INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 14 Control in Semiconductor Manufacturing

CHAPTER 8

An Ion Source Simulator for Cascade Control and Recovery From a Measurement That Stops Meaning Anything

Pages 188–211

Edited by John D. Hedengren

How to cite this chapter. "An Ion Source Simulator for Cascade Control and Recovery From a Measurement That Stops Meaning Anything," in *Control in Semiconductor Manufacturing*, John D. Hedengren, Ed., vol. 14 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 188–211.

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

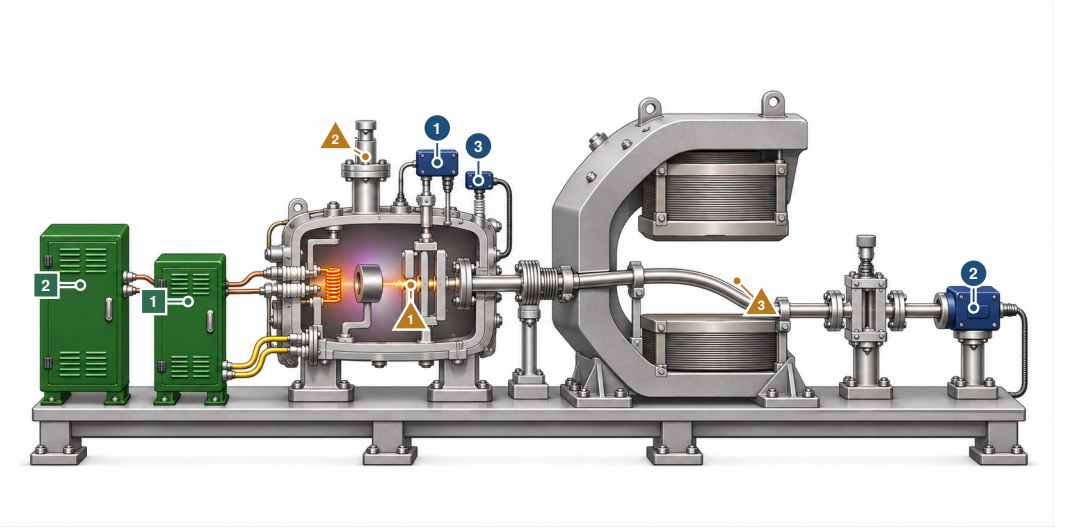
The app for this chapter
https://apmonitor.com/apps/ion-implant-source/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v14-semiconductor/#ion-implant-source

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

CHAPTER 8

An Ion Source Simulator for Cascade Control and Recovery From a Measurement That Stops Meaning Anything



The controller moves 1 bias supply, heats the cathode 2 filament supply

The controller reads 1 extraction current 2 selected beam (Faraday cup) 3 gap voltage and supply current

What disturbs it 1 extraction spark 2 gas flow below its reading 3 beam lost in the magnet

Notice the spark at the extraction gap: for a few milliseconds the beam readings collapse while the hot filament and cathode beside it do not change. Illustration, not to scale.

IN PLAIN TERMS **When a Reading Stops Telling the Truth**

The main idea. Chip makers push boron atoms into silicon wafers with a narrow beam of charged atoms. A hot source makes the beam, and automatic controls keep it steady by turning the source's heating power up or down.

The challenge. Now and then the high voltage that pulls out the beam sparks, and the beam reading drops for a moment while the source is fine. A control that trusts the reading turns the heat up, and the beam overshoots when the voltage returns.

Why it matters. Holding the controls still through the spark, then putting the missed boron back only where it was missed, keeps the wafer good.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 8.1 and 8.2, the fact box of Section 8.3, the try-it and fact boxes of Section 8.6, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 8.6, Section 8.7, Section 8.8 and its open challenge

The control question A spark across the extraction gap halves the beam voltage in half a millisecond while the cathode, with a time constant of seconds, notices nothing. What should a cascade on the beam do with a measurement that has stopped meaning anything, and what does each choice cost in source heat, beam overshoot and missed dose?

Learning objectives

- Name the current each loop of the three-level cascade reads and moves, and explain why the fast loop moves bias current rather than filament current.
- Explain why back-calculation anti-windup does not stop a loop from chasing a spark, from the heat the cathode stores.
- Compute how long supply-current surge and gap collapse coincide, and choose a detection window and a measurement gate from it.
- Compare no repaint, a whole-pass repaint and a local repaint driven by a per-strip dose ledger.
- Judge when a retry limit should turn automatic restarts into a trip.

Prerequisites Sections 1.6.2 (velocity form, anti-windup, bumpless transfer and loop action), 1.4.2 (measurements that stop meaning anything), 1.6.1 (control period, speed and seed) and 1.6.4 (scenarios, scorecard and CSV).

The app https://apmonitor.com/apps/ion-implant-source/

See also Volume 2, Chapter 2 for the velocity-form PID with back-calculation, and Volume 2, Chapter 3 for windup at a true actuator limit. Volume 3, Chapter 2 for matching loop structure to where a disturbance enters, and Volume 3, Chapter 3 for the speed an inner loop must reach. Volume 8, Chapter 13 for a product history kept apart from the instruments that act on it.

8.1 Introduction

Every feedback loop carries an assumption that is almost never written down, which is that the measurement still represents the plant variable the loop is meant to regulate. Saturation does not violate that assumption and neither does noise, and the standard remedies for both leave it intact. A protection event violates it completely while leaving every quantity the controller can see numerically valid.

The clearest place to see it is an ion implanter. A tungsten filament at 2600 K heats a separate cathode at 2200 K by electron bombardment, the cathode fills an arc chamber with boron trifluoride ions, an extraction electrode at 10 kV pulls them out, and an analyzer magnet keeps one mass in five. When that extraction gap sparks, the supply current goes to its limit, the gap voltage halves in about half a millisecond, and the beam goes with it. The two thermal masses notice nothing. A proportional-integral controller on extraction current then reads a thirty-five milliamp error in a variable it is responsible for and does the only thing it knows, which is to raise source power. Measured on the simulator of this chapter, that drives the bias command from 2.493 A to its 4.0 A limit, adds 28.2 K to the cathode, and overshoots the beam by 14.75 % when high voltage returns.

> ⧖ HISTORY **From Curiosity to Production Doping (1950s to 1978)**
>
> Ion beam doping of semiconductors was developed from the early 1950s, at first as an experimental curiosity of solid-state physics. Early workers were drawn to the damage that ions do to a crystal and overlooked the real prize: dopant atoms placed with control and then activated by annealing above 800 °C. Device engineers with manufacturing problems to solve brought it into transistor production, where it gave good control of where a junction sits and let threshold voltages be tailored [1]. By 1978 implanters are considered to have come of age [2]. Some forty years into its development the industry was selling about two billion dollars of implanters a year and had installed thousands of machines [3].

Back-calculation anti-windup does not prevent any of it, and the reason is worth stating precisely. Back-calculation answers the question of whether the actuator has run out of room [4, 5]. During the

dropout it has, for 32 ms, and back-calculation duly holds the integral below 0.23 A. The proportional term alone already asks for the limit, and the cathode stores the heat. The actuator is being spent on an error that does not exist, and no bound on the integrator changes that. The remedy belongs one layer up, in a measurement-quality gate that invalidates the signal on corroborating evidence, an explicit output hold, a frozen rather than merely limited integrator, and an integral state reconstructed from the held output before the loop is handed back [6, 7].

The same structure appears wherever a protection system takes a process temporarily outside its controller's domain. A compressor on anti-surge recycle (Volume 4, Chapter 2), a fired heater in a burner management sequence, a reactor after a trip and a converter in a fault-current limit all produce measurements that keep arriving and stop being evidence. The implanter is the cleanest teaching case because the separation of time scales is extreme. The fault is five milliseconds long and the plant it appears to describe has a four second time constant, so there is no way to confuse the two.

There is a second lesson underneath. Four quantities in an implanter are called current and no two of them are the same thing. At the nominal recipe the source discharge runs at 10.0 A, the extraction supply delivers 35.00 mA, and 7.918 mA of mass-selected $^{11}B^{+}$ reaches the wafer. A twelve percent analyzer detune moves the third by −12.00 % and leaves the second unchanged. One second later the outer loop has raised the second by 0.25 % while the third is still 11.76 % low, so an outer loop reading only the Faraday cup drives the source harder to fix a beamline problem.

This chapter makes measurement-quality gating measurable on an interactive ion source. With the same gains and the same spark the fault-chasing index is 7.46×10^{-2} A s without the gate and 8.28×10^{-3} A s with it, and a local dose repaint returns a wafer to 0.21 % strip non-uniformity and +0.21 % dose where a whole-pass repaint overdoses it by 16.58 %. The chapter follows three threads. The first is a sixteen-state reduced source and beamline model with a switched-conductance high-voltage fault, a coincidence detector at the physics step and a three-level cascade, written closely enough to rebuild. The second separates two judgments that are usually made together, the latching of a fault and the invalidation of one measurement, and measures the timing cost of conflating them in actuator movement. The third is a scenario pair in which one plant, one seed and one spark produce 14.75 % and 0.29 % of beam overshoot under two architectures that differ by one decision. Anti-windup, cascade control and the Child-Langmuir law are all standard. What is particular to the implanter is how they combine.

Section 8.2 places the problem in its domain. Section 8.3 gives the source and beamline model, and Section 8.4 the three loops, the detector and the supervisor as implemented. Section 8.5 separates the fault latch from the measurement gate, Section 8.6 reports the measured outcomes of the twelve scenarios, Section 8.7 names the real sensors, supplies and controller platform, and Section 8.8 states what the model leaves out. The Exercises at the end of the chapter use the scenarios.

8.2 Sources, Sparks and Held Outputs

The ion sources used in implantation descend from the isotope separators of the 1960s [8], and the History box below tells how the first filament source came to dope silicon. Its weakness is that the filament is the plasma electrode, so filament erosion and plasma conditions are the same variable.

The indirectly heated cathode architecture separates them. A filament behind a shield is heated by its own current, its electrons are accelerated across a few hundred volts onto a separate cathode, and that bombardment heats the cathode to emission temperature. The cathode is then the plasma electrode and the filament is a heater, which gives longer life and a cleaner control handle [9]. Langmuir probe measurements inside such a source relate electron temperature and plasma density to the source control parameters that an operator actually sets [10], and the fluorine chemistry of a boron trifluoride feed both feeds the process and attacks the hardware, which is why source lifetime is a live engineering subject rather than a settled one [11].

HISTORY **From Isotope Separators to Chip Doping (1963 to 1992)**

In 1963 a new ion source was described for electromagnetic isotope separators [8], machines that sort the ions of an element by mass to purify one isotope [12]. Its heated filament, a straight tungsten wire placed parallel and close to the extraction slit, drives the arc discharge, and the wear of that wire by evaporation and sputtering limits the life of the source [12]. Thirty years later it was still the most widely used ion source. Its sharply focused milliampere beams of most elements had made it a source of choice for the production doping of silicon devices, and filament lifetime was still a subject of study [13].

The physics that limits extraction is older than either source. The three-halves power dependence of space-charge-limited current on accelerating voltage was established for thermionic diodes in the 1910s [14, 15], and the temperature dependence of thermionic emission itself was settled in the same decade [16]. Both appear in this model in normalized form, because a control study needs the exponent and the bound rather than the absolute current.

On the control side the ingredients are old and the combination is not. Integrator windup and the back-calculation remedy were given their standard treatment in 1989 [4] and placed in a unified framework five years later [5]. Anti-windup, bumpless transfer and reliable switching were later shown to be one problem viewed three ways [6], which is exactly the observation this application is built on. Fault-tolerant control supplies the vocabulary for what happens when a signal is present but no longer informative [7]. What the implanter literature adds is the consequence on the product, because an implant specification is an integral rather than a level and an interruption therefore leaves a spatial defect. Scanned beam uniformity is controlled by the scan itself rather than by the source [17], and a fault handler that did not record where the scanner was cannot repair what the fault cost [18].

8.3 Process Model

Boron trifluoride enters arc chamber AC-101 through mass flow controller FC-101. Tungsten filament FIL-101 is heated by supply PS-F at current I_f. Filament electrons are accelerated onto cathode CA-101 across bias supply PS-B at V_b and current I_b, and the bombardment power $V_b I_b$ is what heats the cathode. Cathode electrons are accelerated into the gas across the source arc potential V_a and confined by axial field B, and repeller RE-101 returns them through the chamber. Extraction electrode EX-101 at V_x and suppression electrode SU-101 at V_s form the beam, analyzer magnet AM-201 bends it, resolving slit RS-201 cuts it, and the scanned wafer is the load. The equipment and every symbol in the equations are drawn in Figure 8.1.

DID YOU KNOW? **Hundreds of Hours on One Cathode**

In 1998 an indirectly heated cathode source ran several times longer than sources with a bare filament cathode. Its cathode lasted from 70 hours at the highest discharge power to over 500 hours in moderate operation, and near 1 kW of discharge power its ions were mostly atomic rather than molecular [9].

Table 8.1 collects the three controlled variables, the three manipulated variables and what acts against them. Time is carried in seconds, currents in amperes for the supplies and milliamperes for the beam, temperatures in kelvin, electrode voltages in kilovolts and electron temperature in electronvolts. Sixteen states are integrated at $\Delta t = 0.2\,\mathrm{ms}$, the fault detection and the dose ledger run at that same step, and the three controllers run at 1 ms, 20 ms and 100 ms.

8.3.1 Two Thermal Masses in Series

The filament heats itself. Its resistive drive is quadratic in current and scaled by a drift factor r_f that represents thinning with age, and the cathode is heated by the bombardment power delivered across the bias potential with a small conduction term from the filament, $k_{fc} = 0.12$, giving Equations (8.1)

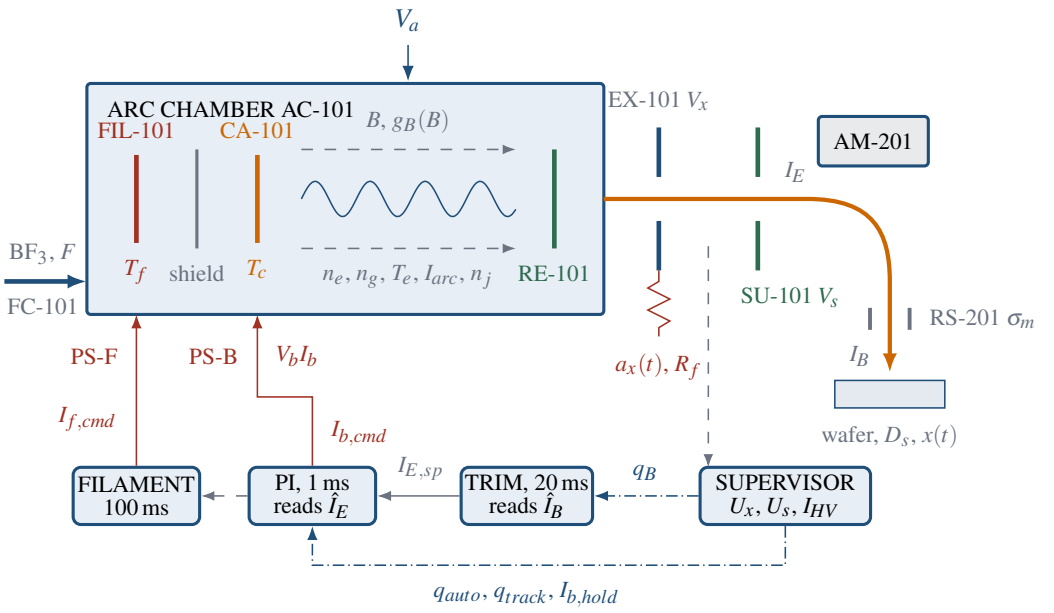


Figure 8.1 The indirectly heated cathode source, the beamline and the control path. Bombardment power V_bI_b heats the cathode, the cathode fills the chamber, and extraction pulls a mixed beam that the analyzer magnet and the resolving slit reduce to one mass. The zigzag at the extraction electrode is the switched conductance $a_x(t)$ through fault resistance R_f that constitutes a spark. The inner loop reads extraction current and sets the bias supply, the trim reads the selected beam and sets the inner setpoint, and the filament supervisor reads the bias command and sets the filament supply. The supervisor latches faults on the high-voltage evidence alone, reads the extraction current only to qualify a recovery, and owns the permissives that decide whether the inner loop and the trim may act on their measurements. In the control path, supply commands are drawn solid, signals read by a controller dashed and permissives dash-dotted.

and (8.2).

$$\tau_f \dot{T}_f = T_w + r_f \left(\frac{I_f}{I_{f,0}} \right)^2 (T_{f,0} - T_w) - T_f \tag{8.1}$$

$$\tau_c \dot{T}_c = T_w + \frac{V_b I_b}{V_b I_{b,0}} \left[1 + k_{fc} \left(\frac{T_f}{T_{f,0}} - 1 \right) \right] (T_{c,0} - T_w) - T_c \tag{8.2}$$

Both are written so that the nominal recipe is a fixed point, which makes $T_{f,0}$ and $T_{c,0}$ operating points rather than fitted constants. Neither state can move measurably in the five milliseconds a spark lasts, and that single fact is what the whole fault-handling design rests on.

8.3.2 Emission, Written as a Ratio

Direct evaluation of Richardson-Dushman emission at two thousand kelvin is numerically stiff and the control problem only ever needs the ratio to the reference, so the model carries the normalized form of Equation (8.3) with $\theta = T_c / T_{c,0}$.

$$\tilde{I}_e = \mathrm{clip}\left[\theta^2 \exp\left(\beta_R \left(1 - \frac{1}{\theta} \right) - \gamma_w w \right), 0, I_{e,max} \right] \tag{8.3}$$

The exponent β_R stands for $\phi / k_B T_{c,0}$ and is an adjustable parameter rather than a measured work function. At the shipped value of 9 the emission is 0.562 at 2090 K, 1.000 at 2200 K, 1.692 at 2310 K

Table 8.1 Controlled, manipulated and disturbance variables

Symbol	Role	Design value	Range
I_E	CV, inner, 1 ms	35.00 mA	10–80 mA
I_B	CV, outer, 20 ms	7.918 mA	1–12 mA
I_b	CV, supervisor, 100 ms	2.500 A	0–4.0 A
$I_{b,cmd}$	MV, inner loop	2.500 A	0–4.0 A
$I_{E,sp}$	MV, outer loop	35.0 mA	10–80 mA
$I_{f,cmd}$	MV, supervisor	155 A	100–190 A
T_f	state, unmeasured	2600 K	—
T_c	state, measured	2200 K	trip 2560 K
n_g	state, unmeasured	1.000	normalized
n_e	state, unmeasured	1.000	normalized
I_{arc}	state, measured	10.0 A	normalized $\times 10$
n_j	state, five species	mix of T_e	normalized
U_x	state, measured	10.0 kV	2–30 kV
U_s	state, measured	2.0 kV	0.5–5 kV
w	state, unmeasured	0 clean	0–1
E_{fault}	state, ledger	0 J	—
g_F	DV, unmeasured	1.000	0.80–1.20
η_{tr}	DV, unmeasured	1.000	0.70–1.05
a_x, a_s	DV, switched	0	$\{0,1\}$
D_s	product, 32 strips	1.000 at end	—

and 2.742 at 2420 K. Five percent of temperature is sixty-nine percent of emission, which is a plant whose gain cannot be quoted without also saying where it is standing. The deposit state w raises the effective work function as the source ages.

The bias supply is limited by what the filament can actually emit, $I_{b,act} = \min(I_{b,cmd}, I_{em,f})$, and both supplies are first order lags of time constants τ_{Ib} and τ_{If} toward their commands.

8.3.3 Gas, Electrons and the Intended Discharge

The normalized gas density follows the indicated flow through a calibration factor, the electron density follows the cathode emission through a bounded confinement factor, and the intended source discharge current is a first order proxy driven by all of them.

$$\tau_g \dot{n}_g = \frac{g_F F}{F_0} - n_g, \qquad \tau_e \dot{n}_e = \tilde{I}_e \, g_B(B) - n_e \tag{8.4}$$

$$\begin{aligned} g_B(B) &= \frac{B/B_0}{1 + B/B_0} \Big/ \frac{1}{2} \\ \tau_{arc} \dot{I}_{arc} &= \frac{V_a}{V_{a,0}} n_e n_g \, g_B(B) - I_{arc} \end{aligned} \tag{8.5}$$

The confinement factor of Equation (8.5) is normalized to unity at the nominal field and saturates by construction, because more field always raises the number of ionizing collisions and never without limit. The axial field bends the electrons into helices and the repeller returns them, and the model represents that as a bounded scalar rather than by solving trajectories.

8.3.4 One Gas, Five Ions

Five lumped species are carried, BF_3^+, BF_2^+, BF^+, B^+ and F^+. Each relaxes on a common time constant toward its share of the ionization rate $n_e n_g$, and the share is set by an effective yield and a

reduced rate coefficient.

$$\tau_i \dot{n}_j = f_j^*(T_e)\, n_e n_g - n_j, \qquad k_j(T_e) = \sqrt{T_e}\, e^{-E_j/T_e}$$
$$f_j^*(T_e) = \frac{y_j b_j k_j(T_e)}{\sum_\ell y_\ell b_\ell k_\ell(T_e)} \tag{8.6}$$

The yields y_j are 0.55, 1.00, 0.45, 11.5 and 0.35 in the order listed, and b_j is one except for B^+, where it is the B^+-fraction disturbance. The thresholds E_j are the gas-phase appearance energies rounded to a tenth of an electronvolt, 15.7 eV for BF_3^+, 15.9 eV for BF_2^+, 24.5 eV for BF^+, 31.0 eV for B^+ and 17.4 eV for F^+. They order the onsets and are not a claim about cross sections. The effective electron temperature is a bounded function of the arc potential, the gas loading and the deposits, with $k_V = 0.035\,\mathrm{eV/V}$, $k_p = 0.9\,\mathrm{eV}$ and $k_w = 1.8\,\mathrm{eV}$,

$$T_e = \mathrm{clip}[T_{e,0} + k_V \Delta V_a - k_p(n_g - 1) - k_w w,\ 2,\ 12] \tag{8.7}$$

in electronvolts, with $\Delta V_a = V_a - V_{a,0}$. At the nominal recipe $T_e = 6.00\,\mathrm{eV}$ and the mix is 19.8 % BF_3^+, 34.8 % BF_2^+, 3.7 % BF^+, 32.3 % B^+ and 9.5 % F^+. Raising the arc potential to 75 V takes T_e to 6.88 eV and the mix to 17.4, 30.7, 4.0, 39.3 and 8.6 %, so the parent falls and B^+ rises fastest, as the thresholds require.

8.3.5 Extraction and Mass Selection

All species share one aperture, so the extracted ion current is the smaller of what the plasma supplies and what the gap will pass.

$$I_{pl} = K_{pl} \sum_j n_j \sqrt{\frac{T_e}{T_{e,0}}}, \qquad f_j = \frac{n_j}{\sum_\ell n_\ell} \tag{8.8}$$

$$I_{CL} = 1.45\, I_{E,0} \left(\frac{U_x}{V_{x,0}}\right)^{3/2} \sqrt{\frac{m_{ref}}{m_{eff}}}, \quad m_{eff}^{-\frac{1}{2}} = \sum_j \frac{f_j}{\sqrt{m_j}} \tag{8.9}$$

$$I_{ion} = \min\left(I_{pl}, I_{CL}\right), \qquad I_E = I_{ion} + I_{par}\frac{U_x}{V_{x,0}} \tag{8.10}$$

The masses m_j are those of the ^{11}B ions, 68, 49, 30, 11 and 19 u, and $m_{ref} = 30\,\mathrm{u}$. The gain K_{pl} is calibrated so that the nominal recipe produces exactly the declared extraction current. At that point $m_{eff} = 25.13\,\mathrm{u}$ and the plasma supplies 33.20 mA against a space-charge limit of 55.45 mA, so extraction is plasma limited and raising the extraction voltage buys essentially no extra current while changing the implant energy. That inequality is the whole argument for recipe-locking V_x and manipulating I_b instead.

Beam optics collapse with the extraction field, and below about a third of the tuned voltage nothing useful is transported, $\kappa = \mathrm{clip}[(U_x/V_{x,0} - 0.30)/0.45, 0, 1]$. Each boron-bearing species splits into an ^{11}B and an ^{10}B component in the ratio f_{11}, the resolving slit passes a Gaussian window in mass-to-charge ratio, and the selected beam is Equation (8.11).

$$S_{jv} = \exp\left[-\frac{(m_{jv}/Z_j - \mu_{set})^2}{2\sigma_m^2}\right]$$
$$I_B = \kappa\, \eta_{tr} \sum_{j,v} S_{jv}\, f_v f_j\, I_{ion} \tag{8.11}$$

At the nominal recipe with $\mu_{set} = 11$ this gives 7.918 mA, which is 22.6 % of the extraction current.

8.3.6 Deposits and the High-Voltage Fault

A single lumped state represents conductive and corrosive deposits, $\dot{w} = k_{dep} n_g s - k_{clean} w$ with s the deposit-rate slider. It raises the effective work function in Equation (8.3) and it raises a stochastic glitch hazard, $\lambda_g = \lambda_0 \exp[\beta_w w + \beta_V (U_x / V_{x,0} - 1)]$. That hazard exists to make faults more likely in degraded scenarios and is never a predicted time to arc.

Each electrode is a first order supply whose time constant is switched by what is happening to it, Equation (8.12).

$$\dot{U}_k = \frac{U_{k,sp}^{eff} - U_k}{\tau_k}, \qquad \tau_k = \begin{cases} \tau_{coll} & a_k = 1 \\ 3\tau_{coll} & \text{HV inhibited} \\ t_r/3 & \text{recovering} \\ \tau_{HV} & \text{otherwise} \end{cases} \tag{8.12}$$

for $k \in \{x, s\}$, with $U_{k,sp}^{eff}$ zero while a fault is conducting or high voltage is inhibited. The supply current is the beam load plus the fault current, limited by the supply,

$$\begin{aligned} I_{HV} &= \min\left[I_{max},\; I_E + \sum_k a_k \frac{U_k}{R_f} \right] \\ \dot{E}_{fault} &= \sum_k a_k \frac{U_k^2}{R_f} \end{aligned} \tag{8.13}$$

which produces the signature the detector is written around. The supply current goes to its limit and the gap voltage collapses, both within a millisecond, while nothing thermal moves at all.

8.3.7 Scan and Dose

The wafer is thirty-two strips crossed by a scan at position $x(t)$ advancing at $1/t_{scan}$ per pass. Two ledgers are kept per strip. The committed dose X_s accrues whether or not the beam is on, and the delivered dose D_s accrues only when it is, Equation (8.14).

$$\dot{X}_s = \frac{\delta_s N_{st}}{n_{pass} t_{scan}}, \qquad \dot{D}_s = b_w(t) \frac{I_B}{I_{B,sp}} \frac{\delta_s N_{st}}{n_{pass} t_{scan}} \tag{8.14}$$

where δ_s is one while the scan is over strip s and zero otherwise, and b_w is the physical beam-enabled indicator, zero while high voltage is inhibited or below 35 % of its setpoint and, during a local repaint, over strips already at their committed dose. The deficit $\max(0, X_s - D_s)$ is what a repaint has to restore, and keeping the two ledgers separate is what makes the difference between putting dose back where it went missing and adding a pass everywhere. The committed ledger does not accrue during a repaint, because a repaint pass delivers dose the ledger already owed. The delivered ledger follows the simulated true beam and the physical beam-enabled indicator and deliberately not the controller's measurement-valid flag, so a dead instrument never erases dose that was actually delivered.

8.3.8 Initialization and Numerics

A warm scenario initializes at the nominal recipe, $T_f = 2600\,\mathrm{K}$, $T_c = 2200\,\mathrm{K}$, $I_f = 155\,\mathrm{A}$, $I_b = 2.5\,\mathrm{A}$, $n_g = n_e = 1$, n_j at the mix of $T_e = 6\,\mathrm{eV}$, $U_x = 10\,\mathrm{kV}$ and $U_s = 2\,\mathrm{kV}$. A cold scenario sets every one of those to zero or to the wall temperature and the supervisor to OFF. Integration is fixed-step fourth order Runge-Kutta at 0.2 ms. The fast task runs after every integration step and owns the beamline evaluation, the fault model, the coincidence detector, the supervisor transitions and the dose ledger, because hardware protection has to act before any supervisory controller can. The historian logs at 100 Hz, which is coarse enough that the sub-millisecond shape of a fault is not on the strip charts and fine enough that a fifteen millisecond event always lands on at least one sample. The application carries a separate fast-event display for exactly that reason. It keeps the last 1250 integration steps, the 250 ms its axis shows.

Table 8.2 Model parameters. All are adjustable in the model dialog unless marked hidden or fixed.

Symbol	Meaning	Value	Unit
$I_{f,0}$	Nominal filament current	155	A
$I_{b,0}$	Nominal bias current	2.5	A
V_b	Bias voltage	350	V
V_a	Source arc voltage	50	V
B	Source magnetic field	1.0	normalized
F_0	Nominal BF_3 flow	3.5	sccm
f_{11}	^{11}B fraction	0.80	—
$V_{x,0}$	Extraction voltage	10	kV
$V_{s,0}$	Suppression magnitude	2.0	kV
$I_{E,0}$	Extraction current, nominal	35	mA
σ_m	Resolving slit window	0.35	u
η_{tr}	Analyzer transmission	0.92	—
I_{par}	Parasitic extraction current	1.8	mA
$T_{f,0}$	Filament temperature	2600	K
$T_{c,0}$	Cathode temperature	2200	K
T_w	Wall temperature (hidden)	500	K
$T_{c,max}$	Overtemperature trip	2560	K
β_R	Richardson exponent	9.0	—
τ_f	Filament thermal	4.0	s
τ_c	Cathode thermal	1.5	s
τ_g	Gas response	0.08	s
τ_e	Electron response	0.002	s
τ_i	Ion species response	0.010	s
τ_{arc}	Discharge response	0.005	s
τ_{Ib}	Bias supply (hidden)	0.005	s
τ_{If}	Filament supply (hidden)	0.05	s
$T_{e,0}$	Electron temperature	6.0	eV
τ_{HV}	Normal supply response	1.0	ms
τ_{coll}	Collapse time constant	0.8	ms
R_f	Fault path resistance	5.0	kΩ
I_{max}	Supply current limit	120	mA
λ_0	Baseline glitch hazard	0.004	1/s
N_{st}	Scan strips (fixed)	32	—
n_{pass}	Passes per wafer	6	—
t_{scan}	Time per pass	2.0	s
e_D^{max}	Mean dose tolerance	1.5	%
U_D^{max}	Strip non-uniformity limit	1.0	%
σ_{I_E}	Extraction measurement noise	0.12	mA
σ_{I_B}	Faraday noise	0.05	mA
τ_{meas}	Measurement filter (hidden)	2.0	ms

8.3.9 Assumptions

This is a lumped control model. It is not a particle-in-cell code, not a Boltzmann solver, not a plasma-chemistry model and not a beam transport simulation, and its purpose is the signs, the nonlinearities, the time-scale separation and the fault logic.

The five species use effective yields and appearance-energy thresholds. The thresholds are published values used to order the onsets, and the yields are chosen so that the nominal recipe produces a plausible $^{11}B^+$ beam. Absolute species fractions are not predicted, and multiply charged ions, neutral fragments, charge exchange and wall recombination are absent.

A glitch is a switched conductance across a supply with a current limit. Real arcs depend on geometry, contamination, pressure, surface state, supply topology, electromagnetic coupling and the

protection hardware itself.

Hardware protection in a real implanter acts locally and far faster than any supervisory controller. The application demonstrates the sequencing and its control consequences, and every threshold, dwell and retry count needs equipment-maker validation before it means anything.

Signal names vary between equipment makers. The cathode shield in Figure 8.1 is passive and protective, and the electrically active electrode opposite the cathode is labeled repeller, which is one of several conventions in use.

The dose map is one dimensional with a single beam-enabled indicator. Two dimensional scan patterns, beam profile, neutralization, wafer charging, angle control and channeling are outside the model, and the twelve second wafer is a compression of a real implant so that faults and wafers happen on comparable time scales in a classroom.

8.4 Controllers as Implemented

Three loops on three clocks, and a supervisor that owns the permissives.

8.4.1 The Inner Loop and Its Two Gates

The inner loop regulates extraction current with cathode bias current on a scan of $\Delta t_c = 1\,\text{ms}$. It is written in the incremental, or velocity, form with a feedforward bias $I_{ff} = I_{b,0}$, and the complete law as coded is Equations (8.15) to (8.20).

$$e_k = I_{E,sp,k} - \hat{I}_{E,k} \tag{8.15}$$

$$s_k = s_{k-1} + K_p\,(e_k - e_{k-1}) \tag{8.16}$$

$$c_k = I_{ff} + s_k + z_{k-1}, \qquad u_k = \text{clip}\left(c_k, 0, I_{b,max}\right) \tag{8.17}$$

$$I_{b,cmd,k} = q_{auto}\,u_k + (1 - q_{auto})\,I_{b,hold} \tag{8.18}$$

Here $\hat{I}_E$ is the extraction-current reading after measurement noise and the first order filter τ_{meas}. Away from the gates this is the velocity form of Section 1.6.2 with no derivative term. The running sum s_k is S_k of Equation (1.13), the demand c_k is v_k of Equation (1.14) with the bias $u_b = I_{ff}$, and u_k is that demand clipped to $I_{b,max} = 4.0\,\text{A}$. Extraction current rises with bias current, so the loop is reverse acting in the convention of that section. While q_{auto} is one, $I_{b,hold}$ follows $I_{b,cmd}$. Equation (8.18) is a switch and not a blend. While the permissive is zero the command *is* the held value, so neither the proportional term nor the integrator can reach the actuator. Blending them with a gain instead would let a large proportional term leak through at exactly the moment the error is meaningless.

The integral state is updated after the output, and it carries three regimes in one recursion, Equation (8.19), after which it is clipped to $\pm 3 I_{b,max}$.

$$\begin{aligned} z_k = z_{k-1} + \Delta t_c \big[& q_{auto}\,(K_i e_k + K_{aw}(u_k - c_k)) \\ & + q_{track} K_t \left(z_{track,k} - z_{k-1}\right) \big] \end{aligned} \tag{8.19}$$

$$z_{track,k} = I_{b,hold} - I_{ff} - s_k \tag{8.20}$$

Running, $q_{auto} = 1$ and $q_{track} = 0$ give ordinary integral action with back-calculation for saturation, Equation (1.15) with $z = I$, $K_i = K_c/\tau_I$ and $K_{aw} = 1/T_t$. The default $K_{aw} = 5/\text{s}$ is a tracking time of 0.2 s, so each scan removes half a percent of the clipped excess. Held, both permissives are zero and the integrator is frozen, which back-calculation alone would not do. Recovering, the tracking permissive drives z with $K_t = 40/\text{s}$ toward the value that reproduces the held output at the present error, so the first automatic output equals the held output to within the measurement noise, 2.4944 A against 2.4930 A in scenario 7, and the transfer back is bumpless by construction rather than by luck [6].

8.4.2 The Outer Trim and the Filament Supervisor

The selected-beam loop trims the extraction setpoint on a 20 ms scan, Equation (8.21).

$$\dot{I}_{E,trim} = q_B \,\mathrm{clip}\left[K_{i,B}\left(I_{B,sp} - I_{B,meas}\right), -r_B, r_B\right] \tag{8.21}$$

with $K_{i,B} = 0.08/\mathrm{s}$, $r_B = 2.0\,\mathrm{mA/s}$, the trim bounded to $\pm I_{trim,max} = \pm 8\,\mathrm{mA}$ and the setpoint $I_{E,sp} = \mathrm{clip}(I_{E,base} + I_{E,trim})$ within 10–80 mA. The permissive q_B requires a valid beam measurement, the NORMAL state, and a ratio of selected to total current between 0.02 and 0.90, which is the cheapest available diagnostic for a beamline fault masquerading as a source fault. Selected beam rises with the extraction setpoint, so the trim is reverse acting like the inner loop.

The filament supervisor runs every 100 ms, acts only in NORMAL and only on a persistent bias offset, $\dot{I}_{f,sp} = K_f(I_{b,cmd} - I_{b,mid})$ with $I_{b,mid} = 2.25\,\mathrm{A}$, is rate limited to 0.5 A per execution and bounded to 100–190 A, and is frozen under the naive architecture and whenever the cathode is within 60 K of its trip. At the shipped gain $K_f = 0.08/\mathrm{s}$ and a bias command of 2.5 A it moves two hundredths of an ampere per second, which is two orders of magnitude slower than the inner loop. Bias current and filament current are two actuators on one variable and the only thing keeping them apart is that factor. A hotter filament lowers the bias command the inner loop needs, so the supervisor, which raises the filament current when the bias command sits above its center, is direct acting in the convention of Section 1.6.2.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.4.3 The Detector and the Supervisor

Detection requires coincidence on either electrode, Equation (8.22).

$$g = \bigvee_{k\in\{x,s\}} \Big(\left[I_{HV} > 0.6 I_{max}\right] \wedge \left[U_k < 0.75 U_{k,sp}\right] \wedge \left[t_{p,k} > t_{detect}\right] \Big) \tag{8.22}$$

Low beam current on its own is never sufficient and is never labeled an arc. The supervisor is a state machine over OFF, WARMUP, NORMAL, QUENCH, RECOVER, QUALIFY, REPAINT and TRIP. On detection it stores the last good command, records the scan coordinate, inhibits high voltage and enters QUENCH for the dwell t_q. It then ramps the affected supplies over t_r, qualifies the recovered extraction current inside eight percent of setpoint for a dwell t_{qual}, and resumes. If the committed dose deficit exceeds the repaint threshold it enters REPAINT first. More than n_{storm} events in a rolling ten seconds, or a cathode overtemperature, trips the source and the restart is an operator action rather than a timer.

Three architectures are selectable and they differ only in the permissives. A naive proportional-integral loop keeps q_{auto} at one throughout and has no outer trim. A cascade adds the trim and the filament supervisor, the structure Volume 3, Chapter 2 matches to where a disturbance enters, and still keeps q_{auto} at one. The arc-aware cascade gates both, and it alone repaints.

8.5 Two Judgments That Are Usually One

The application keeps two judgments apart that are normally made together, and the separation is the most transferable idea in it.

Latching a fault quenches the beam, so a false latch is expensive and the rule that triggers it demands corroborating evidence and a persistence time. Deciding that one measurement has stopped meaning anything is a weaker and much faster question. An extraction-current reading taken while the extraction voltage is not present is meaningless on its own terms, and waiting the coincidence window before saying so costs the controller a scan.

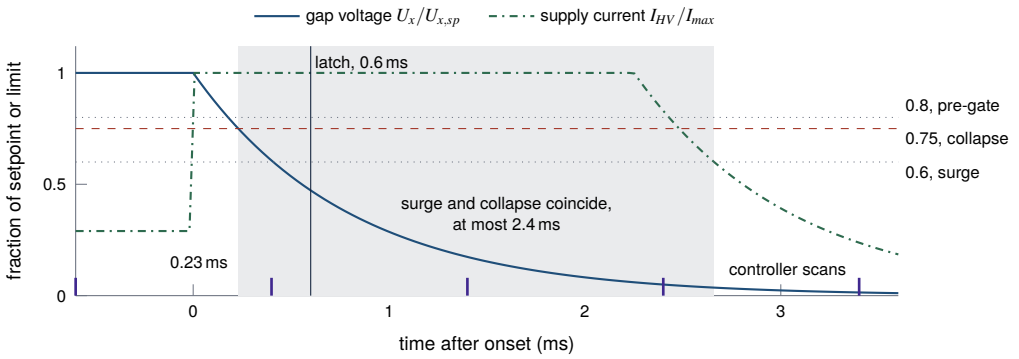


Figure 8.2 A spark against the controller's clock. The curves follow Equations (8.12) and (8.13) with the parameters of Table 8.2, the beam load left out after onset; only the printed times are the chapter's. The gap voltage (solid) passes the dashed 75 % line of the detector 0.23 ms after onset, and the fault latches at 0.6 ms. Notice the scan that falls between the two: it reads a falling extraction current before any latch, which is why the 80 % pre-gate (upper dotted line) needs no persistence. The supply current (dash-dotted) stays on its limit only while the gap still holds voltage, so the shaded coincidence lasts at most 2.4 ms.

That scan is not free. The gap voltage falls through seventy-five percent of its setpoint 0.23 ms after onset, which is the second integration step, and the coincidence rule needs 0.3 ms more, so the fault is latched 0.6 ms after onset. The 1 ms controller scan keeps its own clock. When a spark starts 0.2 or 0.4 ms before a scan, that scan reads a filtered extraction current that is already falling while nothing has been latched, and holding the command in force at the latch would freeze whatever that scan asked for. The application therefore keeps two defenses. A fast pre-gate marks the extraction-current reading invalid whenever the extraction voltage is below eighty percent of its setpoint, which needs no persistence because it latches nothing, and the command is snapshotted at the integration step whenever the electrical evidence is healthy, so that the snapshot, not the command at detection, becomes $I_{b,hold}$. Figure 8.2 draws that sequence against the controller's scans.

The consequence is measurable on a counterfactual copy of the controller with both defenses removed, firing the same fifteen millisecond spark at each of the five integration-step phases of the scan. With the onset 0.4 ms before a scan, the held command is 2.740 A against a pre-fault 2.492 A, the cathode gains 79 K over the following second, and the extraction current climbs to 57 mA, so the 150 ms qualification dwell is never completed and the supervisor is still in QUALIFY nine seconds later. At the other four phases the hold is within 0.018 A of the pre-fault value. Either defense alone, and the application as shipped, hold the command within 0.002 A of its pre-fault value at all five phases, and the cathode moves 0.15 K. The scripted spark of scenario 7 starts 0.2 ms after a scan and is latched before the next one, so on its own it would not expose the difference.

The same separation explains why a dead Faraday cup and a spark are handled differently. Two measurements carry two validity flags. The extraction-current flag is gated on high-voltage health, the selected-beam flag is gated on its own instrument, and a single global fault bit would either freeze the inner loop for no reason or walk the outer one on a dead reading.

8.6 Results

Every number below comes from the application's own engine at the default seed, 20260921 (Section 1.6.1), with the scenario named beside it. Scorecard figures are quoted as the scorecard computes them, and peaks, crossings and held values are read at the 0.2 ms integration step. Table 8.3 lists the measured outcomes. Where noise could matter, the case was rerun with other seeds and their range is given beside the shipped result.

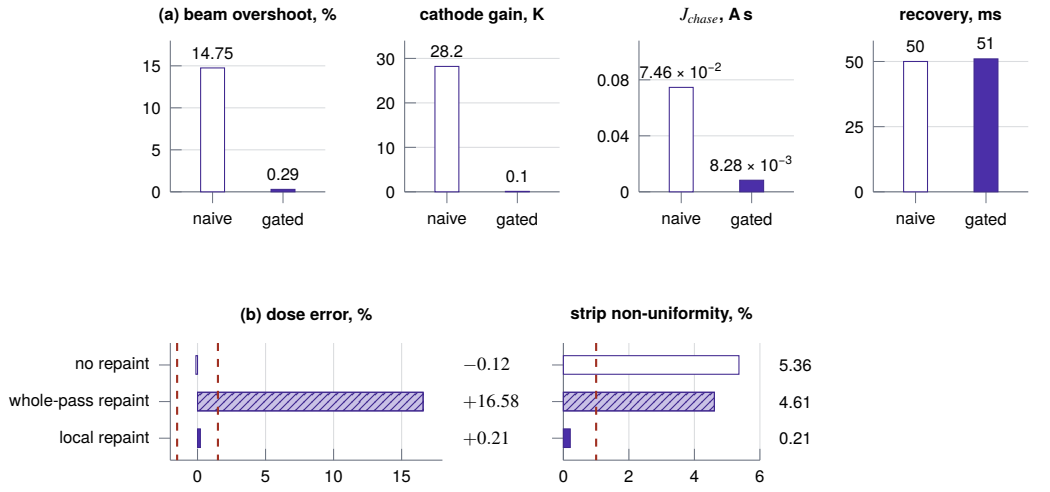


Figure 8.3 Two rows of Table 8.3 as charts. (a) Scenarios 6 and 7, the same fifteen millisecond spark on the naive PI (open bars) and on the gated, arc-aware cascade (solid bars). (b) Scenario 10, no repaint (open), a whole-pass repaint (hatched) and a local repaint (solid), with the mean dose tolerance of 1.5 % and the strip non-uniformity limit of 1.0 % of Table 8.2 dashed. Notice the last panel of (a): the gate removes the overshoot, the heat and the chasing, and leaves recovery where it was.

The twelve scenarios fall into two halves and a challenge. Scenarios 1 through 5 build the source and the cascade, from a cold start through the emission knee, the four currents, the mass spectrum and the outer trim to tuning without faults. Scenarios 6 through 11 are the fault, organized as contrasts rather than as a list of features. Scenario 6 against 7 is the gate, scenario 8 against 9 is specificity, scenario 10 is the product and scenario 11 is the retry limit. Scenario 12 is the graded challenge. Each scenario carries a completion rule in the app, and the Exercises at the end of the chapter are built on several of them. Figure 8.3 draws the two central contrasts of the table, one spark with and without the gate and the three repaint policies.

8.6.1 The Plant Before the Fault

The nominal recipe is a fixed point of every state, and the three currents an operator reads are 10.0 A of source discharge, 35.00 mA of extraction and 7.918 mA of mass-selected beam. The plasma supplies 33.20 mA against a space-charge limit of 55.45 mA, so the extraction voltage is not a current knob.

The nonlinearity that decides the architecture is the emission knee. Between 2090 K and 2420 K the normalized emission runs from 0.562 to 2.742, a factor of 4.9, and its slope with temperature changes by a factor of 3.7 across those three hundred and thirty kelvin, so a loop closed on a thermal variable has no single gain. Closing on extraction current instead puts the nonlinearity inside the loop where the integrator absorbs it.

Scenario 4 is the pair of disturbances that one current measurement cannot separate. Dropping the analyzer transmission by twelve percent moves the selected beam by −12.00 % and the total extraction current not at all. One second later the outer trim has raised the extraction current by 0.25 % and the beam is still 11.76 % low. Shifting the plasma toward fewer B^+ at 22 s moves the beam by a further −9.15 % without moving the total at all. An outer loop with unlimited authority reading only the Faraday cup would drive the source harder in both cases, which is why the trim is rate limited, bounded to ±8 mA and gated on the current ratio. At the shipped trim gain it is also slow. By 22 s it has recovered 2.7 of the twelve points the detune took and has used 1.04 mA of its authority, so at that gain the bound is never reached. The case's rerun raises the trim gain to 0.45/s.

Table 8.3 Measured outcomes at the default seed

Case	Configuration	Metric	Value
2	nominal recipe	I_{arc}, I_E, I_B	10.0 A, 35.00, 7.918 mA
2	—	I_{pl}, I_{CL}	33.20, 55.45 mA
1	T_c 2090 to 2420 K	emission ratio	0.562 to 2.742
3	V_a 50 to 75 V	T_e	6.00 to 6.88 eV
3	—	B^+ fraction	32.3 to 39.3 %
4	η_{tr} 1.00 to 0.88	I_B, I_E at once	−12.00, 0.00 %
4	—	same, 1 s later	−11.76, +0.25 %
6	naive PI, 15 ms	$I_{b,cmd}$	2.493 to 4.0 A
6	—	time at limit	32 ms
6	—	ΔT_c, overshoot	28.2 K, 14.75 %
6	—	J_{chase}, recovery	7.46×10^{-2} A s, 50 ms
7	arc-aware, same spark	$I_{b,cmd}$ held	2.493 A
7	—	ΔT_c, overshoot	0.1 K, 0.29 %
7	—	J_{chase}, recovery	8.28×10^{-3} A s, 51 ms
8	flow to 60 %	I_B dip, $I_{b,cmd}$	6.62 mA, 3.378 A
8	—	fault state entered	none
9	60 ms dropouts	$I_{b,cmd}$, trim	2.489–2.511 A, 0 mA
10	no repaint	dose error, U_D	−0.12, 5.36 %
10	whole-pass repaint	dose error, U_D	+16.58, 4.61 %
10	local repaint	dose error, U_D	+0.21, 0.21 %
11	six events in 4.5 s	trip at, energy	3.6 s, 24.3 J
11	—	energy, settled	32.3 J
11	limit 20	energy, no trip	32.6 J
12	as given	worst I_B error	0.946 mA
12	suggested	worst I_B error	0.177 mA
12	as given, suggested	wafers in spec	0 of 5, 2 of 4

Half of the detune loss is then back 7.6 s after the detune, and after the species shift the trim reaches its 8 mA bound at 33.9 s, leaving the beam 0.50 mA short at 40 s.

8.6.2 Saturation Without Windup

Scenario 6 fires a fifteen millisecond extraction spark on a loop with textbook back-calculation anti-windup, the remedy Volume 2, Chapter 3 applies at a true actuator limit, and no gating. The bias command runs from 2.493 A to its 4.0 A limit, the cathode gains 28.2 K, the selected beam overshoots to 9.065 mA against a 7.9 mA setpoint, and the fault-chasing index $J = \int \max(I_b - I_{b,pre}, 0)\, dt$ over the two seconds after detection is 7.46×10^{-2} A s.

> **TRY IT IN THE APP** **6 · A Naïve PI Chases a Spark**
>
> Load the scenario and press Run. At 3 s a fifteen millisecond spark crosses the extraction gap. Watch the bias-current command climb from 2.493 A to its 4.0 A limit while the extraction voltage is on the floor, and the selected beam overshoot to 9.065 mA when the high voltage returns. Nothing about the source was wrong.

The word windup is slightly wrong for what happens. The proportional term alone asks for 2.80 A above the feedforward once the reading has fallen, so the output sits on its limit for 32 ms, from 3 ms after onset until the high voltage is ramping back. Back-calculation engages for exactly that interval and does what it was designed to do, holding the integral below 0.23 A. The damage is heat. Nearly 4 A of bias for those milliseconds raises the cathode 28.2 K over the next 60 ms, and it is that stored heat, not the integral, that overshoots the beam when the gap returns. The loop has spent its actuator on an error that does not exist.

Pause. Try Exercise 8.2 before reading on.

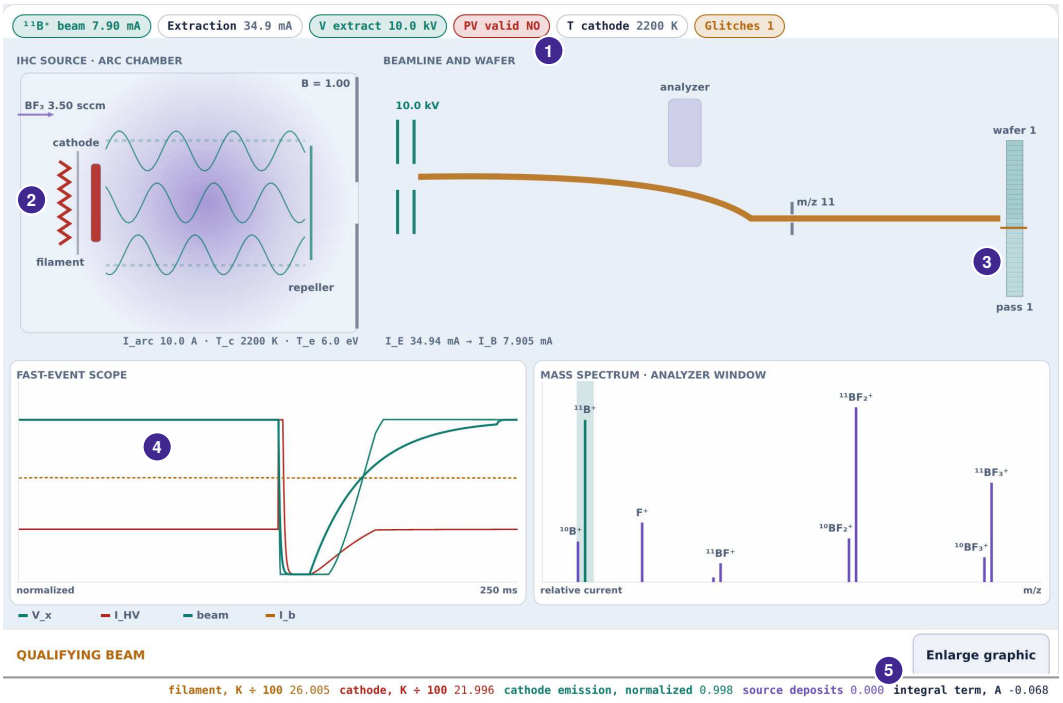


Figure 8.4 The application during scenario 7 at $t = 3.12$ s, 120 ms after a fifteen millisecond extraction spark under the arc-aware architecture, with the supervisor in QUALIFY (the banner under the plant view) and the extraction-current reading still marked invalid (1). The crop keeps the plant view and, under the gray rule, the labels of the thermal strip. The source cutaway (2) carries the filament, the cathode, the repeller and the electron helices whose brightness follows the electron density, and the beamline carries the bent beam into the resolving slit and the scanned wafer (3) with its strips colored by delivered dose. The fast-event display (4) below them holds the last 250 ms at the integration step, which is where the supply current reaching its limit and the gap voltage collapsing are legible while the bias current stays flat at its held 2.493 A. The extraction current reads 34.94 mA, the selected beam 7.905 mA and the fault-energy ledger 8.0 J. The integral term (5) is plotted on the thermal strip on purpose, because that is where a millisecond fault would otherwise be written into a four second plant. Here it reads −0.068 A as the tracking permissive rebuilds it before the loop is handed back.

Scenario 7 is the same gains, the same spark and the same seed with the measurement gated. The bias command is held at 2.493 A from the latch until the beam requalifies, the cathode gains 0.1 K, the overshoot falls to 0.29 % and J falls to 8.28×10^{-3} A s. Over seeds 1 to 5 the gated J runs from 1.3×10^{-3} to 4.8×10^{-3} A s against 6.9×10^{-2} to 7.3×10^{-2} A s ungated, because what remains with the gate is measurement noise rather than the spark. Recovery time is essentially unchanged, 51 ms against 50 ms, because recovery belongs to the high-voltage ramp and not to the controller. That last comparison is the one worth carrying away. Gating the measurement does not make the tool recover faster. It stops the tool from doing damage while it waits.

8.6.3 Specificity Is the Harder Half

A detector is easy to make sensitive. Scenarios 8 and 9 test the other direction. Dropping the real gas flow to sixty percent takes the beam from 7.92 mA to 6.62 mA, the inner loop raises the bias command to 3.378 A to restore the extraction current, and the supervisor never leaves NORMAL.

That is correct behavior, and a fault gate written to include low beam current would have frozen the controller through exactly the disturbance feedback exists to reject. The beam also settles above where it started, at 8.35 mA after thirty seconds, because less gas raises the electron temperature to 6.36 eV and enriches the B^+ fraction, so the same extraction current carries more boron.

Pause. Try Exercise 8.3 before reading on.

Killing the Faraday reading for 60 ms with every electrical quantity normal leaves the bias command inside the band its measurement noise gives it anywhere in the run, 2.489 to 2.511 A during the first dropout against 2.482 to 2.519 A over the whole ten seconds. The outer trim does not move at all, and again there is no quench. The inner loop keeps regulating because its own measurement never became invalid, and the outer trim holds its accumulated value because its measurement is the dead one.

> **DID YOU KNOW? Same Move, Opposite Verdicts**
>
> From the simulator: in Table 8.3 a real gas loss to 60 % drives the bias command to 3.378 A, and that is feedback doing its job. A fifteen millisecond spark drives it to its 4.0 A limit, and that is the fault. The size of the move does not tell them apart. Only the high-voltage evidence does.

8.6.4 Putting the Dose Back Where It Went Missing

An implant specification is an integral, so an interruption is a spatial defect. Scenario 10 fires a fifteen millisecond spark at 4 s, at the start of the third of six passes, and grades the finished wafer three ways. The beam is physically off for 26 ms, which leaves strip 0 at 0.897 of its committed dose. No repaint gives a dose error of −0.12 % and a strip non-uniformity of 5.36 %. A whole-pass repaint gives +16.58 % and 4.61 %. A local repaint gives +0.21 % and 0.21 %.

The whole-pass result is the instructive one. The extra pass starts where the scanner stands when qualification ends, at strip 4, and runs one full pass of travel, so every strip gains a sixth of its dose. Strip 0 ends at 1.064 against about 1.17 for its neighbors, exactly as far below them as before, and the mean ends 16.58 % high. The non-uniformity falls slightly only because it is normalized by a larger mean. Local repaint integrates the deficit of Equation (8.14) strip by strip with the beam blanked over strips already at their committed dose. It needs no record of where the scanner was at detection, because both ledgers are kept per strip and the deficit carries its own location.

8.6.5 Knowing When to Stop

Scenario 11 fires six events in 4.5 s, four on the extraction gap and two on the suppression gap, on a source running dirty. Each one is handled correctly, which is the problem. The tool trips 3.6 s into the storm on the fifth event with 24.3 J on the fault-energy ledger, and the restart is an operator action. A trip inhibits the high voltage but cannot recall a spark already conducting, so the ledger settles at 32.3 J. With the limit raised to 20 the tool restarts through all six and reads 32.6 J, because the only spark after the trip is a 0.32 J suppression spark against 8.0 J for each extraction spark. A trip saves only the sparks that come after it. Raising the limit is not merely permissive, because every restart puts the full extraction voltage back across a gap that has already broken down several times, and the fault-energy column is the honest measure of what that costs. The storm is scripted, so these figures do not change with the seed.

TRY IT IN THE APP **11 · An Arc Storm and a Safe Trip**

Press Run and watch the rolling event count against the Arc-storm trip limit. The source, running dirty, takes six sparks in 4.5 s, and each one is handled correctly. The tool still trips on the fifth, 3.6 s into the storm, and stays tripped. Restarting is an operator action: press Clear the trip and restart.

8.6.6 The Graded Run

The challenge is seventy seconds containing a gas calibration step, an analyzer detune at 12 s, three sparks of two lengths on two electrodes, an arc storm that trips the tool, an operator restart at 34 s and a source cleaning at 35 s, beginning from a single proportional-integral loop with a three millisecond detection window, a nearly dead outer trim at 0.02/s and no repaint. The model parameters and disturbances must stay as given. The run is scored on the worst selected-beam error after 46.2 s, the last third of the run, against a band of ± 0.35 mA around the 7.9 mA setpoint. As given, that error is 0.946 mA, 0.596 mA outside the band, and none of five wafers meets specification. The suggested settings correct the architecture to arc-aware, the detection window to 0.3 ms, the trim gain to 0.45/s and the repaint policy to local, and retune the inner loop to $K_p = 0.10$ A/mA and $K_i = 0.45$ A/(mA s). With them the worst error is 0.177 mA, inside the band with 0.173 mA to spare, and two of four wafers pass. Seeds 1 to 5 give 0.938 to 0.967 mA as given and 0.164 to 0.188 mA with the suggested settings.

The detection window does not do what its three millisecond value suggests. The coincidence of surge and collapse lasts at most 2.4 ms for any spark, because the supply current falls below sixty percent of its limit once the gap voltage has collapsed, so at 3 ms none of the nine sparks in the run is latched, not merely the short ones. Under the arc-aware architecture the fast pre-gate then holds the output only while the extraction voltage is below eighty percent of its setpoint, and the loop chases the tail of the filtered reading to 3.29 A of bias at the 8 s spark, against the limit for the naive loop. What decides the score is the trim gain, because the error left in the scored window is the 12 s detune, which only the selected-beam loop can see. The cascade with the trim gain at 0.45/s and nothing else changed already scores 0.117 mA, while the arc-aware loop with the 0.3 ms window and the shipped trim scores 0.828 mA. The gate, the window and the repaint are visible in the fault-chasing index and the dose ledger instead. Shortening the window alone puts the naive loop's J at 0.31 A s over the run, and the suggested settings bring it to 0.012 A s. Shortening it costs nothing in nuisance trips, because the rule is a coincidence of surge and collapse rather than a threshold on any one signal.

8.7 Technology in Practice

The three loops of Section 8.4 read the current of the extraction supply and a Faraday cup at the end of the beamline, and they move the bias and filament supplies of the source. The detector reads the same supply current beside the gap voltage, and the sequence acts on the extraction supply itself. Figure 8.5 draws them with their tags.

What Measures It

The selected beam, the quantity the wafer is paid for, is read on a Faraday cup.

TECHNOLOGY **Sensor: A Faraday Cup**

A Faraday cup is an insulated metal cup that stops the beam, wired to a current-sensitive amplifier, so it measures the charge the beam carries directly. Every ion that strikes it frees secondary electrons with energies below about 10 eV, and any that escape make the reading wrong by their current. A long cup, a suppression voltage at the entrance or a permanent magnet keeps them in. Because it stops the beam, the cup must also take the beam's energy as heat [19]. In an implanter a cup samples the scanned beam during the implant, a current integrator sums the charge into dose, and the cup is calibrated before the first wafer of a batch. The scan must overscan the cup completely, or the dose varies [20].

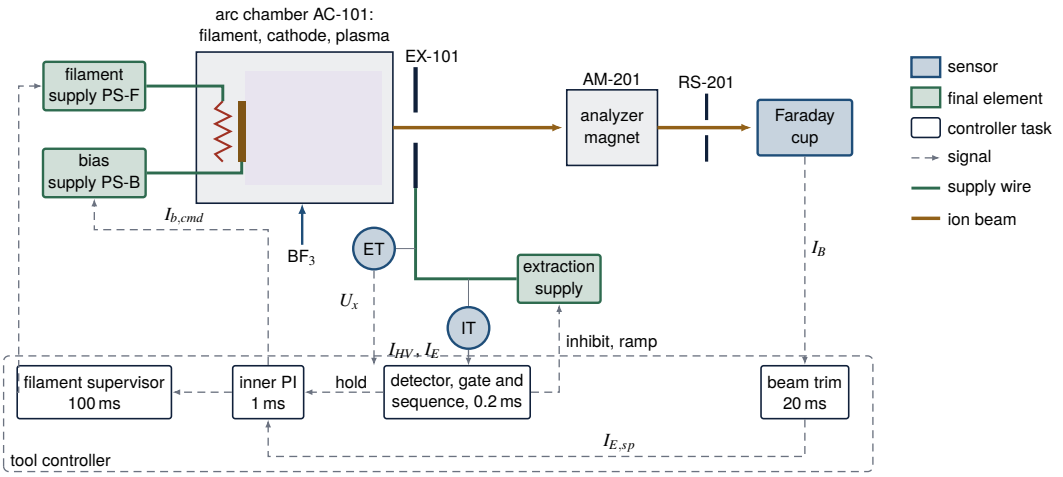


Figure 8.5 The cascade of Figure 8.1 drawn with its hardware. The current transmitter IT on the extraction supply serves twice: the detector reads it as supply current beside the gap voltage ET, and the inner loop receives it as extraction current only through the gate. The Faraday cup, which samples the scanned beam beside the wafer, feeds the beam trim alone. Notice that the detector moves no source supply: it gates a measurement, holds the inner loop and inhibits or ramps the extraction supply.

The model carries the cup as a reading with 0.05 mA of noise (Table 8.2). In scenario 9 it goes to zero for 60 ms while every electrical quantity stays normal, which is how a failed cup or its amplifier looks on a trend, and why the trim has a validity flag of its own.

What Moves It

The inner loop commands the bias supply between 0 and 4.0 A and the supervisor the filament supply between 100 and 190 A.

> **TECHNOLOGY** **Final Element: The Supplies Around the Cathode**
>
> An indirectly heated cathode is driven by three supplies. One sends current through the filament, one biases the filament against the cathode cap so that its electrons bombard and heat the cap, and an arc supply holds the cathode 50 to 150 V negative of the chamber walls. More bias means more energetic electrons striking the cap and more arc current. A design patented in 1996 closed two loops on these supplies. The arc current, compared with a demand the implant controller set from the desired beam current, adjusted the bias supply, and the cathode's emission current adjusted the filament supply. Shielding the filament behind the cap kept it out of the plasma and extended its life [21].

The model keeps the same pair of supplies as first-order lags of 5 ms and 50 ms, with the bias capped by what the filament can emit. Its fast loop closes on extraction current rather than arc current, and its filament supervisor plays the slow loop.

Where the Controller Runs

The loops belong to the implanter's own control system. In a 1996 tool a workstation was linked to a bus controller whose device interfaces drove the sensors and actuators, and recipes set the species, energy, beam current and dose for an automatic set-up [20]. Control loops run with cycle times from milliseconds to minutes [22]. The app runs the detector, the sequence and the dose ledger every 0.2 ms, the inner loop every 1 ms, the trim every 20 ms and the supervisor every 100 ms. Real arc

protection acts in the supply hardware, faster still. The structure of this chapter is a velocity-form PI block, two slow trims and a state machine that owns their permissives.

> **IN PRACTICE** **Counting Glitches Before a Tool Is Accepted**
>
> In a 1996 qualification of a medium-current implanter, the unscanned beam ran for two hours on a chart recorder. A glitch, a sudden drop-out or spike of more than ten percent of the beam current caused by an arc or another upset in the source, could occur at most ten times in those two hours, and fewer than five times an hour in the later qualification tests. Dose uniformity and repeatability were read on monitor wafers by four-point probe. Repeatability was the hardest to reach: filament current, arc current and extraction settings drift as the source ages [20].

What Is Modernizing It

Three developments bear on this source. Deposits cause premature failures of the source, and in-situ chemical cleaning cycles improved beam stability, glitch rate and source lifetime on a production tool compared with a tool without them [23]. Predictive maintenance trained on production data forecasts the breaking of the filament [24]. And because an implanted wafer cannot be reworked, classifiers watch tool data drawn through the equipment's SECS-II interface to detect faults in real time [25]. The deposit state of this chapter, which raises the glitch hazard, is the variable the first of these acts on.

8.8 Limitations

The plasma model is lumped and reduced. It reproduces the ordering of the species onsets and the bounded effect of confinement, and it does not predict absolute species fractions, plasma potential, sheath structure or contamination.

The glitch is a teaching model of a switched conductance. Real protection hardware acts locally and faster than anything in this application, and every threshold, dwell, retry count and ramp order requires validation by the equipment maker before it is used for anything.

The deposit state carries a hazard rate that rises with deposits and extraction voltage. It exists to make faults more likely in degraded scenarios and must never be read as a predicted time to arc.

The historian logs at 100 Hz, so the strip charts carry one or two samples of a fifteen millisecond event and none of a two millisecond one. The sub-millisecond shape is available only on the fast-event display in the scene, and the export does not carry high-rate rows around a fault.

The mass analyzer is an empirical mass window. Beam optics, emittance, neutralization, wafer charging, angle control, channeling, junction depth, activation and crystal damage are all outside the scope, and repaint restores a ledger rather than proving device-level equivalence after a real interruption.

> **OPEN CHALLENGE** **A Source That Lasts on One Dopant**
>
> Replacing the ion source is generally the most frequent maintenance on an implanter [26]. Fabs now expect stable beams with minimal arcing, at least 400 hours of source life and more than 90 % uptime [27]. Life is getting harder to reach. In advanced logic many doping steps have moved into epitaxy, so a tool runs one species for long stretches, which shortens source life. A co-gas with hardware changes doubled or tripled it for some species [11]. Progress would be a source that keeps its life and its arc rate on any single species. The deposit state of this chapter only raises a hazard rate and predicts no arc.

8.9 Conclusion

This chapter studied an indirectly heated cathode boron trifluoride ion source on an interactive simulator in which a three-level cascade regulates a beam whose accelerating gap sparks. Sixteen continuous states carry two thermal masses, two supply currents, the gas and electron densities, the source discharge, five ion species, two electrode voltages, a deposit state and a fault-energy ledger,

and the fault detection and the dose ledger run at the integration step because protection cannot wait for a controller scan.

The measured outcome that organizes the application is what one decision is worth. The same gains and the same fifteen millisecond spark drive the bias command to its 4.0 A limit and the beam 14.75 % past setpoint without a measurement-quality gate, and hold the command where it was and the beam within 0.29 % with one. Recovery time is unchanged at about 50 ms in both cases, because recovery belongs to the high-voltage ramp. The rule that separates them is short. Record the error, because the wafer really did miss dose, and do not feed it to the loop, because the source was never the reason.

KEY POINTS

- A spark collapses the beam reading for milliseconds while the source is fine, so a loop that believes it heats the source for nothing.
- Back-calculation does not prevent it: the proportional term alone drives the bias to its 4.0 A limit, and the cathode stores the heat.
- Gating the measurement, holding the output and freezing the integrator cut the overshoot from 14.75 to 0.29 %, with recovery unchanged near 50 ms.
- Latching a fault needs corroborating evidence: low beam current on its own is never sufficient.
- A local repaint driven by a per-strip ledger returns the wafer to 0.21 % non-uniformity, where a whole-pass repaint overdoses it by 16.58 %.

Where else this applies. A protective override that takes a loop outside its domain in Volume 4, Chapter 2, windup at a true actuator limit in Volume 2, Chapter 3, and a product record kept apart from the instruments in Volume 8, Chapter 13.

References

[1] R. B. Fair, "History of some early developments in ion-implantation technology leading to silicon transistor manufacturing," *Proceedings of the IEEE*, vol. 86, no. 1, pp. 111–137, 1998.

[2] C. B. Yarling, "History of industrial and commercial ion implantation 1906–1978," *Journal of Vacuum Science & Technology A*, vol. 18, no. 4, pp. 1746–1750, 2000.

[3] B. A. MacKinnon and J. P. Ruffell, "Semiconductor ion implanters," in *Application of Accelerators in Research and Industry: Twenty-First International Conference*, ser. AIP Conference Proceedings. AIP, 2011, pp. 46–51.

[4] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *1989 American Control Conference*. IEEE, 1989, pp. 1693–1698.

[5] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, "A unified framework for the study of anti-windup designs," *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

[6] L. Zaccarian and A. R. Teel, "A common framework for anti-windup, bumpless transfer and reliable designs," *Automatica*, vol. 38, no. 10, pp. 1735–1744, 2002.

[7] M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki, *Diagnosis and Fault-Tolerant Control*, 3rd ed. Berlin, Germany: Springer, 2016.

[8] J. H. Freeman, "A new ion source for electromagnetic isotope separators," *Nuclear Instruments and Methods*, vol. 22, pp. 306–316, 1963.

[9] T. N. Horsky, "Indirectly heated cathode arc discharge source for ion implantation of semiconductors," *Review of Scientific Instruments*, vol. 69, no. 2, pp. 840–842, 1998.

[10] J. N. Mefo, B. J. Sealy, E. J. H. Collart, D. G. Armour, and R. M. Gwilliam, "Characterisation of an indirectly heated cathode source using langmuir probe diagnostics," *Nuclear Instruments and Methods in Physics Research Section B*, vol. 240, no. 1-2, pp. 13–17, 2005.

[11] E. Su, A. Chang, K. Wu, and J. Tsai, "Extension of the source lifetime in HC ion implanter with dedicated species," in *2016 21st International Conference on Ion Implantation Technology (IIT)*. IEEE, 2016.

[12] I. N. Draganic, C. P. Leibman, L. J. Rybarcyk, T. M. Bronson, and K. W. Dudeck, "Hot filament performance in a Freeman ion source," in *Proceedings of the 17th International Conference on Ion Sources*, ser. AIP Conference Proceedings, vol. 2011. AIP Publishing, 2018, p. 030005.

[13] D. J. Chivers, "Freeman ion source: An overview (invited)," *Review of Scientific Instruments*, vol. 63, no. 4, pp. 2501–2506, 1992.

[14] C. D. Child, "Discharge from hot CaO," *Physical Review (Series I)*, vol. 32, no. 5, pp. 492–511, 1911.

[15] I. Langmuir, "The effect of space charge and residual gases on thermionic currents in high vacuum," *Physical Review*, vol. 2, no. 6, pp. 450–486, 1913.

[16] S. Dushman, "Electron emission from metals as a function of temperature," *Physical Review*, vol. 21, no. 6, pp. 623–636, 1923.

[17] J. C. Olson, A. Renau, and J. Buff, "Scanned beam uniformity control in the VIISta 810 ion implanter," in *1998 International Conference on Ion Implantation Technology Proceedings*, vol. 1. IEEE, 1999, pp. 169–172.

[18] 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.

[19] P. Forck, "Lecture notes on beam instrumentation and diagnostics," Joint University Accelerator School and GSI, Darmstadt, Germany, Lecture notes, 2011. [Online]. Available: https://www-bd.gsi.de/conf/juas/juas_script.pdf

[20] S. K. Joung, "Qualification of a medium current ion implantation system in a semiconductor production environment," S.M. thesis, Massachusetts Institute of Technology, Cambridge, MA, 1996. [Online]. Available: https://hdl.handle.net/1721.1/40220

[21] P. Sferlazzo, E. K. McIntyre, Jr., W. E. Reynolds, R. M. Cloutier, and T. N. Horsky, "Ion generating source for use in an ion implanter," U.S. Patent 5,497,006, filed 15 November 1994, issued 5 March 1996, 1996. [Online]. Available: https://patents.google.com/patent/US5497006A/en

[22] 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.

[23] T. Romig, D. Newman, M. Mitchell, K. Ditzler, D. Eldridge, S. Bishop, and J. Mayer, "Implanter source life and stability improvement using in-situ chemical cleaning," in *Ion Implantation Technology: 17th International Conference on Ion Implantation Technology*, ser. AIP Conference Proceedings. AIP, 2008, pp. 364–367.

[24] G. A. Susto, A. Schirru, S. Pampuri, D. Pagano, S. McLoone, and A. Beghi, "A predictive maintenance system for integral type faults based on support vector machines: An application to ion implantation," in *2013 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2013, pp. 195–200.

[25] S. C. Horng and Y. L. Lin, "Fuzzy tree classification system for fault diagnosis on ion implanter," *Journal of Physics: Conference Series*, vol. 2020, no. 1, p. 012017, 2021.

[26] A. Uvais, M. Jinguji, Y. Sato, T. Yotsumoto, and A. Botet, "Extending ion source life on high current ion implant tools with in-situ chemical cleaning," in *Ion Implantation Technology 2010: 18th International Conference on Ion Implantation Technology*, ser. AIP Conference Proceedings. AIP, 2011.

[27] T.-J. Hsieh, "Ion implantation ion sources in semiconductor fabrication: source life, beam stability, defects, yield improvements and contamination control," *Journal of Physics: Conference Series*, vol. 3237, no. 1, p. 012076, 2026.

8.10 Exercises

8.1 ⌕ Concept check · 1 min

Four quantities in the implanter are called current. Which pairing describes the inner loop of the cascade, the one that runs on the 1 ms scan?

(a) It reads the selected beam I_B and moves the extraction-current setpoint

(b) It reads the cathode temperature T_c and moves the filament current I_f

(c) It reads the extraction current I_E and moves the cathode-bias current I_b

(d) It reads the bias command $I_{b,cmd}$ and moves the filament current I_f

8.2 ⌕ Concept check · 2 min

In scenario 6 the naive loop runs back-calculation anti-windup with $K_{aw} = 5/\mathrm{s}$, yet the selected beam overshoots by 14.75 % when the high voltage returns. Why does back-calculation not prevent the overshoot?

(a) The naive architecture runs the inner loop with back-calculation switched off
(b) The tracking permissive pulls the integral toward a corrupted held output
(c) The integral is clipped at ± 12 A, far too wide to bound anything
(d) The proportional term alone drives the bias to its limit, and the cathode stores that heat

8.3 Concept check · 2 min

In scenario 8 the real BF_3 flow falls to 60 % and the selected beam dips from 7.92 to 6.62 mA. Why does the arc-aware supervisor stay in NORMAL with the PV-valid flag at one?

(a) The gas loss raises no supply-current surge and collapses no electrode voltage
(b) The 0.3 ms persistence window is too short to see a slow disturbance
(c) The outer trim restores the beam before the detector finishes its count
(d) Low beam current latches a fault only when it outlasts the quench dwell

8.4 Theory and calculation · 6 min

Use the normalized emission of Equation (8.3) with $\beta_R = 9$, $T_{c,0} = 2200$ K and no deposits ($w = 0$).

(a) Compute $\tilde{I}_e$ at 2090, 2310 and 2420 K to three decimals. *Checkpoint:* 0.562, 1.692 and 2.742.

(b) Differentiate $\tilde{I}_e$ with respect to T_c and find the ratio of the slope at 2420 K to the slope at 2090 K, to one decimal. *Checkpoint:* 3.7.

(c) In scenario 6 the cathode gains 28.2 K during the spark. Treat the selected beam as proportional to $\tilde{I}_e$, which holds while extraction is plasma limited, and find the beam rise this predicts, in percent to one decimal. Compare it with the 14.75 % overshoot the chapter measures. *Checkpoint:* 15.0 %, within about a quarter of a point of the measured overshoot.

8.5 Theory and calculation · 8 min

A spark switches the fault conductance across the extraction gap. Use Equations (8.12), (8.13) and (8.22) with Table 8.2: U_x starts at 10 kV and decays toward zero with $\tau_{coll} = 0.8$ ms while the spark conducts, $R_f = 5.0\,\mathrm{k\Omega}$ and $I_{max} = 120$ mA.

(a) Find the time after onset at which U_x falls below 75 % of its setpoint and the time at which it halves, in milliseconds to two decimals. *Checkpoint:* 0.23 ms and 0.55 ms.

(b) Once the gap has collapsed the extraction current is near zero. Find the gap voltage below which the supply current falls under $0.6\,I_{max}$, the time after onset at which that happens, to two decimals, and so how long surge and collapse coincide. Does a 3 ms persistence window latch a 15 ms spark? *Checkpoint:* 0.36 kV at 2.66 ms, so about 2.43 ms of coincidence. No.

(c) Integrate $\dot{E}_{fault}$ over a spark that lasts many τ_{coll} to find the energy of one extraction spark and of one suppression spark at $U_s = 2.0$ kV. Scenario 11 fires four extraction and two suppression sparks. Find the ledger with no trip and compare it with the chapter. *Checkpoint:* 8.0 J, 0.32 J and 32.64 J, against 32.6 J measured.

8.6 Theory and calculation · 7 min

In scenario 4 the analyzer transmission falls from 1.00 to 0.88 at 6 s. The inner loop holds I_E on its setpoint, and the selected beam is proportional to the ion part of the extraction current, $I_E - I_{par}$ at nominal extraction voltage (Equation (8.10)), with 7.918 mA of beam for 33.20 mA of ions before the detune.

(a) Find the incremental gain $g = \partial I_B / \partial I_{E,sp}$ after the detune, in milliamperes of beam per milliampere, to three decimals. *Checkpoint:* 0.210.

(b) Ignoring the rate limit r_B, Equation (8.21) then makes the beam error decay as a first-order lag. Find its time constant at $K_{i,B} = 0.080/\mathrm{s}$ and the fraction of the detune loss recovered 16 s after the detune, in percent to the nearest percent. *Checkpoint:* 59.6 s and 24 %.

(c) Find the time constant and the half-recovery time at $K_{i,B} = 0.45/\mathrm{s}$, to one decimal, and show that the rate limit stays inactive. Compare with the chapter. *Checkpoint:* 10.6 s and 7.3 s, against 7.6 s measured.

8.7 App challenge · 8 min

Load scenario 6, "A naïve PI chases a spark" (default seed), and run it at 1 s/s to 6 s. The scorecard tiles stop changing 2 s after the spark is latched, so any reading after 5.1 s is the same. Then load scenario 7, "The same spark, gated", and run it to 6 s.

(a) Read the J chase, Overshoot and Recovery tiles after each run. *Checkpoint:* 0.0746 A s, 14.8 % and 50 ms. Then 0.0083 A s, 0.3 % and 51 ms.

(b) In scenario 7 press Reset, set Quench dwell to 40 ms and Beam qualification dwell to 500 ms, and run to 6 s again. Read the Recovery tile. Export the CSV and find the first row whose `supervisor_state` reads `repaint`, and the same row in a run with the shipped dwells. *Checkpoint:* 77 ms. REPAINT begins at 3.64 s, against 3.26 s with the shipped dwells.

(c) The dwells grew by 380 ms in all, yet Recovery grew by only about 26 ms. Explain both numbers. *Checkpoint:* Recovery ends when the beam is first back within 2 % of setpoint, during RECOVER and QUALIFY. Only the quench dwell delays that.

8.8 App challenge · 8 min

Load scenario 10, "The dose hole, and what to do about it" (default seed), which fires a 15 ms spark at the start of the third of six passes with Dose repaint at Local. Run it at 1 s/s to 16 s, by which time the first wafer has finished, and export the CSV. Press Reset, set Dose repaint to None, run to 16 s and export again. Press Reset, set it to Whole pass, and repeat.

(a) In each file, find the first row with `wafers_done` equal to 1 and read its time, `metric_dose_error_pct` and `metric_uniformity_pct` to two decimals. *Checkpoint:* Local 13.78 s, +0.21 and 0.21 %. None 12.01 s, −0.12 and 5.36 %. Whole pass 14.01 s, +16.58 and 4.61 %.

(b) Non-uniformity is the strip range divided by twice the mean. Starting from the no-repaint wafer, predict the whole-pass dose error and non-uniformity if the extra pass adds one sixth of the committed dose to every strip, to one decimal. *Checkpoint:* +16.5 and 4.6 %.

(c) From the `supervisor_state` column, how long does each repaint last? Compare it with how much later each wafer finishes than the unrepaired one, and explain why the two agree to within one 10 ms row. *Checkpoint:* Local 1.78 s against 1.77 s, whole pass 2.00 s against 2.00 s. A repaint suspends the production pass instead of consuming it.

8.9 App challenge · 9 min

Load the challenge, "Challenge · a lot through an aging source" (default seed), and run it as given at 3 s/s to 70 s, where the checklist judges it.

(a) Read the checklist note and the J chase tile. *Checkpoint:* Worst 0.946 mA, 0 of 5 wafers in specification, 1 trip. J chase 0.0000 A s.

(b) Press Reset, set Control architecture to Cascade and Selected-beam trim gain to 0.450, and run to 70 s. The item now passes and the note is gone, so export the CSV. Its time column is in minutes. Over the rows from 0.77 to 1.1667 min (46.2 to 70 s), find the largest $|$`beam_meas_mA` − `beam_sp_mA`$|$, and read `wafers_good` and `wafers_done` in the row at 1.1667 min. *Checkpoint:* 0.117 mA, 2 of 5.

(c) Press Reset, set the architecture back to Naïve PI, the trim gain back to 0.020 and Coincidence persistence to 0.30 ms, and run to 70 s. Read the checklist note and the J chase tile, and explain why the tile read zero in part (a) and why the score did not move. *Checkpoint:* Worst 0.946 mA, 0 of 5. J chase 0.3139 A s.

8.10 Application area · 4 min

An implanter loses beam and the cause has to be named before anyone acts.

(a) For an extraction spark, a suppression spark, a loss of BF_3 flow and a dead Faraday cup, say which of these happen before any protection acts: a surge of the high-voltage supply current, a collapse of the extraction voltage, a collapse of the suppression voltage, and a fall of the selected-beam reading. *Checkpoint:* Extraction spark: surge, extraction collapse, beam reading. Suppression spark: surge, suppression collapse. Gas loss: beam reading only. Dead cup: beam reading only.

(b) Give the correct response to each, and say which of them is a protection event. *Checkpoint:* Both sparks: latch, quench, hold and freeze. Gas loss: let feedback act. Dead cup: freeze the trim, keep the inner loop.

8.11 Application area · 3 min

A process engineer asks why the fast loop moves cathode-bias current when the extraction voltage acts on the gap directly and the filament is the heater. Answer from the chapter's numbers for the nominal recipe.

8.12 Review the agent · 5 min

An AI assistant reviewing scenario 6 writes: "The 14.75 % overshoot is integrator windup. Raise the Back-calculation gain K_{aw} from 5 to 20/s so the integrator unwinds four times faster, and the overshoot will fall below 2 %, into the green band of the Overshoot tile." Find the flaw, and say how you would check the claim in the app before anyone changes a setting on a tool.