 35 % of its output and the supervisor holds the lot after the first wafer. A shift lead proposes to keep the tool running by slowing the recipe to 50 °C/s with the slip-aware limit and the hold off. Using the chapter's measurements, what do you advise, and why is holding the lot the right control action?

9.11 Application area · 3 min

A fab introduces a new nitride-capped product. Compare the film-stack table and a calibration wafer as sources of the pyrometers' emissivity: what does each cost, what does each rely on, and when does each go out of date?

9.12 Review the agent · 5 min

An AI assistant reviews case 9 and writes: "The MPC holds the 400 °C/s ramp to a radial difference of 10.9 K, 0.79 of the slip limit, where the decoupled zone loops slip at 1.77 times it. The MPC is therefore the better controller for this chamber. Switch the challenge lot to Zone control MPC with the film-stack table and every wafer will be in specification." Find the flaw and say how you would check the claim in the app.