INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 14 Control in Semiconductor Manufacturing

CHAPTER 10

Copper Plating Bath Simulator Dosed by Amp-Hours and Analyzed Every Two Hours

Pages 240–264

Edited by John D. Hedengren

How to cite this chapter. "Copper Plating Bath Simulator Dosed by Amp-Hours and Analyzed Every Two Hours," in *Control in Semiconductor Manufacturing*, John D. Hedengren, Ed., vol. 14 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 10, pp. 240–264.

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

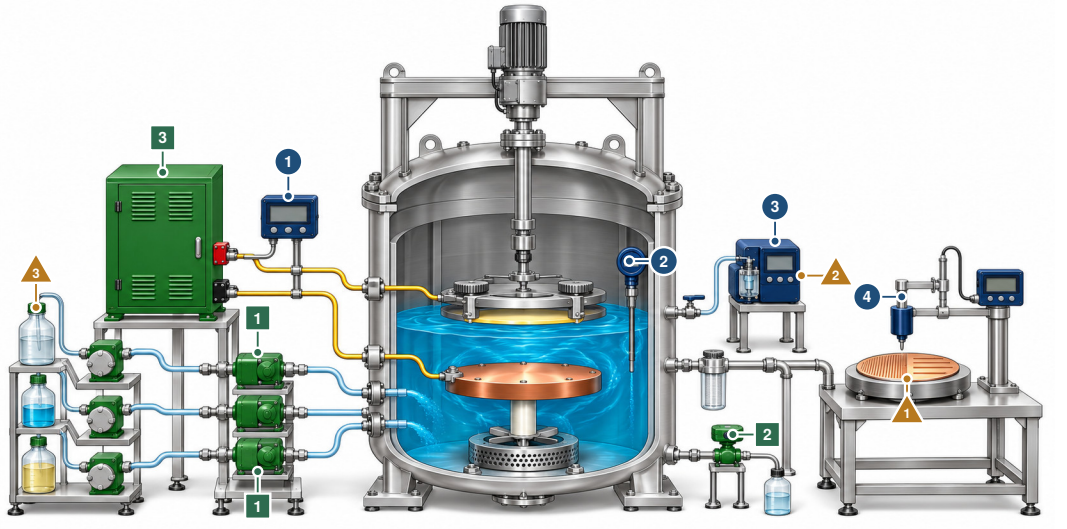
The app for this chapter
https://apmonitor.com/apps/copper-ecd-bath/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v14-semiconductor/#copper-ecd-bath

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

CHAPTER 10

Copper Plating Bath Simulator Dosed by Amp-Hours and Analyzed Every Two Hours



The controller moves 1 three pumps that dose the accelerator, suppressor and leveler
2 a bleed that drains bath to waste, refilled with fresh solution 3 the plating time of each wafer, run by the power supply

The controller reads 1 a counter of the charge the power supply delivers 2 a thermometer in the bath
3 an analyzer that tests a bath sample every two hours 4 a probe that measures each plated wafer's thickness

What disturbs it 1 lots of thick-copper wafers that take twice the charge 2 the analyzer out of service
3 a new lot of additive concentrate that the bath uses faster

Notice that the pumps can follow the charge counter at once, while the analyzer at the right answers every two hours about a sample more than forty minutes old. Illustration, not to scale.

IN PLAIN TERMS **Feeding a Copper Bath by Its Charge**

The main idea. Every copper wire in a chip is grown in a bath of copper and acid, one silicon disk at a time. Three trace chemicals make the copper fill each hole from the bottom up, and the electric charge that plates each disk uses them up.

The challenge. The bath is tested every two hours, and each test shows it as it was over forty minutes before. Worn-out chemicals also build up and shrink the mix in which holes fill, and only fresh solution can dilute them.

Why it matters. Dosing by the charge meets a change of product at once, which a slow test cannot.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 10.2, the fact boxes of the Introduction and of Section 10.7, the try-it boxes of Section 10.7, and the key points
Study	50 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 10.7, Section 10.8, Section 10.9 with its open challenge

The control question A copper plating bath consumes its additives with every amp-hour, while its only analysis arrives every two hours and describes the bath of 42 minutes before. How much of the dosing can such an analysis carry, and what must the amp-hour counter and the bleed do instead?

Learning objectives

- Compute a wafer's charge, bulk step and thickness from Faraday's law, and the additive dose that holds the bath at a given load.
- Explain why a feedforward on the amp-hours answers a lot change at once while a trim on the analysis first sees it hours later.
- Use the report-to-report model to explain why the sampling and the analysis age, not the chemistry, set the gain limit of the trim.
- Price a setpoint of the breakdown products in virgin make-up with the make-up balance, and decide between bleed-and-feed and a dump.
- Judge when an analysis should be rejected against the controller's own model of the bath.

Prerequisites Sections 1.4.3 (run-to-run correction), 1.4.4 (late measurement), 1.4.7 (feedforward), 1.6.2 (velocity-form PI) and 1.6.4 (scenarios). Chapter 6 is a useful first look at a correction made on late results.

The app https://apmonitor.com/apps/copper-ecd-bath/

See also Volume 5, Chapter 3 and Volume 5, Chapter 4 for loops behind slow or late analyzers. Volume 3, Chapter 6 for a dose computed from the load with a slow quality trim. Volume 8, Chapter 9 for corrections made while earlier results are still in flight.

10.1 Introduction

Every copper wire and via of a modern chip is plated. Full copper wiring in a CMOS technology was reported in 1997 [1], and in the damascene process that followed each metal level is etched into the dielectric, lined with a seed layer and filled by electroplating from an acid copper sulfate bath [2]. The features fill without a void only because copper grows faster at their bottoms than at their walls and mouths, and that bottom-up growth is the work of organic additives. A suppressor adsorbs with chloride and inhibits the deposit [3], an accelerator displaces it and gathers at the bottom of a filling feature as the feature's area shrinks [4], and a leveler competes with the accelerator where the current crowds [5]. The filling of trenches and vias in the presence of additives has its own theory [6], the accelerator decomposes in these chemistries [7], and the reliability of finished copper interconnects has been reviewed at length [8].

> ✦ DID YOU KNOW? **Six Levels of Copper in 1997**
>
> The first fully integrated copper wiring of a chip technology, reported in 1997, stacked up to six levels of copper lines. The finest were laid at a pitch of 0.63 µm, above transistors of a technology below 0.25 µm [1].

What the simulator sets out to correct is the belief that a bath whose analyses sit on their setpoints will fill its vias. The additives are analyzed by cyclic voltammetric stripping (CVS) [9], and the application's analyzer takes its sample once every 2 h from a line that holds 2 min of bath and publishes the result 40 min after the draw, so every result describes the bath as it stood at least 42 min before it arrived. In the lot change of case 4, dosed on feedback alone, the result reported at 4:40 still reads 5.96 mL/L, because its sample left the bath just before the thick-copper wafers started at 4:00. The first result to show the doubled consumption arrives at 6:40, from a sample drawn when the accelerator was already 6.5 % low, and 97 of the day's 400 wafers void. The analysis also says nothing about the breakdown products, which set the width of the window the accelerator-to-suppressor ratio must stay inside. In case 2 the dosing holds the accelerator at 5.84 mL/L and the suppressor at 2.000 mL/L all day while the breakdown products climb from 27 to 31.4 mL/L and 234 wafers void.

A transfer function from a dosing pump to an analyzer reading carries neither the age of the reading nor the inventory that closes the window.

Pause. Try Exercise 10.1 before reading on.

The class of problem is general. A composition analyzed seldom and late is also the situation of a fermentation whose key concentrations come from delayed offline assays [10], and of any inferential loop built on slow, multi-rate primary measurements [11]. A species that no actuator removes except a purge is the situation of the inerts and byproducts of a plant with recycle [12]. A sparse analysis that a controller will act on for hours has to be screened for gross errors against what the process should have produced [13], and a recipe corrected between runs by an exponentially weighted average (EWMA) of its errors is the run-to-run controller of semiconductor manufacturing [14], the subject of Section 1.4.3 and of Chapter 6. A reagent consumed in proportion to a measured load is dosed forward in ratio to it, which is textbook practice [15]. The bath puts all of these in one vessel.

However, the electrochemistry literature treats the additives as a mechanism of fill [4, 5] and their analysis as a method [9], and the control literature treats sampled analyses, purges and run-to-run correction as separate structures [11, 12, 14], so neither offers a bath on which a result can be watched arriving 42 minutes late while a window it does not report closes. This chapter works out what a sparse and late analysis can and cannot do for a copper plating bath, on a browser simulator in which every wafer is plated by an exact charge, the analyzer and the laboratory report late and can be wrong, and each wafer's true fill is recorded separately from anything an instrument says. Faraday's law, amp-hour feedforward, the PI trim, the EWMA filters, the gross-error gate and bleed-and-feed are all established practice. Eight cases and a graded day separate a load from a model error, a late result from a wrong one, and a dose from an inventory.

Section 10.2 gives the history of the domain. Section 10.3 sets out the bath as modeled: a 39-state model of seven species, the bath temperature and 31 cumulative ledgers that close to 3×10^{-12}, with discrete wafers, a truth log per wafer, and the framework's shared sampled analyzer and bias update, in the detail an independent rebuild needs. Section 10.4 derives the make-up balance, which divides the breakdown products removed by the concentration at which they leave and shows that in a bath whose byproducts accumulate linearly a dump cycle cannot cost less than bleed-and-feed. Section 10.5 describes the controller. Section 10.6 builds a four-state report-to-report model of the sampled trim, whose spectral radius reaches one at a gain of 210 mL/h per mL/L and whose decay per analysis and period agree with the noise-free engine within 0.003 and 0.9 %. Section 10.7 reports the measured cases, Section 10.8 describes the hardware such a bath runs on, and Sections 10.9 and 10.10 state the limitations and conclusions.

10.2 Brighteners, Stripping Analysis and Damascene Fill

The additives were measured electrochemically two decades before their part in filling a feature was modeled, by the stripping method whose beginnings the history box below tells [9, 16]. The simulated analyzer stands for an instrument of that kind, a sample analyzed in tens of minutes on a schedule of hours. The chemistry behind the signal was mapped in the years that followed. Polyethylene glycol was shown to inhibit copper deposition in the presence of chloride ion [3], the pairing that serves as the suppressor of the damascene bath [5].

> **HISTORY** **Measuring a Brightener by Stripping Copper (1978 to 1985)**
>
> In 1978 a voltammetric stripping method was published for the brighteners and levelers of plating baths. The potential of an inert rotating electrode is cycled in the bath, so that copper is plated on it and stripped off again. The charge needed to strip the copper follows the rate of plating, and so the concentration of the additive. In copper pyrophosphate baths the method reached a precision of about 0.1 parts per million over a range from 0 to 2.5 [16]. A pulse version of 1985 found the absolute concentration of an additive in acid copper baths, fresh and in production, that carried contaminants from circuit board plating [9].

Table 10.1 Variables and States of the Bath, With Setpoints and With the Values Case 4 Begins From

Variable	Role	Design value	Range
S_{CVS}, P_{CVS}, L_{CVS}	CVs, AIC-S, P, L	6.0, 2.0, 3.0 mL/L	3–10, 1–4, 1–6
$\hat{B}$	CV, FIC-BLEED	22 mL/L	4–30
h	CV, R2R	1.00 µm	0.6–1.3
Cu_{lab}, H_{lab}	CVs, AIC-CU	40, 10 g/L	30–50, 5–20
m	quality, true	0.267	void below 0
D_S, D_P, D_L	MVs, dosing pumps	57.4, 2.04, 7.20 mL/h	$\leq$ 150, 12, 40
b	MV, bleed	1.76 L/h	0–4
t_b	MV, bulk step	125.2 s	60–146
w, $\dot{m}_{Cu}$, $\dot{m}_H$	MVs, titration	0	—
thick share	DV, lot mix	0 %	0–100
ΔT_{ch}	DV, chiller drift	0 °C	−3 to 12
$\Delta\eta_a$, $\Delta\eta_c$	DVs, efficiencies	0 %	−5–10, −8–2
CVS down, idle	DVs, switches	off	on, off
outlier	DV, S_{CVS} event	—	+3.0 mL/L
overdose	DV, L event	—	+300 mL
S, P, L, B	states	6.0, 2.0, 3.0, 22 mL/L	—
Cl	state	50 mg/L	—
Cu, H, T	states	40, 10 g/L, 25 °C	—
$z_1 \ldots z_{31}$	states, ledgers	0	—

Damascene wiring gave the additives a task with no tolerance for a void, and the theory of how they fill a feature followed within a few years, as the history box after this paragraph tells [1, 2, 4, 6]. The roles of suppressor, accelerator and leveler in damascene plating were reviewed [5]. The decomposition of the sulfonate accelerator in acid copper chemistries was measured directly [7], and it is the origin of the breakdown products this model carries. A leveling agent also governs the cathode quality of copper electrorefining, which Volume 15, Chapter 8 controls at an inferred limit.

> ⧖ HISTORY **Superfilling and Its Theory (the 1990s to 2005)**
>
> Copper electroplating for on-chip wiring was conceived and developed in the early 1990s. The task was to fill trenches and vias without a void or a seam, and plating turned out to fill them far better than a coating of even thickness could. This superconformal growth, called superfilling, depends on the plating additives [2]. In 2000 simulations showed that a poorly controlled additive concentration can leave a larger void than a bath with no additive at all [6]. By 2005 a model in which the local growth rate follows the accelerator's coverage explained the bottom-up fill: the accelerator gathers on the concave bottom of a feature as it advances [4].

The control problem took its present shape with the production bath. A bath that serves a plating tool for a day of wafers loses additive with every amp-hour and gains byproduct with every molecule it loses, while its analysis stays as sparse as the stripping method allows, and the tools of the introduction, from delayed assays to purges, were developed for other processes and carry over. The simulator places them in one bath and adds what no fab records, the true fill margin of every wafer beside the analyses its controllers saw.

10.3 The Bath as Modeled

The simulated tool plates one 300 mm wafer every 216 s, 400 a day, from a 200 L acid copper bath fed by three dosing pumps and by a drum of virgin make-up solution (VMS), a complete fresh bath at the make-up strength of every species. The bath, the plating cell, the bleed, the instruments and every symbol in the equations appear in Figure 10.1.

Table 10.1 gives each variable's role, its value when case 4 starts and its range. Of the 39 states, eight describe the bath, its seven species and its temperature, and 31 are cumulative ledgers of the charge and of what each species has had dosed, consumed, broken down, bled and fed. Integrating the ledgers with the chemistry in one step lets any balance be checked from the state vector alone.

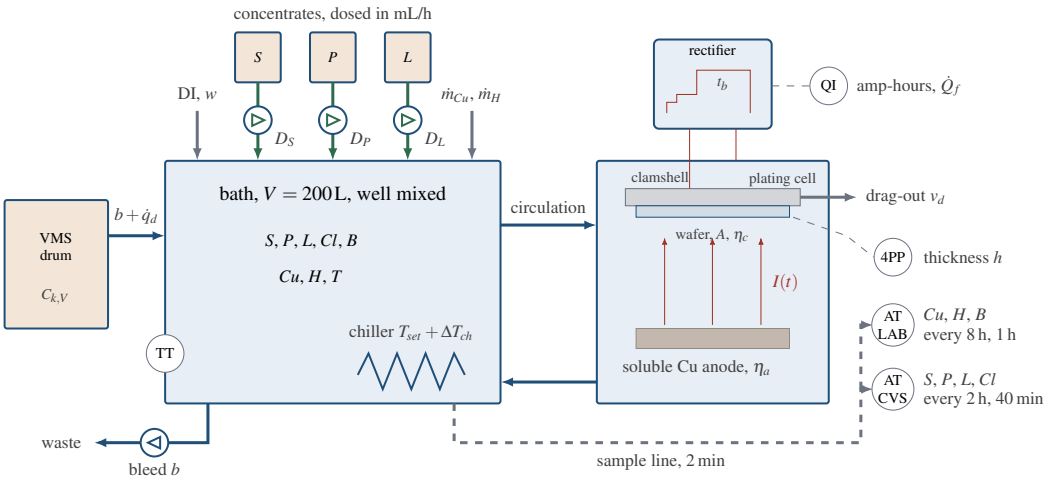


Figure 10.1 The plating bath and the tool it feeds, as modeled. Three pumps dose the accelerator S, the suppressor P and the leveler L as concentrates at D_S, D_P and D_L into a well-mixed bath of volume V, held at $T_{set}+\Delta T_{ch}$ by a chiller and read by TT. The bleed b drains bath to waste, and the VMS drum replaces it and the drag-out v_d of each wafer at the make-up strengths $C_{k,V}$. DI water w and the copper and acid additions $\dot{m}_{Cu}$ and $\dot{m}_H$ are the titration corrections. The plating cell holds the wafer of area A face down in its clamshell over a soluble copper anode, the rectifier drives the three-step current $I(t)$ with a bulk step of length t_b, and the amp-hour counter QI gives the load $\dot{Q}_f$. The CVS analyzer and the laboratory draw from the bath through a 2 min sample line, and the four-point probe 4PP measures each wafer's thickness h. Dashed lines are measurements.

Concentrations are in the units of the tool's panels, mL/L of vendor concentrate for the three additives and their breakdown products, mg/L for chloride and g/L for copper and sulfuric acid, and the engine's time unit is the minute although rates are quoted per hour.

10.3.1 Additives, Chloride and Breakdown Products

Each additive $k \in \{S,P,L\}$ is dosed, consumed at the electrodes in proportion to the charge and to the concentration the surface sees, broken down in the bulk, bled out and fed back with the VMS, which Equation (10.1) writes for the whole bath.

$$\begin{aligned} V\frac{\mathrm{d}C_k}{\mathrm{d}t} &= D_k - c_k I\frac{C_k}{C_{k,V}} - d_k(T)\,V C_k \\ &\quad -(b+w)\,C_k + b\,C_{k,V} \\ d_k(T) &= d_{k,25}\,\mathrm{e}^{E_k(T-25)} \end{aligned} \tag{10.1}$$

The concentration C_k is in mL/L and the dose D_k in mL/h, I is the rectifier current in A, c_k the consumption in mL per A h at the make-up strength $C_{k,V}$, $d_{k,25}$ the breakdown rate at 25 °C in h^{-1} with temperature coefficient E_k in K^{-1} and T in °C, while b is the bleed and w the titration's DI dilution, both in L/h. Chloride is not dosed. It enters only with the VMS and is taken into the deposit at $c_{Cl} = 0.5$ mg per A h, so that $V\,\mathrm{d}Cl/\mathrm{d}t = -c_{Cl}\,I\,Cl/Cl_V - (b+w)\,Cl + b\,Cl_V$. What is consumed or broken down and not built into the deposit stays in the bath as breakdown products, counted in mL/L of the additive they came from, as shown in Equation (10.2).

$$V\frac{\mathrm{d}B}{\mathrm{d}t} = \sum_k y_k\left[c_k I\frac{C_k}{C_{k,V}} + d_k(T)\,V C_k\right] - (b+w)\,B \tag{10.2}$$

The yields are $y_S = 0.8$, $y_P = 1.0$ and $y_L = 0.5$, and no term removes B except bath taken out. At the standard load, a cycle-average current of 9.09 A, the accelerator's 1200 mL are consumed at 45.4 mL/h and broken down at 12.0 mL/h, a turnover of 20.9 h. The suppressor turns over in 8.2 days and the leveler in 3.5, and the bath forms 0.258 mL/L of breakdown products an hour, or 0.463 mL/L under a lot of thick-copper wafers.

10.3.2 Copper, Acid, Temperature and the Exchanges

The soluble anode dissolves copper at the efficiency $\eta_a = 100\,\%$ while the cathode plates it at $\eta_c = 98\,\%$, and the balance of the cathode current makes hydrogen and consumes acid, as shown in Equation (10.3).

$$\begin{aligned} V\frac{\mathrm{d}Cu}{\mathrm{d}t} &= (\eta_a - \eta_c)\,\gamma_{Cu} I + \dot{m}_{Cu} - (b+w)\,Cu + b\,Cu_V \\ V\frac{\mathrm{d}H}{\mathrm{d}t} &= -(1-\eta_c)\,\gamma_H I + \dot{m}_H - (b+w)\,H + b\,H_V \\ \tau_T\frac{\mathrm{d}T}{\mathrm{d}t} &= T_{set} + \Delta T_{ch} - T \end{aligned} \tag{10.3}$$

Here $\gamma_{Cu} = 3600\,M_{Cu}/2F = 1.1855\,\mathrm{g}$ of copper per A h, $\gamma_H = 3600\,M_{H_2SO_4}/2F = 1.8297\,\mathrm{g}$ of acid per A h, $\dot{m}_{Cu}$ and $\dot{m}_H$ are the titration additions in g/h, and $\tau_T = 30\,\mathrm{min}$. At the standard load the anode adds 5.2 g of copper a day to the 8 kg the bath holds and the hydrogen consumes 8.0 g of acid, so in no documented day does either move by more than 0.07 g/L. Two exchanges are discrete. Each wafer leaves with $v_d = 30\,\mathrm{mL}$ of bath, which the feed replaces with VMS, $X \leftarrow X + (v_d/V)(X_V - X)$ for every species X with $B_V = 0$, equivalent to $\dot{q}_d = 0.5\,\mathrm{L/h}$ at the 216 s cycle. A dump-and-remake replaces all 200 L at once and holds the tool for $t_{rm} = 120\,\mathrm{min}$. The tool's DI auto-fill replaces evaporation, so V is fixed.

10.3.3 Wafers and Faraday's Law

Each wafer is a discrete event. It enters at a multiple of the cycle $t_{cyc} = 216\,\mathrm{s}$ and runs a three-step waveform, nucleation for $t_{nuc} = 10\,\mathrm{s}$ at $I_{nuc} = 3\,\mathrm{A}$, fill for $t_{fill} = 30\,\mathrm{s}$ at $I_{fill} = 6\,\mathrm{A}$ and bulk for t_b at I_b, which is 14 A for a standard 1 µm wafer and twice that for a 2 µm thick-copper wafer. Its thickness follows from Faraday's law on the charge it received, Equation (10.4).

$$\begin{aligned} Q_w &= t_{nuc} I_{nuc} + t_{fill} I_{fill} + t_b I_b \\ h &= \frac{\eta_c\, Q_w\, M_{Cu}}{2F\,\rho_{Cu} A} = \frac{\eta_c\, Q_w}{3600\,\kappa_F} \end{aligned} \tag{10.4}$$

The charge Q_w is in A s with the times in s, the 300 mm wafer's area is $A = 706.9\,\mathrm{cm}^2$, copper's density is $\rho_{Cu} = 8.96\,\mathrm{g/cm}^3$, and $\kappa_F = 2F\rho_{Cu}A/(3600\,M_{Cu})$ is 0.534 A h per µm. A standard wafer at 98 % takes 0.545 A h, a bulk step of 125.2 s, and a thick one 1.090 A h and 132.7 s, so the cycle-average current doubles from 9.09 to 18.17 A under a thick lot. The current is held over each physics step at its exact average, so every wafer's charge is exact and the fourth-order Runge-Kutta step integrates a smooth right-hand side. A four-point probe reports each wafer's mean thickness with 0.4 % repeatability as it leaves.

10.3.4 The Fill Window

Whether a wafer's vias fill is read from the true chemistry at the middle of its fill step, through the illustrative window of Equation (10.5).

$$
\begin{aligned}
r &= \frac{S}{P\,g(Cl)}, \qquad g(Cl) = \frac{Cl}{K_{Cl}+Cl}\,\frac{K_{Cl}+Cl_V}{Cl_V} \\
r_c &= r_0\Big(\frac{L}{L_V}\Big)^{\kappa_L}, \qquad w(B) = w_0\Big(1-\frac{B}{B_{cl}}\Big) \\
m &= \min\Big[\frac{w(B)-|\ln(r/r_c)|}{w_0},\ \frac{Pg-P_{min}}{P_V-P_{min}}\Big]
\end{aligned}
\tag{10.5}
$$

The window is a band of the effective accelerator-to-suppressor ratio r, in which chloride weights the suppressor, with $g = 1$ at the make-up's 50 mg/L and $K_{Cl} = 20\,\mathrm{mg/L}$. It is centered at $r_0 = 3.0$ at the make-up leveler and moved up by more leveler with $\kappa_L = 0.5$, and its half-width in natural-log units, $w_0 = 0.35$ in a fresh bath, narrows linearly with the breakdown products until it closes at $B_{cl} = 30\,\mathrm{mL/L}$. The margin m is 1 at the center of a fresh window, 0 on its edge and negative outside it, and below $P_{min} = 0.5\,\mathrm{mL/L}$ of effective suppressor the field is not covered at any ratio. A wafer with $m < 0$ voids, as an underfill with a seam when $r < r_c$ and as a pinch-off void when $r > r_c$ or when the suppressor is short. Past B_{cl} the half-width turns negative, so the margin keeps saying how far past closing the bath is. The window, the rates of Equations (10.1) and (10.2) and the yields are illustrative teaching values rather than a supplier's data, and no feature-scale deposition is computed.

10.3.5 Starting Point, Step and Checks

Table 10.2 gives every parameter with its class. Visible parameters carry a slider in the Show Model panel, hidden ones never appear there, and fixed ones are constants of the code. Every scenario starts with every species at its make-up strength and the breakdown products at 22 mL/L, except the aged baths of cases 2 and 5 at 27 and 32 mL/L, and the controllers start at the doses and bleed that hold that bath. In case 4 they are 57.43 mL/h of accelerator, of which 55.16 mL/h is feedforward, 2.035 mL/h of suppressor, 7.197 mL/h of leveler, a bleed of 1.76 L/h and a bulk step of 125.2 s. The engine of Section 1.6.1 integrates the chemistry and the ledgers together by the fourth-order Runge-Kutta method with a constant step of 1 s, the control layer executes once per 10 s, and the historian stores a row per simulated minute. One seed, 20261102, drives all of the noise, so every analysis and probe result repeats exactly from run to run. Speeds from 1 to 1200 simulated minutes per second of wall time, with 10 as the default, change how fast the trajectory is drawn but not the trajectory. At the default speed a day takes 2.4 min of wall time, and without the interface it computes in about 0.07 s. With the step halved, the challenge's virgin make-up moves by 7×10^{-7} L from 57.0332 L and its worst fill margin by 1×10^{-7} from 0.1178, because the controller and the analyses keep their times and only the integration is refined. The species ledgers close to 3×10^{-12} relative in every case, and the application's 40 verification checks include closure below 1×10^{-9}.

10.3.6 Assumptions

The bath is one well-mixed volume of fixed size. A real tool has a reservoir, several cells and piping with mixing times of minutes, which are short against the hours in which this chemistry moves, and the concentrates' own volume is neglected.

The three additives are lumped species in mL/L of vendor concentrate, consumed per amp-hour and broken down at a first-order rate that grows exponentially with temperature. The coefficients are illustrative and chosen so that a day shows the lessons, which is why the accelerator turns over in 20.9 h. The breakdown products are one lumped species with fixed yields, although the fragments of a real accelerator or suppressor can keep part of its activity.

The fill window is an illustrative correlation. A real window is mapped by cross-sectioning filled test structures for one additive package, one waveform and one feature size, and it also depends on

Table 10.2 Parameters of the Bath and Whether Show Model Exposes Them

Symbol	Meaning	Value	Class
V, τ_T	bath volume, thermal lag	200 L, 30 min	hidden
T_{set}	chiller setpoint	25 °C	visible
S_V, P_V, L_V	additives in the VMS	6.0, 2.0, 3.0 mL/L	visible
Cl_V, Cu_V, H_V	chloride, copper, acid in VMS	50 mg/L, 40, 10 g/L	visible
t_{cyc}	wafer cycle, 400 a day	216 s	hidden
t_{nuc}, I_{nuc}	nucleation step	10 s, 3 A	visible
t_{fill}, I_{fill}	fill step	30 s, 6 A	visible
I_b	bulk current, standard wafer	14 A	visible
—	thick-copper factor, t_b limits	2, 60 s to 146 s	hidden
v_d	drag-out per wafer	30 mL	visible
η_a, η_c	anode, cathode efficiency	100, 98 %	visible
t_{rm}	dump-and-remake downtime	120 min	visible
c_k	consumption per A h, S, P, L	5.0, 0.18, 0.66 mL	visible
c_{Cl}	chloride into the deposit	0.5 mg per A h	hidden
$d_{k,25}$	breakdown at 25 °C, S, P, L	1.0, 0.1, 0.2 %/h	visible
E_k	breakdown temperature slope	0.12, 0.07, 0.07 K^{-1}	hidden
y_k	breakdown-product yields	0.8, 1.0, 0.5	hidden
r_0, w_0	window center, half-width	3.0, 0.35	visible
B_{cl}, κ_L	closing level, leveler slope	30 mL/L, 0.5	visible
K_{Cl}, P_{min}	chloride half-strength, least P	20 mg/L, 0.5 mL/L	hidden
$D_{k,max}$	pump capacities S, P, L	150, 12, 40 mL/h	hidden
b_{max}	bleed capacity	4 L/h	visible
—	CVS interval, analysis time	120, 40 min	visible
—	CVS sample line	2 min	fixed
—	CVS repeatability, S	2 % of S_V	visible
—	CVS repeatability, P, L, Cl	2, 3 %, 1 mg/L	hidden
—	chance of a bad analysis	0	hidden
—	laboratory interval	480 min	visible
—	laboratory delay	60 min	hidden
—	laboratory Cu, H, B	0.2, 0.1 g/L, 0.03 mL/L	hidden
—	probe, thermometer	0.4 %, 0.05 °C	hidden
p_V, p_D	VMS price, downtime cost	$20/L, $1500/h	visible
κ_F, γ_{Cu}, γ_H	Faraday constants	0.534, 1.1855, 1.8297	fixed

Repeatabilities are 1σ. κ_F in A h per µm, γ in g per A h.

rotation, aspect ratio and pattern density, none of which is modeled. The via cross-sections in the scene are drawn from the margin rather than computed.

Faraday's law acts on the flat wafer area with one cathode efficiency, with no edge exclusion, radial profile or terminal effect. Copper and acid follow the efficiency gap, the hydrogen side reaction and the make-up, with no chemical dissolution of the anode and no anode film. The CVS analyzer is a sampled measurement with Gaussian repeatability and without the matrix effects of a real stripping analysis, and the laboratory's analysis of the breakdown products stands in for a separate chromatographic method. All sizes, rates, prices and limits are teaching values, not a design basis.

10.4 A Window Narrowed by an Inventory

Two time scales meet in the ratio. The accelerator turns over in 20.9 h at the standard load and the suppressor in 8.2 days, so a dosing error moves the accelerator within hours while the suppressor barely moves. In case 1, where fixed pumps meet a thick-copper lot at 1 h, the accelerator is at 5.19 mL/L by 6 h and the suppressor still at 1.965 mL/L, and the ratio of 2.65 has left the window through its lower edge at 2.70. The ratio follows the accelerator.

The breakdown products set how far the ratio may wander. The half-width of the window is 41.9 % of the ratio in a fresh bath and shrinks to 9.8 % at the 22 mL/L of most cases, then to 7.3, 4.8 and 3.6 % at 24, 26 and 27 mL/L, until the window closes at 30 mL/L, as shown in Figure 10.2. The leveler moves the center rather than the width. The 300 mL overdose of case 8 raises it from 3.0 to 4.5 mL/L, which moves the center from 3.00 to 3.67 and leaves a bath held at the setpoint ratio of

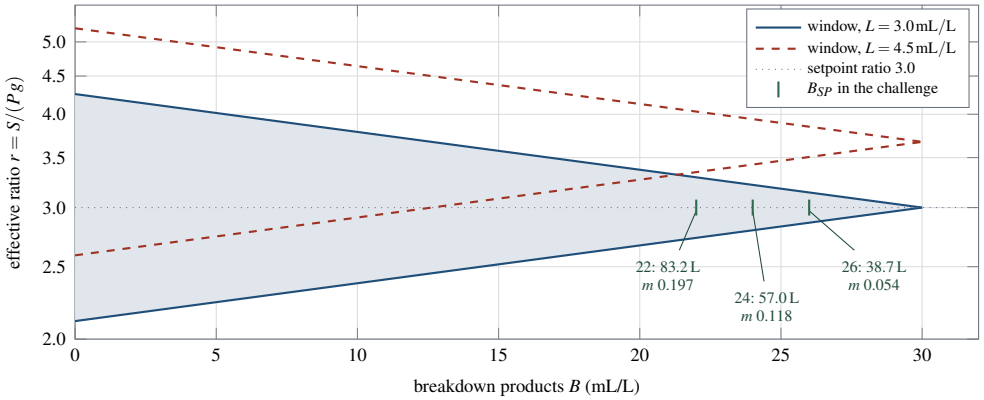


Figure 10.2 Fill window in the effective accelerator-to-suppressor ratio against the breakdown products, Equation (10.5) at the make-up chloride, on a logarithmic ratio axis. The shaded band with solid edges is the window at the make-up leveler, the dashed lines the window after the leveler overdose of case 8, and the dotted line the setpoint ratio. The vertical ticks mark the three setpoints of the breakdown products tried on the challenge day, labeled with the day's virgin make-up and the worst fill margin of its 400 wafers under the suggested dosing.

3.0 below the window at the 22 mL/L of case 8 and at every older age, as the dashed lower edge of Figure 10.2 shows.

Bleed-and-feed is the only exit the breakdown products have, and every liter of bleed is a liter of virgin make-up. Integrating Equation (10.2) over a period with no titration dilution gives the make-up balance of Equation (10.6).

$$V_{VMS} = \int_0^t (b + \dot{q}_d)\,\mathrm{d}t = \frac{\int_0^t GV\,\mathrm{d}t - V[B(t) - B(0)]}{\bar{B}_{out}} \tag{10.6}$$

Here G is the formation rate in mL/L per hour, the right-hand side of Equation (10.2) without its outflow and divided by V, and $\bar{B}_{out}$ is the flow-weighted mean concentration of the bath that left. The numerator is set by the wafers and the bath's age, the denominator by the setpoint. On the challenge day with the suggested dosing the bath forms 1868 mL of breakdown products under every bleed policy. Holding them at 22, 24 and 26 mL/L stores 28, 510 and 897 mL in the bath and removes the rest at mean exit concentrations of 22.12, 23.81 and 25.12 mL/L, which is 83.2, 57.0 and 38.7 L of virgin make-up, while the worst margins of the three days are 0.197, 0.118 and 0.054. The byproducts leave only at the concentration the bath holds them at, so each mL/L of setpoint is make-up saved and window given away.

The same balance settles the choice between bleeding and dumping. A dump cycle at the setpoint B_{SP} replaces the whole volume once every $B_{SP}/\hat{G}$ hours, which is $\hat{G}V/B_{SP}$ per hour, exactly the make-up rate at which bleed-and-feed holds B_{SP}. The application's supervisor prices a day of each policy, $24p_V(b_{hold} + \dot{q}_d)$ for the bleed against $24[(Vp_V + t_{rm}p_D)\hat{G}/B_{SP} + p_V\dot{q}_d]$ for the dump cycle, with b_{hold} the bleed that holds B_{SP}, $p_V = \$20$ per liter and $p_D = \$1500$ per hour of downtime. The dump cycle then carries the bleed's make-up plus the drag-out and the downtime, so it is never the cheaper of the two. Across 69 072 supervisor evaluations in case 2 with the bleed holding its setpoint, spanning downtime costs from 0 to \$10 000 per hour, make-up prices from 1 to \$100 per liter and setpoints from 4 to 30 mL/L, the dump cycle was never cheaper. In a bath whose byproducts accumulate linearly, a dump buys time and a younger bath rather than make-up.

🕮 **Pause.** Try Exercise 10.2 before reading on.

Time is what the supervisor can weigh. The largest bleed, $b_{max} = 4$ L/h, plus the drag-out turns the bath over with $\tau_b = V/(b_{max} + \dot{q}_d) = 44.4$ h toward $B_\infty = \hat{G}V/(b_{max} + \dot{q}_d)$. The time to bring the

Table 10.3 The Loops, Their Actions and Their Shipped Settings

Tag	CV	MV	γ	K_c	τ_I	Output
AIC-S	$\bar{y}_S$	Δ_S	+1	30	6 h	$-F_S$ to $150-F_S$
AIC-P	$\bar{y}_P$	Δ_P	+1	10	24 h	$-F_P$ to $12-F_P$
AIC-L	$\bar{y}_L$	Δ_L	+1	15	12 h	$-F_L$ to $40-F_L$
FIC-BLEED	$\hat{B}$	b	−1	0.5	12 h	0 to 4 L/h
R2R	η_{meas}	t_b	—	—	—	60 to 146 s
AIC-CU	Cu_{lab}, H_{lab}	w, $\dot{m}$	—	—	—	over 30 min

K_c in mL/h per mL/L, and in L/h per mL/L for the bleed. R2R: EWMA, $\lambda_R = 0.3$.
AIC-CU: proportional, $g_T = 0.8$. Gate $g = 0.15$, $\lambda = 0.5$.
Feedforward $\hat{c}_k$ = 4.75, 0.18, 0.66 mL per A h for S, P, L.

estimate $\hat{B}$ back to the level $B_{ok} = B_{cl}(1 - m_{min})$, at which a centered chemistry still keeps the margin m_{min}, is then given by Equation (10.7).

$$t_{rec} = \tau_b \ln \frac{\hat{B} - B_\infty}{B_{ok} - B_\infty} \tag{10.7}$$

In case 5 the bath starts at 32 mL/L, the model's formation rate of 0.248 mL/L an hour gives $B_\infty = 11.0\,\text{mL/L}$, and $m_{min} = 0.15$ gives $B_{ok} = 25.5\,\text{mL/L}$, so that $t_{rec} = 16.5\,\text{h}$ against the 2 h of a remake. The bath of case 2, at 27 mL/L, needs 4.4 h.

10.5 Implemented Controller

Six loops and a supervisor share one controller, called every $\Delta t_c = 10\,\text{s}$ or ten physics steps, although the three dosing trims change only when an analysis arrives, the titration only after a laboratory result and the run-to-run loop once per wafer. The engine step of Figure 1.11 also books the wafer events and exchanges after it integrates. Table 10.3 gives the shipped pairings and settings. Its γ is the action a of Section 1.6.2, with +1 reverse acting and −1 direct acting.

10.5.1 Instruments

No controller reads the bath. The rectifier's amp-hour counter gives the charge, whose rate over each control period passes two first-order lags of 8 min in series, so that the wafer waveform does not ripple the doses, and the result $\dot{Q}_f$ is the feedforward's load. The CVS analyzer and the laboratory are the framework's shared sampled analyzer. The CVS draws a sample at the start of the run and every 2 h after it, through a sample line that holds 2 min of bath, publishes its accelerator, suppressor, leveler and chloride 40 min later and holds them until the next result. Its repeatabilities are 2 % of the make-up strength for the accelerator and the suppressor, 3 % for the leveler and 1 mg/L for chloride. Out of service it draws nothing, loses the analysis under way and holds its last reading, which it flags as stale after two cycles, and a declared bad analysis adds an offset to its next result. The laboratory samples every 8 h and reports 1 h later, copper to 0.2 g/L, acid to 0.1 g/L and the breakdown products to 0.03 mL/L. The probe reports each wafer as it leaves, and the bath thermometer carries 0.05 °C of colored noise.

10.5.2 Amp-Hour Feedforward

Each dosing pump delivers a feedforward and a trim, as shown in Equation (10.8).

$$\begin{aligned} D_k &= \text{clamp}\big(F_k + \Delta_k,\ 0,\ D_{k,max}\big) \\ F_k &= \hat{c}_k \dot{Q}_f + d_{k,25}\, \mathrm{e}^{E_k (T_{set} - 25)}\, V C_{k,SP} \end{aligned} \tag{10.8}$$

The first term of F_k is the amp-hour feedforward with the card's coefficient, and the second is the bulk breakdown the model expects at the chiller setpoint and the concentration setpoint. The suppressor

and leveler coefficients equal the bath's, while the accelerator's ships at 4.75 mL per A h, 5 % below the bath's 5.0, so that the trim has a slow model error to remove, 2.27 mL/h of the 57.43 mL/h the bath needs at the standard load. On feedforward only the trim is zero, and on feedback only there is no feedforward and the trim is the whole dose. A change of strategy moves the trim by the feedforward it adds or removes, so the dose does not jump. A pump cannot take additive out, and a trim asking for less than nothing holds its pump at zero, where it is flagged as saturated. The same split of a dose computed from the measured load and a slow feedback trim on the quality appears in the coagulant dosing of Volume 3, Chapter 6, and Section 1.4.7 places it among the feedforward paths of this volume.

10.5.3 The Sampled PI Trim

Each trim is an instance of the framework's shared PID, written in velocity form and used here without derivative action, and it executes only when an analysis arrives and passes the gate, in the form of Equation (10.9).

$$\begin{aligned} \bar{y}_n &= \lambda y_n + (1-\lambda)\,\bar{y}_{n-1}, \qquad e_n = C_{k,SP} - \bar{y}_n \\ \Sigma_n &= \Sigma_{n-1} + K_c e_n - K_c e_{n-1}, \qquad v_n = \Sigma_n + \mathscr{I}_{n-1} \\ \Delta_n &= \operatorname{clamp}\big(v_n,\ -F_k,\ D_{k,max} - F_k\big) \\ \mathscr{I}_n &= \mathscr{I}_{n-1} + \left[\frac{K_c}{\tau_I}\, e_n + \frac{\Delta_n - v_n}{\tau_I}\right] T_{s,n} \end{aligned} \tag{10.9}$$

In it, y_n denotes the nth accepted result, $\lambda = 0.5$, and $T_{s,n}$ is the time since the previous update, capped at the 120 min cycle. The law is the one of Section 1.6.2: Σ_n is the running sum of Equation (1.13), which telescopes to the positional $K_c e_n$ and is never clamped, $\mathscr{I}_n$ is the integral of Equation (1.15) with back-calculation at the tracking time τ_I, and the clamp acts on the emitted value alone. The derivative and the output rate limit are left unused, and the trim is held between analyses. Because the output is formed before the integral is advanced, the integral learned from result n first reaches the pump with result $n+1$, two hours later. In MANUAL, and on feedforward only, each trim keeps $\mathscr{I} = u_{op} - \Sigma$ at every scan, so its output always equals the value in force and nothing jumps on a return to automatic. A switch to automatic made before the run starts the day from the trims that hold the bath.

10.5.4 The Outlier Gate

Between analyses the controller integrates its own model of each additive, Equation (10.1) with the feedforward coefficients, the doses it commanded, the load $\dot{Q}_f$, the bleed, the drag-out rate and the breakdown at the measured temperature, by explicit Euler at the control period, and it keeps the history. Each result is compared with that model at the moment its sample left the bath. A channel with $|y_n - \hat{C}_k(s_n)| > g\,C_{k,SP}$ and $g = 0.15$ is rejected, and a second rejection on the same side is accepted as a real change. An accepted result re-anchors the model, $\hat{C}_k \leftarrow \hat{C}_k + y_n - \hat{C}_k(s_n)$. The gate is a gross-error test of a sparse analysis against the process model [13], and its confirmation rule costs a real step one analysis. Comparing a late reading with the model of its own time is what Section 1.4.4 asks of any reading that describes an old state of the tool.

10.5.5 Bleed and Feed

The bleed policy is off, a fixed 2.0 L/h, 0.2 L per A h, or the least bleed that holds the estimated breakdown products at B_{SP}. The last is a direct-acting PI ($\gamma = -1$) of the same law on $\hat{B}$, since more bleed is needed when the estimate is above its setpoint, as the application's Show Model note also says. Its output bias is the steady bleed $b_{ff} = \max(0, \hat{G}V/B_{SP} - \dot{q}_d)$, with the tuning of Table 10.3, and it starts at its positional output for the bath it finds. The formation rate $\hat{G}$ is Equation (10.2) evaluated on the controller's model of the additives, B_m integrates it against the bath taken out, and the laboratory

corrects the result every 8 h through the framework's shared bias update of Equation (10.10).

$$\begin{aligned} \hat{B} &= B_m + \beta \\ \beta &\leftarrow \operatorname{clamp}\big(\beta + \lambda_B[B_{lab} - B_m(s) - \beta],\ \pm 20\big) \end{aligned} \tag{10.10}$$

in which $B_m(s)$ is the model's value at the laboratory sample's time and $\lambda_B = 0.5$. Because the accelerator's coefficient is 5 % low, $\hat{G}$ runs 3.5 % low, 0.2488 against 0.2579 mL/L an hour at the reset of case 4, and the bias makes up the difference.

Pause. Try Exercise 10.3 before reading on.

10.5.6 Titration, Run-to-Run and Bath Life

After each laboratory result, copper more than 0.5 g/L above its setpoint is diluted with $g_T Ve/Cu$ liters of DI water over 30 min, copper as far below it gets $g_T Ve$ grams of copper concentrate, and acid more than 0.3 g/L low gets $g_T Ve$ grams of acid, with e the titrated error and $g_T = 0.8$. Copper rises by at most 0.042 g/L and acid falls by at most 0.065 g/L in any documented day, so these corrections never trigger in the shipped cases, and a check that starts the bath 3 g/L high in copper and 2 g/L low in acid confirms their direction. The run-to-run loop infers from each probe result the cathode efficiency that explains it, $\eta_{meas} = 3600\,\kappa_F h_{meas}/Q_w$, corrects the efficiency the recipe assumes through the shared bias update with $\lambda_R = 0.3$ [14], and sets the next bulk step from Equation (10.4), $t_b = (3600\,\kappa_F h_{SP}/\hat{\eta} - t_{nuc}I_{nuc} - t_{fill}I_{fill})/I_b$. With the true efficiency equal to the assumed one, the thickness error is the probe's 0.4 % passed through the average, $0.4\sqrt{\lambda_R/(2-\lambda_R)} = 0.168\,\%$ rms by the noise inflation of Section 1.4.3, against the 0.161 to 0.195 % the engine gives across the cases in AUTO.

The bath-life supervisor computes the margin $\hat{m}$ that the analyses and $\hat{B}$ imply through Equation (10.5), the margin a centered chemistry could still give, $m_c = \max(0, 1 - \hat{B}/B_{cl})$, and the recovery time of Equation (10.7). It advises a dump when $m_c < m_{min}$ and $t_{rec} > t_{rm}$, or when the dump cycle is cheaper and $\hat{B} \geq B_{SP}$, and it reports the reason. Set to automatic, it requests the dump once the advice has stood for 30 min, and the tool finishes the wafer in the cell, drains the bath, remakes it from VMS and stays down for the remake while the controllers' models restart at the make-up strength.

10.5.7 Simulation, Internal and Analysis Models

Section 10.3 is the simulation model. The controllers' internal models are the dosing model with its coefficients and the breakdown at the measured temperature, its formation of breakdown products, Faraday's law with the assumed efficiency and the recovery time of Equation (10.7), and only the supervisor reads the fill window. No control model is identified, and the report-to-report model of the next section is an analysis model of the trim, linearized from the simulation model.

10.6 The Gain a Two-Hour Analysis Can Carry

Around an operating point the accelerator is a first-order bath, $V\,\mathrm{d}x/\mathrm{d}t = -kx + \Delta$, where x is its deviation, Δ the trim, and $k = c_S I/S_V + d_S(T)\,V + b + \dot{q}_d$ is everything that takes accelerator out per unit of concentration. In case 3, after the new lot raises c_S to 6.25 mL per A h, $k = 14.24\,\mathrm{L/h}$ at the standard load and a bleed of 2.28 L/h, the bleed a noise-free run settles to after its first day, which is a time constant $\tau = V/k$ of 14.0 h. The trim changes at report times t_n and holds for $T = 2\,\mathrm{h}$, and result n describes the bath at $s_n = t_n - \theta$ with $\theta = 42\,\mathrm{min}$. Sampling the bath at the sample times, with each trim held until the next report, gives Equation (10.11).

$$\begin{aligned} x_{n+1} &= a\,x_n + \beta_1\Delta_{n-1} + \beta_0\Delta_n, \qquad a = \mathrm{e}^{-T/\tau} \\ \beta_0 &= \frac{1 - \mathrm{e}^{-(T-\theta)/\tau}}{k}, \qquad \beta_1 = \mathrm{e}^{-(T-\theta)/\tau}\,\frac{1 - \mathrm{e}^{-\theta/\tau}}{k} \end{aligned} \tag{10.11}$$

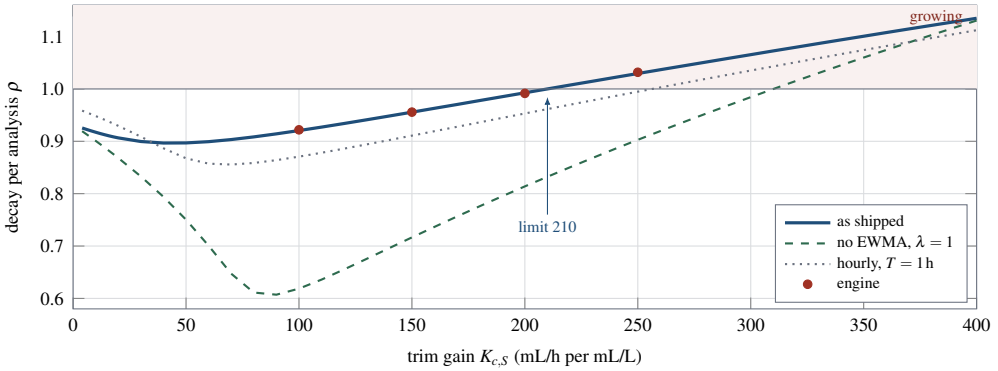


Figure 10.3 Decay per analysis of the accelerator trim against its gain, from the report-to-report recurrence at the operating point of case 3 after the new lot, with $\tau_I = 6\,\text{h}$. The thick solid curve is the trim as shipped, the dashed curve the trim without the average ($\lambda = 1$) and the dotted curve the trim on an hourly analysis. Values above one grow. The round markers are noise-free runs of the engine over 96 h, with the mean half-swing ratio of the cycle-averaged accelerator converted to a factor per analysis.

The trim of Equation (10.9), unclamped and with $T_{s,n} = T$, closes the loop through $\bar{y}_n = \lambda x_n + (1-\lambda)\bar{y}_{n-1}$, $\Delta_n = -K_c\bar{y}_n + \mathscr{I}_{n-1}$ and $\mathscr{I}_n = \mathscr{I}_{n-1} - (K_c T/\tau_I)\bar{y}_n$. The result is a linear recurrence in $(x_n, \Delta_{n-1}, \bar{y}_{n-1}, \mathscr{I}_{n-1})$, and its spectral radius ρ is the factor by which an error shrinks per analysis. Section 1.4.4 derives a boundary of the same kind, Equation (1.7), for a recipe loop whose readings wait in a queue of whole runs. Here the age is a fraction of the update period, and the average and the integral add states.

At the shipped $K_c = 30$ and $\tau_I = 6\,\text{h}$ the recurrence gives $\rho = 0.900$, an error that shrinks by a factor of e in 18.9 h while it swings with a period of 37.6 h, and the fastest decay at this integral time is $\rho = 0.897$ at $K_c = 44$. The radius reaches one at $K_c = 210$, as shown in Figure 10.3. Noise-free runs of case 3 over 96 h confirm the boundary. Averaged over each wafer cycle to remove the plating ripple, the engine's successive half-swings of the accelerator, after the first, shrink or grow by geometric-mean factors of 0.691, 0.847, 0.973 and 1.094 at gains of 100, 150, 200 and 250, which is ρ of 0.922, 0.956, 0.992 and 1.032 against the recurrence's 0.921, 0.956, 0.993 and 1.029, with periods of 18.23, 14.67, 12.86 and 11.46 h against 18.28, 14.72, 12.74 and 11.44 h. At 300 and 400 the growing swing drives the pump to 0 and 150 mL/h and settles into a limit cycle.

The limit belongs to the measurement rather than to the chemistry. With the analysis age cut to the 2 min of the sample line the recurrence allows $K_c = 744$, at 21 min 440, and an hour of age lowers it to 141. The average costs almost a third of the gain, because without it ($\lambda = 1$) the limit is 310, while $\lambda = 0.3$ lowers it to 107. An hourly analysis with the same age raises the limit only to 256, which noise-free runs bracket between 230, where half-swings shrink by 0.913, and 280, where they grow by 1.079. The integral time moves it most, since at $\tau_I = 3\,\text{h}$ the loop is unstable above $K_c = 33$ and at 12 h above 307. Advancing the integral in the same report as its error, rather than one report later, would lower the limit from 210 to 198, because the integral would then add $K_c T/\tau_I$ to the immediate correction. However well the analyzer measures, a trim updated every two hours on a result 42 minutes old needs most of a day to remove an error, while a thick-copper lot doubles the consumption within the few minutes of one wafer. A continuous composition loop loses its margin in the same way behind a lagging online analyzer in Volume 5, Chapter 3.

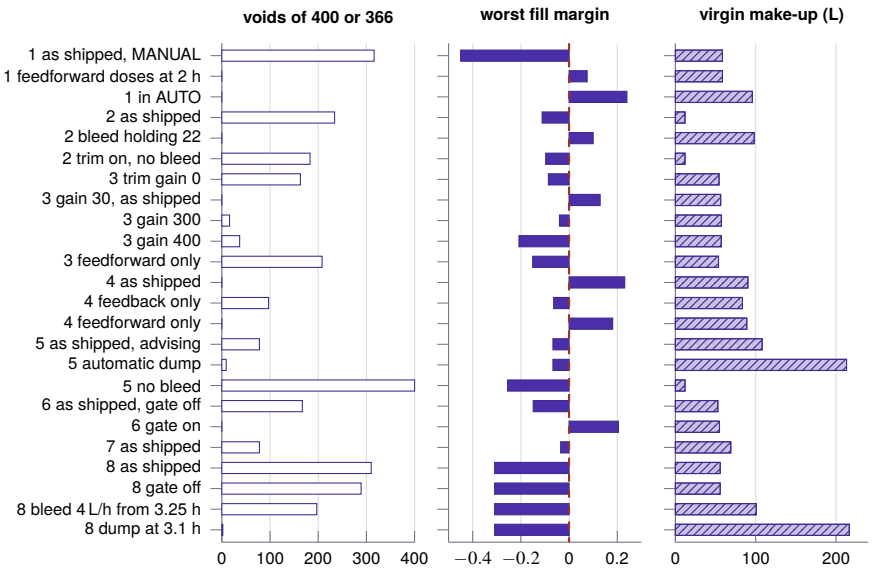


Figure 10.4 The scorecard of Table 10.5 drawn as bars: the voids of each day (open bars), the worst fill margin (solid bars) with the void edge at zero dashed, and the virgin make-up used (hatched bars). Notice that every row with a negative worst margin has voids, that more trim gain in case 3 and the gate in case 8 cost wafers, and that the two dumps buy their few voids with the largest make-up of the table.

Table 10.4 The Case Sequence

#	Scenario and event	Teaching point
1	MANUAL, thick-copper lot at 1 h	dose by the charge
2	Feedforward only, no bleed, $B = 27$	a window closes
3	Accelerator lot, c_S 5.0 to 6.25 at 1 h	the trim's gain limit
4	Thick-copper lot at 4 h	feedforward or feedback
5	$B = 32$, bleed toward 22	past the end of bath life
6	Result 50 % high at 2:40, gate off	a bad analysis
7	CVS down 6–12 h, chiller +10 °C	feedforward by force
8	+300 mL of leveler at 3 h	a dose that stays
9	400 wafers, lot at 8 h, CVS down 14–20 h	graded day

10.7 The Cases as Measured

Eight cases and a graded day make up the sequence of Table 10.4, with the briefings, completion rules and challenge of Section 1.6.4. Every case begins from the reference configuration of Table 10.3, with the feedforward and the trim, the gate, the bleed holding 22 mL/L, the titration and the run-to-run loop on and the supervisor advising, and departs from it only as its row says. Every number that follows comes from the engine at the default seed, over each scenario's 24 h and without the interface, with each setting changed before the run starts unless a time is given. Table 10.5 gathers the scorecard of each case and variant, and Figure 10.4 draws it as bars. A count of voids rests on one realization of the analysis noise wherever the ratio grazes a window edge for hours. Across ten other seeds case 4 on feedback only voids 51 to 117 wafers, case 8 voids 239 to 349 and case 7 none to 125, while the suggested configuration of the graded day plates every wafer at every seed, with worst margins of 0.083 to 0.176 and 56.5 to 57.3 L of make-up.

Table 10.5 Scorecard From the Truth Log, per Case and Variant, Over 24 h

Case and configuration	Voids	First (h)	Worst margin	VMS (L)	Wafers
1 as shipped, MANUAL	316	5.09	−0.450	58.6	400
1 feedforward doses at 2 h	0	—	0.075	58.6	400
1 in AUTO	0	—	0.240	95.8	400
2 as shipped	234	10.01	−0.112	12.0	400
2 bleed holding 22	0	—	0.100	98.1	400
2 trim on, no bleed	183	13.07	−0.097	12.0	400
3 trim gain 0	163	14.27	−0.086	54.4	400
3 gain 30, as shipped	0	—	0.129	56.4	400
3 gain 300	16	22.43	−0.040	57.2	400
3 gain 400	37	21.41	−0.208	57.1	400
3 feedforward only	208	11.57	−0.151	53.7	400
4 as shipped	0	—	0.231	90.3	400
4 feedback only	97	8.09	−0.064	83.5	400
4 feedforward only	0	—	0.181	88.9	400
5 as shipped, advising	78	0.05	−0.067	108.0	400
5 automatic dump	9	0.05	−0.067	213.1	366
5 no bleed	400	0.05	−0.255	12.0	400
6 as shipped, gate off	167	5.69	−0.149	52.9	400
6 gate on	0	—	0.205	54.9	400
7 as shipped	78	12.65	−0.035	69.0	400
8 as shipped	310	3.05	−0.309	55.8	400
8 gate off	289	3.05	−0.309	55.8	400
8 bleed 4 L/h from 3.25 h	197	3.05	−0.309	100.7	400
8 dump at 3.1 h	2	3.05	−0.309	216.5	366

10.7.1 Fixed Pumps Through a Lot Change

Case 1 runs in MANUAL with the pumps at the doses that hold the bath at the standard load, 57.4 mL/h of accelerator, 2.04 of suppressor and 7.20 of leveler, and a bleed of 1.94 L/h. At 1 h a thick-copper lot starts and 383 of the day's 400 wafers take twice the charge. The first void, a seam, comes at 5.1 h in wafer 85, the ratio reaches 2.20 by 24 h, and 316 wafers void. The manual bulk step, set for a standard wafer, also leaves every thick wafer 5.35 % thin, because the nucleation and fill steps do not double with the bulk current. Setting the pumps to the feedforward's doses at 2 h, 93.2 mL/h of accelerator, voids no wafer, although with nothing to remove the feedforward's 5 % error the accelerator drifts to 5.75 mL/L by 24 h and the worst margin is 0.075. In AUTO no wafer voids, the worst margin is 0.240, and the run-to-run loop holds every wafer within 0.65 % of its target. The case's own route to AUTO, a Reset that keeps the manual pumps at the feedforward's doses and a switch before the run, gives the same day.

10.7.2 Feedforward Alone in an Aged Bath

In case 2 the feedforward doses alone with the bleed off, from 27 mL/L of breakdown products. The additives hold, the accelerator at 5.84 mL/L and the suppressor at 2.000 mL/L after 24 h, but the drag-out's 12 L a day is the only exit and the breakdown products climb to 31.4 mL/L. The window's half-width falls from 3.6 % to 1.3 % by 10 h and is gone at 16.3 h. The first void comes at 10.0 h in wafer 167, where the feedforward's drift, a ratio 1.4 % low from the accelerator's coefficient, crosses the narrowed edge, and 234 wafers void, while the supervisor advises a dump from the start. With the bleed holding 22 mL/L it starts at its 4 L/h limit, the breakdown products fall to 21.4 mL/L by 24 h, no wafer voids, and the day uses 98.1 L of virgin make-up instead of 12. With the trim on and no bleed, 183 wafers void from 13.1 h.

10.7.3 How Much Trim a Sampled Loop Takes

At 1 h in case 3 a new lot of accelerator concentrate raises its consumption by 25 %, which the feedforward's coefficient does not know. With no trim the accelerator drifts to 5.34 mL/L and 163 wafers void from 14.3 h, and on feedforward only 208 void from 11.6 h. At the shipped gain of 30 the accelerator stays between 5.76 and 6.00 mL/L, lowest at 10.7 h, the trim reaches 10.5 mL/h by 12 h and 18.7 mL/h by 24 h against the 13.6 mL/h the new lot needs at the setpoint, and the worst margin is 0.129. A gain of 10 is slower, reaching its lowest 5.61 mL/L at 14.7 h. A gain of 100 holds the accelerator between 5.86 and 6.03 mL/L, while 200 swings it from 5.66 to 6.18 mL/L, as the recurrence of Section 10.6 expects of a loop near its limit. At 300 the swing grows into 16 voids from 22.4 h, and at 400 the accelerator swings between 6.53 mL/L at 16.7 h and 5.12 mL/L at 22.7 h and 37 wafers void from 21.4 h.

10.7.4 Feedforward Against Feedback

In case 4 a thick-copper lot starts at 4 h and 333 of the day's wafers are thick. With the feedforward and the trim the accelerator dose rises from 57.4 mL/h at 4 h to 101.8 mL/h at 6 h as the filtered amp-hours double, the accelerator stays between 5.91 and 6.01 mL/L, and no wafer voids, with a worst margin of 0.231. The bleed rises from 1.79 to 3.56 L/h over the same two hours to hold the breakdown products, and the day uses 90.3 L of make-up. On feedback only the result reported at 4:40 reads 5.96 mL/L, and the first to show the lot is reported at 6:40, of a sample drawn at 6:00. The trim then moves the dose from 58.1 to 62.6, 71.6, 81.9 and 90.2 mL/h at successive reports, about 10 mL/h a report once the average has caught up, while the accelerator falls to 5.06 mL/L at 12.6 h, and 97 wafers void from 8.1 h, in wafer 135, to 13.8 h. The run at 11.00 h is shown in Figure 10.5. On feedforward only the accelerator drifts to 5.80 mL/L by 24 h without a void, which is the slow error the trim exists to remove.

> **TRY IT IN THE APP** **4 · A Lot Mix Change**
>
> Load the scenario and press Run. Watch the accelerator pump's dose on the AIC-S card once the thick-copper lot starts at 4 h. The amp-hour feedforward lifts it from 57.4 mL/h to 101.8 mL/h by 6 h, and no wafer of the day voids.

10.7.5 Past the End of Bath Life

Case 5 starts at 32 mL/L of breakdown products, past the closing of the window, with the bleed holding toward 22 mL/L and the supervisor advising. The bleed starts at its 4 L/h maximum, the breakdown products fall to 29.4 mL/L by 6 h and 23.4 by 24 h, the window reopens when they pass 30 mL/L at 4.6 h, and 78 wafers void, pinched off, the last at 4.7 h. The bleed stays within 0.05 L/h of its maximum all day, because the estimate never comes down to its setpoint, so the 108.0 L of make-up are 96.0 L of bleed and 12.0 of drag-out. The supervisor reports from the first minutes that the largest bleed needs 16.5 h to restore the margin. On automatic it requests the dump at 0.50 h and the tool drains the bath at 0.53 h, once its wafer is out, so 9 wafers void before the dump and none after, with 366 wafers plated instead of 400 and 213.1 L of make-up, 200 L of it the remake. The remade bath climbs only to 5.51 mL/L by 24 h, with the bleed shut. With no bleed at all every wafer voids and the breakdown products reach 36.1 mL/L. No setting of the bleed reopens the window in time, and the application says so.

10.7.6 A Bad Analysis and the Gate

In case 6 the result reported at 2:40 reads 9.04 mL/L for a bath at 6.00, with the gate off. The trim takes it at face value and cuts the accelerator dose from 58.2 to 12.3 mL/h. At 4:40 the average still carries half the bad value, 6.56 mL/L against a bath at 5.58, so the dose recovers only to 26.0 mL/h, the accelerator falls to 5.19 mL/L at 10.7 h, and 167 wafers void from 5.7 to 15.6 h. With the gate on,

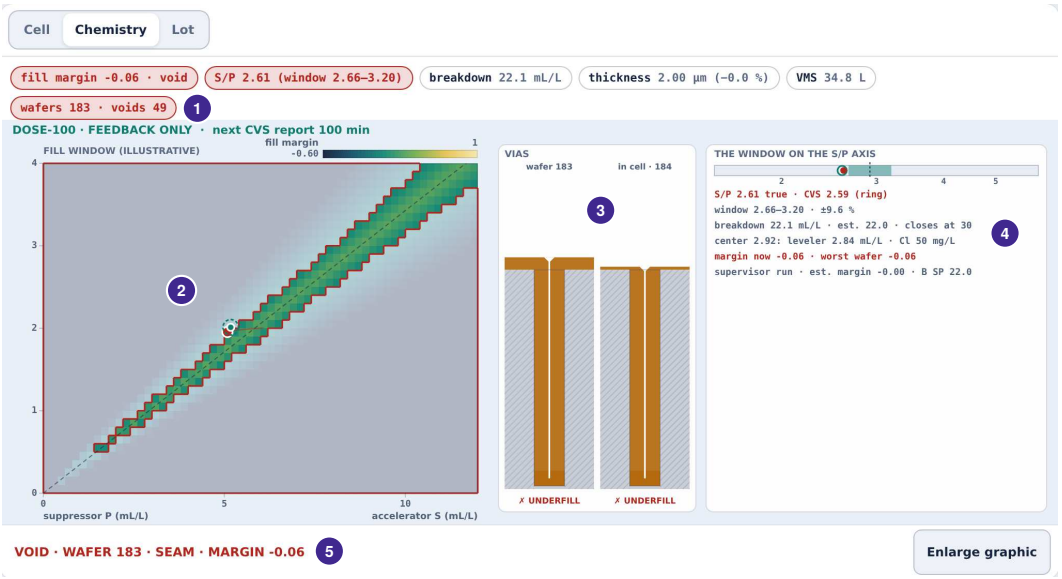


Figure 10.5 The bath in case 4, a lot mix change, at $t = 11.00$ h with the dosing strategy on feedback only, seven hours after the thick-copper lot started, in the Chemistry view, cropped to the plant view. (1) The badges read wafers 183, voids 49 and 34.8 L of virgin make-up used. (2) The fill margin baked over accelerator and suppressor at the bath's leveler, chloride and breakdown products shows the window as a narrow wedge from the origin with a red outline, the true operating point as a filled dot, red because its margin is negative, with its trail just below the wedge, and the last CVS result as a ring. (3) Both via cross-sections, wafer 183 and wafer 184 in the cell, are drawn underfilled with a seam. (4) The window on the S/P axis gives the true ratio 2.61 against a window of 2.66 to 3.20 (±9.6 %) centered at 2.92 by the leveler at 2.84 mL/L, the CVS ratio of 2.59, corrected for the analyzed chloride, as the ring, breakdown products of 22.1 mL/L (estimate 22.0), a margin now of −0.06 and a worst wafer of −0.06. (5) The status line reads VOID · WAFER 183 · SEAM · MARGIN −0.06.

the model predicts 5.96 mL/L for that sample, the result is rejected, and no wafer voids. The gate rejects nothing else in the day's 12 reports.

> **TRY IT IN THE APP** **6 · An Outlier and the Gate**
>
> Press Run with the gate off, as shipped, and watch the accelerator pump after the result of 2:40, which reads 9.04 mL/L for a bath at 6.00. The trim cuts the dose from 58.2 to 12.3 mL/h, and 167 wafers void. Press Reset, set Outlier gate to On on the DOSE-100 card before Run, and run again: no wafer voids.

10.7.7 The Analyzer Out of Service

In case 7 the CVS is out of service from 6 to 12 h, its last result reported at 4:40, and at 6.5 h the chiller drifts so that the bath reaches 34.5 °C by 8 h. At 35 °C the accelerator breaks down 3.32 times as fast as at 25 °C, an extra 2.3 % of it an hour, and the time-based feedforward still assumes the setpoint. The dose stays between 58.9 and 59.5 mL/h and the reading at 5.89 mL/L while the accelerator falls from 6.02 at 6 h to 5.50 mL/L at 12 h. The controller's own model, which uses the measured temperature, follows the fall from 5.93, where the 4:40 result left it, to 5.50 mL/L at 12 h, but no dose is computed from it. The first void comes at 12.6 h in wafer 211, a minute before the returning analyzer's first result at 12:40, of a sample drawn at 12:00. The trim then climbs from 4.0

Table 10.6 Configurations Tried on the Graded Day

Configuration	Voids	First (h)	Worst margin	VMS (L)	B at 24 h
as shipped	226	10.49	−0.374	12.0	28.86
suggested[a]	0	—	0.118	57.0	24.55
suggested, $B_{SP} = 22$	0	—	0.197	83.2	22.14
suggested, $B_{SP} = 26$	0	—	0.054	38.7	26.48
suggested, bleed 4 L/h	0	—	0.260	108.0	20.25
feedforward and trim, no bleed	60	20.45	−0.059	12.0	29.81
feedback only, $B_{SP} = 24$	216	11.09	−0.235	49.5	24.53

[a] feedforward and gated trim, bleed holding $B_{SP} = 24\,\text{mL/L}$

to 39.8 mL/h as it learns the hot bath, the accelerator bottoms at 5.37 mL/L at 14.7 h, and 78 wafers void, the last at 17.3 h.

10.7.8 A Dose That Cannot Be Taken Back

In case 8 the leveler pump puts 300 mL into the bath at 3 h, from 3.0 to 4.5 mL/L, and the window's center moves from 3.00 to 3.67 with it, so the bath's ratio of 3.0 lies below the window and the first void comes at 3.05 h in wafer 51. The gate rejects the first result to show the overdose, reported at 4:40, and accepts the second, at 6:40, on the same side. From then on the leveler pump sits at or near zero against its lower limit, the leveler leaves only by consumption, breakdown and the bath taken out, falling to 3.41 mL/L by 24 h, and 310 wafers void until 21.6 h. With the gate off 289 void until 20.3 h, so the gate's confirmation costs 21 wafers here. A fixed bleed of 4 L/h from 3.25 h ends the voids at 14.8 h with 197 voids and 100.7 L of make-up, and a dump at 3.1 h ends them at once, with 2 voids, 366 wafers and 216.5 L. Raising the accelerator setpoint to 7.2 mL/L from 3.25 h to chase the moved center still voids 251 wafers, 175 of them pinched off, because the window moves back as the leveler falls.

> ✦ DID YOU KNOW? **One Gate, Two Verdicts**
>
> From the simulator: in Table 10.5 the outlier gate takes case 6 from 167 voids to none. In case 8 the same gate lets 310 wafers void against 289 without it, because it waits for a second result before it believes the leveler overdose.

10.7.9 The Graded Day

The challenge runs 24 h from a bath at 22 mL/L, with a thick-copper lot from 8 h to the end and the CVS out of service from 14 to 20 h, and it asks for 400 wafers with no void and at most 65 L of virgin make-up. The day counts only on the bath as given, without a changed model parameter or disturbance. It ships on feedback only with no bleed, and Table 10.6 sets the measured alternatives side by side.

As shipped, 226 wafers void, the first at 10.5 h in wafer 175, the accelerator falls to 4.63 mL/L at 20.7 h, the breakdown products reach 28.9 mL/L, and only the 12 L that replace the drag-out are used. The suggested configuration, the feedforward with the gated trim and the bleed holding 24 mL/L, plates all 400 without a void, with the accelerator between 5.82 and 6.03 mL/L and a worst margin of 0.118, while the breakdown products rise to 24.6 mL/L as the bleed climbs to 2.9 L/h under the thick lot, and it uses 57.0 L, 8.0 L inside the ceiling. Each element pays for itself. The feedforward and trim without a bleed void 60 wafers from 20.4 h, when the breakdown products have reached 28.5 mL/L, and feedback only with the bleed at 24 voids 216 from 11.1 h. The setpoint prices the window as Equation (10.6) accounts for it, 83.2 L at 22, 38.7 at 26 with a worst margin of 0.054, and 108.0 with the bleed at its maximum.

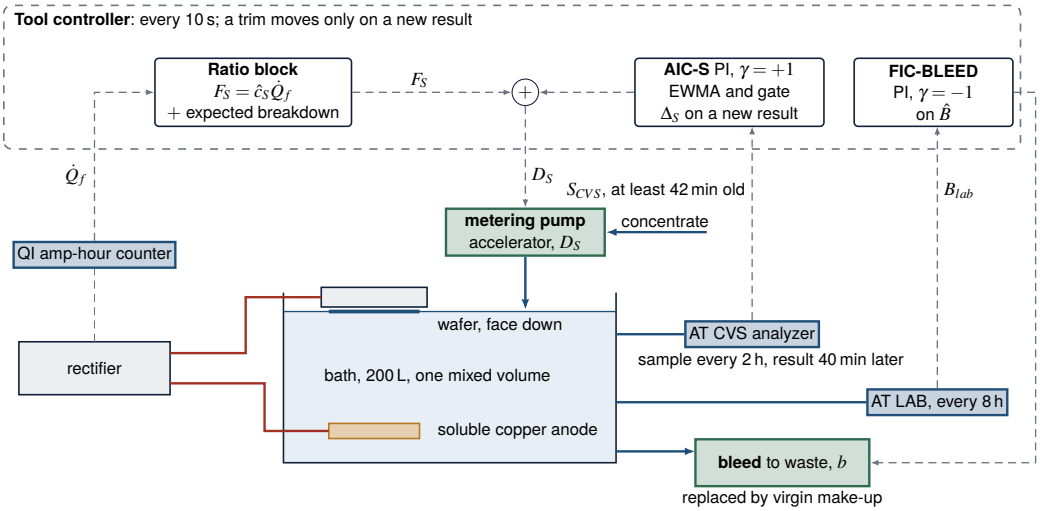


Figure 10.6 The accelerator loop and the bleed of the simulated bath drawn with their hardware. The amp-hour counter QI on the rectifier gives the load, and the ratio block turns it into the feedforward dose. The CVS analyzer AT draws through its sample line every 2 h and reports 40 min later, and AIC-S, a velocity-form PI with $\gamma = +1$, adds its trim only when a result passes the gate. The sum drives the metering pump. The laboratory corrects the model's estimate of the breakdown products every 8 h, and FIC-BLEED, direct acting with $\gamma = -1$, sets the bleed that virgin make-up replaces. Notice that the only signal that reaches the pump within minutes is the charge, not the chemistry.

10.8 Technology in Practice

The loops of Section 10.5 run on a stripping analyzer and a laboratory, the rectifier's amp-hour counter, three metering pumps and a bleed, and the controller of the plating tool. Figure 10.6 draws the accelerator loop and the bleed with that hardware.

What Measures It

The model of Section 10.5 reduces the analyzer to a sample drawn every 2 h through a 2 min line, a result published 40 min later and held, and a repeatability of 2 % of the make-up strength for the accelerator, with no matrix effects.

> **TECHNOLOGY** **Sensor: A Cyclic Voltammetric Stripping Analyzer**
>
> The analyzer sweeps the potential of a rotating inert disk back and forth in a sample of the bath, so that copper is plated on the disk and stripped off again on every cycle. The charge that strips the copper follows the rate of plating, which the additives set, and the method rests on that effect [16]. A result is therefore the strength of the sample's action on the deposit, expressed as a concentration. A direct analysis of the organic compounds of the bath is not commonly available by such electrochemical techniques [17]. A bad analysis shows on the trend as a single step in the reading that the next result does not repeat, as in case 6.

No open, vendor-neutral source gives a production instrument's analysis time or accuracy, so the model's values are teaching values. The one signal of consumption that arrives at once is the rectifier's charge, passed through two lags of 8 min. The laboratory's analysis of the breakdown products stands for a separate chromatographic method.

What Moves It

Each pump is modeled as delivering exactly the dose it is sent, held between zero and 150, 12 and 40 mL/h for the accelerator, suppressor and leveler, and the bleed between 0 and 4 L/h.

TECHNOLOGY Final Element: A Metering Pump on a Current Signal

A metering pump packs its pump mechanism, motor and drive electronics in one unit. A 4 to 20 mA signal sets its flow from zero to full scale, and a knob on the pump sets the flow that full scale means. Many such pumps drive a plunger with a solenoid at a variable rate, so the flow arrives as discrete pulses, a simple, reliable design where an uneven flow is acceptable. In water treatment each chemical is proportioned to the flow of the water it treats, so its pump must be set by the control system, not by hand [18]. A pump that runs away, as in case 8, puts in what no pump can take out.

The model's exact dose stands for the average such a pump delivers, with no pulses, no calibration drift and no failure but the runaway of case 8. Its clamp at zero is the one-sided actuator of that case.

Where the Controller Runs

In the app the six loops and the supervisor share one controller. In a fab a run-to-run correction such as that of the bulk step belongs to advanced process control, which combines model-based control with fault detection and classification [19]. Supervisory and lower-level loops run with cycle times from milliseconds to minutes [20]. The app integrates every 1 s and executes its controller every 10 s, while each trim moves only when an analysis arrives, every 2 h, the bulk step once per wafer, every 216 s, and the estimate of the breakdown products every 8 h.

IN PRACTICE How the Bath Is Analyzed

A review of damascene plating management names bath concentration measurement and control, film metrology, contamination and defects as the areas such a process must measure [21]. In a stripping analysis the suppressor is found by adding the sample step by step to virgin make-up solution until the stripping charge falls to half, the accelerator by adding known standards, and the leveler on a calibration curve. Copper, acid and chloride are titrated or read by absorbance. The analyst holds every solution at one temperature and checks the platinum electrode, which the additives foul over many analyses [22].

In such a controller the chapter's structure is a ratio block on the filtered amp-hour rate, three PI blocks executed on a new analysis, a bath model that gates each result, a bias update and a PI for the bleed, an EWMA on the bulk step, and the supervisor.

What Is Modernizing It

The first development is analysis at the tool. A microelectrode gave the stripping analysis without a rotating disk, for monitoring in the bath itself, in 2007 [23], and in 2025 a portable microfluidic workstation detected the additives online with a mean relative error below 10 %, replacing the laboratory's manual analysis [24]. Chemometrics serve the electroanalytical control of bath concentrations [21]. The second is dosing from a model. Since an exact analysis needs equipment a plating plant typically cannot fully provide, one approach models the electrolyte's composition from data the plant's production systems already record and doses from it, validated on an industrial zinc-nickel line [25]. This chapter's controller already integrates such a model between analyses, but no loop yet doses from it (Section 10.9).

10.9 Limitations

The fill window is an illustrative correlation in the ratio, the leveler and the breakdown products, not a supplier's data or a feature-scale deposition model, and the via cross-sections are drawn from the margin. The breakdown products are one lumped species with fixed yields, the additives' rates are chosen so that a day shows the lessons, the CVS has Gaussian repeatability and no matrix effects, the bath is one well-mixed volume, and Faraday's law acts on a flat wafer with one efficiency.

OPEN CHALLENGE Seeing What the Additives Become

Total organic carbon was once used to follow the decomposition of a plating bath's organics, but it did not correlate strongly with the bath's effective composition. Extraction followed by liquid chromatography has since identified and quantified the byproducts in a sample of a live bath [17]. The breakdown products of this chapter set the width of the fill window, yet its laboratory reports them only every eight hours. Progress would be a measurement of them fast enough for the bleed to act on, in place of an estimate corrected three times a day.

Two parts of the control system are present but idle in these cases. The copper and acid corrections never trigger, because neither species leaves its dead band within a day, and the supervisor's cost comparison cannot choose a dump in a bath whose byproducts accumulate linearly, so every dump the application calls rests on feasibility. The controller's own model tracks the accelerator through an outage, but no loop doses from it, so the drift of case 7 goes uncorrected although the information to correct it exists. All sizes, rates, prices and limits are teaching values.

10.10 Conclusion

An acid copper plating bath is a chemistry dosed against a rare analysis, and the simulator of this chapter logs each wafer's true fill margin beside the results its controllers saw. At the standard load the accelerator turns over in 20.9 h and the suppressor in 8.2 days, so the ratio follows the accelerator, and the breakdown products narrow the window to 9.8 % at 22 mL/L and close it at 30. The amp-hour feedforward answers a doubled load at once, while feedback alone first sees it at 6:40, two hours and forty minutes after the lot, and voids 97 wafers.

A report-to-report model places the gain limit of the sampled trim at 210 mL/h per mL/L, set by the analysis's 42 min age and its average, and the engine's swings decay below it and grow above it as predicted. The make-up a bleed spends is the breakdown products removed divided by the concentration at which they leave, 57.0 L at 24 mL/L on the graded day, and a bath past its life is recovered only by a dump, which ends the voids of case 5 after 9 wafers at a cost of 34 wafers.

The window and the rates are illustrative. A fuller model would add a feature-scale fill model, matrix effects in the analysis, fragments that stay partly active, and a dose computed from the controller's own model between analyses. The bath shows that a sampled analysis can correct a model but cannot carry a load.

KEY POINTS

- A copper bath uses up its additives with every amp-hour, and its analysis arrives every two hours to describe the bath of at least 42 min before.
- The amp-hour feedforward answers a doubled load at once, while feedback alone first sees it at 6:40 and voids 97 wafers.
- The trim's gain limit, 210 mL/h per mL/L, is set by the analysis's age and its average, not by the chemistry.
- The make-up a bleed spends is the breakdown products removed divided by their exit concentration, and only a dump recovers a bath past its life.

Where else this applies. A dose computed from the load with a slow quality trim in Volume 3, Chapter 6, a loop behind a lagging analyzer in Volume 5, Chapter 3, and a recipe corrected on late results in Chapter 6.

References

[1] D. Edelstein, J. Heidenreich, R. Goldblatt, W. Cote, C. Uzoh, N. Lustig, P. Roper, T. McDevitt, W. Motsiff, A. Simon, J. Dukovic, R. Wachnik, H. Rathore, R. Schulz, L. Su, S. Luce, and J. Slattery, "Full copper wiring in a sub-0.25 µm CMOS ULSI technology," in *International Electron Devices Meeting. IEDM Technical Digest*, Washington, DC, USA, 1997, pp. 773–776.

[2] P. C. Andricacos, C. Uzoh, J. O. Dukovic, J. Horkans, and H. Deligianni, "Damascene copper electroplating for chip interconnections," *IBM Journal of Research and Development*, vol. 42, no. 5, pp. 567–574, 1998.

[3] J. P. Healy, D. Pletcher, and M. Goodenough, "The chemistry of the additives in an acid copper electroplating bath," *Journal of Electroanalytical Chemistry*, vol. 338, no. 1-2, pp. 155–165, 1992.

[4] T. P. Moffat, D. Wheeler, M. D. Edelstein, and D. Josell, "Superconformal film growth: Mechanism and quantification," *IBM Journal of Research and Development*, vol. 49, no. 1, pp. 19–36, 2005.

[5] P. M. Vereecken, R. A. Binstead, H. Deligianni, and P. C. Andricacos, "The chemistry of additives in damascene copper plating," *IBM Journal of Research and Development*, vol. 49, no. 1, pp. 3–18, 2005.

[6] A. C. West, "Theory of filling of high-aspect ratio trenches and vias in presence of additives," *Journal of The Electrochemical Society*, vol. 147, no. 1, p. 227, 2000.

[7] A. Frank and A. J. Bard, "The decomposition of the sulfonate additive sulfopropyl sulfonate in acid copper electroplating chemistries," *Journal of The Electrochemical Society*, vol. 150, no. 4, p. C244, 2003.

[8] C.-K. Hu and J. M. E. Harper, "Copper interconnections and reliability," *Materials Chemistry and Physics*, vol. 52, no. 1, pp. 5–16, 1998.

[9] D. Tench and J. White, "Cyclic pulse voltammetric stripping analysis of acid copper plating baths," *Journal of The Electrochemical Society*, vol. 132, no. 4, pp. 831–834, 1985.

[10] R. D. Gudi, S. L. Shah, and M. R. Gray, "Multirate state and parameter estimation in an antibiotic fermentation with delayed measurements," *Biotechnology and Bioengineering*, vol. 44, no. 11, pp. 1271–1278, 1994.

[11] J. H. Lee and M. Morari, "Robust inferential control of multi-rate sampled-data systems," *Chemical Engineering Science*, vol. 47, no. 4, pp. 865–885, 1992.

[12] J. J. Downs and E. F. Vogel, "A plant-wide industrial process control problem," *Computers & Chemical Engineering*, vol. 17, no. 3, pp. 245–255, 1993.

[13] R. S. H. Mah and A. C. Tamhane, "Detection of gross errors in process data," *AIChE Journal*, vol. 28, no. 5, pp. 828–830, 1982.

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

[15] 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.

[16] D. Tench and C. Ogden, "A new voltammetric stripping method applied to the determination of the brightener concentration in copper pyrophosphate plating baths," *Journal of The Electrochemical Society*, vol. 125, no. 2, pp. 194–198, 1978.

[17] L. D'Urzo, H. Wang, A. Pa, and C. Zhi, "SPE-HPLC detection of organic additives in acidic copper plating baths," *Journal of The Electrochemical Society*, vol. 152, no. 4, p. C243, 2005.

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

[19] A. V. Prabhu, "Performance monitoring of run-to-run control systems used in semiconductor manufacturing," Ph.D. dissertation, The University of Texas at Austin, 2008. [Online]. Available: http://hdl.handle.net/2152/17721

[20] 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.

[21] R. Carpio and A. Jaworski, "Review—management of copper damascene plating," *Journal of The Electrochemical Society*, vol. 166, no. 1, pp. D3072–D3096, 2019.

[22] K. Ichiba and K. Watanabe, "Cyclic voltammetric stripping (CVS) for the measurement of additives in electro copper plating baths," *Denki Kagaku*, vol. 88, no. 2, pp. 162–165, 2020, in Japanese.

[23] B.-G. Xie, J.-J. Sun, X.-B. Chen, J.-H. Chen, T.-L. Xiang, and G.-N. Chen, "In situ monitoring of additives in copper plating baths by cyclic voltammetric stripping with a microelectrode," *Journal of The Electrochemical Society*, vol. 154, no. 10, p. D516, 2007.

[24] Y. Zhao, J.-X. Zeng, J.-Q. Yang, T. Song, R. Hu, J.-J. Su, B. Zhang, F.-Z. Yang, D. Zhan, and L. Han, "On-line detection of additive concentrations in acidic copper plating solution for metal interconnection by an electrochemical microfluidic workstation," *Industrial Chemistry & Materials*, vol. 3, no. 5, pp. 607–617, 2025.

[25] A. Leiden, S. Kölle, S. Thiede, K. Schmid, M. Metzner, and C. Herrmann, "Model-based analysis, control and dosing of electroplating electrolytes," *The International Journal of Advanced Manufacturing Technology*, vol. 111, no. 5-6, pp. 1751–1766, 2020.

10.11 Exercises

10.1 Concept check · 2 min

In scenario 4 on feedback only, the thick-copper lot starts at 4:00, yet the accelerator result reported at 4:40 reads 5.96 mL/L, close to its setpoint of 6.00. Why does that result not show the lot?

(a) The outlier gate rejected the low result and held the last good one
(b) Its sample left the bath just before the first thick-copper wafer plated
(c) The EWMA weight of 0.5 hides half of every change in a new result
(d) The accelerator does not fall until the breakdown products start to rise

10.2 Concept check · 1 min

Section 10.4 finds that in this bath a dump cycle at a setpoint B_{SP} never costs less than bleed-and-feed holding the same B_{SP}. Which reason is the one the make-up balance gives?

(a) A dump throws away the whole accelerator inventory, which costs more than the make-up
(b) The supervisor charges every void plated during the remake to the dump
(c) A dump cycle spends the bleed's make-up per hour and adds the downtime
(d) The bleed removes the breakdown products at a higher concentration than a dump does

10.3 Concept check · 2 min

In the convention of Section 1.6.2, how is the PI that sets the bleed from the estimated breakdown products $\hat{B}$ classified?

(a) Reverse acting, action $+1$, because the bleed lowers the breakdown products it controls
(b) Reverse acting, action $+1$, because its output bias is the steady bleed b_{ff}
(c) Neither, because the bleed is a feedforward of $\hat{G}V/B_{SP}$ with no feedback on $\hat{B}$
(d) Direct acting, action -1, because the bleed must rise when $\hat{B}$ is above its setpoint

10.4 Theory and calculation · 7 min

Use Equation (10.4) with $F = 96485\,\mathrm{C/mol}$, $M_{Cu} = 63.546\,\mathrm{g/mol}$, $\rho_{Cu} = 8.96\,\mathrm{g/cm^3}$, $A = \pi(15\,\mathrm{cm})^2$, $\eta_c = 0.98$ and the waveform of Table 10.2.

(a) Find κ_F to three decimals, the charge a 1.00 µm standard wafer needs, and its bulk step at 14 A. *Checkpoint:* 0.534 A h per µm; 0.545 A h; 125.2 s.

(b) A thick-copper wafer of 2.00 µm runs its bulk step at 28 A. Find its bulk step, and the thickness it gets in case 1 if the manual bulk step stays at the standard wafer's unrounded value of part (a). By what percentage is it thin? *Checkpoint:* 132.7 s; 1.893 µm; 5.35 %.

(c) One wafer leaves every 216 s. Find the cycle-average current for standard and for thick-copper wafers. *Checkpoint:* 9.09 A and 18.17 A.

10.5 Theory and calculation · 8 min

Take the standard load of 9.09 A, $V = 200\,\mathrm{L}$, the accelerator at its make-up strength $S_V = 6.0\,\mathrm{mL/L}$, $c_S = 5.0\,\mathrm{mL}$ per A h and $d_{S,25} = 1.0\,\%/\mathrm{h}$ at 25 °C.

(a) From Equation (10.1), find the accelerator consumed at the electrodes and broken down in the bulk, in mL/h. Why do the bleed and its virgin make-up drop out at this state? Find the dose that holds the bath and the turnover time of the 1200 mL inventory. *Checkpoint:* 45.4 mL/h and 12.0 mL/h; 57.4 mL/h; 20.9 h.

(b) The feedforward uses $\hat{c}_S = 4.75\,\mathrm{mL}$ per A h. Find F_S from Equation (10.8) and the trim that makes up the rest. *Checkpoint:* 55.16 mL/h; 2.27 mL/h.

(c) From Equation (10.5), find the window's half-width as a percentage of the ratio, $e^{w(B)} - 1$, at $B = 22$ and 27 mL/L. After the overdose of case 8 the leveler is 4.5 mL/L. Find the new center r_c and the lower edge at $B = 22\,\mathrm{mL/L}$, to two decimals, and say where a bath held at a ratio of 3.0 sits. *Checkpoint:* 9.8 % and 3.6 %; 3.67 and 3.35; below the window.

10.6 Theory and calculation · 8 min

Use the report-to-report model of Section 10.6 at the operating point of case 3 after the new lot: $c_S = 6.25\,\mathrm{mL}$ per A h, $I = 9.09\,\mathrm{A}$, $S_V = 6.0\,\mathrm{mL/L}$, a breakdown of 1.0 %/h, a bleed of 2.28 L/h, a drag-out of 0.5 L/h, $V = 200\,\mathrm{L}$, $T = 2\,\mathrm{h}$ and $\theta = 42\,\mathrm{min}$.

(a) Find k in L/h and the time constant $\tau = V/k$. *Checkpoint:* 14.24 L/h; 14.0 h.

(b) Find a, β_0 and β_1 of Equation (10.11), with the β in h/L to three significant figures, and check that $\beta_0 + \beta_1 = (1-a)/k$. *Checkpoint:* $a = 0.867$; $\beta_0 = 6.21 \times 10^{-3}\,\mathrm{h/L}$, $\beta_1 = 3.11 \times 10^{-3}\,\mathrm{h/L}$; both sides 9.32×10^{-3} h/L.

(c) Find the share $\beta_1/(\beta_0 + \beta_1)$ for $\theta = 42\,\mathrm{min}$ and for $\theta = 60\,\mathrm{min}$, and use it to explain why the limit falls from 210 to 141 as the age grows from 42 min to an hour. *Checkpoint:* 0.334 and 0.482.

10.7 App challenge · 8 min

Open scenario 4, *4 · A lot mix change*, at the default seed, and run each day at a speed of 60 min/s (24 s a day). Read the scorecard tiles and the exported CSV, whose time column is `Time (h)`.

(a) Run the full 24 h as shipped. Read the Voids and Worst fill margin tiles, and the accelerator dose `dose_S_mLh` in the CSV rows at 4.0 and 6.0 h, within 0.3 mL/h. *Checkpoint:* 0 voids, 0.23; 57.4 mL/h and 101.8 mL/h.

(b) Press Reset, set Dosing strategy to FB only on the DOSE-100 card before pressing Run, and run 24 h. Read the two tiles, the CVS reading `S_cvs_mLL` in any row from 4.7 to 6.6 h and from 6.7 to 8.6 h, within 0.01 mL/L, and the first row in which `metric_voids` exceeds zero, within 0.05 h. *Checkpoint:* 97 voids, −0.06; 5.96 mL/L and 5.70 mL/L; 8.1 h.

(c) Explain from the sample times why the dose on feedback only first responds to the lot at 6:40 (the step of under 1 mL/h at 4:40 answers a result from before the lot), and why the voids continue for hours after that. *Checkpoint:* The 4:40 result describes the bath of 3:58; the trim then adds about 10 mL/h a report against the roughly 45 mL/h the lot needs.

10.8 App challenge · 8 min

Open scenario 3, *3 · CVS trim on a sampled, delayed analysis*, at the default seed, and run each day for 24 h at 60 min/s.

(a) Run it as shipped. Then press Reset, set Accelerator trim gain on the AIC-S card to 0 before pressing Run, and run again, and repeat with 300 and with 400. Record the Voids and Worst fill margin tiles of the four days. *Checkpoint:* 0 and 0.13; 163 and −0.09; 16 and −0.04; 37 and −0.21.

(b) In the CSV of the day at 400, find the highest and the lowest true accelerator `S_true_mLL` after 12 h and their times, within 0.02 mL/L and 0.1 h. *Checkpoint:* 6.53 mL/L at 16.7 h; 5.12 mL/L at 22.7 h.

(c) Read ρ at gains of 300 and 400 from the solid curve of Figure 10.3, and explain why the larger gain voids more wafers although it corrects harder. *Checkpoint:* About 1.07 and 1.13, both above one.

10.9 App challenge · 8 min

Open the challenge, *Challenge · a day of 400 wafers*, at the default seed, and run each day for 24 h at 60 min/s. Use Equation (10.6) with the 1868 mL of breakdown products the day forms under every bleed policy.

(a) Run it as shipped and read the Voids, Worst fill margin, Virgin make-up and Day target tiles. *Checkpoint:* 226, −0.37, 12.0 L, not met.

(b) Press Reset. Before Run, set Dosing strategy to FF + trim (the Outlier gate stays On), Bleed policy to Hold B and Breakdown products setpoint to 24 mL/L. Run the day and read the same tiles and the true breakdown products `B_true_mLL` in the last CSV row. Compute the mean concentration at which the breakdown products left, to one decimal. *Checkpoint:* 0, 0.12, 57.0 L, met; 24.55 mL/L; 23.8 mL/L.

(c) Repeat with the setpoint at 26 mL/L, and compute the make-up saved per mL/L of setpoint between 24 and 26. *Checkpoint:* 0, 0.05, 38.7 L, met; 26.48 mL/L, 25.1 mL/L; about 9.2 L per mL/L.

10.10 Application area · 4 min

A process engineer must choose the breakdown-product setpoint for the day of the challenge, with virgin make-up at $20 per liter.

(a) From Table 10.6, price a day of make-up at setpoints of 22, 24 and 26 mL/L. *Checkpoint:* $1664, $1140 and $774.

(b) Which setpoint would you run, and what does the seed spread of Section 10.7 add to the decision? *Checkpoint:* 24: across ten other seeds its worst margin stays between 0.083 and 0.176.

10.11 Application area · 3 min

In case 8 a leveler pump runs away and puts 300 mL into the bath.

(a) Why can the leveler loop not undo the overdose, and how does the leveler leave the bath? *Checkpoint:* A pump can only add; the leveler leaves by consumption, breakdown and bath taken out.

(b) Compare the two recoveries of Section 10.7 in voids, wafers plated and make-up, and name one safeguard that would have prevented the event. *Checkpoint:* Bleed 4 L/h: 197 voids, 400 wafers, 100.7 L; dump: 2 voids, 366 wafers, 216.5 L.

10.12 Review the agent · 5 min

An AI assistant reviews scenario 3 and writes: *"The trim decays by only 0.900 per analysis at its gain of 30, so it takes most of a day to correct the new accelerator lot. The CVS repeats to 2 %, which is precise enough to act on hard. Raise the trim gain to 300: the loop will then correct the lot within one or two analyses, and a thick-copper lot will be caught the same way."* Find the flaw, and say how you would check the claim.