INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 11

Waste Vitrification Melter Simulator for Control Across Four Time Scales

Pages 290–319

Edited by John D. Hedengren

How to cite this chapter. "Waste Vitrification Melter Simulator for Control Across Four Time Scales," in *Control of Nuclear and Fusion Systems*, John D. Hedengren, Ed., vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 290–319.

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

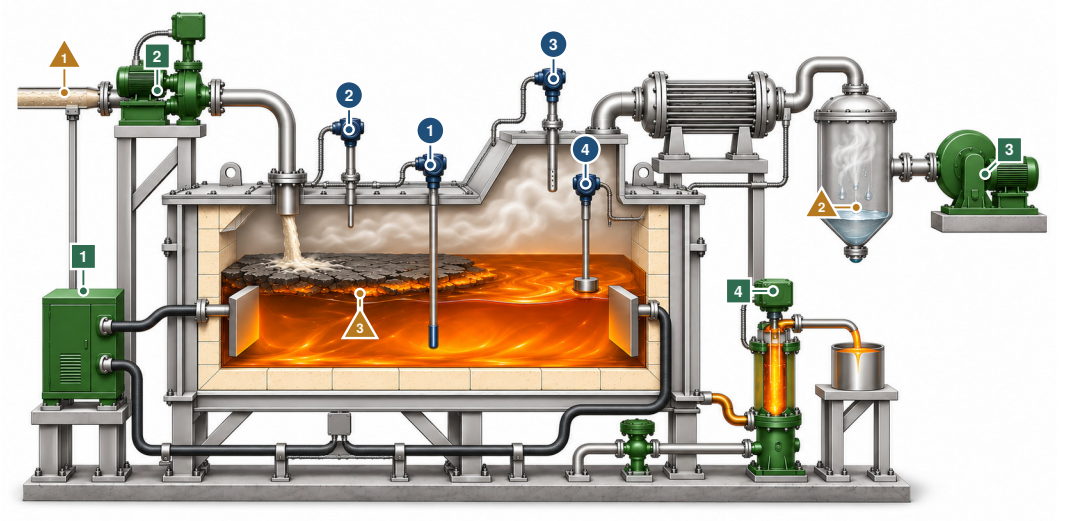
The app for this chapter
https://apmonitor.com/apps/waste-vitrification/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v18-nuclear-fusion/#waste-vitrification

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

CHAPTER 11

Waste Vitrification Melter Simulator for Control Across Four Time Scales



The controller moves 1 electrode power, from the supply cabinet 2 slurry feed pump 3 exhauster that holds the space above the glass under vacuum 4 airlift that pours glass into the canister

The controller reads 1 glass temperature 2 pressure above the glass 3 temperature above the crust, from which its coverage is inferred 4 glass level

What disturbs it 1 a wetter batch of slurry 2 deposits loading the off-gas scrubber 3 foam under the crust

Notice that all the glass is made inside the dark crust, which also blocks the glow of the pool from the space above, so the temperature up there tells how much of the pool the crust covers. Illustration, not to scale.

IN PLAIN TERMS **Making Glass From Waste Under a Floating Crust**

The main idea. A melter turns liquid nuclear waste into glass by spraying it onto a pool of molten glass kept hot by an electric current. The waste floats as a crust while it dries and melts, and all the glass is made inside that crust.

The challenge. The crust also works like a lid that keeps the pool warm. A wetter batch can then cut the output while the pool temperature hardly moves. Four controllers can each do their job well and still leave about a quarter of the output unused.

Why it matters. The glass locks the waste away for a very long time, and a faster melter lets a cleanup finish sooner.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 11.2 and 11.3, the fact boxes of Sections 11.3 and 11.6, the try-it boxes of Section 11.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 11.6, Section 11.7, Section 11.8 and the open challenges

The control question A joule-heated melter makes its glass inside a floating cold cap that also insulates the pool and loads the off-gas train. Why do four correctly tuned loops leave a quarter of the production unused, and what does a feed rate under overrides and feedforward recover?

Learning objectives

- Compute the slurry enthalpy, the glass yield and the evaporation duty of a wetter batch, and explain why power does not set production.
- Compare the electrical gain of a fixed tap with the thermal damping of the pool, and choose the regulation.
- Estimate the time constant of the plenum gas and set it against the dead time of the tap.
- Explain how a low selector with external reset hands the feed between three controllers.
- Measure what coordination, feedforward and MPC each change in production and constraint exposure.

Prerequisites Sections 1.3 (loop action), 1.4.5 (protection layers), 1.6.2 (velocity-form PID), 1.6.3 (predictive control) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/waste-vitrification/

See also Volume 4, Chapter 4 for feed-rate overrides on a low selector, Volume 4, Chapter 6 for a push toward the first binding limit, Volume 4, Chapter 9 for valve-position control. Volume 6, Chapter 7 for the production handle and Volume 6, Chapter 4 for a throughput supervisor.

11.1 Introduction

A joule-heated ceramic melter has no burner and no heating element. Two pairs of Inconel plate electrodes hang in 25.2 t of molten borosilicate at 1150 °C and the current that heats the pool travels through the glass, so the glass is the resistor. Waste slurry mixed with glass formers is sprayed onto the surface, where it does not sink and dissolve. It floats, dries and calcines as a crust two or three centimeters thick called the cold cap, and every kilogram of glass the plant makes is made inside that crust [1, 2]. The crust is therefore the production mechanism. It is also the dominant thermal resistance between the pool and the space above it, because bare glass at 1150 °C radiates 232 kW/m^2 while a crust at 160 °C radiates almost nothing [3]. And it is where every cubic meter of steam and nitrogen oxide is born, so it is the load on the off-gas system as well.

That coincidence is what a block diagram cannot show. A transfer function from electrode power to glass temperature can be written down and tuned against, and it is correct, but it hides the fact that the resistance in its denominator is a state the feed rate creates. The misconception the simulator corrects is measured rather than asserted. Stepping the electrode power up by 140 kW at a fixed feed rate raises the pool by 25.9 °C and leaves glass production at 500 kg/h, approximately where it started, because at steady state the glass rate is bounded by the dry-solids feed multiplied by its oxide yield. What the extra power bought was a smaller crust, coverage falling from 90.8 % to 71.7 %, and a plenum 176 °C hotter carrying the difference away. A second measurement makes the same point from the other side. A feed batch thirteen points wetter costs 37 % of production while the glass temperature moves from 1150.0 °C to 1150.4 °C, because the thicker crust that absorbed the extra evaporation duty of 186 kW also stopped 105 kW of radiation escaping through the bare melt. Glass temperature stays near setpoint while about a third of the production rate is lost.

The class of plant is wider than vitrification. Any process whose dominant resistance is a layer that the control action itself deposits behaves this way, including batch-fed glass furnaces under a batch blanket, rotary kilns under a coating, freeze linings in electric smelting furnaces and fouling layers on any heated surface. Two general consequences follow. The first is that the time constants separate by orders of magnitude, because a gas inventory, a floating crust, a large thermal mass and an accumulating deposit have nothing in common but the plant they sit in. Time-scale separation of that kind is normally exploited to reduce a model [4], and here it has to be exploited structurally instead, because one manipulated variable reaches two loops three decades apart in speed. The second is that no single loop owns the constraint set, so the plant runs against whichever of several limits

is binding today. The standard answer is a production actuator driven by one controller and pulled back by constraint controllers through a selector [5, 6]. The feed-rate loop of the combine harvester in Volume 4, Chapter 4 is built the same way. A melter is an unusually clean example because the constraint that binds first is an actuator demand rather than a measurement.

A student who is handed the four loops of a melter as four independent single-input single-output problems will tune all four correctly and leave a quarter of the plant's output on the table. This chapter models an interactive joule-heated melter in enough detail to be rebuilt, quantifies what coordinating the feed rate against the constraints is worth against independent loops on the same disturbance sequence, and measures the two electrode gains that decide whether the supply is run on power or on voltage. Section 11.2 gives the background of the glass and the machine. Section 11.3 writes out the eleven-state model with its instrument chain, its actuator envelope and its supervisor, to the constant, and lists its assumptions. Section 11.4 measures the four time scales, the two electrode gains, a power step and the pressure loop. Section 11.5 describes the six PI loops, the two classical architectures, the feedforward inversion and the optional predictive controller. Section 11.6 follows the scenario sequence from the electrodes to a graded shift. Its central comparison is 2444 kg of glass against 1968 kg over the same 300 min run, with 0 min above the cold-cap coverage ceiling against 188.3 min, and it shows that the feedforward path is structural rather than cosmetic because it takes the worst pressure excursion from 90.0 Pa to 54.2 Pa. Section 11.7 sets the loops in their plant hardware. Sections 11.8 to 11.10 state the limitations, discuss what transfers to other plants and conclude, and Section 11.11 holds the exercises.

11.2 Vitrification and Its Control Problem

Reprocessing spent fuel for plutonium left the United States with tens of millions of gallons of high-level liquid waste, neutralized with sodium hydroxide and stored in underground carbon steel tanks at Hanford and Savannah River built in the 1940s and 1950s. Tank storage was always an interim answer, and the search for a durable solid began in the 1950s with calcination, with a range of ceramic hosts and with several glass families [7]. Borosilicate glass won for chemical rather than mechanical reasons. It accepts a wide and variable mixture of waste oxides into a single amorphous phase, it is durable in groundwater on a geological time scale, and it tolerates the composition swings a tank farm delivers [8, 9]. Alternatives with better waste loading, including crystalline ceramics and glass-ceramic composites, have never matched it on tolerance to a feed nobody can specify in advance.

The machine that makes the glass took longer to settle than the glass did. Early work used in-can melting and induction-heated metallic melters, both of which are limited by the corrosion of a metal wall by molten borosilicate at 1150 °C. The joule-heated ceramic melter removed the wall from the heating path by making the melt its own resistor, passing alternating current between electrodes immersed in the pool and lining the vessel with refractory rather than metal [7]. The Defense Waste Processing Facility at Savannah River started radioactive operation in 1996 on that design and has run on it since, and the Waste Treatment and Immobilization Plant at Hanford was designed around the same principle at larger scale and with air bubblers added under the melt surface [10]. Bubbling is the single largest step change in throughput the technology has had, and in this simulator it is worth a factor of 2.5 on the melt-to-cap heat transfer coefficient, the parameter that melting-rate correlations fitted to melter feeds and operating conditions treat as the principal lever [11].

Cold-cap behavior became the operating constraint once the melters were large enough that throughput, rather than glass quality, decided whether a cleanup mission would finish. Production is not set by the electrode power, and it is not set by the pool temperature. It is set by the heat that crosses into the crust and by the enthalpy needed to convert what is sitting up there, so the crust is the throughput bottleneck by construction. Laboratory and pilot work has resolved the crust into layers, a wet slurry pool on top, then drying, then calcination, then a reacting foam layer, then the melting front at about 980 °C, and has measured the temperature profile across those few centimeters

directly [2, 12]. The same work identified foaming as the mechanism that destroys melting rate, because a raft of gas-filled glass forms exactly at the interface the heat has to cross [13]. What none of it produced is an instrument. There is no cold-cap sensor on an operating melter, and coverage is inferred from plenum thermocouples through the radiation contrast between bare melt and crust, a relation that computational studies of a pilot-scale melter show to be curved and emissivity dependent [3]. Coverage is consequently an advisory variable on every operating plant, which is precisely the situation that makes it interesting to close a loop on in a simulator.

The electrical side carries its own history. Molten borosilicate conducts by alkali ion migration, so its resistivity follows an Arrhenius law and falls steeply with temperature, and correlations for the viscosity and electrical conductivity of low-activity waste glass melts are fitted over composition and temperature for exactly this reason [14]. A supply holding constant voltage on a falling resistance delivers rising power, which is a positive feedback path around the fastest element in the plant. That is why every operating joule-heated melter regulates power and treats the tap changer as a slave, and Section 11.4.2 measures both gains involved so the design decision can be read off numbers instead of accepted.

> ⧖ HISTORY **Two Melters in Twenty-One Years (1996 to 2017)**
>
> The Defense Waste Processing Facility has vitrified high-level tank waste since 1996 in joule-heated melters whose upper and lower electrodes, each set controlled on its own, pass current through the glass. The first melter served from 1996 to November 2002. The second ran from March 2003 to February 2017 and made over 4.5 million kilograms of glass, twelve years longer than its conservative design life of two years, which rested on experiments that predicted fast wear of the refractory in contact with the glass. Feeding stopped in January 2017 when a pour spout heater failed, and a few days later a lower electrode broke at its stem and leaked cooling water into the glass pool. The stem sat inside the melter shell, so the melter could not be repaired, and a third melter took its place [15].

11.3 Process Model

Melter MLTR-101 is a refractory box with a 10 m^2 melt surface holding 25.2 t of borosilicate glass at a design depth of 0.95 m. Two pairs of plate electrodes on opposite walls carry 1260 kW into the pool at the design point, waste slurry enters through a nozzle in the lid at 1425 kg/h, glass leaves through an airlift into a canister, and the space above the melt is held at −250 Pa gauge by an exhauster pulling through a film cooler and a scrubber train. Figure 11.1 shows the arrangement, the instrument tags and every symbol used in the equations that follow. Reported time and controller periods are in minutes. The energy, mass and pressure balances use seconds before multiplication by 60 in the integrator, while scrubber loading and actuator time constants use minutes. Temperatures are in degrees Celsius, feed and production in kg/h, gas flows in kg/s, pressure in pascals gauge, and electrode quantities in ohms, volts and amperes.

> ✦ DID YOU KNOW? **Ten Square Meters of Molten Glass**
>
> The low-activity waste melters described for Hanford in 2011 have a melt surface of 10 m^2, the size of this chapter's melter, and a design rating of 1.5 t of glass per square meter per day [16]. The chapter's design point makes 1.20 t/m^2 a day.

Table 11.1 classifies what the control system interacts with. Four alarms bound normal operation, a containment limit at −25 Pa above which the plenum can leak outward, a glass temperature floor at 1050 °C, a refractory corrosion ceiling at 1210 °C and a cold-cap coverage ceiling at 97 % beyond which the crust is thickening rather than spreading. None of the four stops the run. Each latches an annunciation and accumulates minutes on the scorecard, so a test that carried the melter outside its limits leaves a record. Eleven states are integrated, seven of them physical and four of them actuator lags.

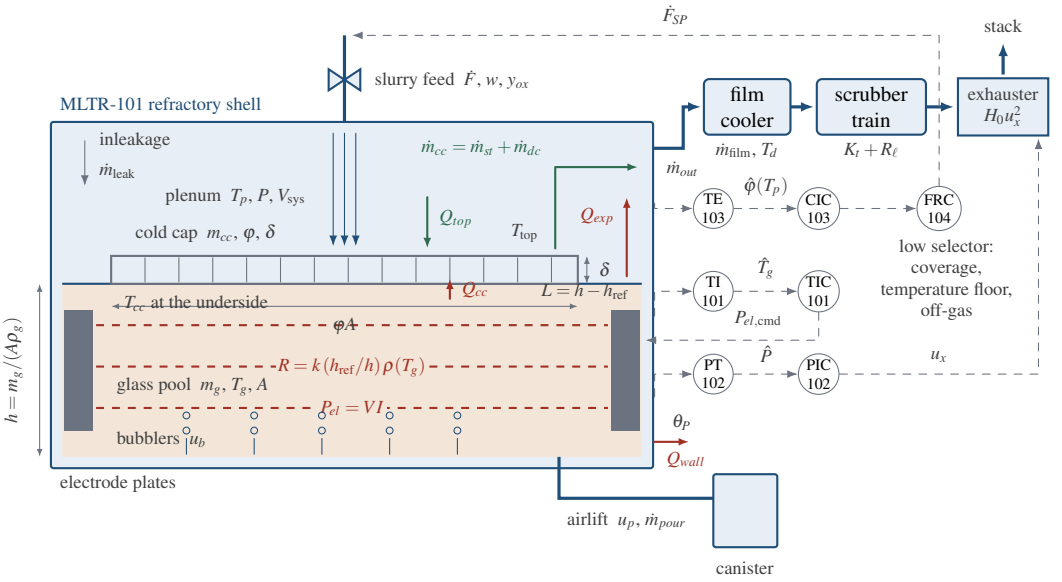


Figure 11.1 Joule-heated ceramic melter MLTR-101 with its off-gas train. Slurry of moisture w enters at rate $\dot{F}$ and forms the cold cap of mass m_{cc}, which covers a fraction φ of the melt area A at mean thickness δ. Current passes between the plate electrodes through the glass itself, so the circuit resistance R carries the pool temperature T_g and the glass depth h. Heat crosses into the cap as Q_{cc}, leaves the uncovered melt as radiation Q_{exp}, returns from the plenum as Q_{top} and leaves the shell as Q_{wall}. The gas the cap makes joins film cooler air and shell inleakage and is drawn through the train resistance $K_t + R_\ell$ by the exhauster at demand u_x. TIC-101 writes electrode power, PIC-102 writes exhauster demand against a pressure tap carrying transport delay θ_P, LIC-105 writes the airlift u_p, and FRC-104 writes the feed rate from the coverage inferred from the plenum temperature.

11.3.1 Glass Pool and the Cold Cap

The pool is one well-mixed inventory heated by the electrodes and cooled by four paths, as shown in Equation (11.1).

$$m_g c_p \frac{\mathrm{d}T_g}{\mathrm{d}t} = P_{el} - Q_{cc} - Q_{exp} - Q_{wall} - Q_{sens} \tag{11.1}$$

where m_g is the glass inventory in kg, c_p is the glass heat capacity in kJ/(kg K), P_{el} is the delivered electrode power in kW, Q_{cc} is the duty crossing into the cold cap, Q_{exp} is radiation from the uncovered melt into the plenum, $Q_{wall} = UA_m(T_g - T_a)$ is the shell loss and Q_{sens} is the sensible heat needed to raise newly made glass from the cap underside to pool temperature. The product $m_g c_p$ is 31.5 MJ/K at the design inventory, and the four loss terms together shed 6.84 kW/K about the design point, which is measured directly in Section 11.4.2.

The cold cap is one lumped mass whose coverage and thickness follow from it, as shown in Equation (11.2).

$$\varphi = 1 - \mathrm{e}^{-m_{cc}/M_{\mathrm{ref}}}, \qquad \delta = \frac{m_{cc}}{A\rho_{cc}\varphi} \tag{11.2}$$

where M_{ref} is the spreading mass in kg, A is the melt area and ρ_{cc} is the bulk density of the crust. The crust spreads before it thickens, which is the choice that gives the layer its characteristic asymmetry. Coverage responds quickly when the crust is small and slowly when it is nearly complete, and Section 11.4 measures the factor of five between the two.

Table 11.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
T_g (TI-101)	CV, state	1150.0 °C	300–1600 °C
P (PT-102)	CV, state	−250 Pa	−6000 to 3000 Pa
φ	CV, inferred	90.8 %	60–96 % SP
L	CV, measured	0 mm	−50 to 50 mm SP
P_{el}	MV	1260 kW	0–1800 kW
V	MV (volt mode)	664 V	0–750 V
u_x	MV	66.6 %	0–100 %
$\dot{F}$	MV or DV	1425 kg/h	0–2600 kg/h
u_p	MV	13.9 %	0–100 %
w	DV, measured	55 %	35–75 %
y_{ox}	DV, unmeasured	0.78	0.60–0.90
u_b bubblers	DV, measured	1×	0–1×
s sulfate	DV, unmeasured	0×	0–1×
η_b exhauster	DV, measured	1×	0.2–1×
η_s supply	DV, measured	1×	0.5–1×
κ train	DV, measured	1×	1–3×
λ inleakage	DV, unmeasured	1×	0.4–3×
m_{cc}	state	215 kg	0–40 000 kg
T_p	state	436 °C	25–1400 °C
m_g	state	25 175 kg	8000–40 000 kg
f foam	state	0	0–1
R_ℓ	state	0 Pa s^2/kg^2	0–200 000

The heat crossing into the crust drives everything downstream of it, and it is written in Equation (11.3).

$$Q_{cc} = h_{cc}\left[\beta_0 + (1-\beta_0)u_b\right]\left(1 - k_f f\right) A\,\varphi\,\max(0, T_g - T_{cc}) \tag{11.3}$$

where h_{cc} is the melt-to-cap heat transfer coefficient in kW/(m^2 K), u_b is the bubbler air fraction, β_0 is the unbubbled fraction of the coefficient, k_f is the loss at full foam, f is the foam index and T_{cc} is the cap underside temperature at which the batch dissolves. Bubbling is therefore worth $1/\beta_0 = 2.5$ on the coefficient and a full foam raft takes 55 % of it back, both of which follow the direction of the effects reported for melter feeds [11, 13].

> ⧖ HISTORY **Four Bubblers for a Bottleneck (2010)**
>
> Until 2010 the same facility had never filled more than 250 canisters of glass in a year, and most years averaged 215. A new contract to close 22 waste tanks in eight years called for 400, a melt rate the melter as built could not reach. In the late summer of 2010 four bubblers were lowered through nozzles in its top head, two of them doubling as the glass thermowell and the level probe. Started on argon on 28 September 2010, they raised the instantaneous melt rate from about 59 to about 91 kg/h. They also changed how the pool behaved. With both electrode sets in manual at fixed output, the first argon flows dropped the lower pool temperature by about 140 °C in less than an hour, so the electrode power is now adjusted whenever a bubbler flow changes [17].

11.3.2 Conversion Inside the Crust

The order of the two conversion steps matters more than either rate. Slurry lands on a crust whose top is already at a couple of hundred degrees, so the water flashes within seconds, and only what is left of the heat afterward calcines and melts the solids. Equation (11.4) writes it that way.

$$\begin{aligned} \dot{m}_{st} &= \min\left(\dot{F}w,\ \frac{Q_{cc}+Q_{top}}{h_w}\right) \\ \dot{m}_{cons} &= \frac{\max(0,\ Q_{cc}+Q_{top}-\dot{m}_{st}h_w)}{h_d} \\ \dot{m}_{glass} &= y_{ox}\,\dot{m}_{cons}, \qquad \dot{m}_{dc} = (1-y_{ox})\,\dot{m}_{cons} \\ \frac{\mathrm{d}m_{cc}}{\mathrm{d}t} &= \dot{F}(1-w) + \left(\dot{F}w - \dot{m}_{st}\right) - \dot{m}_{cons} \end{aligned} \tag{11.4}$$

where $\dot{F}$ is the slurry feed in kg/s, w is the mass fraction of water, h_w is the enthalpy to boil and superheat feed water, h_d is the enthalpy to calcine and melt dry solids, y_{ox} is the oxide yield of the dry solids, $\dot{m}_{st}$ is the steam rate, $\dot{m}_{dc}$ is the carbon dioxide and nitrogen oxide released and $\dot{m}_{glass}$ is the glass production rate. Tying the steam rate to the melting front instead, which is the obvious lumped simplification, makes a wetter batch produce less steam because less material is melting, and that is the opposite of what a melter does. The enthalpy and yield of a kilogram of slurry follow from Equation (11.5).

$$\Delta h_{feed} = w\,h_w + (1-w)\,h_d, \qquad Y = (1-w)\,y_{ox} \tag{11.5}$$

At 55 % water the melter needs 2788 kJ per kilogram of slurry and gets 0.351 kg of glass from it, while at 68 % it needs 3028 kJ and gets 0.250 kg. A wetter slurry costs more energy per kilogram and contains less glass per kilogram, which is why one feed change moves production so far.

11.3.3 Radiation and the Plenum

The two radiation terms trade against each other as coverage moves, and they are given in Equation (11.6).

$$\begin{aligned} Q_{exp} &= A(1-\varphi)\,\varepsilon_m\,\sigma\left[(T_g+T_0)^4-(T_p+T_0)^4\right] \\ Q_{top} &= A\varphi\,\varepsilon_{cc}\,\sigma\left[(T_p+T_0)^4-(T_{\mathrm{top}}+T_0)^4\right] \end{aligned} \tag{11.6}$$

where ε_m is the effective emissivity of the exposed melt, ε_{cc} is the absorptivity of the cap top, T_{top} is the cap top surface temperature, $T_0 = 273.15\,\mathrm{K}$ converts to absolute temperature and σ is the Stefan-Boltzmann constant. Both are in kW after dividing by 1000. The magnitude is the point. Black-body emission at 1150 °C is 232.6 kW/m^2, so a receding crust appears on the plenum thermocouples within a minute, and Q_{top} is roughly a tenth of the melting energy at design coverage, which is why plenum temperature and melting rate are not independent [3]. The plenum itself is a refractory heat capacity

with radiation in, radiation out, gas enthalpy and a shell loss, as shown in Equation (11.7).

$$C_p \frac{\mathrm{d}T_p}{\mathrm{d}t} = Q_{exp} - Q_{top} - UA_p(T_p - T_a) + \dot{m}_{cc} c_{pv}(T_{ex} - T_p) + \dot{m}^{+}_{leak} c_{pa}(T_a - T_p) \tag{11.7}$$

where C_p is the plenum refractory heat capacity in kJ/K, c_{pv} and c_{pa} are the cap off-gas and air heat capacities, T_{ex} is the temperature of the gas leaving the cap, $\dot{m}^{+}_{leak}$ is inleakage when it is inward and UA_p is the lid loss.

11.3.4 Plenum Pressure and the Off-Gas Train

Pressure is a gas inventory worked against an exhauster, and it is the fastest state in the plant, as shown in Equation (11.8).

$$\frac{\mathrm{d}P}{\mathrm{d}t} = \frac{R_g(T_d + T_0)}{V_{\mathrm{sys}}} \left(\dot{m}_{cc} + \dot{m}_{leak} + \dot{m}_{\mathrm{film}} - \dot{m}_{out}\right) \tag{11.8}$$

$$\dot{m}_{out} = \sqrt{\frac{\max(0,\ P + \eta_b H_0 u_x^2/10^4)}{\kappa\,(K_t + R_\ell)}}$$

$$\dot{m}_{leak} = -\operatorname{sign}(P)\,\lambda\, C_{\mathrm{leak}} \sqrt{|P|} \tag{11.9}$$

where R_g is the specific gas constant of the off-gas, T_d is the duct temperature after the film cooler, V_{sys} is the gas volume of the plenum, cooler, duct and scrubber together, H_0 is the exhauster head at full demand, u_x is the exhauster demand in percent, K_t is the clean train resistance, R_ℓ is the accumulated scrubber loading and C_{leak} is the shell and nozzle inleakage coefficient. Three disturbance multipliers appear here and all three default to unity, namely the exhauster availability η_b, the train restriction κ over and above the accumulated loading, and the inleakage multiplier λ. The coefficient $R_g(T_d + T_0)/V_{\mathrm{sys}}$ evaluates to 1524 Pa/kg. Both sides of (11.8) carry units of Pa/s before the conversion to minutes. The sign in (11.9) reverses when the plenum goes positive, which is contamination leaving the melter rather than air entering it, and at the design vacuum the melter draws 541 kg/h of air in through its own seams. That inward flow supplies most of the loop's self-regulation, because $\partial \dot{m}_{leak}/\partial P$ is -3.00×10^{-4} kg/(s Pa) against 9.38×10^{-5} kg/(s Pa) for the exhauster.

Scrubber loading is the slowest state on the plant and it has no measurement of its own, as shown in Equation (11.10).

$$\frac{\mathrm{d}R_\ell}{\mathrm{d}t} = k_{\mathrm{foul}} \frac{\dot{m}_{cc}}{\dot{m}_{cc,0}} \left(1 + k_c f\right) \gamma \tag{11.10}$$

where k_{foul} is the loading rate at design carryover, $\dot{m}_{cc,0} = 0.2083\,\mathrm{kg/s}$ is the design carryover, k_c is the extra carryover at full foam and γ is an operator-set acceleration used so that a week of loading can be watched in an hour. At design carryover and $\gamma = 1$ the train reaches the 6500 Pa s^2/kg^2 of scenario 8 in 6.1 days.

11.3.5 Electrodes and the Power Envelope

The glass is the resistor, and its resistivity is Arrhenius in absolute temperature, as shown in Equation (11.11).

$$\rho(T_g) = \rho_0\, \mathrm{e}^{E_a/R_u(T_g + T_0)}, \qquad R = k\,\frac{h_{\mathrm{ref}}}{h}\,\rho(T_g) \tag{11.11}$$

$$P_{el} \le \min\left(\eta_s P_{\max},\ \frac{(\eta_s V_{\max})^2}{10^3 R},\ \frac{(\eta_s I_{\max})^2 R}{10^3}\right)$$

where ρ_0 is the pre-exponential in Ω m, E_a/R_u is the activation temperature in K, k is the electrode cell constant at the design level in 1/m, h_{ref} is the design glass depth, $h = m_g/(A\rho_g)$ is the actual

depth, and $V_{\max}$, $I_{\max}$ and $P_{\max}$ are the transformer voltage, electrode current and supply ratings, derated together by the supply availability η_s. The cell constant carries h_{ref}/h because the conducting cross-section between plates on opposite walls is proportional to the submerged depth, so a canister pour is an electrical disturbance as well as a hydraulic one. At the design point ρ is 0.0501 Ω m and R is 0.3504 Ω, which places the supply at 664 V and 1896 A for 1260 kW against an envelope of 1605 kW. The Arrhenius form used here is a single lumped law calibrated for a low-activity waste borosilicate, and correlations fitted over composition as well as temperature are available [14].

11.3.6 Inventory, Foam and the Actuators

Glass inventory, the foam raft and the four actuators complete the state vector, as shown in Equation (11.12).

$$\begin{aligned} \frac{\mathrm{d}m_g}{\mathrm{d}t} &= \dot{m}_{glass} - \dot{m}_{pour}, \qquad L = 10^3\left(\frac{m_g}{A\rho_g} - h_{\mathrm{ref}}\right) \\ \tau_f \frac{\mathrm{d}f}{\mathrm{d}t} &= \frac{s}{1+\mathrm{e}^{-(T_g - T_{\mathrm{foam}})/w_{\mathrm{foam}}}} - f \\ \tau_i \frac{\mathrm{d}a_i}{\mathrm{d}t} &= a_{i,\mathrm{cmd}} - a_i, \qquad i \in \{P_{el},\ u_x,\ \dot{F},\ u_p\} \end{aligned} \tag{11.12}$$

where $\dot{m}_{pour}$ is the airlift discharge, L is the level deviation from the design depth in mm, s is the sulfate loading, T_{foam} is the foaming onset temperature, w_{foam} is the onset width, τ_f is the foam time constant and τ_i is the lag of actuator i. The airlift discharge is $\dot{m}_{pour} = (u_p/100)\dot{m}_{pour,max}$, with $\dot{m}_{pour,max} = 1\,\mathrm{kg/s}$ (3600 kg/h). The level is a pure integrator with no self-regulation whatever, and the measured ramp is −0.0222 mm/min of level for every percent of airlift demand.

11.3.7 Instruments

Glass temperature, pressure, coverage, moisture and level reach the controllers. Plenum temperature and glass level are available without measurement noise or an additional instrument lag. The glass thermocouple is a first-order lag of 0.60 min with optional colored noise, and feed moisture comes from an analyzer with a 0.05 min lag. The pressure transmitter is a transport delay of 0.07 min followed by a lag of 0.010 min, and the delay is real, because the tap is an air-purged dip tube through a hot lid with a long impulse line to a transmitter outside the shielded cell. Coverage has no instrument and is inferred from the plenum thermocouple by a soft sensor through the two-point calibration of Equation (11.13).

$$\hat{\varphi} = \mathrm{clamp}\left(\frac{T_{p,\mathrm{bare}} - T_p}{T_{p,\mathrm{bare}} - T_{p,\mathrm{full}}},\ 0,\ 1\right) \tag{11.13}$$

where $T_{p,\mathrm{full}}$ is the plenum temperature under a complete cap and $T_{p,\mathrm{bare}}$ is the temperature over a bare melt. The result is filtered with a time constant of 3.5 min. The underlying relation is curved and this reads it as a straight line, so the soft sensor carries a real bias, which is the reason coverage is advisory on a real melter [3] and the reason a scenario exists in which the inference is fooled.

11.3.8 Initialization, Numerics and Constraints

Table 11.2 lists the constants. The physical states are initialized from a steady-state solve for the scenario's feed, power and disturbances. This avoids marching the slow pool to equilibrium, but does not put every automatic controller at equilibrium. Coordinated control sets its temperature target during the initial scan, which can create an immediate power demand. The steady state is solved as a nested pair. The inner problem is a damped two-by-two Newton iteration on the glass and plenum temperatures with the cap mass held fixed, multi-started from 1150 °C, 1280 °C and 980 °C when the warm start fails, and the outer problem is a bisection in $\log m_{cc}$ until what the melter melts equals what is being fed to it. Consumption is monotone in cap mass, so the bisection cannot fail, and failure at the top of the bracket is itself the answer. The melter is over-fed, no steady state exists, and the

Table 11.2 Model Parameters. Selected Physical and Instrument Constants. Hidden Entries Are Unavailable in the Model Panel.

Symbol	Meaning	Value	Unit	Symbol	Meaning	Value	Unit
Adjustable, shown in the model panel							
A	melt surface area	10.0	m^2	T_{ex}	cap off-gas temperature	600	°C
h_{ref}	design glass depth	0.95	m	V_{sys}	off-gas gas volume	150	m^3
ρ_g	glass density	2650	kg/m^3	T_d	duct temperature	380	°C
c_p	glass heat capacity	1.25	kJ/(kg K)	K_t	clean train resistance	8115	$Pa\,s^2/kg^2$
UA_m	shell heat loss	0.090	kW/K	H_0	exhauster head at full demand	8455	Pa
E_a/R_u	activation temperature	12085	K	$\dot{m}_{film}$	film cooler air	0.25	kg/s
k	electrode cell constant	7.0	1/m	C_{leak}	inleakage coefficient	0.0095	$kg/(s\,Pa^{0.5})$
V_{max}	transformer voltage limit	750	V	C_p	plenum heat capacity	420	kJ/K
I_{max}	electrode current limit	2600	A	UA_p	plenum shell loss	0.120	kW/K
P_{max}	supply rating	1800	kW	k_{foul}	train loading rate	0.60	$Pa\,s^2/(kg^2\,min)$
M_{ref}	cap spreading mass	90	kg	k_c	carryover at full foam	2.5	—
ρ_{cc}	cap bulk density	1100	kg/m^3	$\tau_{P_{el}}$	power supply lag	0.10	min
h_{cc}	melt-to-cap coefficient	0.66	$kW/(m^2\,K)$	τ_{u_x}	damper travel	0.04	min
T_{cc}	cap underside temperature	980	°C	$\tau_{\dot{F}}$	feed transport lag	0.25	min
T_{top}	cap top temperature	160	°C	τ_{u_p}	airlift response	0.30	min
ε_m	exposed-melt emissivity	0.55	—	τ_{T_g}	thermocouple lag	0.60	min
ε_{cc}	cap top absorptivity	0.75	—	θ_P	pressure tap delay	0.07	min
β_0	unbubbled fraction	0.40	—	$\tau_{\hat{\varphi}}$	coverage filter	3.5	min
k_f	heat-transfer loss at full foam	0.55	—	$T_{p,full}$	plenum, complete cap	365	°C
τ_f	foam time constant	12	min	$T_{p,bare}$	plenum, bare melt	1160	°C
T_{foam}	foaming onset temperature	1120	°C	*Hidden, never displayed*			
w_{foam}	foaming onset width	28	°C	ρ_0	resistivity pre-exponential	1.027×10^{-5}	Ω m
h_w	feed water enthalpy	3620	kJ/kg	T_a	ambient temperature	25	°C
h_d	dry solids enthalpy	1770	kJ/kg	R_g	off-gas gas constant	350	J/(kg K)
				c_{pv}	cap off-gas heat capacity	2.10	kJ/(kg K)
				c_{pa}	air heat capacity	1.05	kJ/(kg K)

crust grows until it bridges the plenum. That boundary is the frontier drawn on the operating map of Figure 11.5. The engine of Section 1.6.1 integrates the model with classical fourth-order Runge-Kutta at a fixed step of 0.005 min (0.3 s), the controllers execute every 0.01 min (0.6 s), and the historian declares one sample every seven physics steps (2.1 s), giving approximately 28.6 samples per minute. The default seed is 20 260 118.

The constraints the supervisor and the operating map enforce are a glass temperature floor of 1050 °C, a refractory ceiling of 1210 °C with a guard band of 12 °C below it, coverage between 55 % and 97 %, a plenum temperature limit of 750 °C, an exhauster demand ceiling of 95 % and the supply envelope of (11.11).

11.3.9 Assumptions

Six simplifications carry the model, and each one fails somewhere specific.

One lumped cold cap. A real crust is a layered structure whose relative zone thicknesses set the melting rate [2, 12], and here all of it sits inside one heat transfer coefficient and two enthalpies. The model can therefore show what a thicker cap does, but not why one batch forms a thicker cap than another at the same water content.

Coverage is exponential in cap mass. The shape in (11.2) gives the layer the right asymmetry, fast when small and sluggish when nearly complete, but the constant of 90 kg for 63 % of a 10 m^2 surface is a calibration rather than a measurement.

One Arrhenius resistivity and one electrode pair. A real melter has several electrode pairs on separate transformers, current that leaks into the cold cap and along the refractory, and a resistivity that depends on alkali content as much as on temperature [14]. Equation (11.11) would mispredict the current split on a multi-pair machine entirely.

One resistance and one gas volume for the whole off-gas train. The real train is a film cooler, a submerged-bed scrubber, a wet electrostatic precipitator, mist eliminators, filters and a caustic

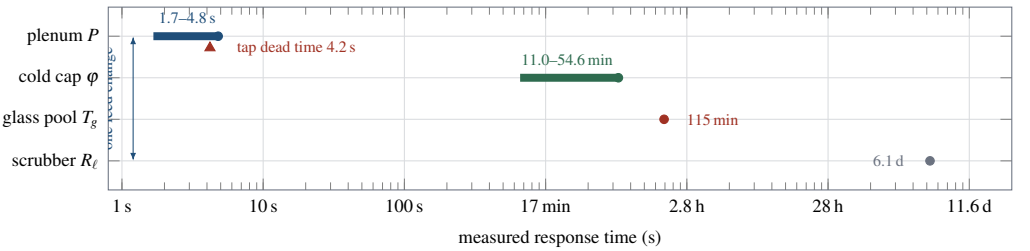


Figure 11.2 Response times measured with damper, feed and power tests on the current model. A moisture disturbance couples to all four states. The scrubber entry is a time to reach the loading of a week-old train rather than a first-order constant, and the cold-cap bar runs between 57.8 and 90.8 % coverage.

scrubber, each with its own volume and condensation. The single R_ℓ of (11.10) reproduces the control-relevant behavior, authority disappearing from the pressure loop over days, and nothing else. Off-gas is also one ideal gas at a fixed duct temperature, whereas the steam fraction really swings from a third to two thirds with feed moisture.

Foaming is one first-order state. Real foaming depends on redox state, sulfate solubility, glass composition and bubbler action together, is strongly hysteretic and is not reliably predictable [13]. The sigmoid in (11.12) reproduces the onset and the insulating consequence and nothing about the chemistry.

Volatilization is absent. An exposed melt on a real machine loses cesium, technetium and halides to the off-gas, which is much of why the coverage floor exists. Here the floor is imposed as a constraint rather than emerging from a mass balance, and the model cannot be used to argue about volatile retention.

11.4 Dynamic Analysis

11.4.1 Four Time Scales on One Plant

The four principal dynamics were measured on the shipped model by stepping one input and reading the 63 % time against a stated late observation. The feed-step baselines were held for 500 min and the responses were observed for another 500 min. Plenum pressure reaches 63 % of an open-loop damper step in 4.8 s, against a gas-inventory time constant of 1.67 s computed from V_{sys} and the self-regulation of (11.8). Cold-cap coverage reaches 63 % of a +50 kg/h feed step in 54.6 min at the design coverage of 90.8 % and, for a +100 kg/h step, 11.0 min at 57.8 % coverage, a factor of 5.0 from the same equation at two operating points. Glass temperature reaches 63 % of a +140 kW power step in 115 min. Scrubber loading reaches the resistance of a loaded train in 6.1 days at design carryover. Figure 11.2 places all four on one logarithmic axis.

The span from 1.7 s to 6.1 days is a factor of 3.1×10^5, and what makes it a control problem rather than a modeling curiosity is that one disturbance excites all four. A batch thirteen points wetter raises the steam rate within one scan, thickens the crust over two hours, changes the pool energy balance over the same two hours and loads the scrubber for the rest of the campaign.

11.4.2 Joule Heating Is Open-Loop Unstable on Voltage Control

The Arrhenius resistivity of (11.11) makes the electrode circuit a feedback path in its own right, and whether that path is positive depends entirely on which quantity the supply regulates. Two gains decide the question, and both were measured on the shipped model rather than estimated.

The first is the electrical gain. Resistance falls by a factor of 1.82 between 1100 °C and 1200 °C, so at a fixed tap of 664 V the delivered power V^2/R changes by 7.52 kW/K about the design point. The second is the thermal gain. Stepping the delivered power by 10 kW on power regulation and

solving for the new steady state moves the pool by 1.462 °C, so the melter sheds 6.84 kW/K. The electrical feedback is larger than the plant's steady-state thermal damping by a factor of 1.10. The fixed-voltage and fixed-power trajectories below demonstrate the resulting difference in open-loop behavior.

The trajectories distinguish the two regulation modes. Taking 34 kW out on power regulation moves the pool from 1150.0 °C to 1145.3 °C and it stops there. Lowering the voltage tap from 664 V to 655 V starts a walk that never stops, and the pool reaches 1129.3 °C after five hours, 1104.1 °C after ten, 1053.3 °C after seventeen and 1020.0 °C after twenty-five, with the delivered power falling from 1258 kW to 338 kW and glass production reaching zero. Nobody would identify that as an instability inside a single shift, which is why the design decision is made structurally instead of by observation. Every operating joule-heated melter regulates power and runs the tap changer as a slave.

🕮 **Pause.** Try Exercise 11.3 before reading on.

The same law has a second consequence that runs the wrong way. The deliverable power of (11.11) is bounded by the transformer voltage at low temperature, where the glass is resistive, and by the electrode current at high temperature, where it is conductive. Because resistance rises as the pool cools, the envelope collapses when the melter needs it most, falling from 1605 kW at 1150 °C to 1178 kW at 1100 °C, 845 kW at 1050 °C and 590 kW at 1000 °C. A cold melter cannot heat itself back up, and the recovery procedure is to stop the feed rather than to turn up the power. The application trends the envelope on the same strip as the power demand for that reason, and reports the margin between them on the scorecard.

11.4.3 Electrode Power Does Not Set Production

The steady-state consequence of (11.4) is worth measuring explicitly because it reverses the intuition a heated reactor teaches. Stepping the electrode power from 1260 kW to 1400 kW at a fixed feed rate of 1425 kg/h raises the pool by 25.9 °C, or 0.185 °C/kW, read 800 min after the step. The native schedule removes the step at minute 400, where the rise is 24.4 °C. The extended test restores 1400 kW immediately after that removal and holds it through minute 820. It is not the unmodified scenario trajectory. The first five minutes run at 0.253 °C/min, which is exactly what 140 kW does to a 31.5 MJ/K pool with nothing else happening, and the remaining hundred minutes are the cold cap catching up. Coverage falls from 90.8 % to 71.7 % and the plenum rises from 436 °C to 612 °C, while glass production moves from 500.2 kg/h to 500.4 kg/h. The transient briefly increases melting as stored cold-cap material is consumed. Once the cap inventory settles, production returns to the rate set by the feed and oxide yield.

Feed rate sets production. Power, temperature and coverage between them decide only how much feed the melter can get away with, and every control decision in Section 11.5 follows from that sentence. In plantwide terms the feed pump is the throughput manipulator, the same choice of production handle that Volume 6, Chapter 7 makes for a reactor and its recycle. It also explains why the temperature setpoint has to move when production is maximized. Melting rate is $h_{cc} A \varphi (T_g - T_{cc})$ by (11.3), so holding both the temperature setpoint and the coverage target fixed pins the melting rate and leaves the feed rate free to decide only how thick the crust gets. At 1150 °C with the cap exactly on its coverage ceiling the absolute ceiling is 1484 kg/h of feed and 521 kg/h of glass, and everything beyond that comes from running hotter.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.4.4 The Pressure Loop and Its Dead Time

Pressure is the opposite extreme and the loop that keeps containment. A closed-loop setpoint step from −250 Pa to −400 Pa moves the damper by 5.01 %, giving a process gain of −29.9 Pa/%, and the pressure reaches 63 % of the step in 9.3 s and settles inside 10 Pa in 38.7 s. The small-signal steady derivative taken directly on the model is −26.8 Pa/%, and the difference between the two readings is the curvature of the square-root characteristic in (11.8). What limits the tuning is neither of these. It

Table 11.3 Implemented Loops. Direction $a = +1$ Is Reverse Acting and $a = -1$ Is Direct Acting. All Six Execute Every 0.01 min With $\tau_D = 0$.

Loop	CV → MV	a	K_c	τ_I (min)
TIC-101	$\hat{T}_g \to P_{el}$	+1	25 kW/°C	45
PIC-102	$\hat{P} \to u_x$	−1	0.015 %/Pa	0.13
CIC-103	$\hat{\varphi} \to \dot{F}$	+1	1800 kg/h/φ	25
LIC-105	$L \to u_p$	−1	4.5 %/mm	40
OVR-T	$T_{g,\min} \to \dot{F}$	−1	60 kg/h/°C	25
OVR-X	$u_{x,\max} \to \dot{F}$	+1	200 kg/h/%	6

is the 4.2 s of transport delay in the purged tap, which is of the same order as the entire closed-loop response, so the high controller gains require verification against this delay and the chosen integral time. Feedback cannot beat a dead time, and Section 11.5.3 shows what feedforward does instead.

11.5 Control Structure

11.5.1 The Six PI Controllers as Implemented

Every loop uses the velocity-form PID of Section 1.6.2. All six have $\tau_D = 0$, so in this chapter's symbols the proportional and integral recursions are Equations (11.14) and (11.15).

$$
\begin{aligned}
e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k)\,, \quad P_k = K_{c,k} e_k \\
s_k &= s_{k-1} + P_k - P_{k-1} \\
c_k &= u_{\mathrm{bias},k} + s_k + I_{k-1} \\
u_k &= \mathrm{clamp}\,(c_k,\ u_{\min},\ u_{\max})
\end{aligned}
\tag{11.14}
$$

$$
I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t} \right] \Delta t_c \tag{11.15}
$$

where $a = \pm 1$ selects the loop direction, K_c is the controller gain in units of actuator per measurement, τ_I is the integral time in minutes, u_{bias} is the feedforward bias described below, T_t is the back-calculation tracking constant defaulting to τ_I, and Δt_c is the 0.01 min controller period. The running sum s_k is S_k of Equation (1.14) and the output before limiting c_k is v_k of Equation (1.15). Back-calculation anti-windup and manual tracking are those of Section 1.6.2. Tracking makes a loop's output agree with its manual station, although a simultaneous change of control architecture can also change the supervisor's target. Table 11.3 lists the six loops with their directions and default settings in the convention of Section 1.3.

The two override directions repay a sentence each, because both are easy to get backwards. OVR-T is direct acting on the feed, because the pool being below its floor has to reduce the feed and the process gain from feed to temperature is negative on a cool pool, where the floor can bind. Near the design point that gain is small and changes sign, as the identified path of Table 11.4 shows. OVR-X is reverse acting, because more feed is more off-gas is more damper, so an exhauster demand above its ceiling has to reduce the feed. Actuator limits are 0 kW to 1800 kW on power, 0 % to 100 % on the damper and the airlift, and 0 kg/h to 2600 kg/h on the feed pump. No loop carries a rate limit, because the four actuator lags of (11.12) supply the physical rate limiting.

11.5.2 Classical Architectures

MANUAL and AUTO decide whether the loops are in service. The architecture selector decides how they are wired to each other once they are. Independent loops and coordinated control remain available alongside the MPC described in Section 11.5.4. In MANUAL every actuator is whatever the operator typed and every controller tracks it, so the transfer either way is bumpless.

Independent loops are what a melter is commissioned with. Temperature holds itself on power, pressure on the exhauster, level on the airlift, and the feed rate is a number an operator typed. It is stable, slow and conservative, because the operator picks a rate that is safe on the worst batch in the tank and the melter then runs at that rate on every other batch as well.

Coordinated control makes the feed rate a manipulated variable. The coverage controller CIC-103 pushes the feed up toward its target, and the two override controllers pull it back down through a low selector, as shown in Equation (11.16).

$$\dot{F}_{SP} = \min\Big(\mathrm{PI}_{\varphi},\ \mathrm{PI}_{T},\ \mathrm{PI}_{x},\ \dot{F}_{\mathrm{op}}\Big) \tag{11.16}$$

where PI_{φ} is the coverage controller output, PI_{T} is the temperature-floor override, PI_{x} is the exhauster-demand override and $\dot{F}_{\mathrm{op}}$ is the operator's feed target, which becomes a hard ceiling rather than a command. The card on screen names which of the four is writing the feed rate at every instant. The same supervisor also walks the temperature setpoint toward the refractory guard band at 0.6 °C/min, clamped between the operator's own setpoint and $T_{g,\max} - 12\,^{\circ}\mathrm{C}$, because Section 11.4 showed that a fixed temperature pins the melting rate. The electrode dryer of Volume 4, Chapter 6 pushes its line speed toward the first binding limit by a similar model-based supervisor. The target it walks toward is the most productive point that satisfies every constraint, found by a 320-point sweep of the steady-state solver over the feed and power plane. That sweep is re-solved at most every 6 min and only when a rounded key built from the measured disturbances and the parameters has changed, and it runs on measured quantities rather than true ones, so it does not know that the exhauster has tripped.

External reset drains the losing controllers' windup through finite back-calculation. After each PI has accounted for its own output clamp, the low selector adds the correction

$$\Delta I_{\mathrm{selector}} = \frac{u_{\mathrm{app}} - u_{\mathrm{PI}}}{T_t}\Delta t, \qquad T_t = \tau_I. \tag{11.17}$$

where u_{PI} is the clamped output u_k of a losing controller and u_{app} is the feed the selector applies. The proportional response remains available while the integral follows the selected feed over a finite time. Manual operation instead tracks each controller to the operator's applied feed for a bumpless return to Auto. OVR-X reads delivered damper travel, not the preceding command, because that command already contains feedforward from the selector's own feed output. Reading the command would close a high-gain loop around the selector and cause scan-to-scan chatter.

11.5.3 Feedforward as an Exact Model Inversion

Feed rate reaches plenum pressure in about two seconds and glass temperature in about two hours, so a supervisor that moves the feed to chase production hands the pressure loop a step disturbance it can only discover after the plenum has told a transmitter 4.2 s downstream of it. The feedforward path is the exact inverse of (11.8) evaluated on measured quantities, as shown in Equation (11.18).

$$\hat{\dot{m}}_{in} = \frac{\dot{F}_{\mathrm{cmd}}}{3600}\Big[\hat{w} + (1-\hat{w})(1-y_{ox})\Big] + \dot{m}^{+}_{leak} + \dot{m}_{\mathrm{film}}$$
$$u_{x,ff} = 100\sqrt{\frac{\hat{K}_t\,\hat{\dot{m}}_{in}^{2} - P_{SP}}{H_0}} \tag{11.18}$$

where $\dot{F}_{\mathrm{cmd}}$ is the feed rate the pump is being asked for, $\hat{w}$ is the moisture the analyzer reports and $\hat{K}_t$ is the train resistance computed from the duct venturi and the differential-pressure cell across the scrubber. The result enters PIC-102 as u_{bias} in (11.14), so it moves the damper in the same scan the feed moves. An equivalent inversion of (11.1) is available as a bias on TIC-101 and is worth much less, because that loop has two hours to notice. Nothing in (11.18) knows that the exhauster

Table 11.4 Nominal Predictive Paths. Time Constants and Delays Are in Minutes. Gains Use Physical Output per kW or per kg/h.

Path	Gain	Lags	Delay[a]	RMS[b]
$P \to \hat{T}_g$	0.144762	77.491	0	0.00350
$F \to \hat{T}_g$	0.003291	28.403	3	0.00218
$P \to \hat{\varphi}$	−0.001498	84.593, 38.547	5	$2.46\,10^{-5}$
$F \to \hat{\varphi}$	0.001043	43.521, 4.625	5	$5.92\,10^{-5}$
$F \to u_x$	0.008380, 0.008236	0.1, 36.901	0	0.00203

[a]Implemented delay. Fitted delays are 0.40, 2.58, 5.20 and 5.15 min. [b]Physical output units. Exhauster paths add in parallel, fitted after the first minute. Other two-lag paths are cascades.

has tripped or that the bubblers have failed, which is deliberate. A feedforward that is right about everything is a simulation rather than a feedforward, and the interesting question is what the feedback has to do when it is wrong.

11.5.4 Model Predictive Control

The optional model predictive controller (MPC) plans electrode power and slurry feed together [18]. It follows the general form of Section 1.6.3, with the plant-specific model, observer and objective below. Its fixed local model is identified from symmetric ±25 kW and ±25 kg/h steps about 1260 kW and 1425 kg/h, following 30 min of settling and 360 min of observation. The outputs are the filtered glass thermocouple, inferred coverage and the pressure controller's exhauster command. Table 11.4 gives the implemented coefficients and independently reproduced fit residuals. The small positive feed-to-temperature gain includes the cap's insulating response. These local coefficients differ from the large-step response times in Section 11.4.

Commands and outputs are normalized as

$$v = \frac{u - u_0}{1000}, \qquad z = D_y^{-1}(y - y_0),$$
$$y_0 = (1150.00648, 0.9102041, 66.6326),$$

where $u = (P_{\mathrm{cmd}}, F_{\mathrm{cmd}})$ is the command pair, $u_0 = (1260, 1425)$ in kW and kg/h, and $y = (\hat{T}_g, \hat{\varphi}, u_x)$ holds the outputs, so the components of y_0 are in °C, coverage fraction and percent. The diagonal scaling D_y has the entries 50 °C, 0.1 and 10 %. Each lag uses $a = \exp(-1/\tau)$ and the sampled recursion $x^+ = ax + (1-a)Kq$, where q is its normalized input and K its scaled gain. The normalized path gain is $1000K/D_{y,ii}$. Cascades use the preceding sample of the upstream lag, with unit gain in the second lag. Two five-sample command histories supply the delays. Eight lag states, ten delay states and three random-walk output biases give 21 states in $x^+ = Ax + Bv$, $z = Cx$. Each output sums its corresponding paths and one bias. Short-delay rounding and cascade sampling add RMS errors of 0.00610 and 0.000245 °C on the two temperature paths, and $3.78\,10^{-5}$ and $5.24\,10^{-5}$ on coverage, relative to the continuous fits.

A steady-state Kalman correction estimates the output biases for offset-free regulation [19]:

$$\bar{x}_k = A\hat{x}_{k-1} + Bv_{k-1}, \qquad \hat{x}_k = \bar{x}_k + L(z_k - C\bar{x}_k). \tag{11.19}$$

The normalized process covariance has diagonal 10^{-8} on the 18 dynamic states and 10^{-4} on the three biases, with measurement covariance $0.0004I_3$. Initialization fills the histories from known commands and puts initial measurement mismatch into the biases. The last sampled applied inputs approximate the preceding minute as a zero-order hold. Envelope clipping or fallback can vary the actual inputs within that minute. This estimator does not recover true coverage from a biased inference.

The decision period is 1 min, while the existing 0.01 min loops retain their scan. A six-hour prediction has 360 samples with block lengths $[1, 1, 2, 4, 8, 16, 24, 36, 48, 60, 72, 88]$ samples. The 24

command variables and three nonnegative horizon-wide slacks minimize

$$J = \sum_{i=1}^{360} \left[(z_{T,i} - r_T)^2 + (z_{\varphi,i} - r_\varphi)^2 - 0.1 v_{F,i} \right] + 5 \sum_{j=1}^{12} \|\Delta v_j\|_2^2 + 10^5 \|\varepsilon\|_2^2, \tag{11.20}$$

where r_T, r_φ are the normalized current supervisory temperature and coverage targets, held over the horizon. All terms are dimensionless. The feed reward is a throughput surrogate, with no energy price or terminal penalty. The shared solver adds a numerical ridge to the QP Hessian. Hard command boxes are $0 \le P_{\text{cmd}} \le P_{\max}$ and $0 \le F_{\text{cmd}} \le \min(F_{\max}, F_{\text{op}})$, with changes at each block boundary limited to 60 kW and 60 kg/h. The first move is anchored to the applied input clipped into the current box. Soft prediction bounds are $T_{g,\min} - 50\varepsilon_T \le T_g \le T_{g,\max} + 50\varepsilon_T$, $\hat{\varphi} \le \varphi_{\max} + 0.1\varepsilon_\varphi$ and $u_x \le u_{x,\max} + 10\varepsilon_x$. Default limits are 1050 and 1210 °C, 0.97 coverage and 95% exhauster demand. Each slack is shared across all 360 samples, including both temperature bounds.

The first solve occurs at initialization. Accepted command pairs hold until the next decision. No future batch information enters either controller. The pressure PI, selected feedforward, level PI, timed pours and physical power envelope remain active. Replaced PIs track the applied commands after clipping. A nonoptimal or invalid plan returns authority to Coordinated until a valid plan is available. Validation checks shapes, finite costs and physical predictions, nonnegative slacks, first-move consistency and every planned command's box and rate limits. The nonlinear plant, its instrument chain and this nominal model remain separate. True foam, cap inventory and future disturbances are evaluation quantities unavailable to the predictive controller.

11.6 Results

The app offers twelve scenarios. They are 1, the glass is the resistor, 2, a power step test, 3, a pressure step test, 4, lined out at the design point, 5, independent loops meet a wetter batch, 6, coordinated control and its feedforward, 7, sulfate foaming, 8, a loaded scrubber, 9, a canister pour, 10, the exhauster trip, 11, MPC and Coordinated on a loaded scrubber, and the eight-hour challenge, listed last. Each fixes the starting mode, the architecture, the six sets of tuning constants, the disturbances and any timed events, and each states its objective before the run. A completion rule judges each one, as Section 1.6.4 describes, and some steps keep an observation across a Reset. A confirmation records an observation that the rule cannot inspect directly. The exercises of Section 11.11 are built on these cases.

11.6.1 The Design Point and What Lining Out Costs

The nominal design initialization uses 1425 kg/h of 55 % slurry and 1260 kW, which lines out at 1150.0 °C of glass and 436 °C of plenum, with the cold cap covering 90.8 % of the 10 m^2 surface at 21.5 mm mean thickness and making 500.2 kg/h of glass at 2.52 kW h/kg, of which 1.58 kW h/kg is spent boiling water that arrived with the waste. The electrode circuit sits at 0.350 Ω, 664 V and 1896 A against an envelope of 1605 kW, and the exhauster holds −250 Pa at 66.6 % demand against 2366 kg/h of total gas, of which the cold cap makes 925 kg/h, the film cooler adds 900 kg/h and the shell leaks 541 kg/h inward. That is 1.20 t/m^2 of glass per day, the number these machines are sold on.

Nothing is misbehaving at that point and a third of the output is missing. Sweeping the feed and power plane against every constraint puts the best feasible point at 662 kg/h of glass, reached at 1885 kg/h of slurry and 1631 kW with the pool at 1204 °C, the cap at 92.2 % coverage and the exhauster at 74.5 %. The binding constraint there is the refractory ceiling, which is why the guard band below it is an engineering decision about melter life rather than a safety factor.

Table 11.5 Measured Outcomes Over the Same 300 min Run With Feed Moisture Stepped to 68 % at $t = 60$ min and Back at $t = 240$ min. Pressure Error Is Accumulated Over $2 \le t \le 300$ min.

	Indep. no FF	Coord. no FF	Coord. FF $\to P$	Coord. FF $\to P, T$
Glass produced (kg)	1968	2444	2444	2454
Above coverage ceiling (min)	188.3	0.0	0.0	0.0
Worst pressure error (Pa)	73.1	90.0	54.2	53.8
Pressure IAE (Pa min)	61.3	136.4	75.0	77.5
Peak glass temperature (°C)	1150.4	1200.7	1200.7	1205.0
Production at 300 min (kg/h)	577	620	620	624
Feed rate at 300 min (kg/h)	1425	1850	1850	1850

11.6.2 The Wet Batch That Nothing Reports

Scenario 5 holds the design point under independent loops and takes the feed from 55 % water to 68 % at $t = 20$ min. Four things move and they move at four different speeds. Within one scan the steam rate goes from 784 kg/h to 969 kg/h, the plenum rises 66.0 Pa before the damper catches it, and the exhauster settles at 70.9 % against 66.6 %. Over the following two hours the crust grows from 90.8 % coverage at 21.5 mm to 99.6 % at 44.9 mm, well past the bridging limit. Production falls from 500.2 kg/h to 314.9 kg/h, a loss of 37 %. And the glass temperature does not move, reading 1150.0 °C before the batch, a maximum of 1150.4 °C during it and 1150.0 °C at the end, with the electrode power going slightly down from 1260 kW to 1242 kW.

> **TRY IT IN THE APP** **5 · Independent Loops Meet a Wetter Batch**
>
> Load scenario 5 and press Run at 10 or 20 times speed. The batch turns wetter early in the run. Watch the glass temperature: it reads 1150.0 °C before the batch and never more than 1150.4 °C during it. Then open the Production tab: glass production has fallen from 500.2 to 314.9 kg/h.

The thermal compensation explains the small temperature excursion. The extra evaporation duty is 186 kW and it does come out of the pool, but the crust that grew to absorb it took the exposed-melt radiation of (11.6) from 110 kW to 5 kW, so the pool lost 105 kW of its own heat leak at the same time. The cold cap absorbed the disturbance and the temperature loop reported nothing wrong. This modeled disturbance illustrates the reason for inferring coverage from the plenum thermocouples and putting it on the screen even though the inference is biased.

Pause. Try Exercise 11.2 before reading on.

11.6.3 Coordination Against Independent Loops

Table 11.5 runs the identical disturbance sequence under four configurations for 300 min. Feed moisture goes from 55 % to 68 % at $t = 60$ min and back at $t = 240$ min. The independent case runs at the design feed rate of 1425 kg/h, and the three coordinated cases are given 2200 kg/h as a ceiling and find their own rate underneath it.

Coordination is worth 2444 kg of glass against 1968 kg, 476 kg or an increase of 24.2 % from the same equipment on the same feed, and it spends 0 min above the cold-cap coverage ceiling where independent control spends 188.3 min. The margin depends on the feed rate chosen for the independent case, and 1425 kg/h is the rate at which this melter makes its nameplate 500 kg/h, so the comparison is against a competent choice. Neither result comes from tuning. The independent loops are correctly tuned and each one holds its own setpoint throughout. What changes is that the coverage controller sheds feed when the crust starts to grow instead of letting a fixed feed rate push it past the bridging limit, and that the supervisor walks the pool to 1198 °C on the dry batch where the operator's setpoint held it at 1150 °C. On a dry batch after twelve hours the coordinated melter

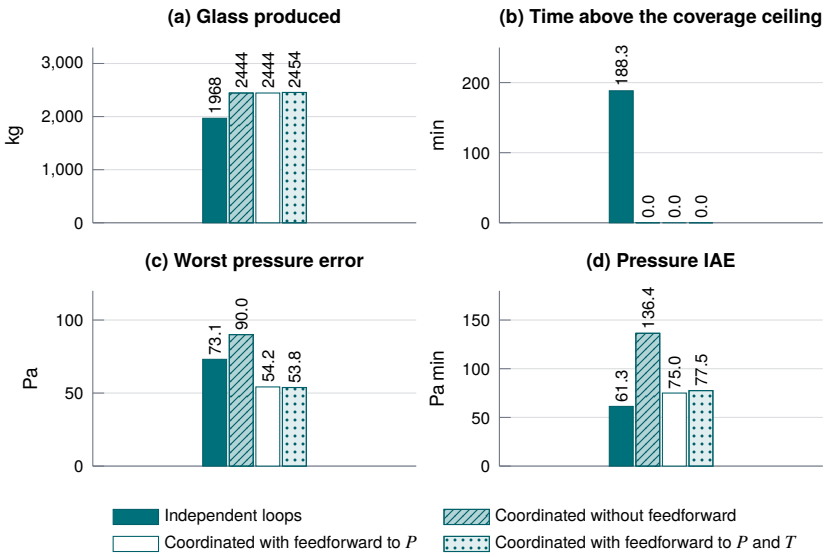


Figure 11.3 The four configurations of Table 11.5 on the same 300 min run with a wetter batch: (a) glass produced, (b) minutes above the cold-cap coverage ceiling, (c) the worst pressure error and (d) the accumulated absolute pressure error. Solid: independent loops. Hatched: coordinated without feedforward. White: coordinated with feedforward to the pressure loop. Dotted: feedforward to both loops. Notice that coordination buys the glass and the coverage, while the pressure feedforward changes only the two pressure panels.

settles at 646.6 kg/h of glass at 1842 kg/h of feed, which is 97.7 % of the best feasible 662 kg/h. The independent column of Table 11.5 reads 577 kg/h at $t = 300$ min rather than 500 kg/h because the oversized crust the wet batch left behind is still melting through faster than the nozzle is feeding it. Figure 11.3 draws four rows of Table 11.5 as bars.

The two middle columns isolate the feedforward. Both produce approximately the same glass, 2444 kg, with differences below one kilogram. The pressure feedforward has only a small indirect effect on the production loops in this comparison. What it changes is the pressure loop, where the worst excursion falls from 90.0 Pa to 54.2 Pa and the accumulated absolute pressure error falls from 136.4 Pa min to 75.0 Pa min. The independent case has the smallest integrated error of the four at 61.3 Pa min, and that is not a point in its favor, because it moved the feed rate exactly once and gave the pressure loop almost nothing to reject. Extending the feedforward to the temperature loop adds 10 kg of glass over five hours with only a small change to the pressure error, which is the expected result when a disturbance path is already two hours slow.

11.6.4 Constraints That Bind at the Actuator

Scenario 8 opens on an off-gas train that has run for a week. Pressure is on setpoint, temperature is on setpoint, production is 500 kg/h and nothing alarms. What changed is the exhauster demand, which reads 88.1 % against the clean train's 66.6 %, because the train now costs 6311 Pa of pressure drop for the same flow instead of 3504 Pa. Twenty-two points of actuator authority have gone and no measurement on the melter reported it. The wet batch that cost 66.0 Pa and took the damper to 70.9 % on a clean train now costs 72.3 Pa and takes it to 92.5 %, within two points of the 95 % ceiling the supervisor is allowed to use and close enough to the stop that a second surge would not be absorbed at all. Turning the feedforward on brings the same excursion back to 45.5 Pa.

A constraint that binds at an actuator has to be relieved at an actuator, and the only actuator that can relieve this one is the feed pump. Holding an actuator position below a ceiling with a slower

manipulated variable is also the idea behind the valve-position controller of Volume 4, Chapter 9. Scenario 10 is the extreme case. The main exhauster trips to a standby unit with 45 % of the head, and the plenum reaches +34.8 Pa and spends 33 s above the −25 Pa containment limit, which no control system could have prevented because the damper travels in two seconds and the head halved. What follows is the point. The exhauster demand crosses its ceiling, the low selector hands the feed rate to OVR-X, and the feed comes down from 1577 kg/h to a minimum of 1052 kg/h. The plenum recovers to −250 Pa without scan-to-scan command chatter. From minutes 100–150 the feed command changes by less than 0.002 kg/h per scan. The standby unit carries the reduced off-gas load, and production at minute 200 is 430 kg/h. The melter produces 1236 kg over the 160 min trip. Under independent loops the feed does not move, the damper sits at 100 %, and the crust bridges completely. That run produces 785 kg, 36.5 % less. Override control lets the plant give up production to relieve the actuator constraint. The coordinated melter still ends at 48.3 % coverage, below the imposed floor, and at a plenum temperature of 788 °C, above its 750 °C limit. The model does not quantify volatilization.

> **TRY IT IN THE APP** **10 · The Exhauster Trips and the Feed Pump Becomes the Pressure Controller**
>
> Load scenario 10 and press Run. The main exhauster trips to a standby unit with 45 % of the head. Watch the plenum pressure jump above its containment limit, then the binding line on the FRC card: the low selector hands the feed rate to OVR-X, the feed falls from 1577 to a minimum of 1052 kg/h, and the plenum recovers to −250 Pa.

11.6.5 Foam, the Pour and the Graded Shift

Scenario 7 sends a sulfate-bearing batch to the melter, which forms a raft between the melt and the cap and cuts the melt-to-cap coefficient of (11.3) by up to 55 %. With independent loops and the feed unchanged, production falls from 500 kg/h to 84 kg/h and the melter produces 542 kg in the 200 min the batch takes to clear, with the electrode power falling too, because a pool that cannot give heat away needs less of it. Two remedies are usually offered and only one works. Dropping the temperature setpoint to 1100 °C does reduce the foam index, but it also takes the driving difference $T_g - T_{cc}$ from 170 °C to 120 °C, and the second effect is larger, so the melter produces 135 kg instead of 542 kg. Cutting the feed target to 900 kg/h produces 1151 kg, and letting the coordinated supervisor find the rate for itself produces 1427 kg, which is 2.6 times what independent loops managed with nobody touching a setpoint.

Scenario 9 requests full airlift for 25 min at minute 20. By minute 50 the cumulative discharge is 1867 kg, including normal level-controlled discharge before and after the request. The level reaches −54.7 mm and resistance rises from 350.4 mΩ to 371.8 mΩ. Power regulation raises the voltage from about 664 V to 684 V. In Auto, voltage regulation also recalculates the tap from the temperature controller's power demand at each scan, so this experiment gives essentially the same delivered power in both modes. It does not hold a fixed voltage. A held tap would lose power as resistance rises, as the Manual experiment in Section 11.4.2 demonstrates. In the added Manual pour comparison, a 664 V tap gives about 1258 kW at minute 20 and 1183 kW at minute 45, while Manual Power holds 1260 kW. The fixed-tap trial also changes temperature, so the difference is not solely a level derivative.

The challenge runs an eight-hour shift with three batches, a partially loaded scrubber accumulating deposits three times as fast as usual, and a view port seal failure at $t = 380$ min. To pass, the melter must produce at least 4100 kg by minute 480 with no time above −25 Pa, less than 5 min more than 100 Pa off the pressure target, less than 10 min below 1050 °C and less than 90 min above the coverage ceiling, and the logged temperature error from $t = 0.66(480)$ min onward must remain within 12 °C. The scenario parameters and disturbances must stay unchanged. The rule judges cumulative glass formation rather than airlift discharge. As delivered, the melter produces 3.880 t and discharges 3.637 t, with 289.4 min above the coverage ceiling and 22.4 min more than 100 Pa off the pressure target.

Table 11.6 Fixed-Setting Comparison Over 480 min. E_T Is Temperature IAE (°C min), $e_{T,\max}$ the Peak Measured Error (°C), t_x Time at or Above 95% Exhauster Demand (min), D_x Its Excess Integral (% min), $x_{\max}$ Peak Demand (%), and t_c True Coverage Above 0.97 (min). n/f Counts Solves/Rejected Plans. Mean Times Are Headless Wall-Clock Observations.

Case	Law	Glass (kg)	Energy (kWh)	E_T	$e_{T,\max}$	t_x	D_x	$x_{\max}$	t_c	n/f	Mean (ms)
Loaded scrubber	C	4038.6	12480	1668.8	40.9	341.275	7.8324	96.01	0.000	0/0	—
	MPC	4013.4	11676	8644.9	40.9	325.735	23.2778	99.98	0.000	480/0	19.05
	C-SIMC	4045.8	12471	1315.3	40.9	358.525	8.8199	96.16	0.000	0/0	—
Wet batch	C	4373.8	12897	1918.3	48.0	0.000	0.0000	77.05	0.000	0/0	—
	MPC	4423.4	12936	2248.6	48.0	0.000	0.0000	78.12	0.000	480/0	18.05
	C-SIMC	4374.3	12883	1732.8	48.0	0.000	0.0000	77.02	0.000	0/0	—
Challenge	C	4413.6	12065	4028.9	48.0	79.125	11.9713	96.69	72.210	0/0	—
	MPC	4425.0	11660	5369.2	48.0	89.500	34.2925	100.00	0.000	480/0	18.34
	C-SIMC	4426.6	12061	2090.4	48.0	70.730	9.9220	96.37	15.120	0/0	—
Sulfate	C	3824.5	10949	2943.7	48.0	0.000	0.0000	68.55	0.000	0/0	—
	MPC	3716.3	9645	14233.0	50.3	0.000	0.0000	68.56	19.185	480/0	18.54
	C-SIMC	3811.3	10957	1437.7	48.0	0.000	0.0000	68.65	0.000	0/0	—
Blower trip	C	4446.4	12405	2527.9	48.0	109.560	10.5789	97.05	0.000	0/0	—
	MPC	4323.8	11391	14794.7	67.5	152.635	76.4851	100.00	0.000	480/0	19.02
	C-SIMC	4450.4	12421	1780.6	48.0	115.710	12.2878	97.30	0.000	0/0	—

Applying suggested settings before Run or applying them and then resetting gives the same trajectory: 4.414 t produced and 4.418 t discharged, with 72.2 min above the coverage ceiling and 0.41 min off pressure. Production rises by 533 kg, or 13.7 %. These settings use power regulation, coordinated control, both feedforward paths, a feed ceiling of 2100 kg/h and a coverage target of 88 %. Both startup paths pass the completion rule. Its temperature test uses the largest logged error from minute 316.8 onward, the final 34 percent of the shift.

11.6.6 Predictive and Selector Control Compared

The added loaded-scrubber case starts with $R_\ell = 6500$ and receives an unannounced moisture change from 55 to 68% at minute 60, returning at minute 240. All five comparisons run 480 min with seed 20260118 and identical instruments. Neither structure receives a batch preview. Coordinated (C) retains shipped tuning, except for the unchanged suggested challenge settings. MPC retains the six-hour settings of (11.20) in every case. A third law, C-SIMC, uses one fixed retuning of the temperature and coverage PIs from the same nominal diagonal paths [20]. The half-rule reduction gives $\tau = \tau_1 + \tau_2/2$ and $\theta = \theta_{\text{fit}} + \tau_2/2$ for coverage. The tuning rule is $K_c = \tau/[K(\lambda+\theta)]$, $\tau_I = \min[\tau, 4(\lambda+\theta)]$. Selecting the highest loaded-scrubber glass total among $\lambda = 5, 10, 20, 40, 80$ min gives $\lambda = 5$ min, with $(K_{c,T}, \tau_{I,T}) = (99.1327, 21.5994)$ and $(K_{c,\varphi}, \tau_{I,\varphi}) = (3528.8971, 45.8335)$ in the units of Table 11.3. This selection is limited to that tested tuning family.

On the loaded train, MPC produces 4013.4 kg against C's 4038.6 kg, while using 11676 against 12480 kWh. Its temperature IAE is 8644.9 against 1668.8 °C min. Exhauster demand is at or above 95% for 325.735 against 341.275 min, but the excess integrals are 23.2778 and 7.8324 % min. The longer C exposure is much shallower. C-SIMC gives 4045.8 kg and an IAE of 1315.3 °C min, with 358.525 min and 8.8199 % min of exhauster excess. During minutes 60–240, C and MPC produce 1144.4 and 1124.6 kg. During minutes 240–480 they produce 2329.6 and 2345.2 kg. Duration, depth, energy and production describe different losses.

MPC also produces less glass on sulfate and blower, at 3716.3 and 4323.8 kg against C at 3824.5 and 4446.4 kg. All 15 runs have zero time below the glass floor, above the glass ceiling or outside the level band. These observations do not establish physical safety. On blower, C's true coverage stays below 0.55 for 142.190 min and its plenum exceeds 780 °C for 137.195 min. MPC avoids both, but containment loss lasts 0.580 against 0.555 min. Containment excess integrals are 15.4941 and 17.0726 Pa min, and the peaks are 59.7903 and 65.6884 Pa beyond the −25 Pa limit. Pressure-error exposure is 1.960 and 4.285 min. On sulfate, MPC instead spends 19.185 min above true coverage

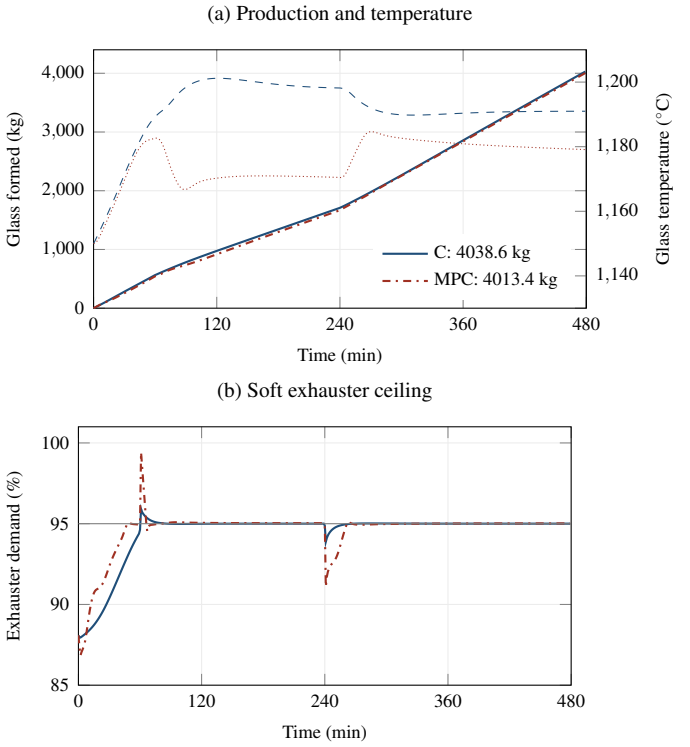


Figure 11.4 Loaded-scrubber comparison with identical axes for both control laws. Coordinated is drawn solid and MPC dash-dotted. Panel (a) plots accumulated glass as thick lines on the left scale and glass temperature as thin lines on the right scale, dashed for Coordinated and dotted for MPC. Panel (b) plots exhauster demand and its 95% ceiling as a thin gray line. The wet batch occupies minutes 60–240. MPC uses less energy but produces slightly less glass, tracks temperature less closely and incurs a larger exhauster-excess integral.

0.97. Time, depth and peak were recorded for every evaluated limit, from physical truth used solely for evaluation. Requested power is clipped for 27.440 and 23.770 min on the loaded train, and 27.440 and 79.780 min on blower. This is unavailable requested power, not a violation of delivered-power capacity.

> ✦ DID YOU KNOW? **The Predictive Controller Wins One Case of Five**
>
> From the simulator: of the five cases of Table 11.6, MPC makes the most glass only on the wet batch, 4423.4 kg against 4374.3 for C-SIMC and 4373.8 for Coordinated. On the loaded scrubber, the challenge, sulfate and the blower trip, a selector law makes more glass than MPC.

Suggested C passes the unchanged challenge with 7.7 °C worst error in its scored interval. MPC produces 4425.0 against 4413.6 kg, but fails at 28.4 °C against the 12 °C band. C-SIMC passes at 2.2 °C. Requested-power travel is 3500 kW for C-SIMC, 2832 for C and 3041 for MPC. The local model, soft output bounds and throughput reward leave losses that output-bias correction cannot remove. These runs do not isolate their individual causes.

Each MPC run accepts 480 plans without fallback. A separate browser measurement gave a warmed mean of 22.10 ms over 219 solves, p95 22.5 ms and maximum 29.0 ms. At the application's 1× rate of one simulated minute per wall second, this is 2.210% of one core. Requested 20× playback

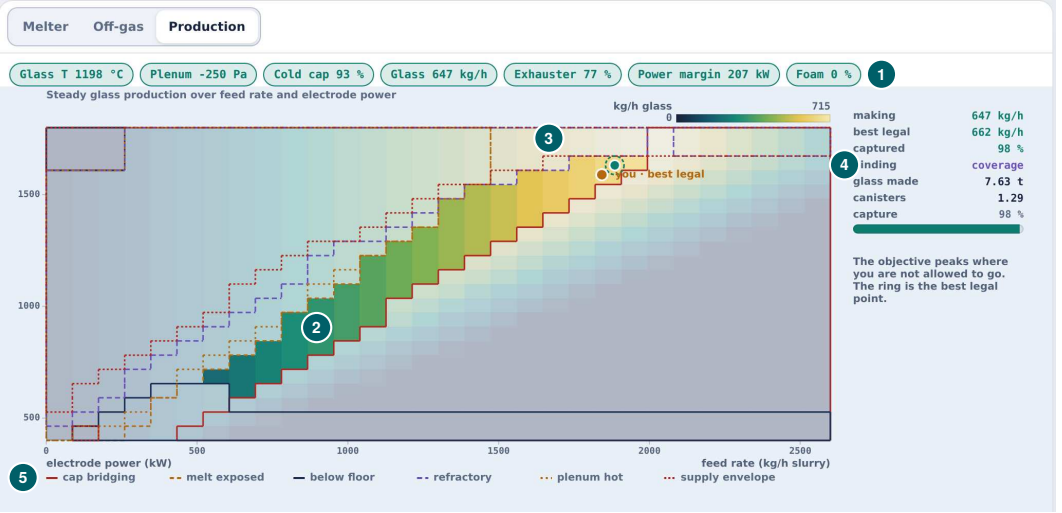


Figure 11.5 Coordinated control at $t = 12.00$ h on the Production view, with the feed ceiling at 2200 kg/h. The operating map plots steady glass production over feed rate and electrode power, with each constraint traced as its own outline, infeasible cells dimmed rather than hidden, the live point trailing behind itself and the best legal point ringed. The melter is making 647 kg/h against a best feasible 662 kg/h, which is 98 % captured, with the coverage controller holding the selection, the pool at 1198 °C, the plenum on setpoint at −250 Pa and the exhauster at 76.6 % of which 10.3 % is feedforward. The view is cut to the Production panel of the application. Callouts: (1) the status line, (2) the operating map with each constraint outline, (3) the live point and the ringed best legal point, (4) the production readings with the binding constraint, (5) the key to the constraint outlines.

measured 18.147 simulated min/s and 49.366 frames/s with zero skipped reached decisions. A 25 ms frame-work guard bounds pending frame work while retaining simulation steps and decisions. A short guard-off observation of 18.238 min/s and 50.332 frames/s does not establish a frame-rate benefit. Trajectory reproducibility excludes wall-clock solver timing.

Figure 11.5 shows the separate dry coordinated run at 720 min, with the objective on the map peaking where the melter is not allowed to go and the best legal point sitting on a corner where two constraints meet.

11.7 Technology in Practice

The tags of Figure 11.1 stand, in a plant, for thermocouples in thermowells, purged dip tubes, transformers on the electrodes, a slurry pump, an exhauster and an airlift. Figure 11.6 draws the loops with that hardware in place of the schematic's symbols.

What Measures It

Glass and vapor-space temperatures are read by thermocouples in thermowells lowered through the top head of the melter, beside the level probe and the feed tube [17]. A thermocouple is a junction of two dissimilar metals whose voltage is roughly proportional to temperature, and the well around it adds a lag [21]. The chapter's glass thermocouple TI-101 is that lag, 0.60 min. The plenum thermocouples carry a heavier job. With no instrument for the cold cap, they are the reading from which coverage is inferred in Equation (11.13), and operating melters set their feed by them [3]. Pressure and level come from purged dip tubes.

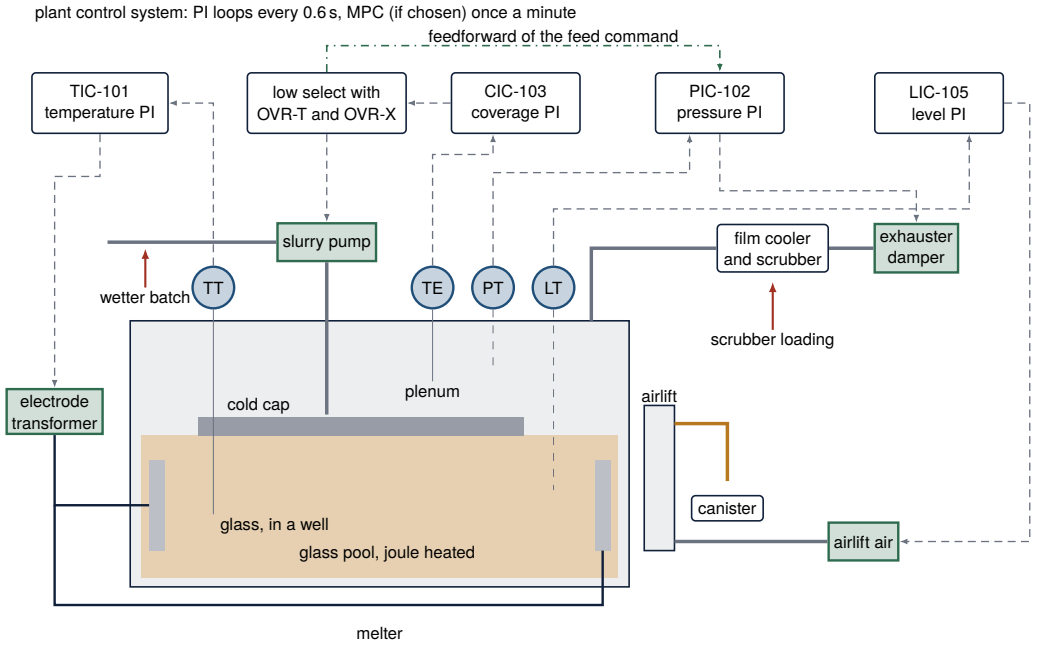


Figure 11.6 The loops of Figure 11.1 with the hardware of a plant. A thermocouple in a well (TT) feeds TIC-101, which sets the power of the electrode transformer. A plenum thermocouple (TE) gives the coverage that CIC-103 holds with the slurry pump, through a low select shared with the overrides OVR-T and OVR-X, where external reset keeps the losing controllers from winding up. Purged dip tubes read the plenum pressure (PT) for PIC-102, which moves the exhauster and receives the feed command as feedforward (dash-dot), and the glass level (LT) for LIC-105, which sets the airlift air. Thick gray lines carry slurry and off-gas, thin dashed arrows signals.

> **TECHNOLOGY** **Sensor: Purged Dip Tubes for Pressure and Level**
>
> A dip tube is a pipe whose open end sits in the process, fed with a purge gas so slowly that only a few bubbles escape. The gas pressure then equals the pressure at the tube's end, and a transmitter anywhere along the line reads it without touching a corrosive or plugging fluid. If the purge stops, the reading is no longer accurate and the tube may plug, too much purge reads high from friction, and each bubble that breaks away leaves a small ripple on the trend [21]. The Savannah River melter reads plenum pressure and glass level from a short and a long pipe of one argon-purged probe, and bubbling moved its indicated level by less than 0.25 cm [17]. The chapter's 4.2 s tap delay stands for this line.

What Moves It

Four final elements carry the six loops. Electrode power comes from a transformer on each electrode set. Slurry enters through feed tubes, and one low-activity design spreads it through six nozzles to give a cap of uniform thickness. Glass leaves by an airlift: air bubbled into the bottom of a riser lifts the glass over a weir into a trough and the canister [16]. Air or argon bubbled into the pool itself raises the heat transfer to the cap [16, 17], which is the multiplier u_b of Equation (11.3). The exhauster at the end of the train holds the plenum below atmospheric pressure, and the chapter models it as a fan whose head grows with the square of its demand u_x.

TECHNOLOGY Final Element: The Electrode Power Supply

Single-phase current reaches the glass through opposing pairs of nickel-alloy plate electrodes, and the glass pool temperature sets the electrode power through a feedback loop [16]. Each electrode set has its own transformer, and the hardware fixes its limits. At Savannah River a change of tap switch position raised the lower transformer's maximum current from 1920 to 2560 A, and the upper controller's current limit went from 71 to 90 % of its 1920 A output, after the allowed current density was doubled. A melter needs both sets: it cannot run on one [15]. The chapter lumps them into one pair with the voltage, current and power envelope of Equation (11.11).

Where the Controller Runs

The loops run in a plant control system, whose loops cycle in times from milliseconds to minutes [22]. The chapter's PI loops execute every 0.01 min (0.6 s) and its predictive controller decides once a minute, both inside that range. The coordinated structure is the override pattern of process control. Several controller outputs meet at a low-select function, only the lowest moves the feed pump, and each de-selected controller must be kept from winding up while it waits [21]. The external reset of Equation (11.17) is that protection, built from one PI block per loop and one selector block.

IN PRACTICE Feeding by the Plenum Temperature

Cameras installed to watch the cold cap at two melters stopped working after very limited time inside, so both plants learned to set the feed rate from the plenum temperature instead. One targeted a plenum between 475 and 525 °C, and the other keeps it at or above 493 °C to control flammable gas. Feeding too fast bridges the crust to the side wall, builds up feed in places, and brings large pressure swings and unstable temperatures [3]. The plenum itself is held slightly below atmospheric pressure to capture the off-gas and keep radiological control [16].

What Is Modernizing It

Two developments address the readings the chapter has to infer. Millimeter-wave sensors view the glass through refractory waveguides from a receiver that can sit outside the biological shield. Their thermal emission signal is linear with temperature, and tests have reached 1500 °C [23]. Computational fluid dynamics models of a pilot melter reproduce its glass and plenum thermocouples on average within 8 and 91 °C, and such models may one day inform the melter's feedback loops [3]. The nonlinear predictive control and moving horizon estimation that Section 11.9 proposes would use the chapter's own equations in the same way.

11.8 Limitations

The model is not a melter design tool. With one lumped cold cap, a single-pair electrode circuit and a one-resistance off-gas train it cannot support a claim about a real machine's throughput, refractory life, off-gas loading or sulfate limit, and the constants that set those answers are calibrations chosen to put the design point at 500 kg/h rather than measurements. Volatilization is absent, so the coverage floor is imposed rather than derived. Glass chemistry does not appear at all, which puts waste loading, durability and the composition constraints that actually size a melter out of reach. The pour is a timed transfer from a well-mixed inventory, with no riser freeze, no canister changeout and no pour spout heating, which are the things that stop a real melter from pouring. Foaming is one first-order state driven by a temperature sigmoid, whereas the real phenomenon is hysteretic and not reliably predictable, which is why sulfate limits on melter feed are set conservatively in practice [13]. Repeatability is itself unrealistic, because freezing the integration step and the controller period reproduces the physical trajectory, excluding wall-clock solver timings, while a distributed control system does no such thing. What survived the simplification is what a controller has to be designed around, namely the self-insulating layer, the four time scales, the positive electrical feedback, the

shrinking power envelope, the biased soft sensor and the actuator constraint, and Section 11.3.9 says where each assumption breaks.

> **OPEN CHALLENGE** **Seeing Inside a Hot, Radioactive Melter**
>
> Waste glass melters lack sophisticated diagnostics because they are hot, corrosive and radioactive. The temperature profiles of the glass, the cold cap and the plenum gas, and the conductivity, density and viscosity of the melt are all wanted, and with them in real time the process could run under feedback control [23]. Readings that disagree with conditions inside are a known operating problem, and the temperature profile matters more as melters grow, because local hot spots have breached refractory [16]. The chapter's one glass temperature and one inferred coverage are that gap in miniature. Progress would be sensors that survive inside the melter and report its profiles.

The predictive model is fixed at the nominal design point. It omits foam, changes in the electrical envelope and future batch composition. Its output-bias estimator corrects measured offsets without identifying those mechanisms or removing the coverage instrument bias. Soft constraints and finite-horizon optimization provide no guarantee of true plant limits or closed-loop stability. The sampled applied-input approximation also omits within-minute variation during clipping and fallback. Halving the physics step from 0.005 to 0.0025 min changes glass by less than 0.005 kg and temperature IAE by less than 0.08 °C min in the loaded-scrubber and suggested-challenge comparisons. Controller scans remain 0.01 min and MPC decisions remain one minute apart.

11.9 Discussion

The melter is a plant where the objective and the constraints live on different actuators. Production comes from the feed pump, the temperature constraint is relieved at the feed pump, the containment constraint is relieved at the damper until the damper runs out and then at the feed pump as well, and the electrode power sets none of it directly. That structure is what the operating map of Figure 11.5 exists to make visible, and reading such a map is a skill that transfers to any plant run against several limits at once.

> **OPEN CHALLENGE** **Coverage From a Thermocouple**
>
> A full-scale melter will rarely let the cold cap be watched, so its feed rate will rest on plenum thermocouples. Simulations of a pilot melter put a swing of up to 606 °C on the plenum from cold-cap coverage and up to 179 °C from the emissivity of the crust alone, and test feeds often build uneven caps that leave plenum temperatures differing from place to place [3]. The soft sensor of Equation (11.13) reads one straight line between two calibration points. Progress would be a reading that tells a lopsided or darker crust from a smaller one, quickly enough for operators to act [3].

The economics change what a controller should optimize. A melter campaign is measured in years, the equipment cannot be inspected or repaired, and the number that decides whether a cleanup mission finishes on schedule is tonnes of glass per day. Every degree of setpoint the supervisor takes toward the refractory limit is throughput bought with melter life, which is why the guard band is adjustable rather than buried. The difference between the operator's 1150 °C setpoint and the supervisor's 1198 °C is 146 kg/h of glass, and the matching cost in refractory corrosion is outside this model entirely.

The model's equations also support methods beyond the implemented PI hierarchy and fixed linear MPC. An eleven-state model with a nonconvex feasible region and an actuator envelope that moves with a state is a natural nonlinear model predictive control problem, and the selector of (11.16) is exactly the structure a constrained optimizer replaces [24]. Moving horizon estimation on the same equations would estimate the cold-cap mass from the plenum thermocouple and the off-gas flow together, which could help distinguish a changing cap from the inleakage disturbance that

biases (11.13) in the final challenge. The model transfers to the Gekko Optimization Suite in Python with the five hidden constants of Table 11.2 recast as estimated parameters [25].

11.10 Conclusion

This chapter studied a joule-heated waste vitrification melter as an interactive laboratory for a plant whose dominant resistance is a state the control action creates. The simulator integrates eleven states whose four principal dynamics span a factor of 3.1×10^5 between the computed gas-inventory time constant of 1.67 s and 6.1 days of scrubber loading. The pressure reaches 63 % of a damper step in 4.8 s, with a cold cap between 11.0 min and 54.6 min depending on coverage and a glass pool at 115 min. Two measured gains decide the electrode regulation, 7.52 kW/K of positive electrical feedback at a fixed tap against 6.84 kW/K of thermal damping, so a 9 V tap change walks the pool from 1150 °C to 1020 °C over twenty-five hours while the same 34 kW removed on power regulation settles at 1145.3 °C. A 140 kW power step raises the pool 25.9 °C, thins the crust from 90.8 % to 71.7 % coverage and leaves production at 500 kg/h, and a batch thirteen points wetter costs 37 % of production while the glass temperature moves 0.4 °C. Making the feed rate a manipulated variable under a coverage controller with temperature-floor and off-gas overrides on a low selector produces 2444 kg of glass in 300 min against 1968 kg for independent loops, with 0 min above the coverage ceiling against 188.3 min, and the feedforward inversion from feed rate to exhauster demand takes the worst pressure excursion from 90.0 Pa to 54.2 Pa with less than a kilogram difference in production. The additional MPC uses less loaded-scrubber energy, 11676 against 12480 kWh, but produces slightly less glass, 4013.4 against 4038.6 kg, with a higher temperature IAE, 8644.9 against 1668.8 °C min, and a larger exhauster-excess integral. No controller wins every constraint. A fuller model would add a layered cold cap, a multi-pair electrode circuit with its own current balance, a volatilization mass balance behind the coverage floor, and a condensing off-gas train. The melter remains a plant for learning to coordinate loops that share one actuator and disagree about time.

KEY POINTS

- A melter makes its glass inside a floating crust, and the same crust insulates the pool and loads the off-gas train.
- Feed rate sets production. Power, temperature and coverage decide only how much feed the melter can get away with.
- On a fixed voltage tap the electrical feedback of 7.52 kW/K exceeds the thermal damping of 6.84 kW/K, so the supply regulates power.
- A coverage controller with two overrides on a low selector makes 2444 kg of glass in 300 min against 1968 kg for independent loops.

Where else this applies. Overrides on a low selector also set the feed of Volume 4, Chapter 4, and a production handle pushed toward the first binding limit drives Volume 4, Chapter 6.

References

[1] R. Pokorny and P. Hrma, “Model for the conversion of nuclear waste melter feed to glass,” *Journal of Nuclear Materials*, vol. 445, no. 1-3, pp. 190–199, 2014.

[2] P. Ferkl, P. Hrma, J. Kloužek, A. A. Kruger, and R. Pokorný, “Cold-cap structure in a slurry-fed electric melter,” *International Journal of Applied Glass Science*, vol. 15, no. 1, pp. 73–87, 2024.

[3] A. W. Abboud, D. P. Guillen, and R. Pokorny, “Effect of cold cap coverage and emissivity on the plenum temperature in a pilot-scale waste vitrification melter,” *International Journal of Applied Glass Science*, vol. 11, no. 2, pp. 357–368, 2020.

[4] P. Kokotović, H. K. Khalil, and J. O’Reilly, *Singular Perturbation Methods in Control: Analysis and Design*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1999.

[5] 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.

[6] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[7] C. M. Jantzen, "Historical development of glass and ceramic waste forms for high level radioactive wastes," in *Handbook of Advanced Radioactive Waste Conditioning Technologies*. Woodhead Publishing, 2011, pp. 159–172.

[8] I. W. Donald, B. L. Metcalfe, and R. N. J. Taylor, "The immobilization of high level radioactive wastes using ceramics and glasses," *Journal of Materials Science*, vol. 32, no. 22, pp. 5851–5887, 1997.

[9] M. I. Ojovan and W. E. Lee, "Glassy wasteforms for nuclear waste immobilization," *Metallurgical and Materials Transactions A*, vol. 42, no. 4, pp. 837–851, 2011.

[10] J. D. Vienna, "Nuclear waste vitrification in the United States: Recent developments and future options," *International Journal of Applied Glass Science*, vol. 1, no. 3, pp. 309–321, 2010.

[11] S. Lee, D. A. Cutforth, D. Mar, J. Klouzek, P. Ferkl, D. R. Dixon, R. Pokorny, M. A. Hall, W. C. Eaton, P. Hrma, and A. A. Kruger, "Melting rate correlation with batch properties and melter operating conditions during conversion of nuclear waste melter feeds to glasses," *International Journal of Applied Glass Science*, vol. 12, no. 3, pp. 398–414, 2021.

[12] D. R. Dixon, M. J. Schweiger, B. J. Riley, R. Pokorny, and P. Hrma, "Temperature distribution within a cold cap during nuclear waste vitrification," *Environmental Science & Technology*, vol. 49, no. 14, pp. 8856–8863, 2015.

[13] S. Lee, P. Hrma, R. Pokorny, J. Klouzek, B. J. VanderVeer, D. R. Dixon, S. A. Luksic, C. P. Rodriguez, J. Chun, M. J. Schweiger, and A. A. Kruger, "Effect of melter feed foaming on heat flux to the cold cap," *Journal of Nuclear Materials*, vol. 496, pp. 54–65, 2017.

[14] P. Ferkl, X. Lu, A. A. Kruger, and J. D. Vienna, "Temperature and composition dependence modeling of viscosity and electrical conductivity of low-activity waste glass melts," *Journal of Non-Crystalline Solids*, vol. 640, p. 123119, 2024.

[15] B. C. Hodges, "Replacement of melter 2 at the defense waste processing facility (DWPF)," in *WM2018 Conference*, Phoenix, AZ, 2018, paper 18608; report SRR-STI-2018-00060, OSTI 1431108. [Online]. Available: https://www.osti.gov/biblio/1431108

[16] S. E. Kelly, "A joule-heated melter technology for the treatment and immobilization of low-activity waste," Washington River Protection Solutions, for the U.S. Department of Energy Office of River Protection, Richland, WA, Tech. Rep. RPP-48935, Rev. 0, Mar. 2011.

[17] M. E. Smith and D. C. Iverson, "Installation of bubblers in the Savannah River site defense waste processing facility melter," in *WM2011 Conference*, Phoenix, AZ, 2011, report SRR-STI-2010-00784, OSTI 1001181. [Online]. Available: https://www.osti.gov/biblio/1001181

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

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

[20] S. Skogestad, "Simple analytic rules for model reduction and PID controller tuning," *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

[21] 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.

[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] P. P. Woskov, J. S. Machuzak, P. Thomas, S. K. Sundaram, and W. E. Daniel, Jr., "Millimeter-wave monitoring of nuclear waste glass melts: An overview," Massachusetts Institute of Technology, Plasma Science and Fusion Center, Cambridge, MA, Tech. Rep. EMSP-81897-2002, Jun. 2002.

[24] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[25] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, "GEKKO optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.

11.11 Exercises

11.1 Concept check · 2 min

The melter runs in Manual at the design point with the feed fixed at 1425 kg/h of slurry at 55 % water. The electrode power is stepped from 1260 to 1400 kW and held until the pool and the cold cap have settled. What does the model predict for the glass production rate?

(a) It returns to about 500 kg/h, the dry-solids feed times its oxide yield

(b) It rises by about 11 %, in proportion to the extra electrode power

(c) It rises, because the melt-to-cap heat flow grows with $T_g - T_{cc}$

(d) It falls, because the thinner cap lets more heat escape as radiation

11.2 Concept check · 2 min

In scenario 5 a batch thirteen points wetter reaches the melter under independent loops at the design feed rate. Production falls by 37 %, yet the glass temperature never moves more than 0.4 °C from 1150.0 °C. Why does the temperature loop see almost nothing?

(a) The temperature loop is tuned fast enough to reject the extra evaporation duty within minutes

(b) The thicker crust cuts the radiation from the bare melt by about as much as the extra duty takes from the pool

(c) The extra water boils in the plenum, so its evaporation never draws heat from the pool

(d) The thermocouple lag of 0.6 min filters the temperature change out of the measurement

11.3 Concept check · 2 min

The melter is switched from power regulation to a fixed voltage tap at its design point, and the tap is then lowered from 664 to 655 V. What happens to the pool, according to the model?

(a) It settles a few degrees cooler, as it does after the same cut on power regulation

(b) It oscillates with growing amplitude, because the supply lag adds phase to the thermal loop

(c) It keeps cooling for many hours, because the power at a fixed tap falls faster with temperature than the losses do

(d) It warms, because a lower voltage drives a larger current through the glass

11.4 Theory and calculation · 6 min

Slurry enters at 1425 kg/h. The feed water needs $h_w = 3620\,\mathrm{kJ/kg}$, the dry solids need $h_d = 1770\,\mathrm{kJ/kg}$, and the dry solids have an oxide yield $y_{ox} = 0.78$. Use Equation (11.5).

(a) Compute the enthalpy per kilogram of slurry, Δh_{feed}, and the glass per kilogram of slurry, Y, at $w = 0.55$ and at $w = 0.68$. *Checkpoint:* 2788 kJ/kg and 0.351; 3028 kJ/kg and 0.250.

(b) Compute the steady glass production at this feed rate for each batch, and the fractional loss. *Checkpoint:* 500.2 kg/h and 355.7 kg/h, a loss of 28.9 %.

(c) Compute the extra evaporation duty of the wetter batch in kilowatts. Scenario 5 reads 314.9 kg/h just before the batch clears, a 37 % loss. Why is the run below the steady value of part (b)? *Checkpoint:* 186 kW.

11.5 Theory and calculation · 7 min

The glass resistivity is $\rho_0 \mathrm{e}^{E_a/R_u(T_g+T_0)}$ with $E_a/R_u = 12085\,\mathrm{K}$. At the design point $T_g = 1150\,^\circ\mathrm{C}$, $P_{el} = 1260\,\mathrm{kW}$ and $R = 0.3504\,\Omega$. The pool holds $m_g c_p = 31.5\,\mathrm{MJ/K}$ and sheds 6.84 kW/K about the design point. The supply limits are 750 V, 2600 A and 1800 kW.

(a) Compute the ratio of the resistance at 1100 °C to that at 1200 °C, at a fixed glass depth. *Checkpoint:* 1.82.

(b) At a fixed tap, $P = V^2/R$. Show that $\mathrm{d}P/\mathrm{d}T_g = P\,(E_a/R_u)/(T_g + T_0)^2$ and evaluate it. Subtract the thermal damping and estimate the time scale $m_g c_p/(\text{net gain})$ on which the pool walks away. *Checkpoint:* 7.52 kW/K; net 0.68 kW/K; about 13 h.

(c) Compute the power envelope of Equation (11.11) at 1150 °C and at 1100 °C, and name the limit that binds. *Checkpoint:* 1605 kW and 1178 kW, both on the voltage limit.

11.6 Theory and calculation · 6 min

The plenum, cooler, duct and scrubber hold $V_{\mathrm{sys}} = 150\,\mathrm{m}^3$ of off-gas with $R_g = 350\,\mathrm{J/(kg\,K)}$ at $T_d = 380\,^\circ\mathrm{C}$. At the design vacuum of −250 Pa the shell inleakage follows Equation (11.9) with $\lambda = 1$ and $C_{\mathrm{leak}} = 0.0095\,\mathrm{kg/(s\,Pa^{0.5})}$, and the exhauster pulls $\dot{m}_{out} = 0.6571\,\mathrm{kg/s}$ through the clean train, $K_t = 8115\,\mathrm{Pa\,s^2/kg^2}$.

(a) Compute the coefficient $R_g(T_d + T_0)/V_{\mathrm{sys}}$, the inleakage in kg/h and its slope $\partial\dot{m}_{leak}/\partial P$. *Checkpoint:* 1524 Pa/kg; 541 kg/h; -3.00×10^{-4} kg/(s Pa).

(b) From Equation (11.8), compute the slope $\partial\dot{m}_{out}/\partial P$ at a fixed damper. *Checkpoint:* 9.38×10^{-5} in kg/(s Pa).

(c) Compute the time constant of the gas inventory, and the share of the self-regulation that the inleakage supplies. Compare the time constant with the 4.2 s transport delay of the pressure tap. *Checkpoint:* 1.66 s; 76 %.

11.7 App challenge · 7 min

Load scenario 1, The glass is the resistor (MANUAL). Before Run, switch to Advanced, set Electrode regulation to Voltage and set Electrode voltage tap to 655 V. Run to $t = 600$ min at 20× (30 seconds) at the default seed and export the CSV file. Its time column is in hours for a run this long.

(a) Read `glass_T_C` at 5 and 10 h, and `power_kW` at 10 h. *Checkpoint:* 1129.3 °C and 1104.1 °C, each within 0.3 °C; 869 kW, within 3 kW.

(b) Reload the scenario, keep Electrode regulation on Power, set Electrode power to 1226 kW before Run and run the same 600 min. Read the glass temperature at the end. *Checkpoint:* 1145.3 °C, within 0.1 °C.

(c) Both runs remove about 34 kW at the start. Explain with the two gains of Section 11.4.2 why one settles and the other does not.

11.8 App challenge · 8 min

Load scenario 6, Coordinated control, and why it needs feedforward, as shipped: Coordinated, Feedforward from feed on To pressure and a feed target of 2200 kg/h. Each run goes to $t = 300$ min at 20× (15 seconds) at the default seed. Read the Glass made tile and the IAE of the melter plenum pressure loop on the scorecard.

(a) Run the case as shipped. *Checkpoint:* 2.44 t; IAE 75.9 Pa min, within 0.5 Pa min.

(b) Pause, press Reset, switch to Advanced and set Feedforward from feed to Off before Run. Run the same 300 min and give the ratio of the two pressure IAE values. *Checkpoint:* 2.44 t; IAE 137.5 Pa min, within 0.5 Pa min; ratio 0.55.

(c) Why does the feedforward change the pressure loop and not the glass? Table 11.5 gives 75.0 and 136.4 Pa min. Why do your readings differ slightly?

11.9 App challenge · 8 min

Load scenario 10, The exhauster trips and the feed pump becomes the pressure controller, as shipped: Coordinated, To pressure feedforward and a feed target of 1800 kg/h. The standby exhauster takes over at $t = 40$ min and the main one returns at 200 min. Run to $t = 200$ min at 20× (10 seconds) at the default seed and export the CSV file. Its time column is in hours.

(a) Read the Containment tile and the peak of `plenum_P_Pa`. *Checkpoint:* 0.6 min; about +35 Pa, within 1 Pa.

(b) From `metric_glass_kg` at the last rows at or before 40 min and 200 min, compute the glass made during the trip, and read the minimum of `feed_kgh` after the trip. *Checkpoint:* 1236 kg, within 5 kg; 1052 kg/h, within 5 kg/h.

(c) Pause, press Reset, set Control architecture to Independent loops before Run and repeat. Give the glass made during the trip and the feed rate. *Checkpoint:* 785 kg, within 5 kg; the feed stays at 1800 kg/h.

11.10 Application area · 3 min

At the end of the exhauster trip of scenario 10 the coordinated melter has held containment and kept producing, but its cold cap covers only 48.3 % of the melt and its plenum is at 788 °C.

(a) What does an exposed melt do on a real machine that this model does not show, and why does the coverage floor exist?

(b) The feed cannot rise, because the exhauster is on its ceiling. Which actuator could an added override use to protect the cap, and what would it cost?

11.11 Application area · 4 min

The supervisor of the coordinated architecture walks the temperature setpoint from the operator's 1150 °C to 1198 °C, 12 °C below the refractory ceiling of 1210 °C, and the chapter puts the difference at 146 kg/h of glass.

(a) Express that difference in tonnes of glass per day and as a share of the 500 kg/h design production. *Checkpoint:* 3.5 t per day; 29 %.

(b) Why is the guard band adjustable rather than fixed in the code, and what must an engineer bring from outside the model to set it?

11.12 Review the agent · 5 min

An AI assistant reviews scenario 11 and writes: "On the loaded scrubber the MPC uses 11 676 kW h against 12 480 kW h for Coordinated, about 6.4 % less. Its exhauster demand is at or above the ceiling for 325.7 min rather than 341.3 min. The predictive controller is better and should replace the selector for all operation of this melter." Find the flaw and explain how to check it in the app.