INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 8

Mission-Profile Load Following and Preview Control of an Aircraft Molten Salt Reactor

Pages 204–228

Edited by John D. Hedengren

How to cite this chapter. "Mission-Profile Load Following and Preview Control of an Aircraft Molten Salt Reactor," in *Control of Nuclear and Fusion Systems*, John D. Hedengren, Ed., vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 204–228.

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

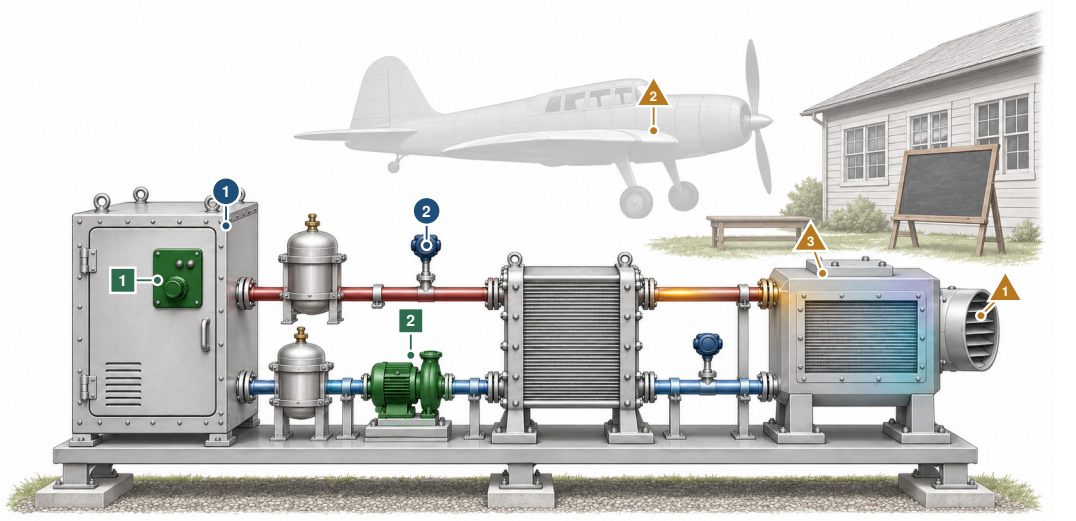
The app for this chapter
https://apmonitor.com/apps/aircraft-msr/

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

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

CHAPTER 8

Mission-Profile Load Following and Preview Control of an Aircraft Molten Salt Reactor



The controller moves 1 control rod at the reactor 2 fuel pump speed

The controller reads 1 reactor power, from a neutron detector 2 core outlet temperature

What disturbs it 1 engine throttle 2 altitude and flight speed 3 fouling of the air heater

Notice that the throttle acts at the far right, on the air heater, while the rod and the pump act at the far left, a whole loop of salt away. Illustration, not to scale.

IN PLAIN TERMS A Reactor That Must Read the Flight Plan

The main idea. This reactor heats the air of two jet engines in place of burning fuel. Its fuel is a hot liquid salt that a pump moves round a loop. A control rod and the pump together set the reactor's power and its temperature.

The challenge. At takeoff the engines ask for many times more heat within a minute. The reactor feels the change only after the cooled salt comes all the way round. The rod moves slowly, and the salt must never get hotter than its metal pipes allow.

Why it matters. A flight follows a plan, so a controller that reads the plan ahead can move before the demand does.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	The first page, the fact box of the Introduction, the history boxes of Section 8.2, the try-it and fact boxes of Section 8.6, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 8.6, Section 8.7, Section 8.8 and the open challenge

The control question A molten salt reactor that heats the air of two turbojets must follow a thermal demand that moves by a factor of fourteen in a takeoff, with a rate-limited rod and a fuel pump that both move power and outlet temperature. What is advance knowledge of the flight profile worth to it, and what does a coordinated predictive controller add on an event the profile does not contain?

Learning objectives

- Compute the demanded duty and the core temperature rise from the model, and explain why a circulating core is critical at positive reactivity.
- Calculate the relative gain of the rod and pump pairing from a measured gain matrix, and explain why it moves from 0.86 at takeoff to 3.58 at idle.
- Distinguish an actuator that has run out of process gain from one that has run out of travel.
- Measure what preview is worth to a feedforward and to a dynamic matrix controller on a scheduled takeoff.
- Explain why two independent loops break the outlet limit in a go-around that the predictive controller survives.

Prerequisites Sections 1.4.2 and 1.4.3 (point kinetics, circulating precursors), 1.6.2 (velocity-form PID), 1.6.3 (predictive control) and 1.6.4 (scenarios and data), and Chapter 7 for a circulating-fuel core with one pump loop.

The app https://apmonitor.com/apps/aircraft-msr/

See also Volume 6, Chapter 2 for pairing and decoupling with the relative gain array. Volume 9, Chapter 5 for feedforward from published orders, Volume 9, Chapter 4 for coordinated two-input two-output predictive control and Volume 9, Chapter 9 for a controlled variable held in a band.

8.1 Introduction

A turbojet raises the temperature of compressed air ahead of the turbine, and nothing in the thermodynamic cycle requires that the temperature rise come from combustion. Between 1946 and 1961 that observation funded the Aircraft Nuclear Propulsion program, whose reactor half produced the Aircraft Reactor Experiment (ARE) at Oak Ridge and whose engine half produced turbojets that ran on reactor heat at the Idaho test site [1, 2]. The program was cancelled for reasons of shielding mass and changing strategic priorities rather than for reasons of control, and the control problem it posed has not been retired. It has become more relevant, because every current proposal for a molten salt reactor serving a process, an isolated grid or a marine plant is a proposal for a reactor that follows a load it does not set [3, 4].

An aircraft is the extreme case of that duty. The thermal demand of the plant studied here moves by a factor of fourteen in forty-five seconds during a takeoff roll and by a factor of nine in twelve seconds on a go-around, against a control rod limited to 0.6 % of travel per second and a fuel salt whose containment alloy sets a hard upper temperature. No grid-connected reactor is asked for anything comparable. What distinguishes the problem from a large load disturbance, however, is not its size. It is that the disturbance is *published*. Throttle, altitude and Mach number twenty seconds in the future are known in the present, because a flight is flown to a plan, while the reactor's own measurement of the same disturbance arrives only after crossing an air heater metal mass, an intermediate salt loop, an exchanger metal mass, a 10 s fuel loop transit and a thermocouple. Figure 8.1 sets the two paths side by side.

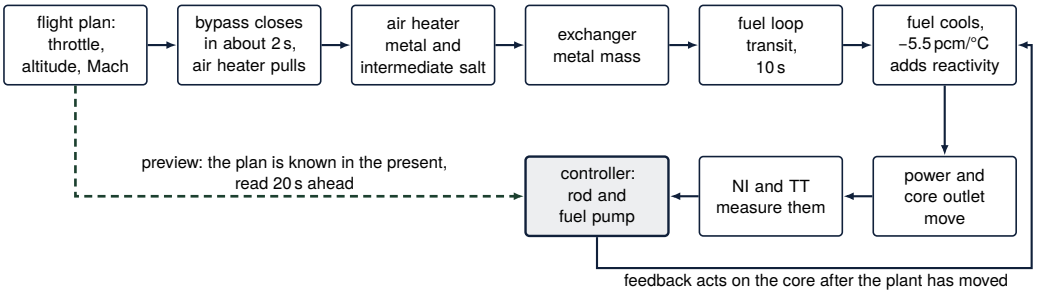


Figure 8.1 Two paths from the same throttle change. Solid arrows: the plant's own path, which crosses the air heater, the intermediate salt loop, the exchanger metal and the fuel loop transit before the fuel temperature coefficient moves the power, and is measured only after that. Dashed arrow: the preview path, by which the flight plan reaches the controller before the throttle moves. Notice that feedback can only start at the end of the solid path, while preview starts at its beginning.

> ✦ **DID YOU KNOW? Five Percent a Minute**
>
> Most current power reactor designs can maneuver between 50 and 100 percent of rated thermal power and ramp as fast as 5 percent of rated power per minute, although these maneuvers are not part of normal daily operation [5]. The takeoff of this chapter moves the demand by a factor of fourteen in forty-five seconds.

Simulation is the natural way to teach that distinction. A block diagram or a transfer function can express a feedforward path, but neither exposes what makes this plant hard, which is that the steady-state gain matrix moves by more than an order of magnitude across one flight and that the pairing chosen at the design point stops being the right pairing before the aircraft lands. The misconception the application is built to correct is that saturation means an undersized actuator. At 6.6 MW the core temperature rise of this plant is about 7 °C at design flow, so the fuel pump cannot move the outlet temperature 23 °C no matter how much travel it has. That is an actuator that has run out of process gain, it looks identical to an undersized valve on a trend, and it is fixed by a different remedy.

The class of problem is broader than nuclear propulsion. Any plant driven by a published trajectory has preview available and most control systems built for steady operation do not use it. Optimal preview tracking has an established theory for known reference models [6], and predictive control in receding-horizon form accepts a setpoint trajectory as naturally as a setpoint [7, 8]. On the plant of this chapter the value of preview can be measured against the same plant, the same actuators and the same limits as a controller that does not use it.

This chapter works with an interactive aircraft molten salt reactor and quantifies what advance knowledge of a flight profile is worth to a reactor that must follow it. Nothing in it is a claim that a nuclear aircraft is practical. Section 8.2 places the problem in its history. Section 8.3 develops a reduced-order model in which the aircraft enters as a scheduled disturbance through two thermodynamic quantities. Section 8.4 states the four control architectures as implemented, among them a two-input two-output dynamic matrix controller whose free response is built from a closed-form steady state rather than from an integrated model, which makes preview control affordable in a browser at about 20 µs of wall-clock time per execution on a desktop processor. Section 8.5 describes the app and its eleven scenarios. Section 8.6 compares the four architectures on one mission under actuator rate limits and a temperature constraint. Section 8.7 names the hardware such loops run on, Sections 8.8 and 8.9 give the limitations and the conclusion, and the exercises follow in Section 8.10.

8.2 Background

The ARE operated at Oak Ridge in November 1954, circulating a molten fluoride fuel salt and rejecting its heat through helium and water rather than through an engine [1, 9], and the History box

below tells how that circulation showed itself.

> ⧖ HISTORY **Watching the Fuel Go Round (1954)**
>
> The Aircraft Reactor Experiment at Oak Ridge pumped its molten fluoride fuel through the reactor and on to two heat exchangers that passed the heat to helium, and the helium gave it to water that was then discharged. At full flow the fuel took about 47 seconds to go round, about 8 of them in the core. The helium blower ran through a magnetic clutch, which gave smooth control at any power where the temperature coefficient was the controlling factor. The circulation showed on the fission chamber recorders as uranium was added from October 30, 1954. Each enriched slug mixed slowly into the carrier salt and made a pip every time it passed through the core, almost exactly 2 minutes apart at minimum pump speed and later 47 seconds apart, which gave the same flow as the flowmeter [10].

Two findings from that experiment shape the model of this chapter. The first is that circulating the fuel carries delayed neutron precursors out of the core before they decay, which modifies the inhour equation without destabilizing the reactor [11]. The second is that the experiment was operated largely by changing the heat removal and allowing the negative temperature coefficient to find the power, which is self-regulation used as an operating procedure rather than as a safety argument [9].

The engine half of the program ran turbojets on reactor heat in its direct-air-cycle Heat Transfer Reactor Experiments [2], as the second History box tells.

> ⧖ HISTORY **Turbojets on Reactor Heat (1956)**
>
> The engine half of the aircraft nuclear propulsion program built a series of Heat Transfer Reactor Experiments, in which compressor air was heated directly in the reactor. In January 1956 the first of them ran a modified turbojet on nuclear power alone. The engine was started on chemical fuel with its air passing through the cold reactor. The control rods were then withdrawn and the reactor brought to power by demanding more neutron flux, and as the reactor heated the air, the chemical fuel valve, sensing the temperature rise, gradually closed. Engine speed was held by the area of the exhaust nozzle. The third experiment ran two modified turbojets and reached full nuclear power in 1959 [12]. The plant of this chapter heats the air through an intermediate salt loop instead.

The arrangement in that work, in which compressed air is heated by the reactor and expanded through the turbine, is the arrangement modeled here, with an intermediate salt loop interposed so that the fuel salt never enters the engine.

Quantitative dynamic modeling of circulating-fuel reactors matured on the Molten Salt Reactor Experiment that followed. The lumped structure used in this chapter, with modified point kinetics, a small number of well-mixed thermal nodes and a pure transport delay around the external loop, is the structure of the MSRE theoretical dynamics analysis [13] and of the stability analysis that preceded it [14], and the operating experience against which such models were compared is documented [15]. The reduction of precursor transport to a single loop transit time with an exponential survival factor remains the standard point-kinetics treatment, and modern work has quantified how much it loses against transport and computational fluid dynamics solutions of the same problem [16, 17] and against Monte Carlo evaluations of the effective delayed fraction [18].

Control of molten salt reactors under load following is an active subject. Core outlet fuel temperature has been treated as the primary controlled variable with fuel circulation as the manipulated variable [3], and load following transients have been examined for natural circulation designs where the flow is not manipulable [4]. What the aircraft case adds to that literature is a second controlled variable that must be held simultaneously, a disturbance an order of magnitude faster, and a disturbance trajectory that is known in advance.

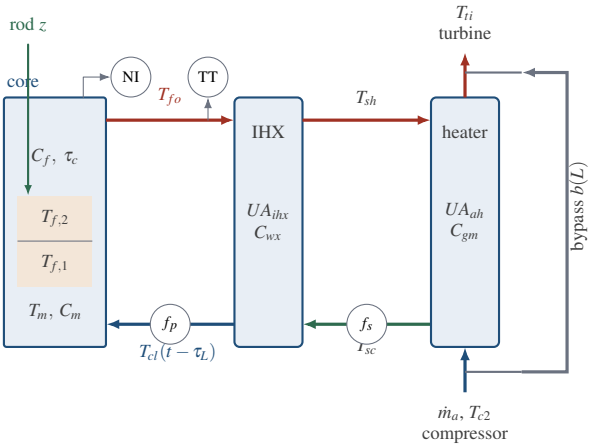


Figure 8.2 Aircraft molten salt reactor propulsion arrangement. Fuel salt circulates through two core nodes at temperatures $T_{f,1}$ and $T_{f,2}$ against a moderator and structure node T_m, leaves at T_{fo}, is cooled in the intermediate heat exchanger, and returns to the core inlet one external loop transit τ_L later. The intermediate salt carries the duty to an air heater, where compressed air at $\dot{m}_a$ and T_{c2} is heated and remixed with the fraction $b(L)$ routed around the heater. The control rod position z and the fuel pump speed f_p are the manipulated variables, and the neutron instrument NI and the outlet thermocouple TT supply the two measurements.

8.3 Aircraft Reactor Model

The plant is a notional molten salt reactor rated at 125 MW thermal driving two turbojets. Fission heats a circulating fluoride fuel salt directly. The salt is pumped through the core, gives its heat to a clean intermediate salt in an intermediate heat exchanger (IHX), and returns. The intermediate salt carries the heat to an air heater in the engine, where it passes to compressed air. Heated air and bypassed air remix ahead of the turbine. Figure 8.2 shows the arrangement and labels every symbol that appears in the equations below.

Table 8.1 classifies the variables the control problem is written in.

8.3.1 Circulating-Fuel Point Kinetics

The fuel is a liquid and it moves, so a delayed neutron precursor born in the core is carried out of it before decaying, as Section 1.4.3 explains and Chapter 7 develops for a ground plant. With a core transit time τ_c and an external loop transit τ_L, the six precursor groups obey

$$\frac{dW_i}{dt} = \beta_i n - \lambda_i W_i - \frac{W_i}{\tau_c} + \frac{W_i(t-\tau_L)\, e^{-\lambda_i \tau_L}}{\tau_c}, \tag{8.1}$$

in which $W_i = \Lambda c_i$ scales the concentration by the neutron generation time, n is power as a fraction of rated, and the last two terms are loss to the loop and delayed return reduced by decay along the way. The prompt neutron mode is removed algebraically by the prompt jump approximation [19] of Equation (1.1), with six groups, which gives

$$n = \frac{\sum_i \lambda_i W_i}{\beta - \rho}. \tag{8.2}$$

The generation time cancels from the model entirely, which is what allows a 0.05 s integration step to carry a thirty-minute flight. The β in the denominator of (8.2) is the *static* delayed fraction,

Table 8.1 Controlled, manipulated and disturbance variables

Symbol	Meaning	Units	Range
Controlled variables			
P	reactor thermal power	% of rated	0 to 150
T_{fo}	core fuel outlet temperature	°C	600 to 900
Manipulated variables			
z	control rod position	% withdrawn	0 to 100
f_p	fuel pump speed	% of design	40 to 125
Disturbance variables			
L	propulsion demand (throttle)	%	10 to 105
h	altitude	ft	0 to 40 000
M	Mach number	—	0 to 0.80
Principal states			
W_i	scaled precursors, $i = 1$ to 6	MW	—
$T_{f,1}, T_{f,2}$	core fuel node means	°C	—
T_m	moderator and structure	°C	—
T_{cl}	fuel cold leg plenum	°C	—
T_{wx}, T_{gm}	exchanger metal masses	°C	—
T_{sh}, T_{sc}	intermediate hot and cold	°C	—
T_{ti}	turbine inlet gas	°C	—
$\dot{m}_a$	engine air mass flow	kg/s	30 to 245
b	air heater bypass fraction	—	0 to 0.92

650 pcm for the six-group thermal uranium-235 data used here, because it comes from production and production does not depend on where the precursor ends up.

Setting dW_i/dt to zero in (8.1) with $W_i(t-\tau_L) = W_i$ gives the effective delayed fraction that the fuel pump controls, which is Equation (1.5) of Chapter 1 with the pump speed written in,

$$\beta_{\text{eff}} = \sum_i \beta_i \frac{\lambda_i \tau_c}{\lambda_i \tau_c + 1 - e^{-\lambda_i \tau_L}}, \qquad \tau_c = \frac{\tau_{c,0}}{f_p}, \quad \tau_L = \frac{\tau_{L,0}}{f_p}. \tag{8.3}$$

At the design flow of 6 s in the core and 10 s outside it, (8.3) gives 390 pcm against a static 650 pcm, so 40.1 % of the delayed neutrons decay outside the core and the rod must hold 260 pcm of permanent excess reactivity to pay for circulation. Applying the same expression to the published ARE transit times of 8.3 s and 47 s gives 378 pcm and a loss of 41.8 %. Because both transit times scale inversely with pump speed, β_{eff} runs from 464 pcm at 40 % flow to 373 pcm at 125 %, and the fuel pump is therefore an actuator that moves the plant's own stability margin every time the temperature loop acts.

8.3.2 Reactivity and Criticality

Reactivity is measured from a design reference state at rated power,

$$\rho = \rho_0 + W(z) - W(z_0) + \alpha_f\left(\bar{T}_f - T_{f,0}\right) + \alpha_m(T_m - T_{m,0}), \tag{8.4}$$

with $\bar{T}_f = (T_{f,1} + T_{f,2})/2$ and $W(z)$ the usual S-curve integral rod worth. A circulating core is critical at *positive* reactivity, because it has to replace the delayed neutrons it is pumping away (Section 1.4.3), so $\rho_0 = \beta - \beta_{\text{eff}}$ at design flow and the steady-state condition is $\rho_{ss} = \beta - \beta_{\text{eff}}$ at whatever flow is in

force. Equation (8.4) is the reactivity ledger of Equation (1.3) without a xenon term, and it pins a weighted core temperature at every power, which is why the mean fuel temperature barely moves across the mission and what a load change actually moves is the core temperature *rise*.

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

8.3.3 Core and Loop Energy Balances

The core is divided into two fuel nodes in series, which is the coarsest division that separates the mean fuel temperature that criticality pins from the outlet temperature that the containment alloy sees,

$$\begin{aligned}\tfrac{C_f}{2}\dot{T}_{f,1} &= \tfrac{1}{2}(1-f_M)Q - 2W_f\left(T_{f,1}-T_{fi}\right)\\ &\quad - \tfrac{UA_{fm}}{2}\left(T_{f,1}-T_m\right), \qquad (8.5)\\ \tfrac{C_f}{2}\dot{T}_{f,2} &= \tfrac{1}{2}(1-f_M)Q - 2W_f\left(T_{f,2}-T_{f,1}^{out}\right)\\ &\quad - \tfrac{UA_{fm}}{2}\left(T_{f,2}-T_m\right), \qquad (8.6)\\ C_m\dot{T}_m &= f_M Q + \tfrac{UA_{fm}}{2}\left(T_{f,1}+T_{f,2}-2T_m\right), \qquad (8.7)\end{aligned}$$

with $Q = P_0 n$ and $W_f = C_f/\tau_c$ the flow capacity rate. Each flowing node carries a linear temperature profile, so its mean is the average of its inlet and outlet and the outlet follows algebraically as $T^{out} = 2T - T^{in}$. The core inlet is the cold leg plenum temperature delayed by one external loop transit, $T_{fi}(t) = T_{cl}(t-\tau_L)$, and the same τ_L that carries the heat carries the precursors in (8.1), because it is the same salt in the same pipe.

8.3.4 Heat Exchange

Both exchangers use the counterflow effectiveness relation rather than a stream exchanging with an isothermal wall. The distinction is not cosmetic on this plant, because the intermediate salt leaves the IHX at 759 °C while the fuel leaves it at 659 °C, and a lumped wall model cannot represent a cold outlet hotter than the hot outlet. With $C_{\min}$ and $C_{\max}$ the smaller and larger capacity rates and $C_r = C_{\min}/C_{\max}$,

$$\varepsilon = \frac{1-e^{-N(1-C_r)}}{1-C_r e^{-N(1-C_r)}}, \quad N = \frac{UA}{C_{\min}}, \quad Q = \varepsilon C_{\min}\Delta T_{in}. \qquad (8.8)$$

Both conductances follow the usual turbulent 0.8 power of mass flow, split as 0.4 per stream, so $UA = UA_0(W_1/W_{1,0})^{0.4}(W_2/W_{2,0})^{0.4}$. The metal mass of each exchanger is carried as a separate capacity that absorbs $UA_m(\bar{T} - T_{metal})$ during a transient and contributes nothing at steady state, which gives the exchanger thermal inertia without also giving it a thermal resistance it does not have.

8.3.5 Aircraft and Engine

The aircraft enters the reactor model through exactly two thermodynamic quantities. International standard atmosphere and ram compression give the total conditions at the compressor face, and an ideal compressor with fixed pressure ratio π_c and isentropic efficiency η_c gives the discharge temperature,

$$T_{t1} = T_{amb}(h)\left(1+0.2M^2\right), \quad T_{c2} = T_{t1}\left[1+\frac{\pi_c^{0.2857}-1}{\eta_c}\right]. \qquad (8.9)$$

Engine air flow follows the corrected flow group and a compressor map written as a power of spool speed,

$$\dot{m}_a = \dot{m}_{max}\, \nu(L)^{2.2}\, \frac{\delta}{\sqrt{\theta}}, \qquad \nu(L) = 0.35 + 0.0065\, L, \tag{8.10}$$

with δ and θ the total pressure and temperature ratios to sea level static. The group $\delta/\sqrt{\theta}$ is 1.04 at sea level and 0.38 at 35 000 ft, while T_{c2} is 248 °C at sea level and 165 °C in the cruise. The two act against each other, so takeoff and high-altitude cruise are physically different disturbances at the same throttle.

Propulsion demand is expressed as the turbine inlet temperature the engine schedules, because thrust at a given air flow is set by how hot the gas is when it reaches the turbine. The demanded reactor duty follows,

$$Q_{dem} = \dot{m}_a c_{p,a} \left(T_{ti}^{SP}(L) - T_{c2} \right), \qquad b = 1 - \left(\tfrac{L}{100} \right)^{\kappa}. \tag{8.11}$$

The second half of (8.11) is the air heater bypass, the fraction of compressor discharge routed around the heater and remixed downstream. It is scheduled open loop on throttle, so the reactor still has to make the heat, and it exists because without a cold air path this plant has no turndown at all. A fixed-geometry air heater with the core held at its design temperature pulls 26 MW at ground idle, 21 % of rated against a demand of 7 %, and the only way to shed the difference would be to cool the salt to 409 °C, which freezes it.

8.3.6 Steady State in Closed Form

Every link in the chain is a fixed thermal resistance at fixed flows, so the whole steady state collapses to two affine relations,

$$\bar{T}_f = T_{c2} + A\, Q, \qquad T_m = T_{c2} + \left(A + \tfrac{f_M}{U A_{fm}} \right) Q, \tag{8.12}$$

with A assembled from the reciprocal conductances of the air heater, the intermediate loop, the IHX and the core. Substituting (8.12) into the criticality condition $\rho_{ss} = \beta - \beta_{\text{eff}}$ leaves one linear equation in Q, so the settled power and every settled temperature at a given rod position, pump speed and flight condition are a closed form and never have to be integrated. That property is what makes both the model feedforward of Section 8.4.4 and the predictive controller's free response of Section 8.4.5 affordable in a browser.

8.3.7 Parameters

Table 8.2 lists the parameters. The column marked *ARE* gives the corresponding published value where one exists, so that a reader can see exactly where the model stops being anchored to an experiment and becomes notional.

8.3.8 Assumptions

The model is a control model. Point kinetics with the prompt jump approximation is silent about anything approaching prompt critical, and the application annunciates $\rho \geq \beta$ rather than integrating through it. The precursors are assumed uniformly mixed in the core, so a precursor about to leave is treated as worth the same as one just born near the flux peak, which transport solutions of the same problem show is an approximation rather than an identity [16, 17]. A pump speed change alters both transit times instantly, so salt already in the loop is treated as having always traveled at the new speed. Reactivity coefficients are constants, there is no burnup, and xenon is not modeled.

On the thermal side there are two fuel nodes, one moderator node and one lumped metal mass per exchanger, so no spatial temperature or flux distribution exists and hot channel factors cannot appear. The whole external transport is applied to the return leg rather than split between a hot leg and a cold leg. The 850 °C limit, the 900 °C trip and the 500 °C freeze alarm are simulation constraints chosen to make the control problem bite rather than the qualification limits of any particular alloy or salt.

Table 8.2 Model parameters, with the published Aircraft Reactor Experiment value where one exists

Symbol	Meaning	ARE	Model	Rationale for the difference
P_0	rated thermal power, MW	2.5	125	two turbojets rather than a helium and water sink; results are reported in percent
$\tau_{c,0}$	core transit at design flow, s	8.3	6.0	compact core sized for an airframe rather than a test loop
$\tau_{L,0}$	external loop transit, s	47	10.0	one exchanger in a nacelle rather than eight in a test cell
β	static delayed fraction, pcm	470	650	six-group thermal uranium-235 data, so β_{eff} can be checked by hand
α_f	fuel temperature coefficient, pcm/°C	−5.5	−5.5	taken directly
α_m	moderator coefficient, pcm/°C	—	−1.0	small and negative by design
f_M	fission energy outside the fuel	0.108	0.108	taken directly
v_{rod}	rod speed limit, %/s	—	0.60	10.5 pcm/s at the design position, the order of the ARE regulating rod
W_{tot}	total rod worth, pcm	—	900	260 pcm is spent permanently on paying for circulation
T_{fo}	core outlet at rated power, °C	828	799	alloy limit set at 850 °C to leave transient margin
T_{fi}	core inlet at rated power, °C	654	659	comparable core temperature rise
C_f	core fuel heat capacity, MW s/°C	—	5.36	gives a 3 s core thermal time constant
C_m	moderator and structure, MW s/°C	—	24	gives a 40 s time constant, the slowest thermal mode
UA_{ihx}	IHX conductance, MW/°C	—	2.53	sized for 125 MW at the design terminal temperatures
UA_{ah}	air heater conductance, MW/°C	—	0.758	sized for 699 °C turbine inlet at takeoff
$\dot{m}_{max}$	engine air, sea level static, kg/s	—	232	low turbine inlet temperature forces a large air flow
π_c	compressor pressure ratio	—	6.0	low on purpose, because a cool discharge leaves room for the reactor
κ	bypass schedule exponent	—	0.67	fits the demanded duty at sea level and is deliberately wrong at altitude

On the aircraft side there is no flight model at all. No lift, no drag, no weight and no equation of motion, so nothing here can say whether the aircraft would climb at the rate the profile asserts. Thrust appears only as the turbine inlet temperature the engine did or did not receive. There is no shielding and no mass accounting, which means the reason the historical program was cancelled is entirely absent from the model.

8.4 Control Architectures

8.4.1 Pairing

Both actuators move both controlled variables. Withdrawing the rod raises the core temperature level, which raises the duty a fixed heat sink pulls, so the rod raises power and outlet temperature together. Speeding the fuel pump narrows the core temperature rise and cools the outlet, and raises the power slightly because a faster flow lowers the thermal resistance A in (8.12). The steady-state gains measured on the application at the sea-level takeoff condition are

$$\begin{bmatrix} \Delta P \\ \Delta T_{fo} \end{bmatrix} = \begin{bmatrix} +0.549 & +0.020 \\ +3.03 & -0.69 \end{bmatrix} \begin{bmatrix} \Delta z \\ \Delta f_p \end{bmatrix}, \tag{8.13}$$

in units of percent of rated power and degrees Celsius per percent of actuator. The relative gain array [20] of the pairing rod to power and pump to temperature is therefore 0.86, which is close to unity and says that the two loops barely interact at the design point. Volume 6, Chapter 2 treats a pairing whose relative gain of one still hides a one-way interaction.

That is not a property of the plant. Repeating the same four measurements at landing idle gives $\partial P/\partial z = 0.023$, a factor of twenty-four weaker, and $\partial T_{fo}/\partial f_p = -0.15$, a factor of four and a half weaker, while $\partial T_{fo}/\partial z$ is between 2.3 and 3.1 °C per percent at every steady condition on the mission except the top of the descent, where the rod sits at 73 % on the flat of its worth curve and the figure is 1.5. The relative gain is 3.58 there. The pairing chosen at the design point is strongly interacting by the time the aircraft is on the ground, and at low power the rod is the better temperature handle while the pump is barely a handle at all.

The intermediate salt circulator is instructive as a negative result. With the rod parked, moving it from 40 % to 130 % of design at cruise changes the settled reactor power by 0.2 points and the core outlet by 0.1 °C, because the air stream through the heater (0.04 MW/°C) is then the smaller capacity rate by a factor of more than seven even at 40 % circulation, and a circulator that only sets the spread of the intermediate loop leaves the duty where it was. The same move at the sea-level design point changes the power by 19.7 points, because at full throttle the heater air stream (0.27 MW/°C) and the intermediate stream at 40 % flow (0.31 MW/°C) are nearly balanced. It is a pump, it costs mass, it is a handle only where the rod is already strong, and it is not a manipulated variable.

8.4.2 Temperature Program

No fixed core outlet setpoint exists that both ends of the mission can reach. With the power loop closed, the fuel pump can place the outlet anywhere between 779 and 950 °C at the top of the takeoff climb and only between 740 and 752 °C at taxi idle, and those bands do not overlap. The setpoint is therefore a program that slides with the demanded power,

$$T_{fo}^{SP} = T_{prog} - k_T\,(100 - P_{SP})\,, \tag{8.14}$$

with $T_{prog} = 795\,^{\circ}\mathrm{C}$ and $k_T = 0.55\,^{\circ}\mathrm{C}$ per percent of power as shipped. Setting k_T to zero restores a fixed setpoint and the pump then spends most of the flight on a stop, which is the subject of scenario 6 in Table 8.3.

8.4.3 Decentralized Feedback

Both loops use the shared velocity-form PID of Section 1.6.2, with the derivative taken on the measurement, executed every $\Delta t_c = 0.25\,\mathrm{s}$ against a 0.05 s integration step. Equation (8.15) writes it in this chapter's symbols, with s_k for the running sum S_k of Equation (1.14) and c_k for the output v_k before limiting of Equation (1.15),

$$\begin{aligned} s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias,k} + s_k + I_{k-1} \\ u_k &= \mathrm{slew}[\mathrm{clamp}(c_k, u_{\min}, u_{\max})\,, u_{k-1}, v\,\Delta t_c] \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - c_k}{\tau_t}\right]\Delta t_c \end{aligned} \tag{8.15}$$

with $P_k = K_c e_k$, $D_k = -K_c \tau_D\, a\,(\mathrm{PV}_k^f - \mathrm{PV}_{k-1}^f)/\Delta t_c$ and $e_k = a(\mathrm{SP}_k - \mathrm{PV}_k)$, where a is the action flag and PV^f is the measurement passed through a first-order filter of 1 s. The output u_k is c_k clipped to the actuator range and then limited to a move of $v\,\Delta t_c$ per scan, and the integral alone absorbs the difference by back-calculation with a tracking time $\tau_t = \tau_I$. Setting τ_I to the top of its range (600 s) removes integral action and with it the back-calculation.

The power loop has a positive process gain, so $a = +1$ and the loop is reverse acting in the convention of Section 1.3. The temperature loop has a negative process gain, because a faster pump cools the outlet, so $a = -1$ and the loop is direct acting. The power loop's slew rate v is the rod drive rate of 0.6 %/s, so a demand the drive cannot execute is never counted as executed, and the temperature loop has no slew limit. The power loop's bias is the rod position at the start of the run, or the feedforward z_{ff} of Section 8.4.4 when that architecture is selected, and the temperature loop's

bias is the 100 % design pump speed. In manual operation and under the predictive controller both loops track the applied actuator positions, so every transfer is bumpless.

8.4.4 Preview Feedforward

The feedforward is solved from the steady state of the plant as declared, with the parameters of Table 8.2 at their defaults, rather than identified from a step test, because the operating point moves by a factor of fourteen in power across one flight and no single identified model covers it. Given the flight condition at $t + t_{lead}$, the demanded duty from (8.11) and the closed form (8.12), one bisection on the monotone rod worth curve returns the rod position that makes the plant critical at that duty. The feedforward is

$$z_{ff}(t) = \frac{1}{\tau_{ff} s + 1} z_{ss}(t + t_{lead}) + K_{ff} \left[P_{SP}(t + t_{lead}) - P_{SP}(t) \right]. \tag{8.16}$$

Two properties of (8.16) are worth stating because both were measured rather than assumed. The first is that the steady-state term contributes almost nothing to a takeoff. The rod position the plant needs is 54.1 % at taxi and 53.5 % at the top of the takeoff climb, so the destination barely moves and the entire job is transient. What a takeoff needs is the rod above its steady position for the length of the ramp, while the fuel, the moderator and two exchanger metal masses are charged to their new temperatures, and the size of that lead is available only from the schedule. That is the second term. Volume 9, Chapter 5 uses published orders the same way, released one travel time ahead at each gate.

The second property is that the steady-state term has to be lagged, and the lag is not a refinement. At low power the rod position this plant needs goes *up* as the throttle comes *down*, because the bypass over-bypasses and a smaller air stream needs a hotter core to drive even a small duty through it. Closing the throttle therefore asks the rod to move one way eventually and the opposite way immediately, and an unlagged steady-state feedforward gets the immediate half exactly backwards. Removing the 45 s lag takes the integrated power error over the takeoff window from 448 to 926 % s.

The feedforward enters the power loop as its bias $u_{bias,k}$ in (8.15), and the fuel pump is left on feedback. A pump feedforward solves for the flow that would hold the outlet on its program at the demanded duty, which during a descent is a steady state the plant will not reach for another minute. Placing that flow in the temperature loop's bias at the same 20 s preview, a variant tested and not shipped, drives the pump from 40 % to 110 % during the descent while the plant is still at the old power. Over the full flight it quadruples the pump travel (481 % against 117 %) and raises the outlet error from 8475 to 9270 °C s, while the power error moves only from 2817 to 2728 % s.

8.4.5 Dynamic Matrix Control

The predictive architecture is a two-input two-output dynamic matrix controller [7], not the quadratic-program controller that most apps offer (Section 1.6.3). It works over a 120 s horizon in twenty-four blocks of 5 s, with four moves per input. The objective is

$$J = \sum_{k=1}^{N_p} \left[w_P e_{P,k}^2 + w_T e_{T,k}^2 \right] + \sum_{j=1}^{N_m} \left[\lambda_z \Delta z_j^2 + \lambda_f \Delta f_{p,j}^2 \right], \tag{8.17}$$

and the unconstrained solution is the single matrix-vector product

$$\Delta u = \left(A^T Q A + \Lambda \right)^{-1} A^T Q e = K e, \tag{8.18}$$

after which the first move of each input is clamped into its actuator box and into the rod rate limit. The gain K is 8×48, eight moves against twenty-four predictions of each output, and is computed when the predictive controller takes over by Gauss-Jordan elimination on an 8×8 system. It is rebuilt

in place when the demand or the power has moved more than eight points of rated, which occurs 22 times over the full flight and ten times during the takeoff alone. Move suppression is normalized against the diagonal of $A^T QA$, so a weight of unity means a move costs as much as the error it removes, without which a value reasonable at takeoff freezes the controller at idle.

Two departures from the textbook form are load-bearing. The first is the free response. Textbook DMC shifts a stored prediction and corrects it with a step bias, which assumes everything unmodeled is a step. On this plant it is a ramp, because the throttle is moving, the bypass valve is moving with it and the reactor's own temperature feedback is dragging the power after them at two points per second, and a step-disturbance prediction starts a takeoff behind and never recovers. The free response here is instead built from the closed form (8.12) evaluated at the previewed flight condition with the actuators where they are,

$$\hat{y}_k = y + (y_k^{ss} + \text{bias} - y)\left(1 - e^{-kT_s/\tau_{free}}\right), \tag{8.19}$$

with $T_s = 5\,\text{s}$ the execution period and τ_{free} equal to 26 s for power and 40 s for the outlet. Each execution therefore costs twenty-five algebraic solves and no differential equations. The measured bias is carried as a level, which is what makes the controller offset free when the air heater fouls. The step responses and the free response use the parameters of Table 8.2 at their defaults, fixed when the application loads, and the prompt jump in the rod-to-power response is scaled by the measured power, so a changed parameter or the fouling reaches the controller only as model mismatch through that bias.

The second departure is that the outlet temperature is a zone of $\pm 12\,^{\circ}\text{C}$ rather than a setpoint, which is standard industrial practice for a variable that is wanted in a range [8] and the way Volume 9, Chapter 9 treats a comfort band. It matters here because the fuel pump saturates for about a third of the flight, so a target the pump cannot reach is a constant term in (8.17) that the optimizer can serve only with the rod, which moves the outlet by 2.3 to 3.1 °C per percent across the mission. At the shipped weights the optimizer barely does so. Closing the zone to zero moves the full-flight power error from 941 to 868 % s, an improvement rather than a cost, while the thrust deficit rises from 10 350 to 11 543 °C s.

The whole execution, free response and matrix-vector product together, is timed by the application whenever the predictive controller takes over and written to the exported file. It costs about 20 µs of wall-clock time on a desktop processor, once every 5 s of plant time.

8.5 Interactive Application

The app runs on the framework of Section 1.6. Twenty states are integrated with fixed-step fourth-order Runge-Kutta at 0.05 s, with the controller executing at 0.25 s and the predictive controller at 5 s. The external loop delay is a ring buffer of 8192 samples pushed every 0.05 s and interpolated linearly, which holds 409 s of history and covers the natural convection floor where τ_L reaches 200 s. Simulation speed is selectable from real time to fifty times, and the thirty-minute flight is therefore about half a minute of wall clock at the fastest setting.

Figure 8.3 shows the interface. Three tabs present the same simulation from three points of view. The mission tab draws the aircraft on its own scheduled altitude trace with the engine beneath it as compressor, air heater and turbine with the bypass duct looping around the heater. The plant tab draws the energy path end to end with a temperature ladder of every terminal temperature and the 850 °C limit across it. The neutronics tab draws the static delayed fraction as one bar split at β_{eff}, both transit times, a plot of β_{eff} against pump speed with the operating point marked, and the reactivity ledger as signed bars in pcm.

Eleven scenarios are provided, listed in Table 8.3 with what each establishes and what completes it. Each pins a flight phase, an architecture, a tuning and a disturbance schedule, and carries the briefing, Learn card and completion rule of Section 1.6.4. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. A finished run exports as a comma



Figure 8.3 The application in scenario 1 at run time 158 s, mission time 218 s, in the initial climb at 2700 ft and Mach 0.29, with both actuators parked, cropped to the mission tab. Power has overshot to 149.4 % of rated against a demand of 93.5 % and the high flux alarm is annunciated (5), while the core outlet is at 818 °C, 32 °C below the alloy limit. The status badges (1) lead with the power deviation. The mission tab shows the aircraft on the scheduled altitude profile (2), the phase name and flight condition, and the engine drawn as compressor, air heater and turbine with the bypass duct (3). The panel leads with the turbine inlet temperature (718 °C against a schedule of 699 °C) (4), which is the quantity that says whether the aircraft received the thrust it asked for.

separated file carrying the flight condition, both loops, every loop temperature, the reactivity ledger, both transit times, the accumulated metrics, every disturbance setting and the parameters the run used, so a single file reproduces the experiment without the application.

8.6 Results

Every number in this section was measured on the application at the tuning stated, with each case selected, its settings changed and then run. Integrated absolute error accumulates from the start of the run and never stops rising, so each figure is read at the end of the run its scenario defines and two figures are comparable only over the same window. Two windows are used throughout, a 300 s takeoff window that opens 30 s before the throttle moves (the case of scenarios 7 and 10), and the full 1800 s flight of scenario 8. Integrated errors are the scorecard's own, accumulated from the 4 Hz log, while peak power errors are taken against the continuous demand at the 0.05 s integration step. Peak deviations, peak temperatures and actuator travel do not have the window problem.

8.6.1 Inherent Load Following and Its Limits

With the rod and the fuel pump parked and nothing in automatic, the reactor follows the takeoff on its own and then overshoots it, the self-regulation that Chapter 2 studies in a solid-fuel core. Power reaches 149.5 % of rated against a peak demand of 100.1 % about fifty seconds after the throttle reaches full, 56.0 points above the demand at that moment. The largest gap between power and demand over the flight is 90.9 points the other way, at the end of the ramp, when the power has barely left its taxi value. The high flux alarm at 118 % annunciates on the way past.

Table 8.3 Scenarios and Completion Requirements

#	Scenario, *kind*	What it establishes	Required to complete
1	Nobody driving *demonstration*	inherent following, and its overshoot	Fly the mission in manual with the rod and the fuel pump untouched to $t = 4$ min, through the takeoff.
2	Hold and measure *experiment*	the gain matrix at two operating points	At the held sea-level, full-throttle condition in manual, step the rod up 2 % (from 54.2 to about 56.2 %) after $t = 5$ s and hold it 4 min, then step the pump up 10 % and hold it 4 min.
3	Takeoff by hand *experiment*	the rod rate limit	In manual, start moving the rod between $t = 1$ s and the end of the throttle ramp at $t = 75$ s, fly to $t = 2$ min with the outlet below 850 °C, then set the rod to the steady position the model gives for the demand.
4	Power loop *experiment*	one loop, and the temperature it ignores	Fly the whole mission to landing in automatic on the power loop alone, with the fuel pump left at 100 %.
5	Two loops *experiment*	the second loop and what it costs	Fly the mission in automatic on two loops until a flight phase pins the fuel pump on a stop, then on to landing.
6	Actuator authority *demonstration*	saturation without an undersized actuator	At ground idle on two loops with the program slope $k_T = 0$, run 5 min with the pump on a stop, then 5 min with $k_T \geq 0.5$ °C per percent. Then set the throttle to 100 %, Reset so the run starts lined out there, and run 5 min without a trip.
7	Preview feedforward *experiment*	lead, lag and anticipation gain	Four takeoff runs on preview feedforward to $t = 5$ min, each setting held throughout: as shipped, $t_{lead} = 0$, $t_{lead} = 45$ s, and $K_{ff} = 0$ at $t_{lead} = 20$ s.
8	Whole flight compared *experiment*	four architectures, no single winner	Four whole missions to landing in automatic, one on each architecture: power loop alone, two loops, preview feedforward and MPC.
9	Go-around *experiment*	the constraint, crossed in a transient	Three runs through the go-around at $t = 100$ s, each held on its architecture for 70 s between $t = 90$ and 165 s: two loops, preview feedforward and MPC.
10	Coordinated MPC *experiment*	the value of the preview switch	Three takeoff runs on MPC to $t = 5$ min: preview on with the ± 12 °C zone (IAE$_P$ below 1500 % s), preview off (above 3000 % s), and preview on with the zone at zero (below 1500 % s).
11	Graded challenge *challenge*	degradations a feedforward cannot see	Fly the whole 1800 s mission with power within ± 1.2 % of demand from $t = 1188$ s, the outlet never above 850 °C and total rod travel under 150 %, with the model parameters and disturbances as given.[a]

[a] With the shipped two-loop tuning ($K_{c,P} = 0.35$, $\tau_{I,P} = 150$ s, $K_{c,T} = 0.5$, $\tau_{I,T} = 180$ s) the worst power error over the judged window is 5.07 %, more than four times the band, although the outlet peaks at only 801 °C and rod travel is 112 %. With the suggested MPC architecture it is 0.92 %, inside the band with 0.28 points to spare, with a peak outlet of 818 °C and 99 % of rod travel.

The mechanism is a delay acting on a stiff feedback. The bypass valve closes in about two seconds and the air heater begins pulling at once, but the cooled salt must cross the air heater metal, the intermediate loop, the exchanger metal and a 10 s fuel loop transit before the core is affected. By the time the fuel temperature has fallen far enough for the −5.5 pcm/°C coefficient to supply the reactivity the ramp needed, the ramp is over and the reactivity is still present. The cost to the aircraft is 48 °C of turbine inlet temperature below schedule at the worst point, integrating to 4383 °C s by landing.

Table 8.4 Measured mission performance. IAE_P in percent-seconds, IAE_T in °C-seconds, thrust deficit is the integral of turbine inlet temperature below schedule in °C-seconds, and peak $|e_P|$ is the largest gap between power and the continuous demand at the 0.05 s integration step

	Takeoff window, 300 s				Full flight, 1800 s			
Architecture	IAE_P	peak $\lvert e_P\rvert$	peak T_{fo}	deficit	IAE_P	IAE_T	peak T_{fo}	deficit
Nobody driving	6235	90.9	819	1781	9990	7723	819	4376
Power loop only	687	10.93	817	64	1312	26525	817	5974
Two loops	641	10.41	810	65	1487	10667	812	7822
Preview feedforward	448	14.8	808	57	2817	8475	809	6791
Dynamic matrix control	729	15.6	818	24	941	20144	818	10350
DMC, preview disabled	6358	99.9	782	6231	10276	32259	786	31385

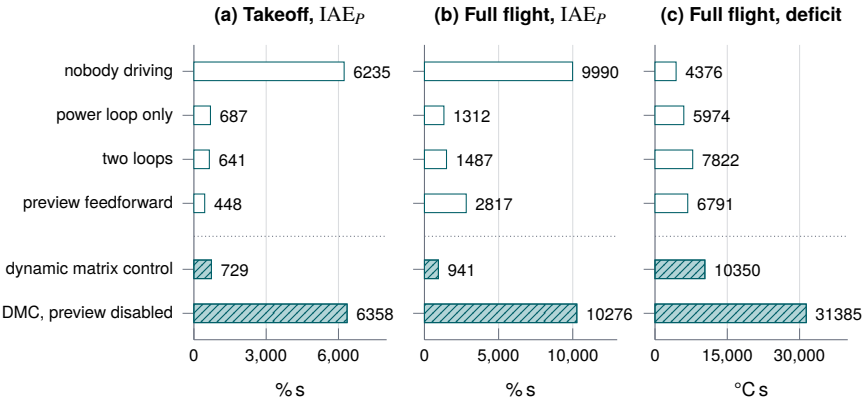


Figure 8.4 Three columns of Table 8.4 drawn as bars. White bars, above the dotted line: the plant with nobody driving and the three architectures without the predictive controller. Hatched bars, below it: the dynamic matrix controller with its preview and with the preview disabled. Notice that the same predictive controller with its preview disabled lands beside the plant with nobody driving in panels (a) and (b), and that the shortest bar belongs to a different row in each panel.

> **TRY IT IN THE APP** **1 · Fly the Mission With Nobody Driving**
>
> Load scenario 1 and press Run, leaving the control rod and the fuel pump alone. Watch reactor power against the demand trace through the takeoff. The power lags at first, then climbs on its own past the demand to 149.5 % of rated, and the high flux alarm at 118 % annunciates on the way past.

8.6.2 Architecture Comparison

Table 8.4 reports the four architectures on both windows at the shipped tuning of $K_{c,P} = 0.9$, $\tau_{I,P} = 45\,\mathrm{s}$, $K_{c,T} = 1.1$ and $\tau_{I,T} = 60\,\mathrm{s}$, with the feedforward at $t_{lead} = 20\,\mathrm{s}$, $K_{ff} = 0.20$ and $\tau_{ff} = 45\,\mathrm{s}$, and the predictive controller at $w_P = 4.0$, $w_T = 0.04$, $\lambda_z = 0.15$, $\lambda_f = 0.30$ and a $\pm 12\,^{\circ}\mathrm{C}$ zone. With the power loop alone the fuel pump stays parked at 100 %.

The takeoff power peaks are read against the engine's step-level demand at each physics scan. Readings against the continuous demand trace give 11.1 and 10.6 points for the first two controlled rows. Figure 8.4 draws three of the columns as bars.

No architecture wins every column, and the rows that appear to lose are the informative ones. The power loop alone removes most of the tracking error, taking IAE_P over the flight from 9990 to

1312 % s and the worst gap from 90.9 points to 10.93, and controls no temperature at all, integrating 26 525 °C s of deviation from the program. Adding the second loop cuts that to 10 667 °C s and costs tracking (1487 % s), because the pump is now moving and every pump move changes the thermal resistance of the whole chain.

Preview feedforward wins the takeoff and loses the flight. Over the 300 s window it integrates 448 % s against 641 for the two loops alone, and over the full flight it integrates 2817, which is worse than either loop architecture. The explanation is in (8.16). The anticipation term acts only while the demand is changing, so it is free during the quiet phases, while the lagged steady-state term still leads the plant during the slow throttle reductions of the descent even at a 45 s lag.

The predictive controller wins the flight by a wide margin, 941 % s against 1487 for the two loops, because it coordinates both actuators against a setpoint trajectory it can see. Its temperature figure of 20 144 °C s is a direct consequence of the $\pm 12\,^{\circ}\mathrm{C}$ zone rather than a failure, and closing the zone to zero returns the temperature figure to 10 358 °C s and the power figure to 868 % s at the cost of the thrust deficit (11 543 against 10 350 °C s).

8.6.3 The Value of Preview

The last row of Table 8.4 is the central experiment. It is the same controller with the same model, the same horizon, the same weights and the same actuators, and the only change is that the setpoint trajectory over the horizon is replaced by the present setpoint held constant, which is what ordinary model predictive control does. Integrated power error over the takeoff goes from 729 to 6358 % s, almost a factor of nine, and the thrust deficit from 24 to 6231 °C s. Over the full flight the tracking error goes from 941 to 10 276, which is slightly worse than the 9990 of the open-loop case.

> ✦ **DID YOU KNOW? Preview Off, Worse Than Nobody Driving**
>
> From the simulator: over the full flight, the dynamic matrix controller with its preview disabled integrates 10 276 % s of power error, slightly more than the 9990 % s of the plant with nobody driving. With the preview it integrates 941 % s (Table 8.4).

The feedforward architecture isolates the same effect in a simpler controller. Setting t_{lead} to zero, which converts (8.16) into an ordinary steady-state feedforward acting on the present flight condition, takes the takeoff error from 448 to 681 % s, which is worse than having no feedforward at all. Setting the anticipation gain K_{ff} to zero and keeping the preview gives 657. Neither half improves on the 641 % s of the two loops alone, and together they are worth 193 % s against having neither.

Preview also buys the thrust deficit directly, and at a price. Extending t_{lead} to 45 s at $K_{ff} = 0.2$ takes the deficit from 57 to 11 °C s with the worst shortfall down from 6 °C to 3, while the integrated power error rises to 930 % s. That trade is physical rather than a tuning artifact. During a ramp the reactor must supply the demanded duty and the energy going into its own thermal capacities at the same time, so tracking the demand exactly and satisfying the engine exactly are different objectives.

8.6.4 Go-Around

The go-around is the transient the mission is built around and the only one that is not in the profile. The aircraft is on approach at about a quarter throttle with the demand at 9 % of rated power and a core that has been cooling for ten minutes, and at $t = 100\,\mathrm{s}$ the throttle goes to 95 % in twelve seconds. Table 8.5 reports it.

The two independent loops exceed the containment limit. Power passes the demand by 144.3 points at 152 s, and the core outlet crosses 850 °C at 147 s, peaks at 859 °C and spends 8.3 s above the limit. Preview feedforward does not help (142.9 points, 859 °C), because the event is not in the profile and there is nothing to read ahead.

🕮 **Pause.** Try Exercise 8.3 before reading on.

Table 8.5 Go-around from the approach, 300 s window. Over and under are the largest excess of power above the demand and shortfall below it, in points of rated power

Architecture	over	under	peak T_{fo}	$t > 850\,°\mathrm{C}$	deficit
Nobody driving	62.6	72.4	792	0 s	2174
Power loop only	14.5	63.0	793	0 s	1354
Two loops	144.3	63.5	859	8.3 s	2228
Preview feedforward	142.9	61.8	859	8.3 s	2150
Dynamic matrix control	51.6	65.8	815	0 s	1417

The overshoot comes from the two loops acting together rather than from the plant or from the power loop. With nobody driving, the power lags the demand by 72.4 points while the throttle opens and passes it by only 42.9 points at 182 s, once the salt cooled by the opened heat sink has crossed the loop and the fuel temperature coefficient has delivered its reactivity. Its larger excess in Table 8.5 comes later, when the throttle returns to the approach schedule. The power loop alone, with the pump parked at 100 %, passes the demand by 7.9 points during the go-around and keeps the outlet at 793 °C. Closing the temperature loop as well takes the pump from 61 % to its 40 % stop as the throttle opens and then to 125 % by 146 s as the outlet climbs, and because the pump also sets β_{eff}, every one of those moves is a reactivity change. The power stalls near 54 % between 130 and 136 s, the power loop keeps withdrawing the rod at its rate limit to 70.3 %, and that stored rod reactivity arrives together with the plant's own delayed insertion.

Of the three architectures that control the outlet, the predictive controller is the only one that stays under the limit. Its worst power error is a 65.8-point shortfall while the throttle is still opening, it passes the demand by 51.6 points, and the outlet peaks at 815 °C. It does so by moving the fuel pump and the rod together against a prediction that already contains the plant's own passive response, so it does not add its own on top. The tell-tale is the fuel pump. The two-loop system reaches 125 % only one second before the outlet crosses the limit, while the predictive controller puts the pump on its 125 % stop at the execution that coincides with the go-around command, before the throttle has moved.

> **TRY IT IN THE APP 9 · Go-Around From the Approach**
>
> Load scenario 9 and run it on Two loops through the go-around at 100 s, watching the core outlet on the Plant tab. It peaks at 859 °C and spends 8.3 s above the 850 °C limit. Pause, press Reset, choose MPC before pressing Run and repeat: the outlet now peaks at 815 °C.

8.6.5 Degradation and the Graded Challenge

The graded scenario runs the full flight from a deliberately sluggish tuning with 0.8 °C of thermocouple noise, 0.5 % of flux channel noise, and three undeclared events: a 20 % loss of air heater conductance at 620 s, a 6 °C thermocouple bias at 900 s and a fuel pump delivery gain of 0.85 at 1180 s. The target is reactor power within ± 1.2 % of demand through the last third of the flight (from $t = 1188$ s) with the outlet below 850 °C and rod travel under 150 %. The sluggish tuning as given misses by more than a factor of four, at 5.07 %.

The predictive controller holds 0.92 % with a peak outlet of 818 °C for 99 % of rod travel and 162 % of pump travel. The two loops at the tuning of Table 8.4 reach 2.68 % for 258 % of rod travel, and tightening them to $K_{c,P} = 1.6$ with $\tau_{I,P} = 30$ s buys 1.55 % for 397 % of rod travel. Preview feedforward at the same tuning is the worst of the four at 8.93 %.

The feedforward result is the one worth stating carefully. A feedforward is an open-loop correction computed from a model, so a plant degradation does not disturb it, it invalidates it, and the feedback

loop underneath then has to clean up a correction that is confidently wrong. The feedforward's model keeps the declared air heater conductance, so the fouling at 620 s leaves it unchanged. The predictive controller survives the same two degradations because its model bias in (8.19) is measured every execution, so the degradation appears in the bias rather than in the controlled variable.

8.6.6 Saturation Without an Undersized Actuator

Holding the aircraft at ground idle (15 % throttle) for five minutes with the program slope k_T set to zero leaves the outlet 23.1 °C below a fixed 795 °C setpoint with the fuel pump pinned on its 40 % stop, and integral action never moves it. The reachable band at that condition is 763 to 773 °C. Restoring the slope slides the setpoint to 743 °C, which is outside the band in the other direction, and the pump pins on its 125 % stop with the outlet 19.5 °C high after a further five minutes. At 100 % throttle, with the run started lined out there by a Reset, both settings work and the pump sits at 101.5 and 103.1 %.

The pump has not run out of travel. At 6.6 MW the core temperature rise is 7.4 °C at design flow, and the pump's whole range only moves it between 6.0 and 18.6 °C, which cannot move an outlet temperature twenty-three degrees. The actuator that can move it there is the rod, worth 2.3 °C per percent at that condition against the pump's 0.15, which is the relative gain of 3.58 written out.

Pause. Try Exercise 8.2 before reading on.

8.7 Technology in Practice

The two measurements of Figure 8.2 are, in a plant, a neutron detector outside the reactor vessel and a thermocouple in the core outlet line, and the two manipulated variables are a regulating rod on its drive and a fuel pump on a speed drive. Figure 8.5 draws the loops with that hardware in place of the schematic's symbols.

What Measures It

Neutron flux is measured by fission chambers, ion chambers and self-powered detectors, and detectors outside the vessel give the flux as a representation of reactor power [21]. In a molten salt reactor the flux reading serves as the fast indication of power for feedback control, while the longer-term power comes from heat balances and from samples of the fuel salt [22]. The core outlet temperature is read by thermocouples in thermowells inserted into the flow [22], and every thermowell adds to the first-order lag of its sensor through its mass and specific heat [23]. The chapter carries both instruments as a first-order lag with colored noise and a constant bias, and the graded scenario adds thermocouple noise, flux channel noise and a thermocouple bias.

> **TECHNOLOGY Sensor: An Ion Chamber Outside the Vessel**
>
> Power-range neutron detectors are typically uncompensated ion chambers, since the neutron signal is large against the gamma [21]. In a molten salt reactor they would sit outside the vessel, should answer a change of power within about a millisecond, and be calibrated against a measurement of thermal power. No high-sensitivity neutron detector that works above 600 °C is commercially available. Heat lowers the resistance of a chamber's ceramic insulators and causes spurious discharges whose pulses cannot be told from neutron pulses [22], so a detector run too hot reports counts that are not neutrons. The Aircraft Reactor Experiment cooled its special fission chambers in the reflector with helium [10]. The chapter's NI is a lag with colored noise and a bias.

What Moves It

The power loop writes the rod drive at no more than 0.6 %/s of travel, and the temperature loop writes the fuel pump between 40 and 125 % of design speed with no slew limit. The rod worth follows an S-curve, and at the top of the descent the rod sits on the flat of that curve, where it moves the outlet temperature least.

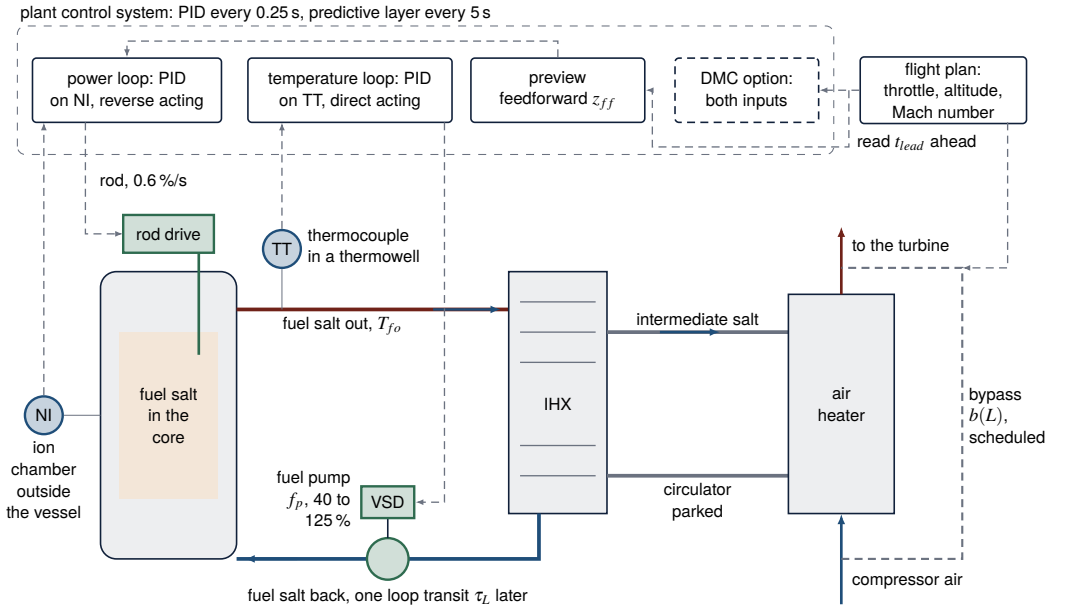


Figure 8.5 The two loops of the chapter with plant hardware in place of the symbols of Figure 8.2. Sensors are blue circles and final elements green: the neutron instrument NI, an ion chamber outside the vessel, the core outlet thermocouple TT in its thermowell, the regulating rod on its rate-limited drive and the fuel pump on a variable-speed drive (VSD). The power loop writes the rod, with the preview feedforward in its bias, and the temperature loop writes the pump. The predictive option replaces both loops and reads the flight plan over its horizon. Notice that the flight plan reaches the controllers directly, while the plant feels the same throttle only through the air heater, the intermediate salt, the exchanger and the fuel loop.

TECHNOLOGY Final Element: A Regulating Rod on a Slow Drive

The Aircraft Reactor Experiment had one stainless-steel regulating rod. It moved 12 inches about the center of the reactor, was worth 0.4 percent in reactivity over that stroke and could be moved at 0.011 percent per second. Its three larger rods of boron carbide and iron each took 13 minutes to withdraw [10]. In power reactors the rod position is indicated twice, by a detector on the drive housing and by counting the signals sent to the drive, and after a dropped or stuck rod the two disagree. Power supply failures and drift have dropped or slipped rods [24]. The chapter's rod moves at 0.6 %/s, 10.5 pcm/s at the design position.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, and a supervisory layer sends setpoints to the field controllers below it [25]. The chapter's two PID blocks execute every 0.25 s over a 0.05 s integration step: a reverse-acting block on the neutron instrument with the rod drive's slew limit on its output and the preview feedforward in its bias, and a direct-acting block on the outlet thermocouple whose setpoint follows the temperature program. The dynamic matrix controller is the supervisory layer. It executes every 5 s, writes both actuators at once and reads the flight plan over its 120 s horizon.

> **IN PRACTICE** **Measuring the Plant While It Runs**
>
> For the reactivity control of a molten salt reactor, the variables to watch are the position of the control element and the response of the neutron flux to each change. Small perturbations may be impressed on the power to evaluate the reactor's dynamics online, as the pseudo-random binary movements of the regulating rod did at the Molten Salt Reactor Experiment. In the United States, general design criterion 26 of 10 CFR Part 50, Appendix A, requires two independent means of bringing and holding the reactor subcritical [22]. Scenario 2 of this chapter measures the gain matrix in the same spirit, by stepping the rod and then the pump at a held condition.

What Is Modernizing It

Three developments reach these loops. Neutron noise, cross-correlated between several detectors outside the vessel, is expected to show how the fuel salt is distributed in the core, for example after a flow blockage or with vessel vibration [22]. Operating plants are moving from analog to digital instrumentation and control through individual system upgrades [21]. And control systems that adopt office computers, operating systems and network protocols have far less isolation from the outside world and a greater need to be secured [25], which matters to a controller that reads a flight plan from outside the reactor.

8.8 Limitations

The reduced-order structure used here is the structure that the Aircraft Reactor Experiment and Molten Salt Reactor Experiment literature supports for circulating-fuel *neutronics* and thermal *time scales* [13, 14]. The thermal system of this aircraft plant is notional. It has not been validated against any experiment, and published work on reduced-order models of this class cautions that transient core temperature predictions from them are less reliable than their power predictions. The temperatures in this application should be read as a control-relevant model rather than as a thermal-hydraulic prediction, and nothing here is evidence about whether a nuclear aircraft could be built.

The predictive controller is unconstrained and then clipped. Move and position limits are imposed by clamping the computed move into the actuator box rather than by solving a constrained quadratic program [7], and the fuel temperature limit enters as a weighted zone rather than as a hard constraint. The go-around result in Table 8.5 is therefore correctly stated as an observation that the outlet stayed below 850 °C and not as a guarantee that it could not exceed it.

The air heater bypass is a schedule rather than a governor, which is a deliberate choice. A real engine control system would close a loop around turbine inlet temperature and hide the reactor behind it, and there would then be no reactor control lesson. The consequence is that the schedule is wrong in ways a real design would not be, particularly at altitude, and the power loop is cleaning that up as well as rejecting the disturbance.

Instruments are a first-order lag plus colored noise plus a constant bias. A thermocouple in a salt thermowell has a nonlinear calibration, drifts, and fails in ways that do not resemble a constant offset. Finally, the fixed step and the deterministic pseudorandom noise make every run reproduce exactly on any machine, which a real plant does not, and that repeatability should be treated as a teaching convenience.

> **OPEN CHALLENGE** **A Neutron Detector for the Hot Core**
>
> Salt reactor cores run near 750 °C, and no high-sensitivity neutron flux detector that works above 600 °C is commercially available. Inside the core, long exposure to heat lowers the resistance of ceramic insulators, causes spurious discharges, outgasses the chamber body into its fill gas and weakens the electrode supports. Detectors outside the vessel avoid this, but they will not give a detailed picture of the flux inside the core [22]. Progress would be a detector that survives years at core temperature, so that power could be read where it is made, not only through the vessel wall.

8.9 Conclusion

This chapter studied a reduced-order model and an interactive app for a notional molten salt reactor supplying heat to aircraft turbojets. The reactor model retains circulating-fuel point kinetics with delayed neutron precursor transport, thermal reactivity feedback and lumped energy balances in the form the Aircraft Reactor Experiment literature supports, and the aircraft enters as a scheduled disturbance through the corrected air flow group and the compressor discharge temperature.

Three results follow from the mission comparison. The plant follows a takeoff without being commanded, lagging it by 90.9 points of rated power at the end of the ramp and then overshooting to 149.5 % of rated, so inherent load following reduces the size of the control problem without removing it. The conventional pairing of control reactivity to power and fuel circulation to core outlet temperature has a relative gain of 0.86 at takeoff and 3.58 at ground idle, so a pairing chosen at the design point is strongly interacting before the aircraft lands, and the fuel pump saturates at low power because its process gain has gone rather than because it is too small. Advance knowledge of the flight profile is worth more than any retuning of the feedback loops, taking the integrated power error over a takeoff from 6358 to 729 % s for a predictive controller that is otherwise identical.

The go-around is the case that separates coordination from feedforward. Two independent loops that hold every steady phase of the flight put the fuel outlet 9 °C above its limit for eight seconds, feedforward does not help because the event is not scheduled, and a coordinated predictive controller peaks at 815 °C. The eleven scenarios of the app, each with a completion rule, and its data export carry these studies of multivariable pairing, feedforward timing and preview control on a plant whose disturbance is a published plan, and the exercises of Section 8.10 build on them.

KEY POINTS

- A reactor whose fuel is a pumped liquid salt follows a takeoff on its own, but it lags the demand and then overshoots to 149.5 % of rated.
- The rod-to-power and pump-to-temperature pairing has a relative gain of 0.86 at takeoff and 3.58 at ground idle.
- At low power the fuel pump saturates because its process gain has gone, not because it is too small.
- Preview of the flight profile takes the takeoff power error of the same predictive controller from 6358 to 729 % s.

Where else this applies. Volume 9, Chapter 5 also acts on published orders ahead of time.

References

[1] E. S. Bettis, R. W. Schroeder, G. A. Cristy, H. W. Savage, R. G. Affel, and L. F. Hemphill, “The Aircraft Reactor Experiment–Design and construction,” *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 804–825, 1957.

[2] S. Comassar, “General electric direct-air-cycle aircraft nuclear propulsion program: Aircraft nuclear propulsion application studies (comprehensive technical report),” General Electric Company, Cincinnati, OH, Tech. Rep. APEX-911, 1962.

[3] W. Zeng, Q. Jiang, J. Xie, and T. Yu, “Core outlet fuel temperature control of liquid molten salt reactor during load following operation,” *Progress in Nuclear Energy*, vol. 121, p. 103214, 2020.

[4] J. P. Carter, R. Christensen, and S. Yoon, “Numerical analysis of dynamic load following response in a natural circulation molten salt pebble bed reactor,” *Nuclear Engineering and Design*, vol. 394, p. 111818, 2022.

[5] International Atomic Energy Agency, *Non-baseload Operation in Nuclear Power Plants: Load Following and Frequency Control Modes of Flexible Operation*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2018, no. NP-T-3.23.

[6] H. A. Pak and R. Shieh, “On MIMO optimal preview tracking control for known trajectory models,” *Optimal Control Applications and Methods*, vol. 12, no. 2, pp. 119–130, 1991.

[7] C. E. Garcia and A. M. Morshedi, “Quadratic programming solution of dynamic matrix control (QDMC),” *Chemical Engineering Communications*, vol. 46, no. 1-3, pp. 73–87, 1986.

[8] 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.

[9] E. S. Bettis, W. B. Cottrell, E. R. Mann, J. L. Meem, and G. D. Whitman, "The Aircraft Reactor Experiment–Operation," *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 841–853, 1957.

[10] W. B. Cottrell, H. E. Hungerford, J. K. Leslie, and J. L. Meem, "Operation of the aircraft reactor experiment," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL-1845, 1955.

[11] W. K. Ergen, A. D. Callihan, C. B. Mills, and D. Scott, "The Aircraft Reactor Experiment–Physics," *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 826–840, 1957.

[12] G. Thornton and A. J. Rothstein, "General electric direct-air-cycle aircraft nuclear propulsion program: Program summary and references (comprehensive technical report)," General Electric Company, Flight Propulsion Laboratory Department, Cincinnati, OH, Tech. Rep. APEX-901, 1962.

[13] T. W. Kerlin, S. J. Ball, and R. C. Steffy, "Theoretical dynamics analysis of the Molten-Salt Reactor Experiment," *Nuclear Technology*, vol. 10, no. 2, pp. 118–132, 1971.

[14] S. J. Ball and T. W. Kerlin, "Stability analysis of the Molten-Salt Reactor Experiment," Oak Ridge National Laboratory, Tech. Rep. ORNL-TM-1070, 1965.

[15] P. N. Haubenreich and J. R. Engel, "Experience with the Molten-Salt Reactor Experiment," *Nuclear Applications and Technology*, vol. 8, no. 2, pp. 118–136, 1970.

[16] J. Shi and M. Fratoni, "GeN-Foam model and benchmark of delayed neutron precursor drift in the Molten Salt Reactor Experiment," *EPJ Web of Conferences*, vol. 247, p. 06040, 2021.

[17] M. Elhareef and Z. Wu, "Benchmarking a point reactor kinetics method with delayed neutron precursors transport using a multiphysics coupled framework," *Annals of Nuclear Energy*, vol. 217, p. 111366, 2025.

[18] K. S. Kupriyanov, O. S. Feinberg, and V. V. Ignatiev, "Effective delayed neutron fraction in a molten salt reactor with circulating fuel," *Physics of Atomic Nuclei*, vol. 85, no. 8, pp. 1391–1399, 2022.

[19] J. J. Duderstadt and L. J. Hamilton, *Nuclear Reactor Analysis*. New York: John Wiley and Sons, 1976.

[20] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[21] International Atomic Energy Agency, *Core Knowledge on Instrumentation and Control Systems in Nuclear Power Plants*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2011, no. NP-T-3.12.

[22] D. E. Holcomb, R. A. Kisner, and S. M. Cetiner, "Instrumentation framework for molten salt reactors," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2018/868, 2018.

[23] 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.

[24] E. Grove, W. Gunther, and K. Sullivan, "Effect of component aging on PWR control rod drive systems," Brookhaven National Laboratory, Upton, NY, Tech. Rep. BNL-NUREG-47470, 1992, aging Research Information Conference, Rockville, MD, March 1992. OSTI 10147987.

[25] 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.

8.10 Exercises

8.1 Concept check · 2 min

At steady power and design flow the control rod of the aircraft reactor holds about 260 pcm of positive reactivity, and Equation (8.4) sets $\rho_0 = \beta - \beta_{\text{eff}}$. Why is a circulating-fuel core critical at positive reactivity?

(a) The negative fuel coefficient has to be paid for at rated power, so the rod holds the power defect

(b) The moderator coefficient is small, so the rod must supply what the moderator does not

(c) The prompt jump approximation removes the prompt neutrons, so the rod must replace them

(d) Precursors that decay in the external loop are lost to the chain, so prompt multiplication makes up the share

8.2 Concept check · 1 min

In scenario 6 at ground idle with the program slope $k_T = 0$, the core outlet settles 23.1 °C below its setpoint with the fuel pump on its 40 % stop, and the integral never moves it. What is the right diagnosis?

(a) The pump has run out of process gain, because the core rise it can redistribute is only a few degrees

(b) The pump is undersized, and a pump with a wider speed range would reach the setpoint

(c) The integral has wound up, and back-calculation is switched off on the temperature loop

(d) The thermocouple lag is too long, so the loop cannot see the outlet move in time

8.3 Concept check · 2 min

In the go-around of scenario 9, preview feedforward leaves the core outlet at 859 °C, no better than the two loops alone. Why does the preview not help here?

(a) The feedforward lag of 45 s is too slow to follow a 12 s throttle change

(b) The go-around is not in the scheduled profile, so the preview has nothing ahead to read

(c) The rod rate limit stops the feedforward term from reaching the rod during the ramp

(d) The feedforward acts only on the rod, so the pump overshoot that drives the outlet is unchanged

8.4 Theory and calculation · 6 min

Equation (8.13) gives the steady-state gains at the sea-level takeoff condition. The app's steady-state model gives, at landing idle, $\partial P/\partial z = 0.02326$, $\partial P/\partial f_p = -0.001116$, $\partial T_{fo}/\partial z = 2.287$ and $\partial T_{fo}/\partial f_p = -0.1523$, in the same units.

(a) Compute the relative gain λ_{11} of the pairing rod to power and pump to temperature at takeoff. *Checkpoint:* 0.862, within 0.005.

(b) Compute λ_{11} at landing idle. *Checkpoint:* 3.58, within 0.02.

(c) With the temperature loop closed and holding the outlet exactly, the steady gain of the rod on power is K_{11}/λ_{11}. Compute it at both conditions and compare it with the open-loop gain. *Checkpoint:* 0.637 at takeoff and 0.0065 at idle, % power per % rod.

8.5 Theory and calculation · 5 min

The core fuel salt has a heat capacity $C_f = 5.357\,\mathrm{MW\,s/°C}$ and a core transit time of 6.0 s at design flow, and both scale as in Equation (8.3). Take the duty as fixed.

(a) Compute the flow capacity rate W_f at design flow and the core temperature rise Q/W_f at the rated 125 MW and at 6.6 MW. *Checkpoint:* 0.893 MW/°C, 140 °C and 7.39 °C.

(b) Compute the rise at 6.6 MW with the pump on its 40 % and 125 % stops. *Checkpoint:* 18.5 °C and 5.91 °C, each within 0.1 °C.

(c) Each core node carries a linear profile, so the mean fuel temperature is very nearly the average of inlet and outlet and $T_{fo} = \bar{T}_f + \Delta T/2$. With the mean held fixed, how far can the whole pump range move the outlet at 6.6 MW? *Checkpoint:* 6.3 °C, within 0.1 °C.

8.6 Theory and calculation · 7 min

At the sea-level takeoff condition the Mach number is 0.25 and the ambient temperature is 288.15 K. The compressor has $\pi_c = 6.0$ and $\eta_c = 0.85$, the air has $c_{p,a} = 1.15\,\mathrm{kJ/(kg\,K)}$, and the engine schedules a turbine inlet temperature of 700 °C at full throttle. Use Equations (8.9) to (8.11) and Table 8.2.

(a) Compute the compressor discharge temperature T_{c2}. *Checkpoint:* 248.0 °C, within 0.5 °C.

(b) At sea level $\delta/\sqrt{\theta} = (1 + 0.2M^2)^3$. Compute it, the engine air flow at full throttle and the capacity rate $\dot{m}_a c_{p,a}$ of the air through the heater with the bypass closed. *Checkpoint:* 1.038, 240.8 kg/s and 0.277 MW/°C.

(c) Compute the demanded duty Q_{dem} and compare it with the rated power. *Checkpoint:* 125.2 MW, within 0.3 MW.

8.7 App challenge · 7 min

Load scenario 1, Fly the mission with nobody driving (MANUAL), and change no setting. Run at 20× to $t = 4$ min at the default seed, so the run takes 12 seconds, and export the CSV file. Its time column is in minutes. The throttle reaches full at $t = 105$ s.

(a) Find the peak of `power_pct` and the time at which it occurs, in seconds. *Checkpoint:* 149.5 %, within 0.1 %, at 157 s, within 1 s.

(b) Find the largest shortfall of `power_pct` below `power_sp_pct` and when it occurs. *Checkpoint:* 90.8 points, within 0.2, at 104.5 s, within 1 s.

(c) Find the largest value of `tti_sp_C` minus `tti_C`, and explain with the chapter's mechanism why the power first lags and then overshoots. *Checkpoint:* 48 °C, within 0.5 °C, at about 113 s.

8.8 App challenge · 8 min

Load scenario 10, Coordinated MPC with the profile in view. Each run goes to $t = 5$ min at 20× (15 seconds) at the default seed. Read the power loop's IAE on the scorecard or in the last row of column `power_iae`, and the thrust deficit in the last row of column `thrust_deficit_Cs`.

(a) Run the case as shipped, with MPC setpoint preview on Preview on. *Checkpoint:* IAE 729 % s and deficit 24 °C s, each within 2.

(b) Pause, Reset, set MPC setpoint preview (Advanced) to Preview off before Run, and run again. Give the ratio of the two power IAE values. *Checkpoint:* IAE 6358 % s, within 5, deficit 6231 °C s, within 5, ratio 8.7.

(c) Reload the scenario, set MPC temperature zone (Advanced) to 0 with the preview on, and run a third time. Does closing the zone cost the power tracking on this window? *Checkpoint:* IAE 638 % s, within 2.

8.9 App challenge · 8 min

Load scenario 9, Go-around from the approach. The go-around is commanded at $t = 100$ s. Run each case to $t = 160$ s at 10× (16 seconds) at the default seed and export the CSV file, whose time column is in minutes.

(a) Run on Two loops as shipped and read the Peak fuel margin tile. *Checkpoint:* −9 °C, an outlet peak of 859 °C.

(b) Pause, Reset, choose MPC before Run, and run again. Read the same tile. *Checkpoint:* 35 °C, an outlet peak of 815 °C.

(c) In each file find the first row at which `pump_cmd_pct` reads 125 %, in seconds, and explain what the difference shows. *Checkpoint:* 146.5 s on two loops and 100.25 s on MPC, each within 0.5 s.

8.10 Application area · 3 min

Without the air heater bypass, a fixed-geometry heater with the core near its design temperature pulls 26 MW at ground idle.

(a) The demand there is 7 % of the rated 125 MW. By what factor does the heater's pull exceed it? *Checkpoint:* About 3, 21 % against 7 %.

(b) The only other way to shed the surplus is to cool the salt to 409 °C. Why is that not available, and what does the bypass cost the power loop instead?

8.11 Application area · 4 min

A safety reviewer wants to cite Table 8.5 as evidence that the dynamic matrix controller keeps the fuel outlet below 850 °C in a go-around.

(a) Give two reasons from the chapter why the table does not support that claim.

(b) Name one change to the controller that would make the limit a constraint it enforces, and say what would still not be proven.

8.12 Review the agent · 5 min

An AI assistant writes: “The relative gain of rod to power and pump to temperature is 0.86 at takeoff, so the pairing is nearly decoupled. Tune the two loops one at a time at the sea-level design point. Because the relative gain is a property of the plant, the same pairing and tuning will behave the same way for the whole mission.” Find the flaw and say how you would check it in the app.