INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 17 Control of Energy Conversion and Storage

CHAPTER 3

Solid Oxide Fuel Cell Current That Waits for Its Fuel

Pages 49–79

Edited by John D. Hedengren

How to cite this chapter. "Solid Oxide Fuel Cell Current That Waits for Its Fuel," in *Control of Energy Conversion and Storage*, John D. Hedengren, Ed., vol. 17 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 49–79.

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

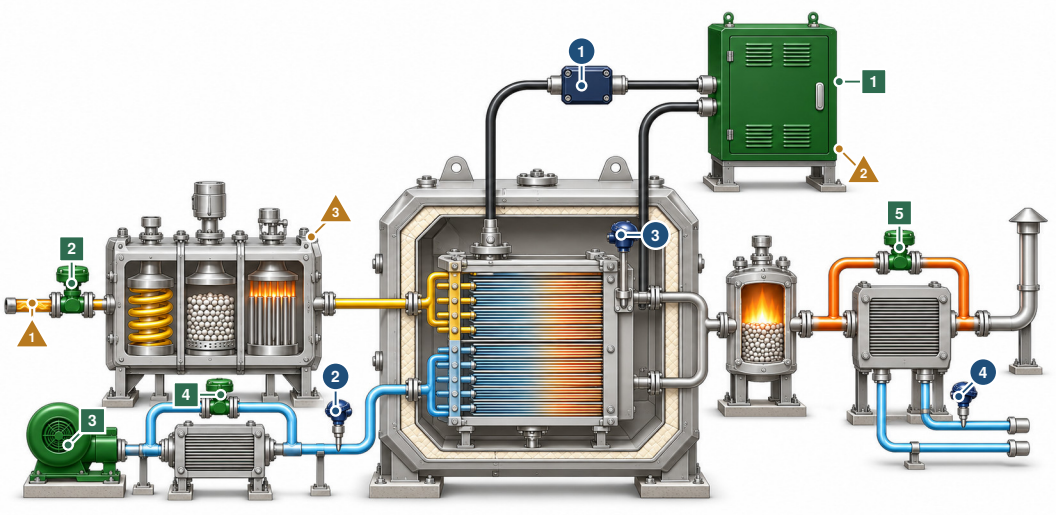
The app for this chapter
https://apmonitor.com/apps/sofc-chp/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v17-energy/#sofc-chp

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

CHAPTER 3

Solid Oxide Fuel Cell Current That Waits for Its Fuel



The controller moves 1 inverter that sets the stack current 2 fuel valve 3 air blower 4 valve that lets air skip the preheater 5 valve that lets exhaust skip the water heater

The controller reads 1 stack current and voltage 2 air entering the stack 3 hot end of the stack 4 hot water to the customer

What disturbs it 1 quality of the fuel gas 2 power the data center asks for 3 a fuel processor that slows down

Notice that the inverter can change the stack current at once, while the fuel must first pass the three compartments of the fuel processor before it reaches the stack. The data center and its battery are not drawn. Illustration, not to scale.

IN PLAIN TERMS A Current That Must Wait for Its Fuel

The main idea. A solid oxide fuel cell makes power for a data center from gas without burning it. Its current can change in a blink, but the hydrogen it uses comes later, after a slow trip through a fuel processor.

The challenge. If the current runs ahead of its fuel, the cell starves and its voltage sags, so the simple control asks for more current still. A limit called a governor holds the current back until the fuel is there, and a battery carries the load while it waits.

Why it matters. A starved cell can be hurt, and a data center cannot pause while its power catches up.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of the Introduction, the history box of Section 3.2, the try-it box of Section 3.5, the fact and try-it boxes of Section 3.5.3, and the key points
Study	55 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 3.5, Section 3.6, Section 3.7 and its open challenges

The control question A solid oxide fuel cell can change its current within a fraction of a second, but the hydrogen that current consumes reaches the anode only after tens of seconds in a fuel processor. How does a controller make the current wait for its fuel, and what carries the load while it waits?

Learning objectives

- Compute the hydrogen a current reacts, its fuel utilization and the current a governor admits.
- Explain the positive feedback by which load-to-current control runs the anode short of hydrogen.
- Predict the delay of the hydrogen through a fuel processor held as stages in series.
- Tell a current limited by arriving fuel from one limited by heat, by the inverter rating or by missing capacity.
- Evaluate how a wrong estimate of gas quality or holdup erodes the governor's margin.

Prerequisites Sections 1.4.2 (time scales), 1.4.4 (a store as an inventory), 1.4.5 (heat and power from one fuel), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios), with Chapter 2 for the shared bus and UPS. Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/sofc-chp/

See also Volume 9, Chapter 4 for an offset-free controller that carries a fuel bias and Volume 9, Chapter 8 for the data hall. Volume 5, Chapter 2 for fuel that acts after a transport delay and Volume 5, Chapter 3 for a slow analyzer.

3.1 Introduction

Solid oxide fuel cells are one of four fuel cell types used for stationary combined heat and power (CHP), and a catalog prepared for the U.S. Environmental Protection Agency gives them projected electric efficiencies of 45 to 60 %, with smaller units typically near 50 % [1]. In the simulated plant, 1000 equivalent cells of 1.2 m^2 at 790 °C carry an islanded data center of 2.8 MW of IT load and 0.635 MW of cooling and auxiliaries at 3.435 MW net, on 7.22 MW of methane on its lower heating value (LHV), and return 1.10 MW of heat as hot water. With every actuator held, a five percent step of the stack current moves the cell voltage from 0.7479 to 0.7397 V within 0.2 s, the anode's fuel margin completes 63 percent of its fall in 1.38 s, and the hydrogen arriving does not move, while a step of the fuel valve needs 17.3 s to deliver 63 percent of its hydrogen to the anode.

+ DID YOU KNOW? A Small Installed Fleet

The EPA catalog of 2017 counted about 83.6 MW of fuel cell CHP installed in the United States, with California first at just under 45 MW. A typical stack joins several hundred cells in electrical series [1].

The simulator is built against a common expectation, that a fuel cell follows its load as quickly as its power electronics do. A load-to-current law that divides the power the site needs by the stack voltage takes the fuel utilization to 0.993, 3.2 s after a 0.35 MW workload step, and the anode's fuel margin to 0.157 against a limit of 0.25. The thinning hydrogen lowers the voltage, the law then asks for more current, and only the inverter's floor on its input voltage stops the runaway. The current can move before the fuel is ready, so something has to make it wait and something else has to carry the load while it does.

The control problem belongs to a general class. A fast actuator stands in front of a slow supply, and a constraint on an internal state (here the hydrogen in the anode) must hold at every instant while the reference changes. The reference governor answers that class by modifying the command so that the constrained response stays admissible, a theory that grew from maximal output admissible sets [2] into a body of work with applications across many plants [3]. Model predictive control holds such constraints over a receding horizon [4], and for tubular solid oxide stacks it has been studied with multivariable models [5] and with constraints against starvation and thermal cracking in load following [6]. The same shape appears when a proton exchange membrane stack must not draw

current faster than its air compressor supplies oxygen [7], and at a fuel cell powered data center whose fuel delivery cannot follow a workload surge, where storage or power capping covers the shortfall [8].

However, the published material treats these pieces apart. Dynamic stack and plant models were written for stack design and for power system simulation [9–11], system studies showed that fuel flow set in proportion to current is effective and that fuel delivery delay limits the transient response [12], control studies chose the controlled and manipulated variables for load following [13, 14], and the data center work sizes storage at rack scale [8]. A student who wants to switch a governor off and watch the anode run short of hydrogen needs a plant in which the current, the processor's holdup, the anode inventory, the stack temperatures and a finite UPS are settings with consequences in amperes, moles and kilowatt-hours. This chapter weighs what a reference governor, a current-to-fuel feedforward and a reserve loop each add to naive load-to-current control on a simulated islanded solid oxide fuel cell plant at a data center, where the estimator's commissioning assumptions erode the governor's margin, and where no control setting replaces missing capacity.

The methods are established ones: the Nernst and Butler-Volmer relations, a holdup modeled as stages in series, the velocity-form proportional-integral-derivative (PID) controller with back-calculation and a low selector of constraint ceilings. The results are matched comparisons on one plant at one seed, and two of them correct the app's first account of its own cases. The current held back after a burst comes from the utilization trim reading the fuel still in the processor as a surplus, not from the UPS recharge, and 17.7 of the 26.1 kWh the UPS gives in the challenge are demand above the inverter's rating, which no governor setting moves. The plant has 27 physical states and 50 ledger states whose energy and atom balances close to 5×10^{-7} of their throughput. It evaluates no logarithm outside its domain, and its controllers are confined to a declared packet of measurements. A predictive controller offered as an option is compared with the governor by ablation, and ten graded cases carry the lessons into the app.

Section 3.2 traces the domain from ceramic electrolytes to fuel-limited load following. Section 3.3 gives the plant equation by equation, from the fuel processor's holdup through the cell voltage and the two stack halves to the bus, the UPS and the data hall, and says where its simplifications stop holding. Section 3.4 gives the controllers: the load-to-current law, the reference governor and its ceilings, the estimator of the arriving hydrogen, the fuel, air and heat loops, the reserve loop, the workload cap and the optional predictive controller of Section 3.4.6. Section 3.5 reports the measurements: three matched comparisons (Section 3.5.1), the recharge that was not the cause (Section 3.5.2), the combined challenge (Section 3.5.3), the predictive controller against the governor (Section 3.5.4) and the checks of the implementation (Section 3.5.5). Section 3.6 describes the hardware of real plants, Section 3.7 states the limitations, Section 3.8 draws the conclusions and Section 3.9 holds the exercises.

3.2 From Ceramic Electrolytes to Fuel-Limited Load Following

The fuel cell is older than the engines it now competes with, as the history box in this section tells. The solid oxide cell's electrolyte is an oxide-ion conductor, usually zirconia, with a nickel and zirconia cermet anode and a lanthanum manganite cathode, and it works above 600 °C [15]. The technology's progress was reviewed at the turn of the century [16].

High temperature is what makes the cell a CHP prime mover and also what slows it. The waste heat is hot enough for cogeneration or a bottoming cycle [17], and the catalog warns that long heat-up and cool-down periods restrict the cells' use in cyclic service [1]. The fuel side sets a second limit. Because the reactant activities fall as utilization rises and the cell voltage cannot exceed the lowest local Nernst potential, which in co-flow sits at the outlet, a utilization of 90 % at conditions typical of these cells costs more than 200 mV of Nernst voltage [17]. Running an anode short of hydrogen is also a risk to the cell, since a nickel anode that oxidizes and is reduced again expands, by the order of

1 % on its first complete oxidation in several studies, and the strain accumulates until it can disrupt the anode or the cell [18].

> ⧖ HISTORY **From a London Demonstration to Ceramic Stacks (1830s to 2007)**
>
> William Grove demonstrated a hydrogen fuel cell in London in the 1830s, and the idea waited more than a century for a practical use. Fuel cells returned to the laboratory in the 1950s, when the United States space program needed power systems with low or no air emissions [1]. The solid oxide cell has the longest continuous development of any type, from the late 1950s, and its solid electrolyte can be cast as tubes, flat plates or monoliths [17]. Early electrolytes conducted well enough only near 1000 °C. Planar cells for 700 to 850 °C instead of 900 to 1000 °C followed, and the lower temperature admits less expensive materials of construction [17]. Test units ran for more than ten years, yet in 2007 these cells were not yet shipping commercially [1].

Dynamic models arrived in two lines. Stack designers simulated planar stacks in three dimensions and in time [9], and power engineers built integrated plant models for system simulation [10] and for the load-following performance of fuel cells beside microturbines [11]. A dynamic model of a 25 kW tubular system put the voltage's response to a change of anode hydrogen on a scale of seconds and the temperature's on a scale of hours, showed that changing the fuel flow in proportion to the current is highly effective, and showed that a fuel flow delay imposes major limitations on the transient response [12]. Model-based control studies of internal-reforming cells followed [13], and a study of four load-following schemes took the DC power, the operating temperature, the fuel utilization and the steam-to-carbon ratio as controlled variables, the current density and the air, fuel and steam flows as manipulated ones, and the cathode outlet air temperature as the measured stand-in for a cell temperature that is hard to measure directly [14].

The answer that limits the current itself came from the low-temperature cell. For a proton exchange membrane stack a linear filter on the current demand had been proposed against oxygen starvation, and a load governor published in 2005 regulated the current drawn so that pointwise-in-time constraints on the oxygen excess ratio and the cathode's oxygen were strictly enforced, and compared its result with that filter's [7]. Data center designers then met the fuel-delivery version of the problem. In a rack-level study the fuel cell matched a power ramp 1.5 minutes after it began, drawing first on the fuel already in the stack, and an energy storage device or server power capping covered the shortfall in between, the storage recharging as a constant load on the fuel cell [8]. Published grid benchmarks pose load following with storage as dynamic optimization with perfect foresight [19], and the plant of this chapter turns that coupling into causal feedback.

3.3 An Islanded Fuel Cell Plant and Its Data Hall

The plant in Figure 3.1 is a teaching design with synthetic values throughout, drawn as the application's Plant view draws it. Fuel gas, methane with an inert share, passes a valve and a steam ratio station into a fuel processor whose external reformer and shift turn each methane molecule and two of steam into four of hydrogen and one of carbon dioxide. The reformate feeds the anodes of a stack of N_s series cells, preheated air from a blower feeds the cathodes, and both off-gases burn in a burner whose exhaust heats the reformer, the air, the steam and finally a hot-water loop. An inverter delivers the stack's power to an islanded bus that also carries the data hall, its chillers and a UPS. There are 27 physical states and 50 ledger states that integrate every boundary flux, and the variables are classified in Table 3.1. Internally the model works in seconds, watts, joules, moles, kelvin and pascals.

3.3.1 Fuel Valve, Steam Ratio and Processor Holdup

Every species enthalpy is a formation enthalpy plus one constant mean heat capacity, $h_i(T) = \Delta h_{f,i} + c_{p,i}(T - T^\circ)$ with $T^\circ = 298.15\,\mathrm{K}$, so the processor, the stack, the burner and the exchangers close on one basis and no heating value or reaction heat is ever added on top. The valve follows its command through a 1 s lag up to 12 mol/s, the ratio station meters steam at $S/C = 2.5$ on the metered

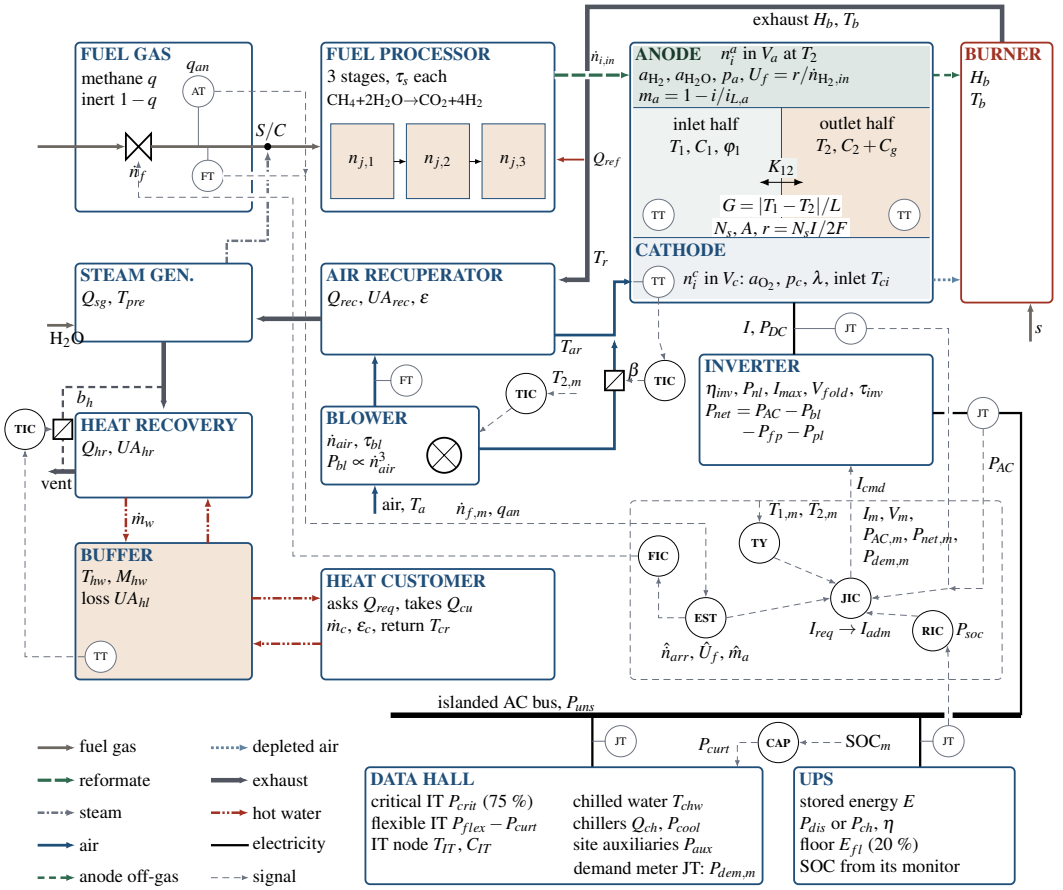


Figure 3.1 The simulated plant as the Plant view draws it. Fuel gas $\dot{n}_f$ of methane fraction q passes the valve, joins steam at the ratio S/C and travels through three processor stages $n_{j,k}$ to the anode, whose inventory n_i^a sets a_{H_2} and a_{H_2O}. The blower's air $\dot{n}_{air}$ leaves the recuperator, mixes with the bypassed share β at T_{ci} and feeds the cathode inventory n_i^c. The two stack halves at T_1 and T_2 exchange $K_{12}(T_2 - T_1)$, the current I makes P_{DC}, and the off-gases burn to H_b with supplementary fuel s. The exhaust gives Q_{ref}, Q_{rec}, Q_{sg} and Q_{hr} in that order before the vent, and the bypass b_h sends part of it around the heat-recovery exchanger. Dashed lines are measurements and commands: JIC (load-to-current and the reference governor) sets the inverter from the meters, the estimator EST, the thermal limiter TY and the recharge request P_{soc} of RIC, FIC sets the valve, TIC loops set the blower, β and b_h, and CAP curtails flexible IT. A process model, not a piping diagram, and not to scale.

gas as if it were pure methane, and the processor holds the feed in three equal stages, as shown in

Table 3.1 Variable Roles and Values at the Design Point (25 °C, 2.8 MW of IT, 1.1 MW of Heat Demand)

Variable	Role	Design	Range
P_{net}	CV, net power	3.435 MW	SP
U_f, m_a	CV, utilization, margin	0.750, 0.733	< 0.85, > 0.25
T_2, T_{ci}	CV, outlet node, air inlet	790, 620 °C	SP
T_{hw}, SOC	CV, hot water, UPS charge	80 °C, 70 %	SP
I	MV, stack current	4996 A	0–7200
$\dot{n}_f$	MV, fuel gas	8.629 mol/s	2–12
$\dot{n}_{air}$	MV, air	457.6 mol/s	150–650
β, b_h	MV, two bypasses	9.6, 22.9 %	0–60, 0–100
P_{IT}	DV, IT load, 75 % critical	2.80 MW	1.5–3.6
Q_{req}, q	DV, heat demand, methane	1.10 MW, 1.00	0–2, 0.7–1
$\Delta\tau_{fp}$, f_{air}	DV, slowdown, blower	0 s, 100 %	0–90, 30–100
$n_{j,k}$	9 states, processor stages	15 s of feed	—
n_i^a, n_i^c	4 + 2 states, electrodes	73.6, 254 mol	—
T_1, T_2	states, stack halves	724.8, 790 °C	—
T_{hw}, T_{IT}, T_{chw}	states, water, hall	80, 23.3, 14	—
Q_{ch}	state, chiller duty	2.85 MW	—
E	state, UPS energy	0.175 MWh	20–95 %
5 states	actuator lags	0.05–5 s	—

Table 3.2 Symbols of the Equations Not Defined in Tables 3.1 and 3.3

Symbol	Meaning	Symbol	Meaning
r, i	H_2 reacted, I/A	E°, E_{rev}	standard, Nernst
a_i, x_i	activity, mole fraction	p_a, N_a	anode pressure, moles
$\dot{n}_{a,out}$	anode outflow	ρ, ρ_{eff}	resistance (ASR)
$i_{L,a}$, $i_{L,c}$	limiting currents	φ_1, φ_2	current shares
C_g	gas heat capacity	G_i, B_i	streams leaving node 1
G	gradient proxy	$Q_{l,j}$	enclosure loss
H_b, s	burner, extra fuel	T_r, T_{x2}	gas temperatures
Q_{ref}, Q_{sg}	reformer, steam	Q_{rec}, Q_{hr}	recuperator, recovery
T_{ar}	recuperated air	Q_{cu}, H_{vent}	heat taken, vent
P_{DC}, P_{AC}	stack, inverter output	P_{bl}, P_{fp}, P_{pl}	blower, processor, plant
P_{gen}, P_{par}	generation, parasitic	P_{dis}, P_{ch}	UPS out, in
P_{uns}, P_{curt}	unserved, curtailed	P_{dem}, P_{soc}	demand, recharge
I_{req}, I_{adm}	request, admitted	I_{ss}, I_{tr}	fuel ceilings
I_V, I_{th}	guard, thermal	I_{air}, I_{inv}	air, inverter
$\hat{n}_{feed}$, $\hat{n}_{arr}$	H_2 fed, arriving	$\hat{U}_f$, $\hat{m}_a$	estimates
I_{ff}, $\dot{n}_{ff}$	feedforward	$\hat{U}_{ff}$, x_{th}	trim PV, excess

Equation (3.1).

$$
\begin{aligned}
\frac{\mathrm{d}n_{j,k}}{\mathrm{d}t} &= \dot{n}_{j,k-1} - \frac{n_{j,k}}{\tau_s}, \qquad k = 1,2,3 \\
\tau_s &= (\tau_{fp} + \Delta\tau_{fp})/3 \\
\dot{n}_{j,0} &= \left(q\dot{n}_f,\ (S/C)\,\dot{n}_f,\ (1-q)\dot{n}_f\right) \\
\dot{n}_{\mathrm{H_2},in} &= 4n_{\mathrm{CH_4},3}/\tau_s \\
\dot{n}_{\mathrm{H_2O},in} &= (n_{\mathrm{H_2O},3} - 2n_{\mathrm{CH_4},3})/\tau_s \\
Q_{sg} &= H_{feed}(T_{pre}) - H_{raw}(T_a) \\
Q_{ref} &= H_{a,in}(T_{fp}) - H_3(T_{pre})
\end{aligned}
\tag{3.1}
$$

where j runs over methane, steam and inert, the stages leave at their contents over τ_s (at least 0.5 s), the outlet also carries $n_{\mathrm{CH_4},3}/\tau_s$ of carbon dioxide and $n_{inert,3}/\tau_s$ of inert, Q_{sg} raises and preheats the liquid water and the feed to $T_{pre} = 150\,^\circ\mathrm{C}$, and Q_{ref} takes the products to $T_{fp} = 750\,^\circ\mathrm{C}$. Reforming and shift are complete at the outlet, so there is no carbon monoxide and its electro-oxidation, internal reforming and anode recycle are absent. The steam ratio per carbon is $2.5/q$, 2.94 at 85 % methane, never below the 2 that complete reforming needs. A reformer with its own furnace, equilibrium

and steam supply is the plant of Volume 11, Chapter 12, where the steam must likewise never fall behind the carbon fed. At design $Q_{ref} = 2.421\,\mathrm{MW}$ and $Q_{sg} = 1.114\,\mathrm{MW}$. Lengthening the holdup also delays the feed already in transit, so a slower processor briefly delivers less.

3.3.2 Electrode Inventories and Faraday's Law

The anode holds hydrogen, steam, carbon dioxide and inert in one mixed volume V_a, the cathode oxygen and nitrogen in V_c, both at the outlet node's temperature, and each drains through a linear conductance into the burner's back pressure p_b, as shown in Equation (3.2).

$$\begin{aligned} r &= \frac{N_s I}{2F}, \qquad \frac{\mathrm{d}n_i^a}{\mathrm{d}t} = \dot{n}_{i,in} - x_i^a \dot{n}_{a,out} + \nu_i r \\ \frac{\mathrm{d}n_{\mathrm{O_2}}^c}{\mathrm{d}t} &= x_{\mathrm{O_2}} \dot{n}_{air} - x_{\mathrm{O_2}}^c \dot{n}_{c,out} - \tfrac{1}{2} r \\ \dot{n}_{a,out} &= k_a \big(p_a - p_b\big)^+, \qquad p_a = N_a R T_2 / V_a \\ a_i &= n_i R T_2 / (V p^\circ), \qquad U_f = r / \dot{n}_{\mathrm{H_2},in} \end{aligned} \tag{3.2}$$

where $\nu = (-1,+1,0,0)$ for hydrogen, steam, carbon dioxide and inert, N_a is the anode's total moles, $x_i^a = n_i^a / N_a$, $p^\circ = 1\,\mathrm{atm}$, $p_b = 104.3\,\mathrm{kPa}$ and $x_{\mathrm{O_2}} = 0.21$. The nitrogen balance of the cathode has the same form without the reaction. The utilization is undefined below 1×10^{-3} mol/s of hydrogen. The anode holds 73.64 mol at 108.5 kPa, about 1.5 s of its throughput, which is why its margin answers a current step on a clock of about a second. Because one equivalent string carries the whole array, fuel and electrical quantities scale together.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.3.3 Cell Voltage and the Domain of Each Logarithm

The voltage is the Nernst potential at the outlet composition less activation, ohmic and concentration losses, with the standard potential computed from the same heat capacities as the energy balance, as shown in Equation (3.3).

$$\begin{aligned} E^\circ(T) &= -\big[\Delta h_r(T) - T\Delta s_r(T)\big]/(2F) \\ E_{rev} &= E^\circ(T_2) + \frac{RT_2}{2F} \ln \frac{a_{\mathrm{H_2}} \sqrt{a_{\mathrm{O_2}}}}{a_{\mathrm{H_2O}}} \\ V &= E_{rev} - \sum_{e=a,c} \frac{R\bar{T}}{F} \sinh^{-1} \frac{i}{2 i_{0,e}(\bar{T})} - i\rho_{eff} \\ &\quad + \frac{RT_2}{2F} \ln\left(1 - \frac{i}{i_{L,a}}\right) + \frac{RT_2}{4F} \ln\left(1 - \frac{i}{i_{L,c}}\right) \\ \rho_{eff} &= \frac{2}{1/\rho(T_1) + 1/\rho(T_2)} \\ \rho(T) &= \rho_0 \exp\big[(E_\Omega/R)(1/T - 1/T_{800})\big] \\ i_{L,a} &= k_{L,a} a_{\mathrm{H_2}}, \qquad i_{L,c} = k_{L,c} a_{\mathrm{O_2}} \\ m_a &= 1 - i/i_{L,a} \end{aligned} \tag{3.3}$$

where $i = I/A$ in $\mathrm{A/cm^2}$, Δh_r and Δs_r carry the constant Δc_p of the reaction, $\bar{T} = (T_1 + T_2)/2$, the exchange current densities are Arrhenius in $\bar{T}$ with a symmetric transfer coefficient of one half, ρ is the area-specific resistance (ASR) with $T_{800} = 1073.15\,\mathrm{K}$, the two halves conduct in parallel and node 1 carries the share $\varphi_1 = \rho(T_2)/[\rho(T_1) + \rho(T_2)]$ of the current, and m_a is the anode fuel margin, the share of the anode's limiting current still unused. At design $E_{rev} = 0.8916\,\mathrm{V}$ less 0.0645, 0.0608 and 0.0184 V gives 0.7479 V at $0.416\,\mathrm{A/cm^2}$, and $\varphi_1 = 0.337$.

Every logarithm has a physical domain, positive activities and $i < i_L$. The functions return an undefined value outside it and never clip an argument. The inverter keeps the plant inside the domain, because it cannot hold its DC input below $V_{fold} = 0.60\,\mathrm{V}$ per cell. Its current reference follows the command through a 50 ms lag, and when the voltage at that reference would fall below the floor, or the bus could not take the power, it draws instead the largest current that satisfies both, found by 40 bisections on the monotone voltage curve with an undefined voltage counted as infeasible. The current that flows is therefore always one at which the cell equation is defined, and refused generation is never created, because less current flows.

3.3.4 Two Stack Halves and Their Heat

Both gas streams enter node 1 and leave node 2, each node's gas leaves at its temperature, and the conduction between the halves is equal and opposite, as shown in Equation (3.4).

$$\begin{aligned} C_1 \frac{\mathrm{d}T_1}{\mathrm{d}t} &= \textstyle\sum_i \dot{n}_{i,in} c_{p,i} (T_{fp} - T_1) \\ &\quad + \textstyle\sum_i \dot{n}_i^{air} c_{p,i} (T_{ci} - T_1) + K_{12}(T_2 - T_1) \\ &\quad - \varphi_1 \left[r \Delta h_r(T_1) + P_{DC} \right] - Q_{l,1} \\ (C_2 + C_g) \frac{\mathrm{d}T_2}{\mathrm{d}t} &= \textstyle\sum_i (G_i + B_i) c_{p,i} (T_1 - T_2) \\ &\quad + R T_2 \frac{\mathrm{d}N}{\mathrm{d}t} - \varphi_2 \left[r \Delta h_r(T_2) + P_{DC} \right] \\ &\quad - K_{12}(T_2 - T_1) - Q_{l,2} \end{aligned} \tag{3.4}$$

where G_i and B_i are the anode and cathode streams leaving node 1 after its share of the reaction, $C_g = \sum n_i (c_{p,i} - R)$ is the heat capacity of the electrode gas inventories, $RT_2\,\mathrm{d}N/\mathrm{d}t$ their flow work, $Q_{l,j} = \frac{1}{2} U A_l (T_j - T_a)$ the enclosure loss and $\varphi_2 = 1 - \varphi_1$. The gradient proxy is $G = |T_1 - T_2|/L$ with $L = 0.10\,\mathrm{m}$, 651.8 K/m at design, a coarse stress indicator and not a prediction of a hot spot or of lifetime. The stack makes 2.67 MW of heat for its 3.736 MW of DC power, 0.715 watts per watt, and the cathode air carries most of it away at $\lambda = 7.42$.

3.3.5 Burner, Recuperator and the Heat Customer

The off-gas burns completely, and supplementary fuel gas s is fired only when the reformer's heating gas would otherwise leave below $T_{fp} + \Delta T_{ref} = 790\,^{\circ}\mathrm{C}$, as shown in Equation (3.5).

$$\begin{aligned} H_b &= H_{a,out}(T_2) + H_{c,out}(T_2) + s\, h_{fuel}(T_a) \\ s &= \frac{\left[Q_{ref} + H_b'(T_{fp} + \Delta T_{ref}) - H_{b,0} \right]^+}{h_{fuel} - \Delta h_p} \\ Q_{rec} &= \varepsilon C_{min} (T_r - T_a) \\ T_{ci} &= (1 - \beta) T_{ar} + \beta T_a \\ Q_{hr} &= \varepsilon' C'_{min} (T_{x2} - T_{hw}) \\ H_{vent} &= H_b - Q_{ref} - Q_{rec} - Q_{sg} - Q_{hr} \end{aligned} \tag{3.5}$$

where $H_{b,0}$ is the off-gas enthalpy, H_b' the products' enthalpy flow, Δh_p the enthalpy of the products of one mole of fuel gas at 790 °C, and every temperature is read from the enthalpy flow on the products' own composition, $T = T^{\circ} + (H - \sum \dot{n}_i \Delta h_{f,i}) / \sum \dot{n}_i c_{p,i}$. The recuperator heats the unbypassed air $(1 - \beta)$ against the gas leaving the reformer at T_r with the counterflow effectiveness ε of $UA_{rec} = 55\,\mathrm{kW/K}$ and delivers it at T_{ar}, the steam generator takes Q_{sg}, and the heat-recovery exchanger heats the water loop with the gas not bypassed, $(1 - b_h)$, which arrives at T_{x2}. The hot water and the

customer close as

$$\begin{aligned} M_{hw}c_w\frac{\mathrm{d}T_{hw}}{\mathrm{d}t} &= Q_{hr} - Q_{cu} - UA_{hl}(T_{hw} - T_a) \\ Q_{cu} &= \min\!\left(Q_{req},\ \varepsilon_c \dot{m}_c c_w (T_{hw} - T_{cr})^{+}\right) \end{aligned} \tag{3.6}$$

At design the burner reaches 942.6 °C, the supplementary fuel is 0.374 mol/s, and the train gives 2.421 MW to the reformer, 8.631 MW to the air, 1.114 MW to the steam and 1.122 MW to the water, venting at 105.0 °C. Burned off-gas is fuel the stack did not consume, so it is never counted as stack fuel, and the efficiencies on one LHV boundary are 47.55 % electric and 62.78 % CHP.

3.3.6 Inverter, Bus and the Stored-Energy Rule

The bus balance and the UPS come from the framework's shared site module `core/cogen.js`, written for the gas engine plant of Chapter 2, moved to the framework with the turbine trigeneration plant of Volume 4, Chapter 9 and imported here unchanged, as shown in Equation (3.7).

$$\begin{aligned} P_{net} &= \eta_{inv} N_s V I - P_{nl} - P_{bl} - P_{fp} - P_{pl} \\ P_{gen} + P_{dis} + P_{uns} &= P_{dem} + P_{ch} \\ P_{dem} &= P_{crit} + P_{flex} - P_{curt} + P_{cool} + P_{aux} + P_{par} \\ \frac{\mathrm{d}E}{\mathrm{d}t} &= \eta_{ch} P_{ch} - P_{dis}/\eta_{dis} \\ P_{dis} &\le \min\!\left(P_{dis,max},\ (E - E_{fl})\eta_{dis}/\Delta t\right) \end{aligned} \tag{3.7}$$

where $P_{bl} = 110\,\mathrm{kW}\,(\dot{n}_{air}/450\,\mathrm{mol/s})^3$, $P_{fp} = 15\,\mathrm{kW} + 1\,\mathrm{kW}$ per mol/s of fuel gas, $P_{pl} = 40\,\mathrm{kW}$, $P_{gen} = P_{net}^{+}$, and a negative net output appears as P_{par} on the demand side. All terms are nonnegative, and each step checks the balance one term at a time. Islanded, the UPS takes the imbalance inside the step's limits, which the energy held when the step begins sets and all four Runge-Kutta stages keep. Those stages carry nonnegative weights, so the integrated energy can at most land on the floor, nothing is clipped afterward and no step needs splitting. Charge is limited the same way toward the 95 % ceiling and by the 0.5 MW rating, and the inverter's current limit includes the charge limit, so a surplus the bus cannot take is never generated. A deficit beyond the UPS is unserved, and this application replaces the module's priority order with a proportional one, so an unplanned shortfall falls on every feeder in the share P_{uns}/P_{dem} and the workload cap is the only planned shedding of the flexible 25 % of the IT load. From its 70 % start the 1 MW, 0.25 MWh UPS holds $0.25 \times 0.50 \times 0.95 = 0.118\,75\,\mathrm{MWh}$ above its 20 % floor, 7.1 min at its rating.

3.3.7 Data Hall, Instruments, Events and Numerics

The data hall is an equipment node and a chilled-water node, as shown in Equation (3.8).

$$\begin{aligned} C_{IT}\frac{\mathrm{d}T_{IT}}{\mathrm{d}t} &= P_{IT,del} - UA_{IT}(T_{IT} - T_{chw}) \\ C_{chw}\frac{\mathrm{d}T_{chw}}{\mathrm{d}t} &= UA_{IT}(T_{IT} - T_{chw}) + Q_x - s_c Q_{ch} \\ \tau_{ch}\frac{\mathrm{d}Q_{ch}}{\mathrm{d}t} &= \mathrm{clip}_{[0,\,Q_{ch,max}]}(Q_{cmd}) - Q_{ch} \\ Q_{cmd} &= UA_{IT}(T_{IT} - T_{chw}) + Q_x \\ &\quad + K_{ch}(T_{chw} - T_{chw,SP}) \end{aligned} \tag{3.8}$$

where $P_{IT,del}$ is the IT power actually served, Q_{cmd} the chiller plant's own duty command, Q_x the half of the site auxiliary load that ends in the chilled water, s_c the served share of the cooling feeder, and the chillers draw $P_{cool} = Q_{ch}/\mathrm{COP} + P_{pump}$ with a COP of 6 and 60 kW of pumps and fans. T_{IT} is a lumped thermal state of the IT equipment, not a measured chip or inlet temperature.

Table 3.3 Model Parameters. Class Is the Source of Each Value (L Checked Against the Literature, F Fitted, T Synthetic Teaching Value) With Its Visibility in the Parameters Panel (V Shown, H Hidden)

Symbol	Meaning	Value Unit	Class	Symbol	Meaning	Value Unit	Class
Stack and electrochemistry				*Air, burner and heat recovery*			
N_s, A	cells, area each	1000, 1.2 —, m^2	T, H	$\dot{n}_{air,max}$, $\dot{n}_{air,min}$	blower range	650, 150 mol/s	T, V/H
$E^\circ(800\,^\circ\mathrm{C})$	from Δh, Δs, Δc_p	0.982 V	L	τ_{bl}, P_{bl}	lag, power at 450	2, 110 s, kW	T, H
ρ_0, E_Ω	ASR, its activation	0.10, 91.4 $\Omega\,cm^2$, kJ/mol	T, F	UA_{rec}, β_{max}	recuperator, bypass	55, 60 kW/K, %	T, V/H
$i_{0,a}$, $E_{a,a}$	anode exchange current	1.5, 80 A/cm^2, kJ/mol	T, H	τ_β, τ_{b_h}	damper lags	3, 5 s	T, H
$i_{0,c}$, $E_{a,c}$	cathode exchange current	0.6, 100 A/cm^2, kJ/mol	T, H	ΔT_{ref}	reformer approach	40 K	T, H
$k_{L,a}$, $k_{L,c}$	limiting currents	8, 12 A/cm^2	T, H	UA_{hr}, $\dot{m}_w$	recovery, water	45, 15 kW/K, kg/s	T, V/H
V_a, V_c	gas volumes	6, 20 m^3	T, V/H	M_{hw}, UA_{hl}	buffer, loss	20, 0.4 t, kW/K	T, V/H
k_a, k_c	outlet conductances	11.4, 56 mol/(s MPa)	T, H	$\dot{m}_c$, T_{cr}, ε_c	customer loop	20, 55, 0.8 kg/s, °C, —	T, H
C_1, C_2	stack halves	3, 3 MJ/K	T, V	*Power, site and UPS*			
K_{12}, UA_l, L	conduction, loss, distance	10, 0.2, 0.10 kW/K, kW/K, m	T, H	η_{inv}, P_{nl}	efficiency, no-load	0.97, 10 —, kW	T, V/H
Fuel processor				I_{max}, $P_{AC,max}$, V_{fold}	limits, foldback	7200, 4.3, 0.60 A, MW, V	T, V/H
τ_{fp}, stages	processor holdup	15, 3 s, —	T, V	τ_{inv}	current lag	50 ms	T, H
S/C, q_{nom}	steam ratio, assumed q	2.5, 1 —	T, V/H	P_{crit}/P_{IT}	critical share	75 %	T, V
T_{pre}, T_{fp}	feed, reformer outlet	150, 750 °C	T, H	C_{IT}, UA_{IT}	IT node	20, 0.3 MJ/K, MW/K	T, H
$\dot{n}_{f,max}$, τ_{fv}	valve open, lag	12, 1 mol/s, s	T, V/H	C_{chw}, τ_{ch}, K_{ch}	chilled water, loading	150, 60, 1 MJ/K, s, MW/K	T, H
Protection (teaching thresholds)				$Q_{ch,max}$, COP, P_{pump}	chillers, pumps	4, 6, 60 MW, —, kW	T, H
$U_{f,max}$	utilization ceiling	0.85 —	T, V	P_{aux}	site auxiliaries	0.10 MW	T, H
m_{lim}, m_{trip}	margin limit, trip	0.25, 0.10 —	T, H/V	$P_{dis,max}$, $P_{ch,max}$	UPS ratings	1.0, 0.5 MW	T, V/H
T_{lo}, T_{hi}, T_{trip}	node limits, trip	650, 830, 860 °C	T, H/V/H	E_{cap}, η	energy, efficiency	0.25, 0.95 MWh, —	T, V/H
G_{max}, R_{max}	gradient, node rate	1000, 5 K/m, K/min	T, V/H	floor, ceiling, SOC_0	UPS charge	20, 95, 70 %	T, V/H/V

Controllers read only a measurement packet. Each analog reading is the truth through a first-order transmitter lag, discretized at the scan with the value at the scan instant, with noise added after the lag: current and voltage (0.02 s, 1 A and 0.2 V on about 750 V), the power meters (0.1 s, 2 kW), the flow meters (0.5 s), the stack thermocouples (5 s, 0.3 K), the cathode inlet and hot water (3 s) and the UPS output (0.05 s), with the battery monitor reporting the charge once a second. A sampled analyzer reports the methane fraction from a sample drawn every 60 s through a 30 s line and analyzed for 60 s, so a reading is 90 to 150 s old. No instrument in the control system reads the anode's hydrogen. The site demand meter reads the load request, so shed load still registers. Every scan draws the same count of random numbers whatever the settings.

Events happen once, at a step boundary, with a mark on the trends. The stack trips and latches at a true anode margin of 0.10, read from an anode-outlet analyzer in the protection chain (a declared teaching simplification), or when a measured node reaches 860 °C, and a trip zeroes the current and puts the fuel to its 2.0 mol/s hot standby in any mode. Each case starts in equilibrium for its own settings, solved by bisection on the current for the site's demand, Newton iterations on the two node balances with the air, and bisection on both bypasses, and its derivatives there are below 2.4×10^{-11}. A fixed step of 0.02 s of the classical fourth-order Runge-Kutta scheme carries the states, the fastest retained modes being the cathode's 35 ms and the anode's 57 ms pressure relaxation and the inverter's 50 ms lag, and the historian logs one row a second. The integration step, the scan and the log are the three clocks of the engine of Section 1.6.1. Table 3.3 lists the parameters with the class of each value.

3.3.8 Where the Simplifications Stop Holding

The plant is hot and operating, so cold start, thermal cycling and degradation chemistry are outside it. One mixed volume per electrode puts the whole anode at its outlet composition, the operative one for the voltage in co-flow [17], but it cannot show a local depletion upstream. The design point, 0.31 W/cm^2 at 0.748 V, sits inside the handbook's 300 to 400 mW/cm^2 for planar stacks on reformate [17] and its 47.55 % inside the catalog's range [1], but no polarization constant was fitted to a cell. The constant mean Δc_p puts E° 0.006 to 0.014 V below the handbook's two-digit values above 650 °C, 0.982 V against 0.99 V at 800 °C. The 3 MJ/K per half lets the stack's temperatures move within a fifteen-minute case, while 1200 m^2 of planar cells with metal interconnects would hold roughly 15 to 50 MJ/K, so the thermal lessons come early (Section 3.5.5). The proportional shortfall is an averaged energy account, since a real islanded bus without shedding would drop whole

Table 3.4 Loops, Scans, Default Tuning and Action (+1 Raises the Output When the Measurement Is Below Its Setpoint), All PI

Loop	PV → MV	Scan	K_c	τ_I	Act.	Output range
trim	$\hat{U}_{ff} \to \dot{n}_f$	1 s	5 mol/s	15 s	−1	$[2, 12] - \dot{n}_{ff}$
air	$T_{2,m} \to \dot{n}_{air}$	1 s	6 mol/(s K)	120 s	−1	$[150,\ \dot{n}_{air,m} + 52]$
preheat	$T_{ci,m} \to \beta$	1 s	1 %/K	30 s	−1	0–60 %
hot water	$T_{hw,m} \to b_h$	1 s	30 %/K	90 s	−1	0–100 %
thermal	$x_{th} \to I_{th}$	1 s	60 A/K	20 s	+1[a]	$[1080,\ \bar{I}]$
guard	$V_m/N_s \to I_V$	0.1 s	20 kA/V	2 s	−1	$[1080,\ \bar{I}]$
reserve	$\mathrm{SOC}_m \to P_{soc}$	5 s	3 MW	300 s	+1	0–0.3 MW

[a] On the negative excess, so an excess lowers the ceiling. $\bar{I} = \min(I_{max},\ \max(I_{req}, I_{adm}) + 300\,\mathrm{A})$. Setpoints $U_{f,SP} = 0.75$, $T_{2,SP} = 790\,^\circ\mathrm{C}$, $T_{ci,SP} = 620\,^\circ\mathrm{C}$, $T_{hw,SP} = 80\,^\circ\mathrm{C}$, $V_g = 0.68\,\mathrm{V}$, $\mathrm{SOC}_{SP} = 70\,\%$.

feeders or collapse. Real limits come from the fuel cell, inverter and UPS manufacturers' data and the applicable codes.

3.4 Admitting Current: Governor, Estimator and Loops

One framework scan of $\Delta t_c = 0.1\,\mathrm{s}$ runs the instruments, the estimator, the load-to-current law, the voltage guard, the governor, the inverter command and the workload cap. Integer counters inside it run the fuel, air, bypass, hot-water and thermal-limiter loops every 1 s and the reserve loop every 5 s, each integrating over the real interval between its own runs, and counters, filters, PID states and the estimator are created again at every reset. Table 3.4 lists the loops. Their action follows Section 1.3: +1 is reverse acting, as in the reserve loop, where more recharge raises the charge, and −1 is direct acting, as in the air loop, where more air cools the outlet node.

3.4.1 One PID Code for Every Loop

Every loop runs the framework's velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13), with the back-calculation time $T_t = \tau_I$ and the loop's own elapsed interval Δt_k in place of the control period. Every loop here is PI. A loop that does not own its output calls `track()`, which back-solves the integral for the value actually applied, and no PID is updated and tracked in one scan, so every transfer is bumpless. The plant signs were measured from the design point: 100 A more lowers the voltage from 0.7479 to 0.7447 V in 0.2 s, 0.4 mol/s more fuel lowers the utilization from 0.750 to 0.719 after 30 s, and 20 mol/s more air cools the outlet node by 1.06 K in 60 s.

3.4.2 Load-to-Current and the Reference Governor

The current request is the power asked divided by the measured voltage, and the governor admits the smallest of six ceilings, as shown in Equation (3.9). A low selector over limits on one actuator is also the structure of the jet engine fuel control of Volume 4, Chapter 3.

$$
\begin{aligned}
I_{req} &= \frac{P_{dem,m} + P_{soc} + (P_{AC,m} - P_{net,m})^{+} + P_{nl}}{\eta_{inv} V_m} \\
I_{ss} &= \frac{2 F U_{gov} \hat{\dot{n}}_{feed}}{N_s}, \qquad I_{tr} = \frac{2 F U_{gov}}{N_s} \min_{[t,\,t+H]} \hat{\dot{n}}_{arr} \\
I_{air} &= \frac{4 F x_{\mathrm{O}_2} \dot{n}_{air,m}}{N_s \lambda_{min}} \\
I_{inv} &= \min\left(I_{max},\ (P_{AC,max} + P_{nl})/(\eta_{inv} V_m)\right) \\
I_{adm} &= \min(I_{req}, I_{ss}, I_{tr}, I_V, I_{air}, I_{th}, I_{inv})
\end{aligned}
\tag{3.9}
$$

where $P_{AC,m} - P_{net,m}$ is the plant's own consumption from two meters, $U_{gov} = 0.80$ is the declared margin below the 0.85 ceiling, $\lambda_{min} = 2$, and $H = 5\,\mathrm{s}$. I_{ss} is the plan's steady ceiling on what the metered feed will make, 5329 A at design, and I_{tr} the transient anode constraint on the smallest

arrival the estimator predicts. The voltage guard I_V is a PI on the measured cell voltage with a 0.68 V setpoint, the backstop for a shortfall the estimator cannot see. The thermal limiter I_{th} is a PI on the excess

$$x_{th} = \max\big(T_{max,m} - T_{lim},\ (G_m - G_{lim})L,\ (\dot{T}_m - R_{max}) \cdot 60\,\mathrm{s}\big) \tag{3.10}$$

with $T_{lim} = 820\,^{\circ}\mathrm{C}$, $G_{lim} = 900\,\mathrm{K/m}$ and the heating rate taken from thermocouples filtered over 10 s and differenced over 60 s, so their noise reads as about 0.1 K/min against the 5 K/min limit. The guard and the limiter wait at $\bar{I}$, 300 A above the larger of the requested and the admitted current, tracking that ceiling while they have nothing to correct, and report no ceiling while they wait, so the gap never acts as a ramp limit and neither can wind up. A switch to Manual starts the operator's current at I_{adm}, where the fuel, air and inverter ceilings still cap it while the guard acts only in Auto, and with the governor off only I_{th} and I_{inv} remain, both protective in either mode. The inverter command moves toward I_{adm} at most 3000 A/s.

3.4.3 An Estimate of the Arriving Hydrogen

The governor and the trim need the anode's hydrogen, which no instrument reads, so an estimator keeps a declared model driven by the packet, as shown in Equation (3.11).

$$\begin{aligned}
\frac{\mathrm{d}\hat{c}_k}{\mathrm{d}t} &= (\hat{c}_{k-1} - \hat{c}_k)/\hat{\tau}_s, \qquad \hat{c}_0/\hat{\tau}_s = \hat{q}\dot{n}_{f,m} \\
\hat{n}_{arr} &= 4\hat{c}_3/\hat{\tau}_s, \quad \hat{n}_{feed} = 4\hat{q}\dot{n}_{f,m}, \quad \hat{\tau}_s = \hat{\tau}_{fp}/3 \\
\frac{\mathrm{d}\hat{n}_{\mathrm{H_2}}}{\mathrm{d}t} &= \hat{n}_{arr} - \frac{\hat{n}_{\mathrm{H_2}}}{\hat{N}_a}\hat{n}_{out} - \frac{N_s I_m}{2F} \\
\hat{N}_a &= \frac{p_{nom} V_a}{R T_{2,m}}, \qquad \hat{U}_f = \frac{N_s I_m}{2F\hat{n}_{arr}} \\
\hat{m}_a &= 1 - \frac{I_m/A}{k_{L,a}\hat{a}_{\mathrm{H_2}}}
\end{aligned} \tag{3.11}$$

where the inputs are the metered fuel flow $\dot{n}_{f,m}$ (0.5 s lag), the methane fraction $\hat{q}$ (the operator's entry, or the analyzer's report 90 to 150 s old), the measured current I_m and the outlet thermocouple $T_{2,m}$. The stages use the commissioning holdup $\hat{\tau}_{fp} = 15\,\mathrm{s}$ and are integrated by explicit Euler at 0.05 s substeps, the anode is one volume at its nominal pressure $p_{nom} = p_b + 4\,\mathrm{kPa}$, its outflow is the arriving gas per methane, $4 + (2.5/\hat{q} - 2)^+ + 1 + (1-\hat{q})/\hat{q}$ moles, and $\hat{a}_{\mathrm{H_2}} = (\hat{n}_{\mathrm{H_2}}/\hat{N}_a)(p_{nom}/p^{\circ})$. The horizon of I_{tr} moves the valve toward its present command through its 1 s lag and the stages forward in 0.25 s steps for 5 s, assuming the measured inputs stay constant and knowing no scheduled event. The utilization estimate is undefined below 0.5 mol/s of arriving hydrogen, and the margin estimate is set to $-\infty$ rather than evaluated when $\hat{a}_{\mathrm{H_2}} \le 0$, so the estimator takes no logarithm. At the priming call it starts at its own steady state for the first packet. Its bias is $\hat{q} - q$ and $\hat{\tau}_{fp} - \tau_{fp}$, the only two quantities it can get wrong, because its anode model takes the plant's declared volume, area, limiting current and steam ratio as commissioning values. Neither the estimator nor the governor reads simulator truth.

3.4.4 Fuel, Air, Preheat and Hot Water

The feedforward sends fuel for the current at the utilization setpoint as soon as it is asked for, and the trim corrects the utilization of that same current, as shown in Equation (3.12).

$$\begin{aligned}
I_{ff} &= \min(I_{req},\ I_{air},\ I_{th},\ I_{inv}) \\
\dot{n}_{ff} &= \frac{N_s I_{ff}}{8F\hat{q}U_{f,SP}}, \qquad \hat{U}_{ff} = \frac{N_s I_{ff}}{2F\hat{n}_{arr}} \\
\dot{n}_{f,cmd} &= \mathrm{clip}_{[2,12]}\big(\dot{n}_{ff} + \mathrm{PI}_u(U_{f,SP} - \hat{U}_{ff})\big)
\end{aligned} \tag{3.12}$$

The fuel ceilings are left out of I_{ff} on purpose, so the fuel starts moving at once and the admitted current climbs as the hydrogen arrives. A trim on the utilization of the admitted current would read the held current's low utilization as a fuel surplus, cut the feed and deadlock with the governor, which the first build measured as an emptied UPS in the governed challenge. Switching the feedforward moves its old contribution into the trim's integral, so the valve does not jump. The air loop moves the blower to hold the outlet node, with its ceiling following the metered flow so that a restricted blower cannot wind it up, the preheat loop holds the cathode inlet with the bypass, and the hot-water loop opens the heat-recovery bypass when the buffer runs hot.

3.4.5 Reserve Loop and Workload Cap

The reserve loop asks the fuel cell for extra power P_{soc} to restore the UPS toward 70 %. While the governor holds the current it does not integrate but tracks what the admitted current actually leaves above the demand, so it cannot wind up and the governor keeps priority. The workload cap, when armed, sheds the whole flexible share while the reported charge is below 40 % or the UPS gives more than 0.9 MW, and releases it once the charge is above 45 %, the UPS below 0.5 MW, the current unheld and 30 s have passed. Three models are kept apart. The plant is Section 3.3, nothing is fitted to data, and the internal model of the governor, the feedforward and the trim is the estimator of Equation (3.11) with its two commissioning values, while Section 3.4.6 gives the predictive controller's own.

3.4.6 Model Predictive Control of Current and Fuel

A structure switch offers a model predictive controller (MPC) in place of the governor's two fuel ceilings and the feedforward. It is never the default, and case 10 alone loads it. Every $T_s = 1$ s, or ten scans, a quadratic program plans the current I and the fuel command $\dot{n}_f$ over $N = 60$ steps and its first step is applied, the receding-horizon scheme of industrial predictive control [4] that Section 1.6.3 describes for every app. Seen from this controller, the governor of Equation (3.9) is a one-constraint relative that holds the utilization on the hydrogen it expects by cutting the current alone, while the feedforward sends the steady fuel for the current asked now. The plan holds that limit, the margin and the cell voltage over a minute with the valve as a second handle, so fuel can be sent ahead of a current that waits for it.

The model is written from Section 3.3 with the declared defaults of Table 3.3, never with the run's parameters, so a change made to the plant is model mismatch. Its states are the valve, the estimator's three stages and the anode's hydrogen, as shown in Equation (3.13).

$$\begin{aligned}
\frac{\mathrm{d}v}{\mathrm{d}t} &= \frac{\dot{n}_f - v}{\tau_{fv}}, \qquad \frac{\mathrm{d}c_1}{\mathrm{d}t} = \hat{q}v - \frac{c_1}{\hat{\tau}_s} \\
\frac{\mathrm{d}c_k}{\mathrm{d}t} &= \frac{c_{k-1} - c_k}{\hat{\tau}_s}, \qquad k = 2,3 \\
\frac{\mathrm{d}n}{\mathrm{d}t} &= \frac{4\hat{\theta}c_3}{\hat{\tau}_s} - \frac{n}{N_a}\dot{n}_{out} - \frac{N_s I}{2F} \\
\dot{n}_{out} &= \frac{c_3}{\hat{\tau}_s}\,\frac{2\hat{\theta}\hat{q} + (S/C)q_{nom} + 1}{\hat{q}} \\
N_a &= \frac{p_a V_a}{R T_{2,m}}, \qquad p_a = p_b + \dot{n}_{out}/k_a
\end{aligned} \tag{3.13}$$

Here v is the delivered flow, c_k are stage contents in moles of methane as in Equation (3.11) with the estimator's commissioning $\hat{\tau}_s = \hat{\tau}_{fp}/3$ and $\hat{q}$, n is the anode's hydrogen, $\hat{\theta}$ is the quality ratio below, $q_{nom} = 1$, and N_a is the anode's total moles at the measured outlet temperature. The cell voltage and the margin follow from Equation (3.3) with the declared cell constants and the measured temperatures and air held over the horizon. The free response is integrated by RK4 with both inputs held, and the

moves act through a linearization at each decision, sampled with an exact zero-order hold and joined to the free response by an affine term per step. Run open loop for 60 s after a step of fuel or current on the lined-out plant with the noise off, the model's RMS errors are at most 0.0021 in utilization, 0.0001 in margin and 0.0010 V in cell voltage, and its Jacobian matches one derived by hand to 2.8×10^{-7}.

The quality ratio $\hat{\theta}$, the hydrogen really arriving over the hydrogen the estimator credits, makes the plan offset free as an input disturbance [20], the disturbance state of Section 1.6.3. The offset-free predictive controller of Volume 9, Chapter 4 carries a fuel bias of the same kind. Together with an anode inventory $\hat{n}$ it is predicted every 0.1 s scan with $\hat{\theta}$ held and corrected by the measured cell voltage, as shown in Equation (3.14).

$$\begin{bmatrix} \hat{n} \\ \hat{\theta} \end{bmatrix} = \begin{bmatrix} \hat{n}^{-} \\ \hat{\theta}^{-} \end{bmatrix} + L \left[\frac{V_m}{N_s} - V\left(I_m, \hat{n}^{-}, \hat{\theta}^{-}\right) \right] \tag{3.14}$$

The minus marks the prediction, $V(\cdot)$ is the model's cell voltage, and L is the steady-state Kalman gain at the design point, with the entries 53.5536 mol/V for $\hat{n}$ and 1.881 78 V^{-1} for $\hat{\theta}$. It is computed for process variances of 1×10^{-3} mol^2 and 2×10^{-6} per scan, a random walk on $\hat{\theta}$, against 4×10^{-7} V^2 on the voltage. It matches an independent Riccati iteration to 4.4×10^{-11}, the error decays with time constants of 0.65 and 0.39 s, and lined out with the instruments' noise $\hat{\theta}$ has a standard deviation of 0.0007 within its bounds of 0.5 and $1.01/\hat{q}$. Leaner gas shows once it reaches the anode, within 27.6 s of the change in case 10 against the analyzer's 90 to 150 s, but a late processor or a drifted stack is read as a change of gas too (Section 3.7).

The plan minimizes Equation (3.15) subject to Equation (3.16).

$$\begin{aligned} J = {} & q_I \sum_{k=0}^{N-1} \left(\frac{I_k - I_{tgt}}{I_{ref}} \right)^2 + r_f \sum_{k=0}^{N-1} \left(\dot{n}_{f,k} - \dot{n}_{f,ref} \right)^2 \\ & + \sum_{j=1}^{12} \left[s_I \left(\frac{\Delta I_j}{I_{ref}} \right)^2 + s_f \, \Delta \dot{n}_{f,j}^2 \right] + \sum_{r=1}^{3} \left(\rho_r \sigma_r^2 + \ell_r \sigma_r \right) \end{aligned} \tag{3.15}$$

$$\begin{aligned} & 0 \le I_k \le I_{box}, \qquad 2 \le \dot{n}_{f,k} \le \dot{n}_{f,max} \\ & \frac{N_s I_k}{2F} \le \frac{4 U_{gov} \hat{\theta} c_{3,k}}{\hat{\tau}_s} + \sigma_1 \\ & m_{a,k+1} \ge m_{lim} - \sigma_2, \qquad V_{k+1} + b_V \ge V_g - \sigma_3 \end{aligned} \tag{3.16}$$

Here $I_{ref} = 1000$ A, $q_I = 100$, $s_I = 0.01$, $r_f = s_f = 0.05$ per $(\mathrm{mol/s})^2$, and Δ is the change of a block from the one before, the first from the commands applied now, over twelve blocks (four of one step, two of two, then 3, 4, 6 and 9 steps and two of 15), which with the three slacks σ_r gives 27 decision variables and 231 rows. The box $I_{box} = \min(I_{max}, I_{inv}, I_{air}, I_{th})$ and the valve's range are hard. The soft rows are the governor's own limit on the hydrogen arriving when each step's current applies, with $c_{3,k}$ the third stage at the step's start, the warning margin $m_{lim} = 0.25$, and the guard's setpoint V_g with the bias b_V of the measured voltage over the model's. Each row carries a single slack shared by all 60 steps, priced at $\rho_r = 10^4$, 10^3 and 10^2 and $\ell_r = 10^3$, 10^2 and 10^2 in mol/s of hydrogen, tenths of margin and hundredths of a volt. The target I_{tgt} is the request capped by the box, by the current that full fuel carries at the limit and by the current at which the steady voltage at full fuel meets V_g, and $\dot{n}_{f,ref}$ is the feedforward of Equation (3.12) corrected by $\hat{\theta}$. The weights are design values kept on every case.

Between plans the current follows the request inside the first step's allowance, the largest current that keeps the first step's rows on their bounds, because holding the planned move left the current below the noisy request for 495.2 s of case 3 instead of 10.2 s. The admitted current never passes a protective ceiling and never takes the utilization row above U_{gov}, and the voltage guard, the air ceiling, the thermal limiter and the inverter still low-select it as in Equation (3.9). A plan is accepted

only on an optimal status, the expected array shapes, every planned command inside its box and finite costs, slacks and predictions, and a test injects four faults that must each be rejected. A rejected plan returns the current and the fuel to the governor, feedforward and trim, crediting hydrogen scaled by $\hat{\theta}$. Forced rejections at twelve instants of case 10, with the noise off, admitted exactly the current the MPC would have admitted, and thirty in a row after the gas change kept the true utilization at 0.802 or below. The MPC reads only the measurement packet, the estimator's stages, the applied commands, the settings and the run's limits, and fed one run's packets over a different hidden plant it issued identical commands for 8999 scans. A plan cost 0.38 ms on average in one browser run (largest 4.5 ms), 3.8 % of one core at 100 times real time.

3.5 Experiments on One Plant

A script loads each case as the case menu does, applies any setting named in the text before time starts (so it is a starting condition), steps the shipped application at its 0.02 s step on the default seed, and exports the CSV record from which the figures are drawn, as Section 1.6.4 describes for every app. On the simulator's truth the power error is $e = P_{dem} + P_{soc} - P_{net}$, the power loop's own setpoint less its output. IAE is $\int |e|\,\mathrm{d}t$ over the stated window in kilowatt-hours, the largest error and the UPS peak are taken from 0.2 s after the event (inside the first scan a step's jump is caught at the end of one physics step, which depends on the step size), recovery is the last moment in the window with $|e| > 20\,\mathrm{kW}$, less the event time, with the rest of the window as its dwell, held time counts the seconds with the admitted current below the request, and travel is the total variation of the fuel valve state. The ten cases are listed in Table 3.5, where case 1 shows the three clocks of the Introduction before any loop is closed.

> **TRY IT IN THE APP** **1 · A Manual Current Step: What Moves, and When**
>
> Load scenario 1 on the Fuel and current view and press Run in Manual. After 30 s type 5250 into Stack current and watch for two minutes. The cell voltage falls from 0.7479 to 0.7397 V within 0.2 s, the anode fuel margin falls from 0.733 to 0.671 by 5 s, and the fuel stays at 8.63 mol/s.

3.5.1 Three Matched Comparisons

Each group of Table 3.6 starts from one initial state with the same seed, meters, limits and timed disturbance and changes only the settings named, and Figure 3.2 shows the trajectories.

Group A is nominal disturbance rejection on the identical step. Naive control asks for the current at once, the utilization reaches 0.993 and the margin 0.157, and for 20.0 s the inverter's foldback, not a controller, stops the current, while the UPS gives up to 0.465 MW. The twenty seconds of high current at low voltage heat the outlet node at up to 8.0 K/min on the filtered thermocouples (3.7 K/min governed) against the limiter's 5 K/min, so the limiter holds the current for 48.9 s and naive control recovers last, at 94.1 s. The feedforward alone moves the fuel at once but cannot make it arrive sooner, so the margin still falls to 0.158 and the foldback lasts 11.2 s. The governor alone keeps the margin at 0.614 and the utilization at 0.800, at the price of 30.5 s of held current and 1.16 kWh from the UPS, and with both the current waits only 13.5 s, the UPS gives 0.65 kWh, the IAE falls from 2.638 to 0.797 kWh and the margin stays at 0.623 or above, as shown in Figure 3.2a and b. Every run ends within 0.7 points of its 70 % start, so no hidden discharge of the store is rewarded, and across seeds 1 to 5 the naive margin is 0.157 and the governed 0.622 to 0.623.

Pause. Try Exercise 3.1 before reading on.

Group B shows a capacity shortfall that more control cannot remove. A 1.2 MW burst for ten minutes takes the demand to 4.83 MW, the fuel cell reaches its inverter's AC rating at 4.10 MW of net output, and the UPS, from 30 %, holds only $0.25 \times 0.10 \times 0.95 = 0.024\,\mathrm{MWh}$ above its floor. It reaches the floor at 201.3 s, and from then on the shortfall falls on every feeder, so 68.0 kWh of critical IT is lost among 109.2 kWh unserved. A 0.5 MW recharge limit still leaves 68.0 kWh of

Table 3.5 The Ten Cases With One Measured Result as Shipped and the Completion Rule of Each

#	Case and events	Outcome as shipped	Required to complete (kind)
1	Manual, current 4996 → 5250 A after 30 s	voltage 0.7479 → 0.7397 V in 0.2 s, margin 0.733 → 0.671 by 5 s, fuel unchanged at 8.63 mol/s, U_f 0.750 → 0.788	step to ≥ 5200 A after 30 s, run 120 s, confirm (experiment)
2	Governor and feedforward off, IT +0.35 MW at 60 s	U_f 0.993 at 3.2 s, margin 0.157 at 5.4 s, foldback 20.0 s, UPS up to 0.465 MW, thermal limiter 107–157 s	see the margin below 0.25, confirm (demonstration)
3	Case 2 with feedforward and governor	carries the site by itself 13.2 s after the step, UPS 0.64 kWh, margin ≥ 0.623, $U_f \leq 0.800$	see the fuel limit, then the site carried, confirm (experiment)
4	UPS at 35 %, recharge 0.1 MW, IT +0.9 MW 360–720 s	38.8 % at 360 s, floor at 705 s, 2.6 kWh unserved, 1.6 kWh critical	see the floor, recharge limit ≥ 0.3 MW, Reset, run 900 s, confirm (experiment)
5	Processor 55 s, estimate 15 s, IT +0.5 MW at 60, methane 85 % at 420 s	at 600 s $\hat{U}_f$ 0.685 against 0.805, guard holds 5974 of 6614 A, UPS 0.42 MW, 46.8 % at 900 s	see the guard bind, type 85 %, run 240 s, confirm (demonstration)
6	Blower to 50 % at 60, IT +0.4 MW at 420 s	outlet node 813.7 °C at 420 s, thermal override from 503 s, peak 825.0 °C, UPS up to 1.0 MW, 41.6 % at 900 s	see the override after 420 s, confirm (demonstration)
7	Heat demand 1.1 → 0.2 MW at 60 s	hot water peaks 80.97 °C, recovered heat 1.12 → 0.22 MW, nodes move under 0.1 K	bypass over 70 %, run to 600 s, confirm (exploration)
8	UPS at 30 %, IT +1.2 MW 60–660 s, cap off	fuel cell at most 4.10 MW of 4.83 MW, floor at 201 s, 68 kWh critical of 109 kWh unserved	see critical loss, arm the cap, Reset, run 720 s, confirm (demonstration)
9	Challenge, Table 3.7	utilization 1.048 and margin 0.158 at the first burst, fails	pass the shared verdict at 1200 s (challenge)
10	MPC, IT +0.5 MW 60–300 and 660–900 s, methane 0.90 at 420 s with the estimate at 1.00	peak U_f 0.807 and margin ≥ 0.568 with 5.58 kWh from the UPS, against 0.877, 53.2 s above 0.85 and 19.14 kWh under the governor	run 1200 s on the MPC, switch to the governor, run 1200 s, confirm (experiment)

critical IT unserved, and raising the governor's utilization limit to 0.85 leaves 67.5 kWh, because neither adds a watt of capacity. Only less load helps. With the cap armed the flexible quarter is shed at once by plan, the demand never passes 3.68 MW, nothing goes unserved, the UPS keeps 28.1 % or more and ends at 43.3 %, and 225.0 kWh of flexible work is deferred instead, as shown in Figure 3.2c and d. The capped run serves less and burns 1.896 MWh against 2.194 MWh, which is not a saving.

🕮 **Pause.** Try Exercise 3.3 before reading on.

Group C is a model mismatch. The processor holds 55 s of feed against the estimator's 15 s, and at 420 s the gas falls to 85 % methane while the estimator assumes 100 %. At 600 s the estimate reads a utilization of 0.685 while the truth is 0.805, and only the voltage guard holds the stack, at 0.680 V per cell and 5974 A against 6614 A requested, so the UPS gives about 0.4 MW to the end and falls to 46.8 %. With 85 % typed at 480 s the estimate meets the truth (0.7885 against 0.7902 at 600 s), the feedforward opens the valve to its 12 mol/s limit and the current rises to 6209 A, and the UPS still gives 0.279 MW because twelve moles a second of this gas cannot carry the site at a safe utilization. The analyzer, used from the start, reaches the same state a couple of minutes later (0.794 against 0.776 at 600 s), and the UPS energy in the window falls from 48.37 to 34.65 and 36.50 kWh. Correcting only the processor estimate to 55 s helps only before the gas change, where the true utilization peaks at 0.816 instead of 0.887.

3.5.2 The Recharge That Was Not the Cause

Case 4 was built to show that restoring the reserve fast makes a second load step. It does, but a different mechanism dominates. With the suggested recharge (0.4 MW, 6 MW per unit) the UPS is at 49.1 % when the burst arrives, falls no lower than 28.4 %, and after the burst ends at 720 s the current waits for fuel for 28.9 s of the next three minutes. Setting the recharge limit to zero at 720 s leaves

Table 3.6 Three Matched Comparisons on the Simulator's Truth. No Run Tripped or Exceeded the Thermal Limits

Group, window	Setting	IAE (kWh)	Max e (MW)	Rec. (s)	Low m_a	Peak U_f	Held (s)	Fold (s)	Travel (mol/s)	UPS (MW)	UPS (kWh)	Unserved, crit., curt. (kWh)	SOC (%)
A. Rejection, case 2,	naive (shipped)	2.638	0.466	94.1	0.157	0.993	48.9	20.0	4.79	0.465	2.42	0, 0, 0	70 → 69.3
0–400 s, IT +0.35 MW	feedforward only	1.115	0.457	17.0	0.158	0.992	0	11.2	13.17	0.457	0.91	0, 0, 0	70 → 70.0
at 60 s	governor only	1.306	0.241	27.6	0.614	0.800	30.5	0	2.02	0.241	1.16	0, 0, 0	70 → 69.7
	both (suggested)	0.797	0.237	13.2	0.623	0.800	13.5	0	5.96	0.237	0.65	0, 0, 0	70 → 69.8
B. Capacity, case 8,	as shipped	136.5	1.081	695.1	0.584	0.801	632.7	0	16.89	0.798	26.52	109.2, 68.0, 0	30 → 25.4
0–900 s, IT +1.2 MW	recharge 0.5 MW	136.8	1.079	725.1	0.584	0.801	634.9	0	19.30	0.792	26.58	109.1, 68.0, 0	30 → 27.7
60–660 s	$U_{gov} = 0.85$	136.3	1.058	695.1	0.440	0.851	607.6	0	21.73	0.804	27.23	108.3, 67.5, 0	30 → 25.6
	workload cap	12.7	0.956	647.6	0.656	0.800	302.8	0	31.56	0.956	11.51	0, 0, 225.0	30 → 43.3
C. Mismatch, case 5,	$\hat{q} = 1.0$ (shipped)	48.37	0.426	n.r.	0.529	0.813	480.0	0	2.27	0.426	48.37	0, 0, 0	67.2 → 46.8
420–900 s, methane	$\hat{q} = 0.85$ typed at 480 s	34.65	0.402	n.r.	0.532	0.814	480.0	0	1.63	0.402	34.65	0, 0, 0	67.2 → 52.6
1.00 → 0.85 at 420 s	analyzer from the start	36.50	0.427	n.r.	0.542	0.811	480.0	0	1.89	0.427	36.50	0, 0, 0	67.2 → 51.8

Rec. is the recovery time from the event with a 20 kW band, its dwell the rest of the window (A: 245.9, 323.0, 312.4, 326.8 s, B: 144.9, 114.9, 144.9, 192.4 s), and n.r. means not inside the band at the window's end. Held counts the seconds with the admitted current below the request (A naive: the thermal limiter, A governed: the arriving-fuel ceiling, B: mostly the AC rating, C: the voltage guard), Fold the inverter's foldback, Travel the fuel valve state's. The current's own travel is mostly the noise of the measured voltage the law divides by (33 A/s before the step, 13.2 kA of the 14.3 to 18.6 kA in A). Fuel on LHV over the window: 0.9217, 0.9205, 0.9117, 0.9060 MWh (A), 2.194, 2.213, 2.204, 1.896 MWh (B), 1.163, 1.205, 1.200 MWh (C), with the same heat delivered in every run of a group. At 0.01 s no entry moves by more than one unit beyond the last digit shown except B's $U_{gov} = 0.85$ maximum error (1.027).

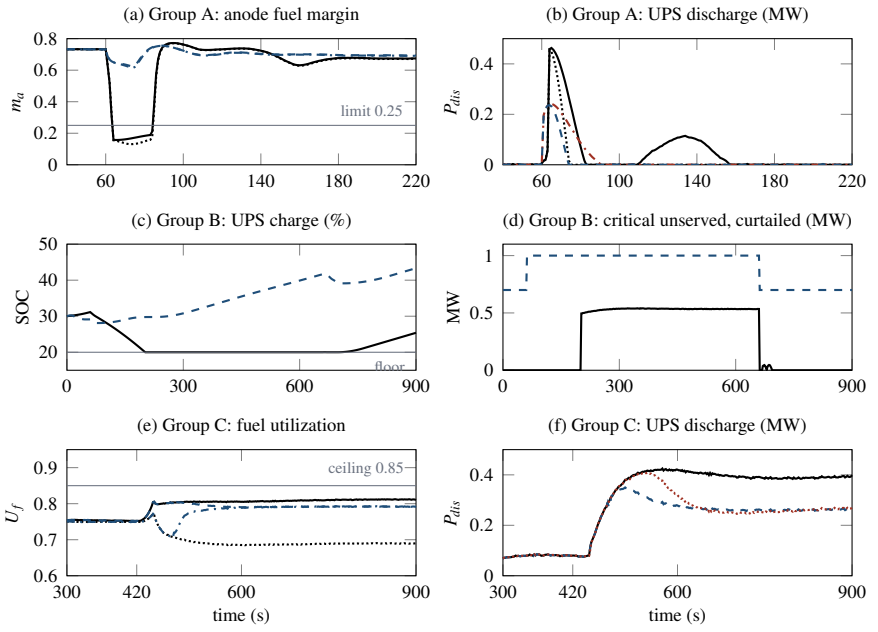


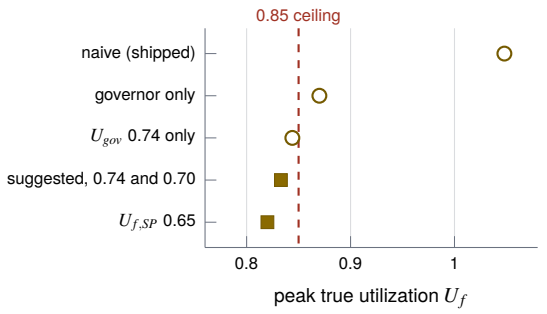
Figure 3.2 Matched comparisons A to C, drawn from the exported CSV records. (a) Case 2 after the 0.35 MW step at 60 s: true anode margin naive (solid) and with feedforward and governor (dashed), with the estimator's margin of each (dotted, dash-dotted), against the 0.25 limit. (b) UPS discharge of the four runs of group A: naive (solid), feedforward only (dotted), governor only (dash-dotted), both (dashed). (c) and (d) Case 8, as shipped (solid) and with the workload cap (dashed): true UPS charge against the 20 % floor, and the critical IT unserved as shipped with the flexible IT curtailed by the cap (dashed). (e) and (f) Case 5 after the gas changes at 420 s: true (solid) and estimated (dotted) utilization as shipped, the same pair with 85 % typed at 480 s (dashed, dash-dotted), and the UPS discharge as shipped (solid), typed (dashed) and with the analyzer (dotted).

the same 28.9 s, while setting the trim's gain to zero instead leaves 5.0 s. As the load falls, the trim's measurement, the utilization of the requested current over the estimated arriving hydrogen, reads the fuel still in the processor as a surplus and cuts the feed from 10.96 to 8.43 mol/s (9.27 mol/s with the trim's gain at zero), and the stack then waits for fuel to come back. The recharge's own cost appears with an aggressive setting (0.5 MW, 20 MW per unit, 60 s), which waits 41.6 s and lowers the margin to 0.575 against 0.621, and keeps the fuel cell at 3.93 MW against a 3.56 MW demand at 780 s. This is a property of any trim on a transport-delayed estimate, and the app's texts now state it.

3.5.3 The Combined Challenge

The challenge runs 1200 s from the design point with the UPS at 70 %. A 0.6 MW burst lasts four minutes from 120 s, the processor slows by 15 s over five minutes from 300 s while the estimator keeps 15 s, a 0.8 MW burst lasts four minutes from 600 s, and the heat demand falls to 0.2 MW at 840 s. One verdict function, shared by the scorecard, the completion rule and the tests, asks over the whole window for a true utilization below 0.85, an anode margin above 0.25, both nodes within 650 to 830 °C with the gradient within 1000 K/m, no trip, no critical IT unserved, the UPS never at its floor (lowest charge above 20.5 %) and at least 50 % at the end. Table 3.7 lists five policies run from the same initial state with the same meters and events, Figure 3.3 draws them as a chart, and Figure 3.4 compares three of them.

As shipped the first burst runs the anode short at once, the utilization reaches 1.048 at 122 s and the margin 0.158, and between the bursts the naive loop cycles on its own as the processor slows, its



U_f	low m_a	UPS (kWh)	SOC low, end	verdict
1.048	0.158	43.4	53.1, 60.3 %	fails
0.870	0.412	31.2	57.4, 61.9 %	fails
0.844	0.466	55.6	46.6, 46.6 %	fails
0.833	0.492	26.1	59.4, 63.2 %	passes
0.820	0.517	20.7	61.6, 65.2 %	passes

Figure 3.3 The five policies of Table 3.7 as a chart: the peak true utilization over the 1200 s (marks and first column), with the lowest anode margin, the energy the UPS gave, its lowest and final charge and the verdict in the other columns. The suggested policy is the governor on with $U_{gov} = 0.74$ and $U_{f,SP} = 0.70$, the last row the same with $U_{f,SP} = 0.65$. Open circles fail, filled squares pass. Notice that the governor alone still crosses the dashed ceiling, and that lowering its limit without the setpoint stays under the ceiling but takes the most energy from the UPS.

Table 3.7 The Challenge Over 1200 s Under Five Policies. No Run Tripped, Left Critical IT Unserved or Left the Node and Gradient Limits

Policy	Peak U_f	Low m_a	UPS (MW)	UPS (kWh)	SOC low, end	Fuel (MWh)	Verdict
naive (shipped)	1.048	0.158	0.847	43.4	53.1, 60.3	2.863	fails
governor only	0.870	0.412	0.759	31.2	57.4, 61.9	2.778	fails
U_{gov} 0.74 only	0.844	0.466	0.873	55.6	46.6, 46.6	2.636	fails
suggested	0.833	0.492	0.772	26.1	59.4, 63.2	2.745	passes
$U_{f,SP}$ 0.65	0.820	0.517	0.784	20.7	61.6, 65.2	2.846	passes

The suggested policy is the governor on with $U_{gov} = 0.74$ and $U_{f,SP} = 0.70$, the last row the same with $U_{f,SP} = 0.65$. Nodes 704.7 to 801.5 °C and gradients 738 to 859 K/m over all five. The UPS starts at 70 % in every run and the heat delivered is 0.277 MWh in each.

margin falling below 0.25 at 387, 436, 490 and 550 s with the cell at the inverter's floor before the late fuel arrives. The governor alone fixes the bursts but not the slower processor, since at the second burst the estimator still assumes 15 s and admits current a little before the hydrogen is there, and the true utilization reaches 0.870 while the estimate reads 0.799. Lowering only the governor's limit to 0.74 leaves the trim's 0.75 setpoint above it, so the governor holds the current below the request for 1121 s of the 1200 s and the UPS ends at 46.6 %.

> **✦ DID YOU KNOW? A Safer Limit Alone Drained the UPS**
>
> From the simulator: in the challenge, lowering only the governor's limit to 0.74 drew 55.6 kWh from the UPS, more than the 43.4 kWh of naive control, and left it at 46.6 %, so the run failed (Table 3.7). The limit and the trim's setpoint have to move together.

Lowering the limit to 0.74 and the setpoint to 0.70 together passes, with the utilization at 0.833 at worst, the margin above 0.49 and the UPS ending at 63.2 %. The suggested run uses less UPS energy than the governor alone, 26.1 against 31.2 kWh, and 20.0 of its 26.1 kWh come in the second burst, whose 4.37 MW demand at its peak exceeds what the inverter's 4.3 MW AC rating delivers after the plant's own consumption. Of that energy 17.7 kWh lie above the rating, against 20.1 with the governor alone and 24.5 as shipped, and no governor setting can move it to the fuel cell. Over 31 seeds, the default and seeds 0 to 29, the shipped policy passes 0 times and the suggested one 31 times, at both 0.02 s and 0.01 s, with its peak utilization 0.832 to 0.833. Its state at 611 s is shown in Figure 3.5.

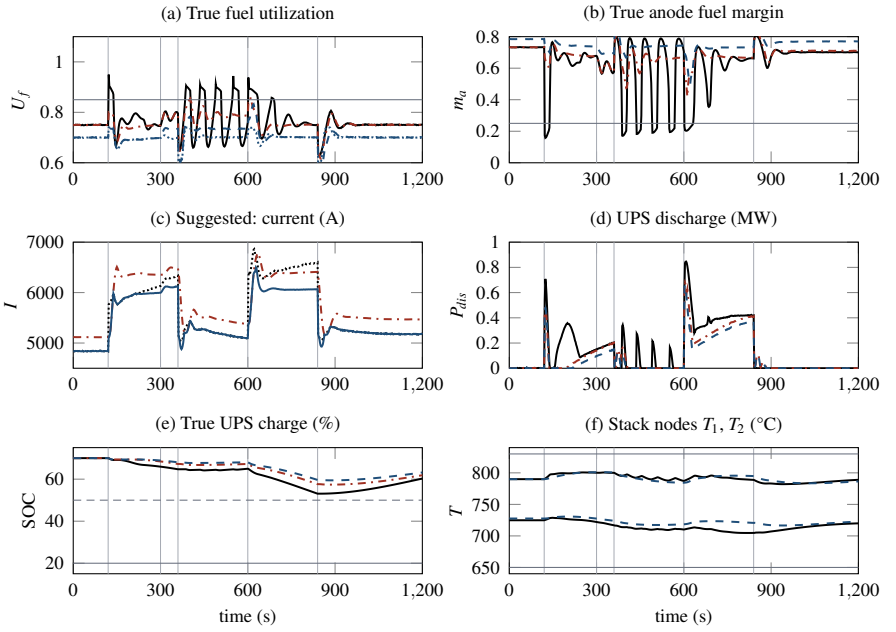


Figure 3.4 The challenge, from the exported records: naive as shipped (solid), the governor alone (dash-dotted) and the suggested policy (dashed), with the events at 120, 300, 360, 600 and 840 s marked (bursts begin, the processor slows, the first burst ends, the second begins, the second ends as the heat demand falls). (a) True utilization against the 0.85 ceiling, with the suggested run's estimate (dotted). (b) True anode margin against the 0.25 limit. (c) Suggested run: requested current (dotted), the arriving-fuel ceiling (dash-dotted) and the flowing current (solid). (d) UPS discharge. (e) True UPS charge against the 20 % floor and the 50 % end target. (f) Stack nodes T_1 and T_2 of the shipped (solid) and suggested (dashed) runs, inside the 650 to 830 °C limits.

> **TRY IT IN THE APP** **9 · Challenge: Workload Bursts, a Slower Fuel Processor and a Heat-Demand Drop**
>
> Load scenario 9 and run it as shipped at 100 times speed. At the first burst the utilization reaches 1.048 and the margin 0.158, and the verdict fails. Turn the Reference governor on, set the Governor utilization limit to 0.74 and the Fuel utilization setpoint to 0.70, then Reset and run the 1200 s. The utilization peaks at 0.833 and the verdict passes.

The margin to the ceiling is thin, 0.0175 of utilization. One factor at a time, with no retuning, the suggested setting fails for a true processor holdup above 20.75 s (the estimator keeping 15 s), a challenge slowdown above 20.05 s against its 15 s, or methane below 0.9626 with the estimate at 1.0, the direction in which an estimate errs dangerously, since gas leaner than assumed makes less hydrogen than the governor credits. A utilization setpoint of 0.65 under the same limit passes with 0.030 of margin for 3.7 % more fuel, 2.846 against 2.745 MWh. Two plan targets are read by the verdict as the plant allows. The plan asks for a positive anode fuel margin, but zero is unreachable, because the inverter folds back at 0.60 V per cell, naive control bottoms out at 0.157 and the stack trips at 0.10, so the verdict asks for 0.25, which lies beyond the 0.85 ceiling at steady state (0.49 there). A charge below 20 % is impossible by construction, so the verdict asks that the UPS never reach its floor.

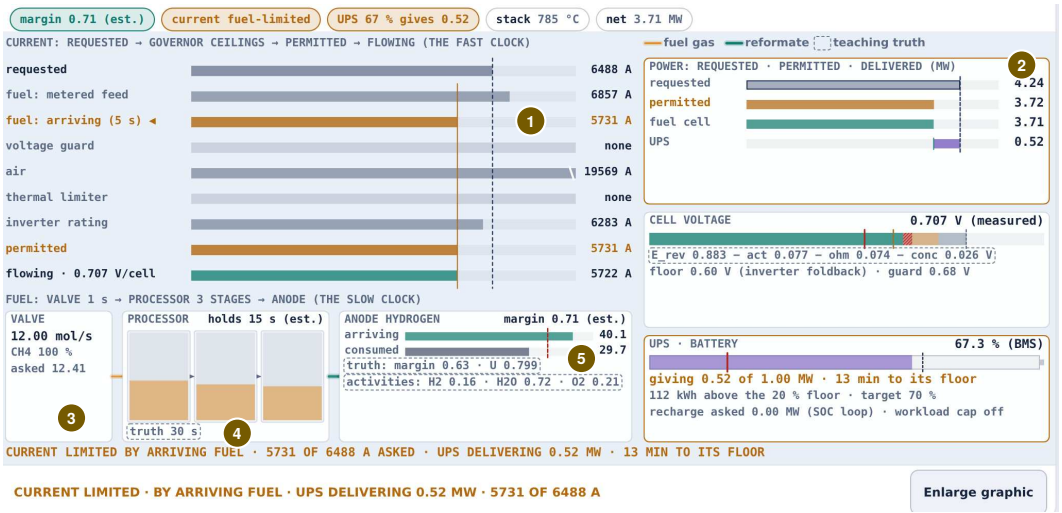


Figure 3.5 Suggested tuning in the challenge at $t = 611$ s on the Fuel and current view, eleven seconds into the second burst, cropped to the plant. (1) The current ladder shows 6488 A requested against ceilings of 6857 A on the metered feed, 5731 A on the arriving fuel, which binds, and 6283 A on the inverter's rating, with 5722 A flowing at 0.707 V per cell. (2) The fuel cell delivers 3.71 MW of the 4.24 MW asked and the UPS 0.52 MW from 67.3 %, 13 min from its floor. (3) The valve is at its 12.00 mol/s limit, (4) the processor's true holdup is 30 s against the estimator's 15 s, the estimator credits 40.1 mol/s of arriving hydrogen against 29.7 mol/s consumed, and (5) the anode's teaching truth reads a margin of 0.63 and a utilization of 0.799 against the estimate's 0.71 and 0.739.

3.5.4 The Predictive Controller Against the Governor

Table 3.8 compares the two structures on cases 3 to 6, 8 and 10 and the challenge over 1200 s, which holds every case's events, on the default seed, with identical instruments, noise and events and no preview for either structure. The governor keeps its shipped tuning ($U_{gov} = 0.80$, $U_{f,SP} = 0.75$), the challenge also runs with the governor on and with its suggested setting, and the MPC keeps case 10's settings everywhere. Ablations divide the credit between the MPC's information and its planning. They are the governor given $\hat{\theta}$ (a counterfactual built outside the application), the MPC with $\hat{\theta}$ off, and the MPC with its fuel capped at the larger of its steady reference and the fuel applied, so that none is sent ahead. The hindsight governor is the best of 96 fixed settings on each case (U_{gov} from 0.74 to 0.83, $U_{f,SP}$ from 0.65 to 0.75, two transient horizons, the feedforward on or off and two recharge loops) by the least UPS plus unserved energy with $U_f < 0.85$, $m_a > 0.25$ and no trip. Every MPC run solved 1200 plans with none rejected, and only the runs with $\hat{\theta}$ off paid a slack (11, 12 and 20 plans in cases 5, 10 and the challenge).

In case 10, two 0.5 MW bursts run from 60 to 300 and from 660 to 900 s, and at 420 s the gas falls to 90 % methane while the estimate stays at 100 %. At the first burst the plan holds the valve at 12 mol/s for 7.0 s, and the current reaches the request in 13.0 against 16.0 s under the governor, with 1.85 against 2.73 kWh from the UPS between 60 and 360 s. After the gas change $\hat{\theta}$ falls below 0.91 within 27.6 s and settles at 0.900, and the true utilization stays at 0.807 or below. The governor still credits pure methane, the thinner anode lowers the voltage, and the same power draws more current, as shown in Figure 3.6, the positive feedback of the Introduction opposed only by the voltage guard. The governor's utilization spends 53.2 s above 0.85, all after the gas change, and in the second burst the guard holds the current for 280.3 s while the UPS gives 14.04 against 3.58 kWh. Over the case the UPS gives 19.14 against 5.58 kWh on 2.7716 against 2.7010 MWh of fuel.

Table 3.8 The Governor Against the Predictive Controller Over 1200 s on the Default Seed, With the Rows That Separate the Estimate From the Plan. No Run Tripped, and Only the Naive Challenge Took the True Margin Below 0.25

Case	Structure	Peak U_f	Above 0.85 (s)	$\int (U_f - 0.85)^+ dt$ (s)	Low m_a	UPS (kWh)	Unserved (kWh)	Critical (kWh)	Fuel (MWh)
10 · methane 0.90	governor	0.877	53.2	0.384	0.421	19.14	0	0	2.7716
	MPC	0.807	0	0	0.568	5.58	0	0	2.7010
	MPC, $\hat{\theta}$ off	0.880	36.9	0.219	0.403	10.91	0	0	2.7684
	MPC, no fuel sent ahead	0.807	0	0	0.564	6.63	0	0	2.7014
	governor + $\hat{\theta}$	0.805	0	0	0.559	6.45	0	0	2.6996
	governor, $\hat{q} = 0.90$ at 425 s	0.800	0	0	0.568	6.32	0	0	2.7004
	governor, lead-lag 15/8 s	0.888	80.1	0.842	0.365	23.87	0	0	2.7588
	hindsight / MPC (0.74, 0.65)	0.822 / 0.744	0	0	0.513 / 0.664	2.52 / 2.66	0	0	2.7102 / 2.7546
5 · processor 55 s	governor	0.887	10.1	0.113	0.382	88.51	0	0	2.9158
	MPC	0.825	0	0	0.549	62.21	0	0	2.9886
	MPC, $\hat{\theta}$ off	0.870	11.8	0.103	0.430	62.90	0	0	2.9900
	governor + $\hat{\theta}$	0.825	0	0	0.538	64.64	0	0	2.9899
	hindsight / MPC (0.74, 0.65)	0.841 / 0.763	0	0	0.488 / 0.661	51.71 / 54.14	0	0	3.0230 / 2.9173
9 · challenge	naive (shipped)	1.048	146.5	5.651	0.158	43.41	0	0	2.8630
	governor on	0.870	7.6	0.047	0.412	31.22	0	0	2.7776
	suggested (0.74, 0.70)	0.833	0	0	0.492	26.10	0	0	2.7447
	MPC	0.814	0	0	0.559	28.16	0	0	2.7680
	MPC, $\hat{\theta}$ off	0.854	1.6	0.003	0.474	29.55	0	0	2.7835
	governor on + $\hat{\theta}$	0.815	0	0	0.537	30.04	0	0	2.7594
	hindsight / MPC (0.77, 0.65)	0.844 / 0.783	0	0	0.473 / 0.617	20.55 / 19.41	0	0	2.8470 / 2.8696
3 · step	governor	0.800	0	0	0.623	0.70	0	0	2.8323
	MPC	0.801	0	0	0.629	0.59	0	0	2.8344
	MPC, no fuel sent ahead	0.801	0	0	0.618	0.84	0	0	2.8364
	governor, lead-lag 15/2 s	0.800	0	0	0.626	0.54	0	0	2.8325
	hindsight / MPC (0.77, 0.65)	0.761 / 0.758	0	0	0.694 / 0.699	0.26 / 0.26	0	0	2.9103 / 2.9163
4 · reserve	governor	0.800	0	0	0.584	44.69	2.62	1.60	2.6794
	MPC	0.801	0	0	0.589	44.66	2.54	1.56	2.6815
	hindsight / MPC (0.77, 0.65)	0.770 / 0.774	0	0	0.646 / 0.646	40.02 / 39.96	0	0	2.7956 / 2.8014
6 · blower 50 %	governor	0.800	0	0	0.637	118.87	0.39	0.24	2.1875
	MPC	0.802	0	0	0.644	117.50	0	0	2.1956
	hindsight / MPC (0.74, 0.65)	0.740 / 0.744	0	0	0.724 / 0.726	53.08 / 52.80	0	0	2.5674 / 2.5697
8 · low reserve	governor	0.801	0	0	0.584	26.52	109.20	68.00	2.8844
	MPC	0.802	0	0	0.587	26.82	108.93	67.85	2.8959
	hindsight / MPC (0.77, 0.65)	0.768 / 0.770	0	0	0.649 / 0.644	28.25 / 28.25	97.08 / 97.27	60.41 / 60.54	3.0157 / 3.0169

Pairs in parentheses are $(U_{gov}, U_{f,SP})$, and a slash separates the hindsight governor from the MPC run at the same pair, a single entry applying to both. The hindsight governors keep the recharge loop as shipped and a transient horizon of 5 s, with the feedforward off in cases 3 and 6. The lead-lag is a 15 s lead over the best of 2, 3, 5 and 8 s lags, and $\hat{q}$ typed is the operator's estimate set 5 s after the gas change. Utilization and margin are the simulator's truth at every physics step, and the energies come from the plant's ledgers.

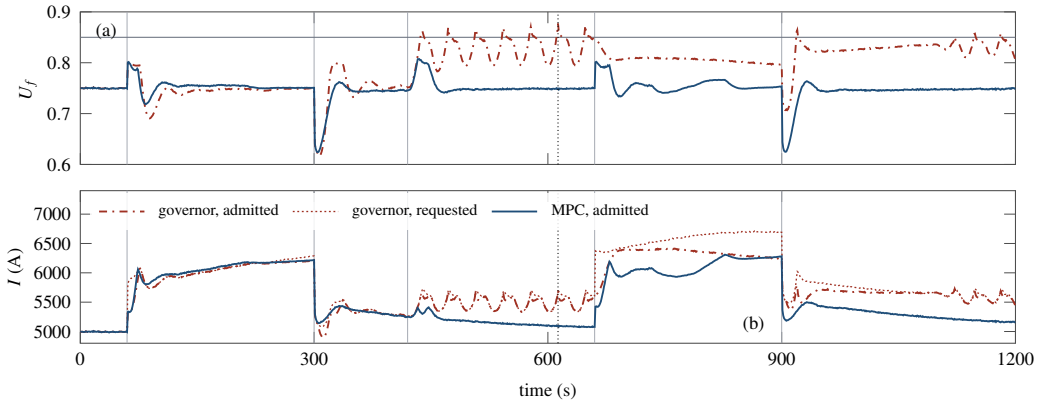


Figure 3.6 Case 10 under the governor (dash-dotted) and the predictive controller (solid), from the exported records. (a) True fuel utilization against the 0.85 ceiling. (b) Admitted current, with the governor's request dotted where its voltage guard holds the current below it. Thin lines mark the bursts at 60 and 660 s, their ends at 300 and 900 s and the gas change at 420 s. At the dotted line, 613 s, both serve the same 3.43 MW demand, the governor with 5689 A at 0.670 V per cell and a utilization of 0.877, the MPC with 5097 A at 0.736 V and 0.749.

The ablations assign the result. With $\hat{\theta}$ off the MPC reaches 0.880, while the governor given $\hat{\theta}$ holds 0.805, and 0.800 with its estimate typed to 90 % after the change. The utilization result belongs to the estimate read from the cell voltage, which the governor could read as well. The plan's own gain is the fuel sent ahead, 5.58 kWh against 6.63 kWh without it and 6.45 kWh for the governor given $\hat{\theta}$. A lead-lag feedforward also sends fuel ahead and beats the plan on case 3's clean step, 0.54 against 0.59 kWh, but without $\hat{\theta}$ it leaves case 10 at 0.888. In case 5, whose processor holds 55 s against 15 s, the governor peaks at 0.887 in the first burst and the MPC at 0.825, because $\hat{\theta}$ reads the late hydrogen as leaner gas, and in the challenge with the governor on the peaks are 0.870 and 0.814. Where the plant matches the estimate (cases 3, 4, 6 and 8) the two structures peak within 0.002 of each other and use UPS energy within 1.4 kWh. The rankings hold for seeds 11 to 13 and a 0.01 s step, the MPC peaking at 0.807 to 0.808, 0.825 and 0.814 against 0.877 to 0.878, 0.887 to 0.888 and 0.870 to 0.871 in cases 10, 5 and the challenge.

The challenge's rule is unchanged, and the MPC at its shipped settings meets it, with a peak of 0.814, a margin of at least 0.559 and 63.6 % left in the UPS. Two limits the verdict does not score are passed under it as under the suggested governor. The node heating rate on the limiter's filtered reading reaches 7.25 K/min (7.22) against a 5 K/min alarm, and the AC output reaches 4.3048 MW (4.3051 MW) against the 4.3 MW rating, which the inverter applies on the measured voltage and which holds the MPC's current for 96.3 and 133.3 s in the two bursts.

On stored energy the MPC loses to a tuned governor. The hindsight governor needs less UPS plus unserved energy than the MPC in every case, 0.26 against 0.59 (case 3), 40.02 against 47.20 (4), 51.71 against 62.21 (5), 53.08 against 117.50 (6), 125.33 against 135.75 (8), 2.52 against 5.58 (10) and 20.55 against 28.16 kWh (the challenge), in all seven with a utilization setpoint of 0.65 and 0.3 to 16.9 % more fuel. Held to no more fuel than the MPC, it still wins cases 3, 4, 5 and 8 and the challenge (0.27, 43.07, 61.19, 130.64 and 26.10 kWh) and loses cases 6 and 10 (118.41 and 34.56 kWh). Run at the hindsight limits, the MPC comes within 6 % of that governor, losing cases 5 and 10 (54.14 against 51.71 and 2.66 against 2.52 kWh) and winning the challenge (19.41 against 20.55 kWh), while its true peak stays lower wherever the plant departs from the estimate (0.763, 0.744 and 0.783 against 0.841, 0.822 and 0.844). The operating limits decide the stored energy more than the structure does.

3.5.5 Verification, Reference Data and Cost

These figures check the implementation and claim nothing about control performance. A script written apart from the application recomputes every boundary flux at each step's start and end and integrates it by the trapezoid rule over nine disturbed runs. The plant's energy closes to 4.5×10^{-7} of the fuel's LHV, the atoms to 3.2×10^{-7} of their throughput, the hot water and the data hall to 3.5×10^{-7} and 8.5×10^{-8} and the UPS's stored energy to 4×10^{-5}, all trapezoid error, and the bus residual is 0 W at every step with no negative term and no simultaneous charge and discharge. No step evaluated an undefined voltage, and the lowest hydrogen activity was 0.0746. With the stack tripped at 120 s the UPS at its 1 MW rating reaches its floor at 545.32 s against the analytic $120 + (E(120) - 0.2E_{cap})\eta_{dis}/1\,\mathrm{MW} = 545.31\,\mathrm{s}$. The design point recomputed by hand agrees with the application to 2.1×10^{-12}, halving the step changes no challenge verdict, and all 89 of the application's checks pass, 22 of them added with the predictive controller. One 1200 s challenge takes 0.46 s headless and 12.0 s at the browser's top speed of 100 times real time.

Two sensitivities bound the lessons. Case 6's thermal override comes at 503.1 s with 3 MJ/K per half, at 695.1 s with 6 MJ/K and not within 900 s with 15 MJ/K, and case 2's heating-rate hold of 48.9 s disappears from 6 MJ/K, while the three challenge verdicts are the same from 1.5 to 15 MJ/K. With $\rho_0 = 0.112\,\Omega\,\mathrm{cm}^2$, an aged stack, the voltage guard binds early and the governor alone passes, and with $0.08\,\Omega\,\mathrm{cm}^2$ the suggested setting keeps only 0.0066 of margin. Agreement with reference data is the weaker claim of the E°, power-density and efficiency checks above, the model has not met measurements of a stack, and no learner study has been made.

3.6 Technology in Practice

The governor and the loops of Section 3.4 act through an inverter, a fuel valve, a blower and two dampers, and they read meters, thermocouples and an analyzer. Figure 3.7 draws the current and fuel loops with that hardware in place of the schematic's symbols.

What Measures It

No instrument reads the anode's hydrogen, so the loops read what surrounds it: the stack current and voltage, the metered fuel, the gas quality and the stack temperatures. The temperatures come from thermocouples. The most common failure of a thermocouple is to fail open, or burn out, and an open input can pick up electrical noise that reads as a wildly varying temperature, so the instrument is set to read high or low on a burnout, whichever is the safer side. Where too much heat is the danger, a high reading drives the controller toward the safe state [21]. The chapter's stack thermocouples lag 5 s with 0.3 K of noise, and the gas quality comes from a sampled analyzer, whose principle Figure 3.8 shows.

> **TECHNOLOGY** **Sensor: A Process Gas Chromatograph for the Gas Quality**
>
> The methane fraction that the estimator needs comes from an online analyzer, and the most common one in continuous process analysis is the gas chromatograph. A sample valve injects a precise quantity of gas into a long, narrow column, a carrier gas washes it through, and each species leaves the column at its own time, so a simple detector sees a train of peaks [21]. The analyzer's computer can turn the concentrations into the heating value of natural gas. Time is the price. Cycle times of minutes are common even for online chromatographs, and enough dead time makes a feedback loop oscillate [21]. The chapter's analyzer draws a sample every 60 s through a 30 s line and analyzes it for 60 s, so each reading is 90 to 150 s old.

What Moves It

The inverter sets the current, the fuel valve sets the gas to the processor, and the blower and the two dampers set the air and the heat. Between the valve and the anode the processor is itself a chain of lags, and enough lags in series make a response that looks dead at first [21], which is why the chapter's three stages need 17.3 s to deliver 63 percent of a fuel step's hydrogen.

plant controller: EST-312, JIC-300 and JY-300 every 0.1 s, FIC-310 and TY-320 every 1 s. The inverter follows I_{adm} through a 50 ms lag and folds back at 0.60 V per cell.

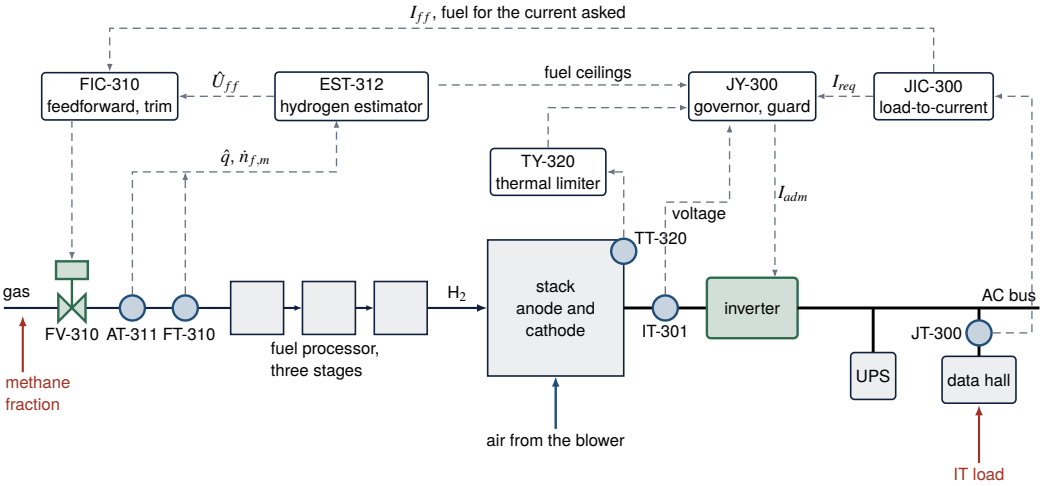
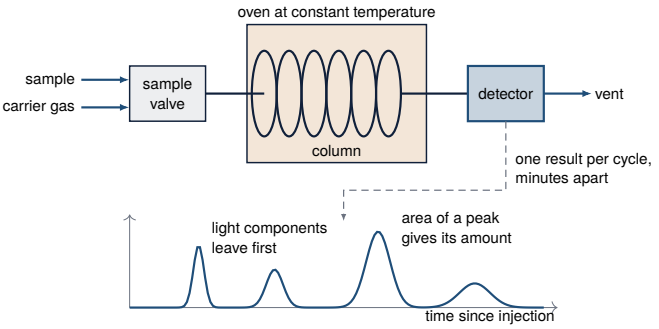


Figure 3.7 The current and fuel loops drawn with their hardware. The site demand meter JT-300 feeds the load-to-current law JIC-300, which asks for I_{req}. The estimator EST-312 turns the metered fuel FT-310 and the gas quality of the analyzer AT-311 into the hydrogen arriving, and the governor JY-300 admits the least of its fuel ceilings, the voltage guard on the stack meter IT-301, the thermal limiter TY-320 on the thermocouple TT-320 and the request. The inverter draws the admitted current. FIC-310 opens the fuel valve FV-310 for the current asked and trims it on the estimated utilization. The air, preheat and hot-water loops, the reserve loop and the workload cap are left out. Notice that the inverter acts on the fast clock while every fuel ceiling waits for the three processor stages.



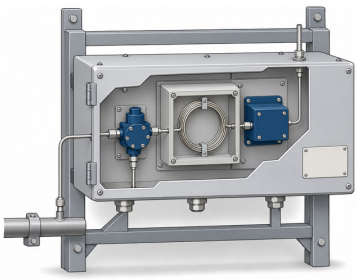


Figure 3.8 Principle of a process gas chromatograph, one kind of online analyzer for the gas quality. A valve injects a sample into carrier gas, the column holds back each species for its own time, and the detector sees them leave one by one. Notice that a result exists only at the end of each cycle, which is why the chapter's analyzer reports a methane fraction 90 to 150 s old.

TECHNOLOGY **Final Element: The Inverter That Sets the Stack Current**

A fuel cell stack is an unregulated voltage source whose voltage falls as its load rises [1]. A direct-current converter boosts that voltage and is the device that actually draws the current from the stack, and an inverter then makes alternating current by pulse width modulation at high frequency [1, 17]. The conversion is typically 92 to 96 % efficient. A specification for a small stand-alone unit asks the inverter to shut down when its input voltage dips below a minimum and names IEEE Standard 929 [17]. The chapter's inverter follows its command through a 50 ms lag and folds back at 0.60 V per cell.

Where the Controller Runs

A fuel cell system has three primary subsystems: the stack, the fuel processor and the power conditioner [1]. Programmable logic controllers (PLCs) manage local processes by feedback under a supervisory scheme, and loops at both levels run at cycle times from milliseconds to minutes [22]. The chapter's controller runs the estimator, the load-to-current law and the governor every 0.1 s, the other loops every 1 s and the reserve loop every 5 s. Its governor is a low selector, a function that passes the least of several signals, and a control selector block can report its selection back so that a de-selected controller stops integrating [21], the role of `track()` in Section 3.4.

> **IN PRACTICE** **Keeping a Fuel Cell Plant in Service**
>
> A fuel cell plant has few moving parts, but its catalysts and its stack wear out on a schedule. Major overhauls replace the shift catalyst every 3 to 5 years, the reformer catalyst every 5 years and the stack every 5 to 10 years, and filters are changed every 2,000 to 4,000 hours [1]. Early commercial units were watched by remote monitoring with an allowance for predictive maintenance. Fuel quality is watched as well. A fuel of lower energy content gives a less concentrated hydrogen-rich gas at the anode and some loss of stack performance, and sulfur and chlorides have to be removed before the processor's catalyst [1], the leaner gas of Section 3.5.1.

What Is Modernizing It

Three developments touch the loops of this chapter. Internally reforming cells, which join the endothermic reformer and the exothermic cell, are harder to design and control but are expected to take a significant market share [1]. The chapter's plant reforms outside the stack, and its Conclusion names internal reforming for a fuller model. Process chromatographs send several concentrations over a digital bus, and temperature programming of the column shortens their dead time [21]. Ethernet, Internet Protocol and wireless devices replacing older proprietary networks raise the likelihood of cybersecurity incidents in control systems [22]. The predictive controller of Section 3.4.6, which reads the gas quality from the cell voltage within 27.6 s where the analyzer needs 90 to 150 s, is a fourth development, and the chapter measures both its gains and its costs.

3.7 Limitations

All sizes and thresholds are synthetic teaching values set against the NETL handbook and the EPA catalog [1, 17] and fitted to no stack beyond the two conductivities behind E_{Ω}. The processor has no heat balance of its own, the burner has no temperature limit (after a trip it briefly burns the feed still in transit, to about 1340 °C on the truth), the two-node stack shows a gradient but no hot spot, and the light stack brings case 6's override and case 2's heating-rate hold earlier than a real stack would. The protection reads the true margin while the controllers do not, the estimator knows the plant's declared anode parameters, and the trim's reading of processor holdup as surplus remains in the shipped tuning. No forecast is present, the predictive controller's quadratic program is the only optimizer, no case outcome qualifies equipment, and none is a savings or reliability estimate.

> **OPEN CHALLENGE** **Stack Life Against Load Following**
>
> The high temperature that lets a solid oxide cell reform its own fuel also adds materials and mechanical design difficulties that shorten stack life and raise cost. The recommended stack rebuild every 5 to 10 years is expensive [1]. Mismatched thermal expansion and hard sealing between flat cells limit thermal cycling and stack life [17], and long heat-up and cool-down periods keep these cells out of many cyclic duties [1]. The chapter's gradient proxy is a coarse stress indicator, not a prediction of lifetime. Progress would be a controller that follows the load while it counts the wear each transient costs the stack.

The predictive controller's model takes the processor's holdup and the gas quality from the estimator's commissioning values and corrects them only through $\hat{\theta}$, which reads every unexplained

voltage deficit as fuel. In case 10, a stack whose area-specific resistance is 0.12 or 0.14 $\Omega\,\mathrm{cm}^2$ instead of 0.10 reads as leaner gas, on the safe side, with $\hat{\theta}$ up to 0.070 and 0.135 below the true ratio and peaks of 0.773 and 0.751 on 2.7690 and 2.9115 MWh of fuel against 2.7010. A stack at 0.085 or 0.07 $\Omega\,\mathrm{cm}^2$ reads as richer gas, with $\hat{\theta}$ up to 0.048 and 0.097 high, and the MPC peaks at 0.830 and 0.850, the second above the ceiling for 0.6 s, against 0.888 and 0.889 under the governor. A processor 20 or 40 s slower than commissioned is absorbed in part, with peaks of 0.818 and 0.825 against 0.878 and 0.888. The model also holds the stack temperatures and the air over the horizon, leaves out the UPS and the anode's pressure dynamics, and shares one slack per soft row across all 60 steps.

> ⚠ **OPEN CHALLENGE** **Storage Sized for the Worst Surge**
>
> Fuel delivery keeps fuel cells slow to follow a sudden rise in data center demand, and the gap becomes a temporary power shortfall. Both answers cost something. Capping the servers' power limits their performance and can break service agreements, while storage alone has to be oversized for the largest surge, which greatly raises the total cost of ownership. For a fuel cell it is the ramp rate that needs capping, not the peak, since the cell eventually meets the peak demand [8]. The chapter's workload cap acts on fixed thresholds of the UPS charge and output. Progress would be one supervisor that sets the store, the cap and the fuel limits together against the measured surges of a real hall.

3.8 Conclusion

This chapter has described a solid oxide fuel cell plant for an islanded data center whose current moves on the inverter's 50 ms, whose anode answers in 1.38 s and whose fuel arrives in 17.3 s. On an identical 0.35 MW step, naive load-to-current control takes the margin to 0.157 and leaves the inverter's foldback to stop the current, the feedforward alone changes nothing at the anode, and the governor with the feedforward keeps the margin at 0.623 and cuts the power IAE over 400 s from 2.638 to 0.797 kWh. When the estimator believes the gas is pure methane, only the voltage guard holds the stack, and when a burst exceeds the inverter's rating from a low reserve, only shedding load protects the critical IT. An added predictive controller keeps case 10's utilization at 0.807 where the governor reaches 0.877, through an estimate of gas quality read from the cell voltage that holds the governor at 0.805 as well, and its plan adds the fuel sent ahead, 5.58 against 6.63 kWh from the UPS without it. In the challenge, a governor limit of 0.74 with a setpoint of 0.70 passes on 31 of 31 seeds with 0.0175 of utilization to spare. Its stack is light, its processor has no heat balance of its own and its thresholds are synthetic. A fuller model would add internal reforming, a processor with its own heat balance, a distributed anode that can deplete locally, a heavier stack, and a predictive governor that knows the processor's holdup. The simulator shows that the current of a fuel cell has to wait for its fuel, and that the wait has to be paid for from storage.

KEY POINTS

- A fuel cell's current can change in a fraction of a second, but its fuel arrives many seconds later, so the current has to wait and storage carries the load.
- On an identical 0.35 MW step, the governor with the feedforward keeps the anode margin at 0.623 where naive control takes it to 0.157.
- When the estimator believes the gas is pure methane, only the voltage guard holds the stack, and when a burst exceeds the inverter's rating from a low reserve, only shedding load protects the critical IT.
- In the challenge, a governor limit of 0.74 with a setpoint of 0.70 passes on 31 of 31 seeds.

Where else this applies. A low selector of limits on one actuator also guards the fuel of the jet engine of Volume 4, Chapter 3, and the same site bus and UPS carry the gas engines of Chapter 2.

References

[1] U.S. Environmental Protection Agency, Combined Heat and Power Partnership, "Catalog of CHP technologies," U.S. Environmental Protection Agency, Washington, DC, Full Report, Sep. 2017. [Online]. Available: https://www.epa.gov/sites/default/files/2015-07/documents/catalog_of_chp_technologies.pdf

[2] E. G. Gilbert and K. T. Tan, "Linear systems with state and control constraints: the theory and application of maximal output admissible sets," *IEEE Transactions on Automatic Control*, vol. 36, no. 9, pp. 1008–1020, 1991.

[3] E. Garone, S. Di Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, 2017.

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

[5] B. J. Spivey and T. F. Edgar, "Dynamic modeling, simulation, and MIMO predictive control of a tubular solid oxide fuel cell," *Journal of Process Control*, vol. 22, no. 8, pp. 1502–1520, 2012.

[6] B. J. Spivey, J. D. Hedengren, and T. F. Edgar, "Constrained control and optimization of tubular solid oxide fuel cells for extending cell lifetime," in *2012 American Control Conference (ACC)*. IEEE, 2012, pp. 1356–1361.

[7] J. Sun and I. V. Kolmanovsky, "Load governor for fuel cell oxygen starvation protection: a robust nonlinear reference governor approach," *IEEE Transactions on Control Systems Technology*, vol. 13, no. 6, pp. 911–920, 2005.

[8] Y. Li, D. Wang, S. Ghose, J. Liu, S. Govindan, S. James, E. Peterson, J. Siegler, R. Ausavarungnirun, and O. Mutlu, "Techniques for efficiently handling power surges in fuel cell powered data centers: Modeling, analysis, results," arXiv preprint arXiv:1806.04506, 2018. [Online]. Available: https://arxiv.org/abs/1806.04506

[9] E. Achenbach, "Three-dimensional and time-dependent simulation of a planar solid oxide fuel cell stack," *Journal of Power Sources*, vol. 49, no. 1–3, pp. 333–348, 1994.

[10] J. Padullés, G. W. Ault, and J. R. McDonald, "An integrated SOFC plant dynamic model for power systems simulation," *Journal of Power Sources*, vol. 86, no. 1–2, pp. 495–500, 2000.

[11] Y. Zhu and K. Tomsovic, "Development of models for analyzing the load-following performance of microturbines and fuel cells," *Electric Power Systems Research*, vol. 62, no. 1, pp. 1–11, 2002.

[12] F. Mueller, J. Brouwer, F. Jabbari, and S. Samuelsen, "Dynamic simulation of an integrated solid oxide fuel cell system including current-based fuel flow control," *Journal of Fuel Cell Science and Technology*, vol. 3, no. 2, pp. 144–154, 2006.

[13] P. Aguiar, C. S. Adjiman, and N. P. Brandon, "Anode-supported intermediate-temperature direct internal reforming solid oxide fuel cell. II. Model-based dynamic performance and control," *Journal of Power Sources*, vol. 147, no. 1–2, pp. 136–147, 2005.

[14] Y. Komatsu, S. Kimijima, and J. S. Szmyd, "Numerical analysis on dynamic behavior of solid oxide fuel cell with power output control scheme," *Journal of Power Sources*, vol. 223, pp. 232–245, 2013.

[15] N. Q. Minh, "Ceramic fuel cells," *Journal of the American Ceramic Society*, vol. 76, no. 3, pp. 563–588, 1993.

[16] S. Singhal, "Advances in solid oxide fuel cell technology," *Solid State Ionics*, vol. 135, no. 1–4, pp. 305–313, 2000.

[17] U.S. Department of Energy, National Energy Technology Laboratory, "Fuel cell handbook," U.S. Department of Energy, Office of Fossil Energy, National Energy Technology Laboratory, Morgantown, WV, Report, Seventh Edition, Contract DE-AM26-99FT40575, Nov. 2004. [Online]. Available: https://netl.doe.gov/sites/default/files/netl-file/FCHandbook7.pdf

[18] D. Sarantaridis and A. Atkinson, "Redox cycling of Ni-based solid oxide fuel cell anodes: A review," *Fuel Cells*, vol. 7, no. 3, pp. 246–258, 2007.

[19] N. S. Gates, D. C. Hill, B. W. Billings, K. M. Powell, and J. D. Hedengren, "Benchmarks for grid energy management with Python Gekko," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 4868–4874.

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

3.9 Exercises

3.1 Concept check · 2 min

Under naive load-to-current control, with the governor and the feedforward off, the IT load steps up by 0.35 MW. The current rises at once, and the anode fuel margin then goes on falling for several seconds. Which mechanism drives the continued fall?

(a) The utilization trim cuts the fuel valve as the utilization rises, so less hydrogen is fed

(b) The stack cools under the higher current, so its resistance rises and the voltage falls

(c) The UPS takes the load, and the inverter lowers the current to leave room for its recharge

(d) The thinner hydrogen lowers the cell voltage, and the law then asks for still more current

3.2 Concept check · 1 min

Case 1 runs in Manual with the fuel valve and the air held. The stack current is stepped by about 5 %. In what order do the cell voltage, the anode fuel margin and the hydrogen arriving at the anode respond?

(a) The voltage within 0.2 s, the margin within seconds, the arriving hydrogen not at all

(b) The arriving hydrogen first, then the margin, and the voltage last of the three

(c) The margin within 0.2 s, the voltage within seconds, the arriving hydrogen after 17 s

(d) All three together, on the processor's clock of about 17 s to 63 percent

3.3 Concept check · 2 min

In case 8 the UPS starts at 30 % and a 1.2 MW burst takes the demand above what the inverter can deliver for ten minutes. Which change keeps all of the critical IT load served?

(a) Raising the recharge power limit to 0.5 MW, so the UPS refills faster

(b) Arming the workload cap, so the flexible quarter of the IT load is shed

(c) Raising the governor's utilization limit to 0.85, so more current is admitted

(d) Switching to the predictive controller, so fuel is sent ahead of the burst

3.4 Theory and calculation · 5 min

The stack has $N_s = 1000$ cells in series and carries $I = 4996\,\mathrm{A}$ at the design point. The valve delivers 8.63 mol/s of pure methane, and the processor makes four moles of hydrogen of each. Take $F = 96\,485\,\mathrm{C/mol}$.

(a) Find the hydrogen the current reacts and the fuel utilization. *Checkpoint:* 25.89 mol/s and 0.750.

(b) Find the steady ceiling I_{ss} of Equation (3.9) at $U_{gov} = 0.80$, and the utilization if the current were stepped to 5250 A with the fuel held, as in case 1. *Checkpoint:* 5329 A and 0.788.

(c) Find the fuel the feedforward of Equation (3.12) sends for 4996 A at $U_{f,SP} = 0.75$. *Checkpoint:* 8.63 mol/s.

3.5 Theory and calculation · 7 min

The valve follows its command through a 1 s first-order lag, and the processor holds the feed in three equal stages of $\tau_s = \tau_{fp}/3$ each. The hydrogen arriving at the anode therefore answers a valve step through four first-order lags in series.

(a) Find the mean delay of the chain at the design holdup $\tau_{fp} = 15\,\text{s}$. *Checkpoint:* 16 s.

(b) For the three stages alone, the share of a step delivered at time t is $1 - e^{-x}(1 + x + x^2/2)$ with $x = t/\tau_s$. Find the x at which it reaches 63.2 %, the time that gives at $\tau_s = 5\,\text{s}$, and an estimate for the whole chain made by adding the valve's 1 s. *Checkpoint:* $x \simeq 3.26$, 16.3 s, about 17.3 s.

(c) In case 5 the processor holds 55 s of feed while the estimator keeps 15 s. Estimate the mean delay and the 63 percent time the same way. *Checkpoint:* 56 s and about 61 s.

3.6 Theory and calculation · 5 min

The UPS stores 0.25 MWh, keeps a floor at 20 % of it, discharges at up to 1 MW and delivers with an efficiency of 0.95.

(a) From its 70 % start, find the energy it can deliver above the floor and how long that lasts at its rating. *Checkpoint:* 0.119 MWh and 7.1 min.

(b) In case 8 it starts at 30 %. Find the energy it can deliver, and how long that covers the gap between the 4.83 MW demand and the 4.10 MW the fuel cell delivers at its inverter's rating. *Checkpoint:* 0.0238 MWh and about 117 s.

(c) The run of case 8 reaches the floor at 201.3 s, 141 s after the burst begins at 60 s. Give one reason it outlasts the estimate of part (b). *Checkpoint:* The UPS charges before the burst, and the demand builds up.

3.7 App challenge · 8 min

Load scenario 2, Naive load-to-current: a workload step and the fuel margin, and change no setting. Run to 400 s at a speed of 20 with the default seed, export the CSV file, and read the scorecard tiles Lowest anode fuel margin and Fuel utilization peak.

(a) Report both tiles, and the column `metric_ups_kWh` in the row at 400 s. *Checkpoint:* 0.16, 0.993 and 2.42 kWh, within 0.02 kWh.

(b) Press Reveal suggested tuning, which turns the feedforward and the governor on, press Reset, run to 400 s again and report the same three values. *Checkpoint:* 0.62, 0.800 and 0.65 kWh, within 0.02 kWh.

(c) The naive run puts no limit of its own on the current. Explain why it nonetheless draws more energy from the UPS than the governed run. *Checkpoint:* The foldback and the thermal limiter hold its current.

3.8 App challenge · 9 min

Load scenario 5, A slow fuel processor and lower-quality gas: current limitation and UPS energy, and change no setting. The processor holds 55 s of feed while the estimator keeps 15 s, and the gas falls to 85 % methane at 420 s. Run to 900 s at a speed of 20 with the default seed and export the CSV file.

(a) In the row at 600 s, read `Uf_est` and `truth_Uf`, and `I_req_A` and `I_meas_A`. Divide the estimate by the truth. *Checkpoint:* 0.685 against 0.805, ratio 0.85; 6614 A against 5973 A, within 10 A.

(b) Name the ceiling that holds the current, from the current card or the column `limit`, and explain the ratio of part (a). *Checkpoint:* The voltage guard. The estimator credits pure methane.

(c) Press the reload button, run to a time between 470 and 490 s, type 85 % into Fuel quality estimate (methane) and run on to 900 s. Read the tile UPS lowest · now and compare it with the run as shipped. *Checkpoint:* 52.6 %, within 0.3 points, against 46.8 %.

3.9 App challenge · 9 min

Load scenario 9, the challenge, which runs 1200 s with two workload bursts, a slowing processor and a drop of the heat demand. Run each part to 1200 s at a speed of 100 with the default seed, pressing the reload button before each part, and read the tiles Fuel utilization peak, UPS lowest · now and Challenge target.

(a) Run as shipped, with the reference governor off. *Checkpoint:* 1.048, 53.1 and 60.3 %, not met.

(b) Switch Reference governor On and set Governor utilization limit to 0.74, leaving Fuel utilization setpoint at 0.75. *Checkpoint:* 0.844, 46.6 and 46.6 %, not met.

(c) Switch the governor On with the limit at 0.74 and the setpoint at 0.70. Explain why part (b) fails while this run passes. *Checkpoint:* 0.833, 59.4 and 63.2 %, met.

3.10 Application area · 3 min

The protection trips the stack at a true anode fuel margin of 0.10, and the challenge verdict asks for a margin above 0.25.

(a) Why must a plant keep a solid oxide anode from running short of hydrogen, apart from the power it loses? *Checkpoint:* The nickel anode oxidizes and expands.

(b) Which reading does the simulated trip use, and which readings must a real control system work from instead? *Checkpoint:* The true margin; the voltage, the current and the fuel flow.

3.11 Application area · 4 min

At the design point the plant takes 7.22 MW of methane on its lower heating value, delivers 3.435 MW net and returns 1.10 MW of heat as hot water.

(a) Compute the electric and the combined heat and power efficiencies. *Checkpoint:* 47.6 % and 62.8 %.

(b) The model's stack has 3 MJ/K per half, while a real stack of this area would hold roughly 15 to 50 MJ/K. Which lessons of the cases would come later on a real stack, and which would not change? *Checkpoint:* Case 6's override and case 2's heating-rate hold come later; the challenge verdicts do not change.

3.12 Review the agent · 5 min

In case 4, with the suggested recharge (0.4 MW, 6 MW per unit), the current waits for fuel for about half a minute after the burst ends at 720 s. An AI assistant explains: "The recharge loop causes this. As soon as the burst ends it asks the fuel cell for up to 0.4 MW to refill the UPS, and that new load outruns the fuel. Set Recharge power limit to zero at 720 s and the wait disappears." Find the flaw and say how you would check it in the app.