INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 4 Override, Split-Range, and Constraint Control

CHAPTER 10

A Supercritical CO_2 Brayton Simulator for Three-Layer Mid-Ranging Control

Pages 243–270

Edited by John D. Hedengren

How to cite this chapter. "A Supercritical CO_2 Brayton Simulator for Three-Layer Mid-Ranging Control," in *Override, Split-Range, and Constraint Control*, John D. Hedengren, Ed., vol. 4 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 10, pp. 243–270.

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

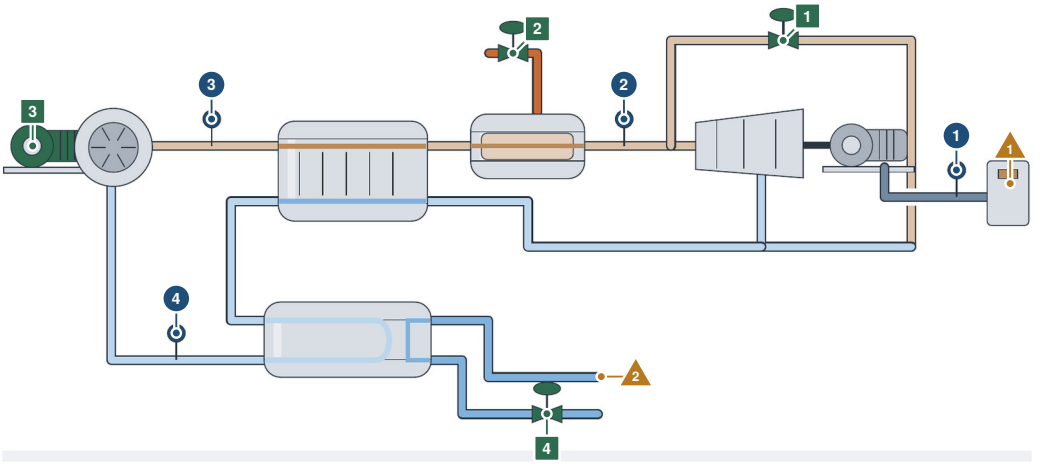
The app for this chapter
https://apmonitor.com/apps/sco2-brayton/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v04-override-constraint/#sco2-brayton

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

CHAPTER 10

A Supercritical CO_2 Brayton Simulator for Three-Layer Mid-Ranging Control



The controller moves 1 bypass valve around the turbine 2 fuel to the heater 3 compressor speed 4 cooling water valve

The controller reads 1 net electric power 2 turbine inlet temperature 3 compressor outlet pressure 4 compressor inlet temperature

What disturbs it 1 power the grid asks for 2 warm cooling water

Notice that three handles, the bypass, the heater and the compressor speed, all move one power output, while the cooler on the cold side sets how much work the compressor needs. Illustration, not to scale.

IN PLAIN TERMS **Three Handles on One Power Plant**

The main idea. This power plant pushes carbon dioxide around a closed loop. The gas is squeezed until it is almost a liquid, so the compressor needs little work. Three handles change the power: a fast valve that wastes energy, a compressor speed that costs nothing, and a slow fire that burns fuel.

The challenge. The controller uses the fast handle first and then passes the work down to the slower, cheaper ones. When the cooling water warms on a hot afternoon, the fluid swells and the power falls, and no controller can win it back.

Why it matters. Done well, the plant follows the grid quickly while it burns less fuel.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on density (Section 10.1), the two history segments (Section 10.2), the try-it segments and the fact on the dispatch cut (Section 10.6) and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 10.6), Technology in Practice (Section 10.7), Limitations (Section 10.8) and the open challenge

The control question A supercritical carbon dioxide power block has one variable the grid asks for, net power, and three actuators that order one way by speed and another by cost. Which structure delivers a dispatch that the obvious architectures cannot reach at all, and what does a third mid-ranging layer buy that no tuning can?

Learning objectives

- Rank the turbine bypass, the compressor speed and the firing rate by response time and by what each costs, from their measured gains.
- Compute compressor work and mass flow near the critical point and explain why the compressor inlet loop is worth more than the hierarchy.
- Explain why a valve position controller reads the unsaturated demand of the loop below it, and size the error it then sees.
- Configure three-layer mid-ranging and judge the resting bypass position as a purchase of responsiveness with fuel.
- Attribute the gain of a model predictive controller to tuning, to its third handle and to its model of the inner loops.

Prerequisites Sections 1.4.7 (mid-ranging), 1.4.5 (saturation), 1.4.4 (external reset), 1.6.2 (velocity form) and 1.6.4 (scenarios). Section 1.6.3 is needed only for the MPC.

The app https://apmonitor.com/apps/sco2-brayton/

See also Volume 3, Chapter 2 for structure chosen before tuning, Volume 3, Chapter 3 for the inner-loop speed a cascade needs and Volume 3, Chapter 5 for a megawatt power cycle on one fluid inventory. Volume 17, Chapter 4 for one variable handed between two actuators, Volume 17, Chapter 2 for an exhaust bypass on two clocks and Volume 17, Chapter 6 for a heat source with a lag of its own.

10.1 Introduction

A closed supercritical carbon dioxide Brayton cycle is attractive because the fluid is compressed close enough to its critical point to behave almost like a liquid, so the compressor absorbs a small fraction of the turbine work and the cycle reaches steam-plant efficiency in a machine an order of magnitude smaller [1, 2]. That same proximity to the critical point is the source of every control problem the application poses.

> **DID YOU KNOW? Three Degrees That Halve the Density**
>
> Just above the critical point, warming carbon dioxide from 32 to 35 °C at the same pressure halves its density. In a model of a recompression cycle for a sodium-cooled reactor, raising the pre-cooler outlet from 32.0 to 33.2 °C lowered the thermal efficiency from 39.32 to 38.68 % [3].

The grid asks for net electrical power, and three actuators move it. The turbine bypass diverts flow around the turbine, so the power falls in about eight seconds and every watt of the reduction is discarded as heat into the recuperator. The compressor speed changes the mass flow and the discharge pressure together, so it raises the cycle efficiency as it raises the power, and it costs nothing. The firing rate is the actuator that adds heat to the cycle, and it reaches sixty-three percent of a turbine inlet temperature change in 16.6 min. The fastest actuator is therefore the most expensive one and the slowest is the one that buys the fuel, which is the ordering that makes mid-ranging the natural structure [4]. Section 10.6 measures what each single-actuator structure fails to do, and the structure that works applies mid-ranging three times in series. A model predictive controller (MPC) is offered as well, moving the bypass and the two inner-loop setpoints as one plan, a method with a long industrial record [5] that has also been analyzed for a supercritical carbon dioxide power converter [3]. Section 10.6.8 measures it against the hierarchy as shipped and as retuned.

The plant brings a specific list of process dynamics and control principles together. Time-scale separation across two orders of magnitude, cascade structure and the condition a secondary loop

must satisfy, valve position control and the choice of a resting position, actuator saturation with anti-windup, the handover of a saturated signal between controllers, constraint reordering rather than constraint removal, and the difference between a controlled variable that is held and a plant that is healthy. Each is standard material in a first course [6–8], and this plant supplies a setting in which they arrive together. The physical fundamentals underneath the model are equally standard, because the cycle is an energy balance around six stations, a compressor and a turbine characteristic, three exchangers with thermal mass and one strongly nonlinear property variation near the critical point. Industrially the behavior matters because a power block that follows dispatch by throwing work away earns less than one that follows it by moving pressure, and because the loop with the largest return on this plant is the one that is usually specified by a mechanical engineer and never tuned again.

This chapter works with a supercritical carbon dioxide power block simulator in which one controlled variable is reached by three actuators of very different speed and cost. It quantifies what each additional layer of the mid-ranging hierarchy buys and what it cannot buy. Section 10.2 places the model against the cycle and control literature, and Section 10.3 gives the equations and constants needed to rebuild the simulator. Section 10.4 measures the three actuators and the compressor inlet loop in open loop. Section 10.5 describes the seven PID loops in three layers, the handover of a saturated demand and the MPC option. Section 10.6 reports the measured scenarios, from the architecture comparison to the graded dispatch day and the MPC against the hierarchy. Section 10.7 describes the hardware of such a plant. Sections 10.8 to 10.10 state the limitations, the three results that carry past this plant and the conclusion.

10.2 Background

Supercritical carbon dioxide cycles have been studied since the 1960s and returned to attention as compact power conversion for nuclear, concentrating solar and waste heat duty. Reviews of the technology and of the cycle layouts that have been proposed agree on the source of the advantage, which is that compression near the critical point costs a small fraction of what compression of a gas costs, and on the price of it, which is that the same region has property gradients steep enough to make both design and control difficult [1, 2]. The reference properties themselves come from the Span and Wagner equation of state [9], and the specific heat of carbon dioxide varies by an order of magnitude across the cooler and the cold end of the recuperator.

> ⧖ HISTORY **A Test Loop at the Critical Point, 2008 to 2010**
>
> Earlier studies of the supercritical carbon dioxide Brayton cycle had been largely analytical. A United States national laboratory then built a small test loop on a skid of about 2 by 3 m, with a motor-driven radial compressor of about 50 kW, to study the stability and control of the cycle near the critical point. Assembly finished in May 2008, and commissioning ran through the 2009 fiscal year. The machine reached 65 000 rpm, flows above 4.1 kg/s and pressure ratios just over 1.65, at inlet densities near 70 % of that of water. Its compressor ran on both sides of the saturation curve, near choke and even in surge. The report named the control problem this chapter models: near the critical point a very large heat flow is needed to move the temperature by as little as 1 K [10].

Dynamic studies of these cycles concentrate on load following, because the recuperator makes the plant slow and the applications are ones where the heat source is not dispatchable. Dynamic models coupled to thermal storage have been built for solar-driven layouts [11], transient models of printed circuit recuperators show behavior a single lumped wall cannot reproduce [12], and wide-range load following strategies for recompression layouts treat the recompression split fraction as the principal handle [13, 14]. The application implements a simple recuperated cycle with no recompression branch, so the handles available to it are firing rate, compressor speed, coolant flow and a turbine bypass. An organic Rankine cycle of the same megawatt scale, with its own coupling through one fluid inventory, is the subject of Volume 3, Chapter 5.

Pause. Try Exercise 10.1 before reading on.

Mid-ranging, also called valve position control, pairs a fast actuator of limited range with a slow actuator of large range and returns the fast one toward a resting position so that its authority is available again (Section 1.4.7). The structure is old in the process industries and its tuning has been given a model-based treatment [4, 15]. Chapter 9 applies it to chilled water from a gas turbine. What this plant adds is a third layer and an explicit cost ordering, because on this plant the fastest actuator is also the only one that destroys work, and the slowest is not the cheapest. Predictive control is the usual alternative to such a ladder, because one optimization can spread a demand over several handles by speed and by cost at once [5].

> **HISTORY A Ladder of Handles for Load Following, 2011**
>
> In 2011 a national laboratory report on supercritical carbon dioxide cycles for fast reactors set out a load-following strategy that ranks its handles much as this chapter does. Moving carbon dioxide out to inventory tanks gave the highest efficiency at reduced load, but only over the range the tank volume allowed, from 50 to 90 % load. The turbine bypass was the simplest and fastest handle, almost instantaneous because it acts through pressure rather than temperature, so it made the final, fine adjustment of generator output. It was also the least efficient, because the heat entering the cycle did not fall with the output. The report therefore added a throttle valve at the turbine inlet below 50 % load, with its settings chosen so that the bypass action stays small [16].

10.3 Process Model

The cycle and the control hierarchy are shown in Figure 10.1. The model carries a specific enthalpy for each of the six cycle stations, two metal wall temperatures, a cooler effectiveness, a heat source temperature, a firing-rate actuator state and three actuator positions, which is fourteen states in all. The clock is in minutes, temperatures in °C, pressures in MPa and specific enthalpies and specific work in kJ/kg.

10.3.1 The Design Point

At design the compressor runs at 85 % speed against a 15 °C coolant supply, the firing rate is unity and the bypass is shut. That gives 5.499 kg/s of carbon dioxide circulating between 6.78 and 20.58 MPa, with station temperatures of 20.50, 38.62, 273.86, 593.10, 384.25 and 64.42 °C around the loop. The turbine produces 201.74 kJ/kg against 20.84 kJ/kg of compression, the recuperator transfers 2040.6 kW at an effectiveness of 0.9254, the primary heater supplies 2269.0 kW, and the net electrical output is 929.0 kW at 40.94 % thermal efficiency.

The recuperator carries 47.4 % of the heat the cycle needs. That is the number that makes the cycle efficient and the number that makes Section 10.6.6 uncomfortable, because a controller that holds the power hides the loss of it completely.

10.3.2 The Compressor and the Near-Critical Region

Compression work is proportional to the pressure ratio less one and to a penalty that rises steeply as the inlet temperature leaves the near-critical region, as shown in Equations (10.1) and (10.2).

$$w_c = K_{c0}\left[1 + (\Pi_{\max} - 1)\frac{u}{1+u}\right]\left(\frac{P_2}{P_1} - 1\right), \tag{10.1}$$

$$u = \left(\frac{\max(0,\ T_1 - T_c)}{\Delta T_c}\right)^{n_c} \tag{10.2}$$

where T_c is the temperature above which the fluid begins to expand, ΔT_c is the rise that produces half of the maximum penalty and $\Pi_{\max}$ is that maximum. The penalty is written as a saturating function rather than a bare power law so that a transient excursion cannot become a runaway that no real

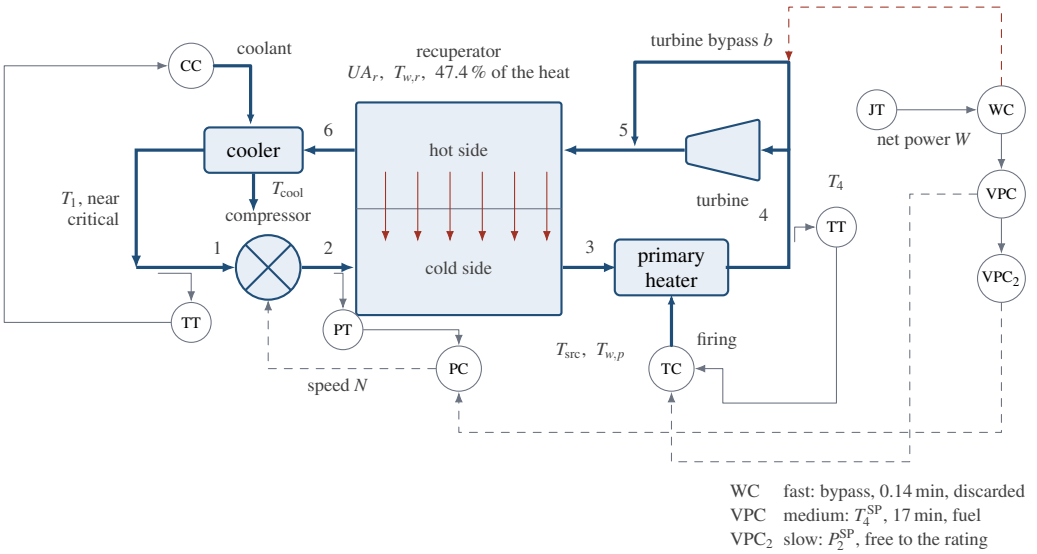


Figure 10.1 The simulated cycle and its control hierarchy. Net power is taken by the turbine bypass in seconds (WC), a valve position controller returns the bypass toward its resting position by moving the turbine inlet temperature setpoint (VPC), and a second valve position controller returns the turbine inlet temperature to its reference by raising the compressor discharge pressure setpoint (VPC$_2$). Those two setpoints go to the firing controller (TC) and the speed controller (PC). The compressor inlet temperature loop (CC) is not part of the hierarchy and is worth more than all of it.

machine has. Mass flow falls in the same variable through a density ratio, and pressure ratio rises with speed, as shown in Equations (10.3) and (10.4).

$$\frac{P_2}{P_1} = \Pi_0 \left(\frac{N}{N_0} \right)^{n_\Pi}, \tag{10.3}$$

$$\dot{m} = \dot{m}_0 \, f_i \, \frac{N}{N_0} \, \mathrm{e}^{-k_\rho (T_1 - T_{1,0})} \tag{10.4}$$

The exponential density ratio in (10.4) is clipped to 0.40 to 1.55. Speed therefore raises the mass flow and the pressure ratio together, which is why its power gain is large and its efficiency gain is positive. Inlet temperature moves (10.1) and (10.4) in opposite directions at the same time, which is why the compressor inlet loop is worth what Section 10.6.5 measures.

Turbine specific work is an ideal expansion with an isentropic efficiency, as shown in Equation (10.5).

$$w_t = f_t \, \eta_t \, c_{p,p} \, (T_4 + 273.15) \left(1 - \Pi_t^{-\gamma} \right) \tag{10.5}$$

The work is proportional to the absolute inlet temperature, which is why the turbine inlet temperature is the plant's other power handle and why the metal limit is what binds that handle.

10.3.3 Stations, Duties and the Exchanger Walls

Pressures around the loop follow from the compressor characteristic and four fixed pressure drops, as shown in Equation (10.6).

$$\begin{aligned} P_2 &= P_1 \, \Pi(N), \quad & P_3 &= P_2 - \Delta P_{23}, \quad & P_4 &= P_3 - \Delta P_{34}, \\ P_5 &= P_1 + \Delta P_{51}, \quad & P_6 &= P_5 - \Delta P_{56} \end{aligned} \tag{10.6}$$

Table 10.1 Controlled, manipulated and disturbance variables with the states that carry them. The six station enthalpies are numbered around the loop from the compressor inlet.

Symbol	Meaning	Unit	Range
Controlled variables			
W	net electrical power	kW	300–1300
T_4	turbine inlet temperature	°C	470–618
P_2	compressor discharge pressure	MPa	13.5–23.8
T_1	compressor inlet temperature	°C	15–32
Manipulated variables			
φ	firing rate	—	0.45–1.40
N	compressor speed	%	55–102
u_c	coolant valve demand	%	0–100
b	turbine bypass position	%	0–78
Disturbances			
T_{cool}	coolant supply temperature	°C	5–34
f_r	recuperator condition	—	0.40–1.05
f_p	primary heater condition	—	0.50–1.10
f_w	cooler condition	—	0.50–1.10
f_t	turbine condition	—	0.80–1.03
f_i	carbon dioxide inventory	—	0.80–1.10
ΔT_{src}	heat source temperature bias	°C	−80 to 60
States			
$h_1 \dots h_6$	station specific enthalpies	kJ/kg	50–2200
$T_{w,r}$	recuperator wall temperature	°C	—
$T_{w,p}$	heater wall temperature	°C	—
ε	cooler effectiveness	—	0.02–0.995
T_{src}	heat source temperature	°C	—
φ_a	firing rate after the actuator lag	—	0–2
N, u_c, b	actuator positions	—	—

The compressor inlet pressure P_1 is set by the carbon dioxide inventory and is treated as a constant. Station temperatures are recovered from the enthalpies by a constant specific heat per segment, as shown in Equation (10.7).

$$\begin{aligned} T_1 &= T_{1,n} + \frac{h_1 - h_{1,n}}{c_{p,w}}, & T_2 &= T_1 + \frac{h_2 - h_1}{c_{p,c}}, \\ T_3 &= T_2 + \frac{h_3 - h_2}{c_{p,rc}}, & T_4 &= T_3 + \frac{h_4 - h_3}{c_{p,p}}, \\ T_5 &= T_{5,n} + \frac{h_5 - h_{5,n}}{c_{p,rh}}, & T_6 &= T_5 - \frac{h_5 - h_6}{c_{p,rh}} \end{aligned} \tag{10.7}$$

The three exchangers are metal walls with thermal mass rather than algebraic effectiveness relations, because their inertia is a large part of what the controller is fighting. The recuperator wall is heated by the mean of the two hot-side temperatures and cooled by the mean of the two cold-side temperatures, and the heater wall is heated by the source and cooled by the cold stream, as shown in Equations (10.8) and (10.9).

$$\tau_{w,r}\dot{T}_{w,r} = \tfrac{1}{2}(T_5 + T_6) - 2T_{w,r} + \tfrac{1}{2}(T_2 + T_3), \tag{10.8}$$

$$\tau_{w,p}\dot{T}_{w,p} = (T_{\text{src}} - T_{w,p}) - g_p(T_{w,p} - T_3) \tag{10.9}$$

The primary-heater conductance slider is limited to 1.3 kJ/(kg °C) to 4 kJ/(kg °C). The lower limit allows for the heater-fouling multiplier, which can halve the conductance. A minimum of 0.7 would still permit sub-zero cycle temperatures in this surrogate model. At the 1.3 minimum, with fouling

factors of 0.5, 0.75 and 1, all twelve cases stay above 17.2 °C, whether the slider is set before Run or moved during the run. The limit bounds the slider's calibrated range and changes neither the nominal design nor the shipped scenarios.

The duties each side of those walls are conductances times the driving difference, as shown in Equation (10.10).

$$\begin{aligned} q_r &= UA_r f_r\left(T_{w,r} - \tfrac{1}{2}(T_2 + T_3)\right), \\ q_p &= UA_p f_p (T_{w,p} - T_3) \end{aligned} \tag{10.10}$$

The recuperator adds to the inertia of the firing path rather than creating it. Part of the added heat returns through the recuperator and has to cross its wall again, but the firing actuator, the source, the heater wall and the heater fluid already form a chain of lags of 2, 4, 5 and 3 min. In a counterfactual run with the recuperator wall lag at its 0.5 min minimum, the time for net power to reach sixty-three percent of a firing step falls only from 17.0 to 15.8 min.

The cooler is written differently. Its state is an effectiveness rather than a temperature, and its duty is bounded so that the outlet approaches the coolant supply temperature and can never pass it, as shown in Equation (10.11).

$$q_w = \varepsilon\,(h_6 - h_{\text{cool}}), \qquad h_{\text{cool}} = h_{1,n} + c_{p,w}(T_{\text{cool}} - T_{1,n}) \tag{10.11}$$

The coolant valve is characterized. A cooler's effectiveness saturates hard in coolant flow and this one is already 91 % effective at design, so an uncharacterized valve would hand the controller a loop gain that varies roughly twenty to one across its travel. The trim is chosen so that the controller output is the effectiveness on a linear scale and the loop gain is constant, as shown in Equation (10.12), with the physical flow and therefore the cubic pump power recovered afterward.

$$\varepsilon_v = f_w\left[\varepsilon_{\text{lo}} + (\varepsilon_{\text{hi}} - \varepsilon_{\text{lo}})\frac{u_c}{100}\right], \qquad \frac{\dot{m}_c}{\dot{m}_{c,\text{ref}}} = \frac{-\ln(1-\varepsilon)}{\text{NTU}_0} \tag{10.12}$$

That choice is not visible in the animation and it is the difference between a loop that can be tuned once and a loop that oscillates at one end of its range and crawls at the other.

10.3.4 State Equations

Each station enthalpy is a first-order tank driven by the work or the duty applied to it, as shown in Equation (10.13).

$$\begin{aligned} \tau_{\text{whx}}\dot{h}_1 &= (h_6 - q_w) - h_1, & \tau_{\text{comp}}\dot{h}_2 &= (h_1 + w_c) - h_2, \\ \tau_{rc}\dot{h}_3 &= (h_2 + q_r) - h_3, & \tau_{\text{phx}}\dot{h}_4 &= (h_3 + q_p) - h_4, \\ \tau_{\text{turb}}\dot{h}_5 &= h_4 - (1-b)w_t - h_5, & \tau_{rh}\dot{h}_6 &= (h_5 - q_r) - h_6 \end{aligned} \tag{10.13}$$

The heat source, the firing actuator, the cooler and the three actuator positions are first-order lags toward their commands, as shown in Equation (10.14).

$$\begin{aligned} \tau_{\text{fuel}}\dot{\varphi}_a &= \varphi_{\text{cmd}} - \varphi_a, \\ \tau_{\text{src}}\dot{T}_{\text{src}} &= T_{\text{src},0} + s_g(\varphi_a - 1) + \Delta T_{\text{src}} - T_{\text{src}}, \\ \tau_{w,c}\dot{\varepsilon} &= \varepsilon_v(u_c) - \varepsilon, \\ \tau_N\dot{N} &= N_{\text{cmd}} - N, \\ \tau_c\dot{u}_c &= u_{c,\text{cmd}} - u_c, \\ \tau_b\dot{b} &= b_{\text{cmd}} - b \end{aligned} \tag{10.14}$$

Electrical output is the turbine work that is not bypassed, less the compressor work and the coolant pump power, as shown in Equation (10.15).

$$W = \dot{m}\,\eta_g w_t (1-b) - \frac{\dot{m}\,w_c}{\eta_m} - P_{\text{aux}}, \qquad \eta = \frac{W}{\dot{m}\,q_p} \tag{10.15}$$

with $P_{\text{aux}} = P_{\text{aux},0}(\dot{m}_c/\dot{m}_{c,\text{ref}})^3$ from (10.12), where $\dot{m}_c$ is the physical coolant flow. After every step the six enthalpies are clipped to 50 kJ/kg to 2200 kJ/kg, the firing state to 0 to 2, the speed to 20 % to 110 %, the coolant demand to 0 % to 100 %, the bypass to 0 % to 100 % and the cooler effectiveness to 0.02 to 0.995.

10.3.5 Measurements

Four measurements close the loops, and each is a first-order filter on the true value with optional noise, updated at the controller scan. The turbine inlet thermocouple has a time constant of 0.35 min, the compressor inlet thermocouple 0.25 min, the pressure transmitter 0.05 min and the power transducer 0.10 min. Noise on the two thermocouples and the power transducer is band-limited with correlation times of 0.5, 0.4 and 0.2 min, the compressor inlet carrying 15 % of the thermocouple amplitude, and it is switched off by default and in every scenario. Nothing on the plant measures the mass flow, the recuperator effectiveness or the thermal efficiency, all three of which are computed for display from the states.

Pause. Try Exercise 10.3 before reading on.

10.3.6 Parameters

Table 10.2 lists the constants needed to rebuild the simulator. Those marked hidden are calibration references and are not offered for adjustment, while the five specific heat surrogates can be changed with the other constants.

10.3.7 Initialization and Numerics

Initialization solves a three-by-three damped Newton system in the compressor inlet, recuperator cold outlet and turbine inlet enthalpies, with both wall temperatures and the cooler effectiveness eliminated in closed form. Successive substitution fails here, because the recuperator returns about four times what a first pass puts into it and the iteration diverges. The reason is worth stating because it is also why the plant is slow. In the dynamic model the same feedback is broken by the wall states having their own time constants, which is what makes the plant stable and the arithmetic unstable. Three inverse calculations use the same solve, returning the firing rate that holds a chosen turbine inlet temperature, the coolant demand that holds a chosen compressor inlet temperature and the firing rate that produces a chosen net power. Manual mode offers the first two as one-press helpers.

The engine of Section 1.6.1 integrates the model by fourth-order Runge-Kutta at a fixed 0.02 min (1.2 s). All seven controllers execute on a 0.1 min (6 s) scan, and the historian records ten samples per simulated minute.

10.3.8 Assumptions

Carbon dioxide properties are constant specific heats per segment, with the near-critical behavior carried entirely by (10.1) and (10.4) as fitted functions rather than derived from an equation of state [9]. There is no recompression branch and no inventory control loop, so the flow split and the charge are parameters rather than actuators. Turbomachinery is quasi-steady, with no shaft dynamics, no surge line and no windage. Each exchanger is a single lumped wall. The heat source is a first-order lag toward a temperature proportional to firing rate, which suits a fired heater. The compressor inlet pressure is fixed by the inventory and does not move with the loop temperature. Pressure drops are constants rather than functions of flow. Section 10.8 states where each of these fails.

10.4 Actuator Characteristics and Open-Loop Response

10.4.1 Three Actuators for One Variable

Table 10.3 reports the steady-state gain and the open-loop response of net power to each actuator, measured from the design point in manual with the other actuators fixed. Response times are counted

Table 10.2 Model constants with their default values.

Symbol	Meaning	Value	Unit
$\dot{m}_0$	design mass flow	5.5	kg/s
N_0	design compressor speed (hidden)	85	%
Π_0	design pressure ratio	3.0354	—
n_Π	ratio exponent on speed	1.35	—
K_{c0}	work per unit pressure ratio	10.2388	kJ/kg
T_c	near-critical threshold	22	°C
$\Pi_{\max}$	largest compression penalty	2.6	—
ΔT_c	half-penalty rise	8.0	°C
n_c	penalty exponent	1.8	—
k_ρ	density loss per degree	0.024	1/°C
η_t	turbine efficiency	0.88	—
γ'	expansion exponent	0.21875	—
η_g, η_m	generator, motor efficiency	0.96, 0.95	—
$P_{\text{aux},0}$	pump power at design flow	15	kW
P_1	compressor inlet pressure	6.78	MPa
ΔP_{23}, ΔP_{34}	cold side, heater drop	0.20, 0.20	MPa
ΔP_{51}, ΔP_{56}	exhaust rise, hot side drop	0.30, 0.20	MPa
$T_{\text{src},0}$	source temperature at design	700	°C
s_g	source temperature per unit firing	280	°C
UA_r	recuperator conductance	10.8986	kJ/(kg °C)
UA_p	heater conductance	1.3478	kJ/(kg °C)
g_p	heater wall to fluid coupling	0.3921	—
ε_{lo}, ε_{hi}	cooler effectiveness, shut, open	0.55, 0.985	—
NTU_0	cooler NTU at design flow (hidden)	2.3772	—
$\dot{m}_{c,\text{ref}}$	coolant flow at design (hidden)	70	%
τ_{fuel}, τ_{src}	firing actuator, source lag	2.0, 4.0	min
τ_{rc}, τ_{rh}, $\tau_{w,r}$	recuperator cold, hot, wall lag	2.0, 2.0, 4.0	min
τ_{phx}, $\tau_{w,p}$	heater fluid, wall	3.0, 5.0	min
τ_{whx}, $\tau_{w,c}$	cooler fluid, wall	3.0, 4.0	min
τ_{comp}, τ_{turb}	compressor, turbine outlet	0.50, 0.20	min
τ_N, τ_c, τ_b	speed, coolant, bypass lag	0.35, 0.40, 0.15	min
$c_{p,w}$	cooler surrogate	4.1982	kJ/(kg °C)
$c_{p,c}$	compressor surrogate	1.1501	kJ/(kg °C)
$c_{p,rc}$, $c_{p,rh}$	recuperator surrogates	1.5774, 1.1602	kJ/(kg °C)
$c_{p,p}$	heater surrogate	1.2924	kJ/(kg °C)
$T_{4,\max}$	turbine metal limit	620	°C
$T_{4,\text{trip}}$	overtemperature alarm	645	°C
$P_{2,\max}$	high-pressure rating	24.0	MPa
$T_{1,\max}$	compressor inlet alarm	32.0	°C
Defaults for the operator-set inputs			
W^{SP}	dispatch setpoint	930	kW
T_4^{SP}, T_4^{ref}	turbine inlet setpoint, reference	593, 593	°C
T_1^{SP}	compressor inlet setpoint	20.5	°C
P_2^{SP}	discharge pressure setpoint	20.58	MPa
b^{ref}	bypass resting position	3	%
φ, N, u_c, b	manual actuator positions	1.0, 85, 82.6, 0	—

Table 10.3 Measured open-loop response of net power to each actuator, from the design point with the other actuators fixed.

Actuator	Steady gain	t_{63} (min)	Cost
Turbine bypass	−10.2 kW/%	0.14	all of it
Compressor speed	+22.1 kW/%	0.33	none
Firing rate	+305 kW per unit	17.02	fuel

from the controller scan that applies the step, so they exclude the up to 0.1 min a change waits for that scan.

Table 10.3 is the whole design problem. The fast actuator is the one that throws work away, because opening the turbine bypass diverts flow around the turbine and the reduction in power is discarded as heat into the recuperator rather than saved. The efficient actuator is the compressor

Table 10.4 Measured open-loop gains and response times of the four measurements, from the design point in manual.

Pairing	Gain	t_{63} (min)	t_{95} (min)
Bypass → net power	−10.2 kW/%	0.14	0.36
Speed → net power	+22.1 kW/%	0.33	0.85
Speed → discharge pressure	0.328 MPa/%	0.35	1.05
Firing → net power	305 kW per unit	17.02	35.85
Firing → turbine inlet	258 °C per unit	16.61	34.87
Coolant → compressor inlet	−0.270 °C/%	8.04	17.21

Table 10.5 Closed-loop steady state at four coolant supply temperatures with the compressor inlet controller in service. The coolant valve saturates between 19 and 22 °C.

Coolant (°C)	T_1 (°C)	Valve (%)	$\dot{m}$ (kg/s)	w_c (kJ/kg)	Net (kW)
15	20.50	82.6	5.500	20.84	929.0
19	20.50	97.4	5.500	20.84	890.9
22	22.82	100.0	5.202	21.38	807.6
29	29.80	100.0	4.400	37.13	597.4

speed, which moves the mass flow and the discharge pressure together through (10.3) and (10.4) and therefore raises cycle efficiency as it raises power, from 38.78 % at 78 % speed to 42.61 % at 92. The patient actuator is the firing rate, which reaches sixty-three percent of a power change in 17.0 min and ninety-five percent in 35.9 min. The ratio of the slowest to the fastest is 122.

Table 10.4 sets the loop hierarchy. In open loop the discharge pressure responds nearly fifty times faster than the turbine inlet temperature, which is the separation a cascade needs and the reason the two supervisory layers can be tuned independently of the inner loops they write to. The firing path contains no transport delay. Its apparent dead time comes from the chain of first-order elements between the fuel valve and the turbine inlet, which is the actuator lag, the source lag, the heater wall and the heater fluid in series. The first two percent of the turbine inlet response appears 3.3 min after the step, and a two-point first-order fit gives an apparent dead time of 5.9 min with a time constant of 10.8 min.

10.4.2 The Compressor Inlet Loop

Compression work and mass flow move in opposite directions in the same variable, so the compressor inlet temperature is the plant's most valuable number. Table 10.5 reports the closed-loop steady state at four coolant supply temperatures with the compressor inlet controller in service and everything else at design.

Fourteen degrees of coolant supply temperature, which is the difference between a spring morning and a summer afternoon on the same cooling tower, nearly doubles the compressor work and removes 331.6 kW, or 35.7 %, of the plant's output. The first four degrees cost 38.1 kW with the compressor inlet held exactly on setpoint, because holding it costs coolant flow and the pump power of (10.15) is cubic in flow. The remaining ten cost 293.5 kW, because the valve has no travel left and the inlet temperature then follows the supply. Nothing in the control hierarchy recovers the loss, because it is not a setpoint error. The cycle has become a different cycle.

10.5 Controller Design

10.5.1 Velocity Form PID Algorithm

Every loop on the plant runs the velocity-form PID law of Section 1.6.2 [17]. Derivative action is available and is set to zero in every loop, because no measurement on this plant is fast enough to

Table 10.6 The seven controllers, their actions and the tunings that Scenarios 3 to 10 load and the challenge suggests. No loop uses derivative action.

Controlled	Manipulated	a	K_c	τ_I (min)
Inner loops				
T_4	firing rate	+1	0.006 per °C	12
P_2	compressor speed	+1	1.2 %/MPa	1.0
T_1	coolant valve	−1	3 %/°C	8
Supervisory layer				
W	bypass	−1	0.09 % per kW	1.2
W	T_4^{SP} (no fast path)	+1	0.60 °C per kW	30
b	T_4^{SP}	+1	3.0 °C/%	30
T_4^{SP}	P_2^{SP}	−1	0.035 MPa/°C	60

reward it. The error and the increments are given by Equations (10.16) and (10.17).

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{10.16}$$

$$\Delta u_k^P = K_{c,k}\,e_k - K_{c,k-1}\,e_{k-1}, \qquad \Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \tag{10.17}$$

where $\Delta t_c = 0.1$ min (6 s) is the scan period common to every loop and a is the action. The increments are combined and the demand is clipped to the actuator span, as shown in Equation (10.18).

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P \\ c_k &= s_k + I_{k-1} \\ u_k &= \min(u_{\max},\ \max(u_{\min},\ c_k)) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{10.18}$$

Here s_k is the running sum S_k of Equation (1.13), c_k is the output before limiting, v_k of Equation (1.14), with the bias folded into the integral, and the integral I used at scan k was accumulated through scan $k-1$. At reset I is seeded with the starting output, which is the manual actuator position for the inner loops and the bypass loop and the present turbine inlet temperature or discharge pressure for the supervisors, so every loop starts where the plant is. Anti-windup is back-calculation with a tracking time $T_t = \tau_I$. While the output is clipped, the last term removes $\Delta t_c/\tau_I$ of the clipped excess each scan, which is $1/12$ for the bypass loop and $1/600$ for the pressure supervisor, while the running sum s_k is never truncated. Transfer between Manual and Auto is bumpless as Section 1.6.2 describes. A supervisor that the chosen architecture leaves out of service runs the same tracking routine against the value in force, so a change of architecture is bumpless too. Two quantities are needed from the algorithm rather than from the actuator. The applied output u_k drives the plant, and the unsaturated demand c_k is what one controller hands to another, for the reason Section 10.5.3 gives.

10.5.2 Seven Loops in Three Layers

Table 10.6 lists every controller, its action and its tuning. The three inner loops hold a turbine inlet temperature, a compressor discharge pressure and a compressor inlet temperature, and closed they reach sixty-three percent of a setpoint move in 10.4, 2.3 and 8.3 min. Closing the loops shrinks the ratio of the first two to 4.5, against nearly fifty in open loop (Table 10.4). The supervisors above them carry integral times of 30 and 60 min, so each is slower than the loop it writes to by a factor of nearly three or more, which is the only structural requirement a cascade makes [6]. Volume 3, Chapter 3 measures what happens when an inner loop is too slow for that requirement.

The supervisory layer is three mid-ranging controllers in series, as shown in Equations (10.19) to (10.21).

$$b = \mathrm{PID}\left(W^{\mathrm{SP}}, W\right), \tag{10.19}$$

$$T_4^{\mathrm{SP}} = \mathrm{PID}\left(b^{\mathrm{ref}}, b^{\mathrm{raw}}\right), \tag{10.20}$$

$$P_2^{\mathrm{SP}} = \mathrm{PID}\left(T_4^{\mathrm{ref}}, T_4^{\mathrm{SP}}\right) \tag{10.21}$$

The signs repay attention. A bypass demand above its resting position means the thermal side is making more power than is wanted and discarding the excess, so the turbine inlet temperature should come down, and $a = +1$ in (10.20) delivers that. A turbine inlet setpoint above its own reference means the temperature is carrying load the pressure should be carrying, so the pressure setpoint should rise, and $a = -1$ in (10.21) delivers that. Both layers are switchable, so the application offers the three architectures that Section 10.6.1 compares.

10.5.3 Handing a Saturated Demand Upward

Equation (10.20) watches b^{raw}, which is the bypass controller's unsaturated demand c_k from (10.18), and not the applied bypass position u_k, as Section 1.4.7 recommends. The distinction is the difference between a hierarchy that works and one that works in only one direction.

The bypass is clamped shut at zero. When the plant is short of power the fast controller demands a negative bypass, which it cannot have, so the applied position sits at zero and reports a shortfall of exactly b^{ref} no matter how much power is missing. A valve position controller watching that signal receives the same constant error for a 5 kW deficit and a 500 kW deficit, and it walks the turbine inlet temperature up at the same slow rate in both cases. The unsaturated demand goes properly negative and states how much power the fast path wishes it could make. While the bypass is held shut against a steady shortfall, the back-calculation of (10.18) lets the integral decay and the demand settles at the proportional term $K_c(W - W^{\mathrm{SP}})$, which is –9 % for a 100 kW shortfall, so the handover carries magnitude information in both directions. An earlier build of the application handed over the clamped position, and the symptom was a plant that shed load in seconds and picked it up in hours.

Pause. Try Exercise 10.2 before reading on.

10.5.4 Model Predictive Control

Beside manual operation, the inner loops alone and the cascade, the supervisory layer can be a model predictive controller (MPC), which is never the default and which only scenario 11 loads. Section 1.6.3 describes the option in general. It replaces the bypass loop and both valve position controllers of Equations (10.19) to (10.21). Every $T_s = 0.5$ min, or five scans, it plans the bypass command b, the turbine inlet setpoint T_4^{SP} and the discharge pressure setpoint P_2^{SP} together over $N = 60$ steps (30 min) and applies the plan's first move alone [5]. A dispatch change between decisions is answered on the scan it arrives. The three inner loops of Table 10.6 stay in service and carry the setpoints out, the commands are bounded by the clamps the cascade uses (bypass 0 % to 78 %, inlet setpoint 470 °C to 618 °C and at most $T_{4,\max} - 2$ °C, pressure setpoint 13.5 MPa to 23.8 MPa), and every alarm stays on the run's own limits above both structures.

The controller's model describes the closed inner loops rather than the plant. Step tests of ±10 °C and ±0.5 MPa on scenario 4, with the tuned loops of Table 10.6 and the noise off, give the turbine inlet loop as a damped second-order response behind a lag and the pressure loop as a lead-lag, as shown in Equation (10.22).

$$\begin{aligned} T_4 &= \frac{\omega^2\, T_4^{\mathrm{SP}}}{(s^2 + 2\zeta\omega s + \omega^2)(\tau_3 s + 1)} + G_P(s)\, P_2^{\mathrm{SP}} + G_b(s)\, b \\ P_2 &= \frac{(as+1)\, P_2^{\mathrm{SP}}}{(\tau_1 s + 1)(\tau_2 s + 1)}, \qquad \tau_b \dot{x}_b = b - x_b \end{aligned} \tag{10.22}$$

Here s is the Laplace variable in 1/min, $\omega = 0.19668$ 1/min, $\zeta = 0.28393$, $\tau_3 = 4.4749$ min, $a = 1.0561$ min, $\tau_1 = 3.2096$ min, $\tau_2 = 0.4062$ min, and x_b is the bypass position behind the declared valve lag $\tau_b = 0.15$ min. The cross terms G_P and G_b have the form $s(\beta_2 s + \beta_1)/[(s^2 + 2\zeta_x \omega_x s + \omega_x^2)(s + 1/\tau_x)]$ with their own fitted constants, and carry the disturbance that the pressure setpoint and the bypass make on the turbine inlet before its loop removes it. Each loop has unit gain because each has integral action. The model reaches 63 % of an inlet setpoint step at 10.80 min, as the plant does, and overshoots by 26.7 % at 21.2 min against 26.3 % at 20.9 min, with an RMS residual of 0.61 % of the peak. The pressure loop reaches 63 % at 2.40 min with a residual of 0.26 %, and the cross terms fit to 1.51 and 1.78 % of their peaks.

Net power is Equation (10.15) itself, written with the declared defaults of Table 10.2, the design turbine condition and inventory, the speed that gives P_2 through (10.3) and the coolant pump's design power, never with the run's values. Its slopes change more than threefold across the handles' ranges (−3.6 to −13.7 kW/%, 0.27 to 1.54 kW/°C and 16.7 to 75.5 kW/MPa), so the map is relinearized at every decision, as shown in Equation (10.23).

$$\hat{W}_k = \bar{W} + g_b(x_{b,k} - \bar{x}_b) + g_T(\hat{T}_{4,k} - \bar{T}_4) + g_P(\hat{P}_{2,k} - \bar{P}_2) \tag{10.23}$$

The bars mark the estimates at the decision, $\bar{W}$ includes the power bias below, the compressor inlet temperature is held at its measured value over the horizon, and g_b, g_T and g_P are the partial derivatives of the map at that point. At design they are −10.65 kW/%, 1.230 kW/°C and 70.10 kW/MPa, against steady changes of −10.66, 1.223 and 70.04 in the same units on the plant with its inner loops closed.

The estimator runs every scan whichever structure is in charge, which lets the MPC take over from a current estimate. The model is driven by the commands applied, and an output bias on each of the three readings the shipped loops use is updated as shown in Equation (10.24).

$$d_{y,k} = d_{y,k-1} + \left(1 - e^{-\Delta t_c/\tau_y}\right)\left(y_k - \hat{y}_k - d_{y,k-1}\right) \tag{10.24}$$

Here y is the turbine inlet thermocouple, the pressure transmitter or the power transducer, $\hat{y}_k$ is the model's prediction of that reading for the scan (for power, through the transducer's 0.10 min lag), and τ_y is 1.0, 0.5 and 0.5 min in that order. A constant bias on each measured output removes the steady offset of a model that is wrong in a declared way [18], which is how a fouled recuperator, a worn turbine or a warm afternoon reaches the plan without being announced. The update is an output-bias filter, not a Kalman filter, and while the inner loops are open the model is reset onto the readings every scan.

A decision solves the quadratic program of Equation (10.25) on the model sampled at T_s, with the three commands held over 12 blocks of 1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15 and 15 steps, which with the slack makes 37 decision variables.

$$\begin{aligned}
\min_{\mathbf{u},s} \; & \sum_{k=1}^{N} \left(\hat{W}_k - W^*\right)^2 + \ell_b \sum_{k=1}^{N} b_k + \rho s^2 + \sigma s \\
& + \lambda \sum_{j=1}^{12} \left[\left(g_b \Delta b_j\right)^2 + \left(g_T \Delta T_{4,j}^{\mathrm{SP}}\right)^2 + \left(g_P \Delta P_{2,j}^{\mathrm{SP}}\right)^2\right] \\
& + r_T\, g_T^2 \sum_{k=1}^{N} \left(T_{4,k}^{\mathrm{SP}} - T_4^{\mathrm{ref}}\right)^2 \\
\text{s.t. } & \mathbf{u}_{\min} \le \mathbf{u}_j \le \mathbf{u}_{\max}, \qquad s \ge 0, \\
& \hat{T}_{4,k} \le T_{4,\max} - \delta_T + o\, e^{-kT_s/\tau_f} + s
\end{aligned} \tag{10.25}$$

Here $\mathbf{u}_j$ holds b, T_4^{SP} and P_2^{SP} over block j, Δ is the change from the block before (from the commands applied now for the first), $\lambda = 0.05$ prices a move by the square of the power it is worth, $r_T = 0.02$ returns the inlet setpoint toward the reference T_4^{ref} of (10.21), and $\ell_b = 2|g_b|\, e^*$ per percent prices

the bypass so that, at the margin, it opens only for the part of a power error above $e^* = 10\,\text{kW}$. The metal row holds the predicted inlet $\delta_T = 1\,^\circ\text{C}$ under the run's limit $T_{4,\max}$, o is the amount by which the estimated inlet already exceeds $T_{4,\max} - \delta_T$ and decays through a funnel with $\tau_f = 5\,\text{min}$, and $\rho = 10^5$ per $^\circ\text{C}^2$ and $\sigma = 10^4$ per $^\circ\text{C}$ price the slack s. The bounds are hard and the metal row is soft, sharing a single slack across all 60 steps, since the shared solver allots one to each soft output. The target W^* is the dispatch clipped to the steady powers the model reaches with every handle at its bound (the bypass shut, the inlet setpoint at the lower of its clamp and $T_{4,\max} - \delta_T$ and the pressure setpoint at its clamp, and the reverse), so a dispatch beyond reach does not pull the plan into the slack. With e^* fixed at 10 kW, under half the ±25 kW band, the horizon and λ were chosen once on scenario 11 (horizons of 20, 30 and 45 min give 795, 781 and 784 kW min), and nothing is retuned per case. The compressor inlet carries no row, because the three handles move it by at most 0.0064 °C per unit with the cooler saturated and no compared run came within 4.25 °C of its 32 °C alarm.

The controller keeps a plan only when the solver's status is optimal, which is tested even when the call succeeds, when every array has the expected size, when each planned command stays within its bounds and the first matches the move, and when commands, predictions, cost and slack are all finite. A specification test injects five such faults and requires each to be rejected. Otherwise the cascade takes the handles with the fast path the learner has set, because its bypass loop, power loop and both valve position controllers have tracked the applied commands with no preload, the historian marks the fallback, and the next valid plan takes the handles back. On scenario 11 without noise, forcing a rejection at 40 min moves the inlet setpoint by 5.0×10^{-4} °C and the pressure setpoint by 1.8×10^{-5} MPa on the hand-over scan. The first plan is made at $t = 0$, and none of the 5970 plans in the nine MPC runs of Table 10.11 was rejected. In Chrome a plan takes 2.05 to 2.21 ms on average over three cases, which is 0.41 to 0.44 % of one core when a simulated minute passes each second and 8.2 to 8.8 % at the top speed of twenty such minutes, and a run played at that speed logs the same rows as one stepped without the browser.

Three models stay separate, the plant on the run's parameters, the four filtered instruments, and the controller's model built from the declared defaults. The MPC reads the instruments, the commands applied, the dispatch and the operator settings, and of the run's parameters only the metal limit, which it treats as a protection limit in the row, the clip and the setpoint bound. A test replays the same instrument readings over a different hidden plant and requires the same estimate and plan. A learner's change to the plant is therefore mismatch whose steady part the biases remove, while a change to the inner-loop tunings is mismatch in the dynamics that nothing corrects, because the model assumes the loops at their identification values (K_c of 0.006 per °C and τ_I of 12 min for the turbine inlet, 1.2 %/MPa and 1.0 min for the pressure).

10.6 Results

Every number in this section was produced by driving the application's own model at its shipped constants and seed. The architecture comparisons of Sections 10.6.1 and 10.6.2 use one protocol. Each structure is selected before Run on Scenario 4, which starts from the design state and has no scripted events, run for 240 min at a dispatch of 930 kW so that the supervisory integrators are close to rest, and then stepped. Integral error is the integral of the absolute deviation of net power from the dispatch setpoint over the 185 min that follow, evaluated at the 0.02 min integration step, and the band is ±25 kW.

10.6.1 Three Architectures on a Dispatch Increase

The central experiment steps the dispatch from 930 to 1050 kW and runs the same plant under four configurations.

The first three rows of Table 10.7 do not merely track badly. They never arrive. Turbine inlet temperature is bounded by the 618 °C clamp on its own setpoint, which at the design speed and pressure is worth 959.5 kW, or 90 kW short of the request, and adding a bypass fast path in front

Table 10.7 Dispatch increase from 930 to 1050 kW. The band is ±25 kW, the turbine inlet setpoint is clamped at 618 °C and the turbine metal limit is 620 °C.

Structure	Rest (kW)	In band (min)	IAE (kW min)	Peak T_4 (°C)	Net (kW)
T_4 only, bypass shut	930.0	never	16968	624.4	959.5
T_4 only, bypass at 3 %	926.6	never	22827	618.0	926.6
Bypass and T_4	930.0	never	16751	618.3	959.5
Three deep	930.0	**6.8**	**453**	617.0	**1050.0**

Table 10.8 Dispatch reduction from 930 to 800 kW, same protocol.

Structure	In band at (min)	IAE (kW min)	Bypass loss (kWh)	Efficiency (%)
T_4 only, bypass shut	44.8	4998	0.0	38.70
T_4 only, bypass at 3 %	44.0	4999	91.8	38.43
Bypass and T_4	2.0	255	242.9	38.15
Three deep	**1.3**	**197**	107.0	**38.77**

of it changes the destination not at all, because a bypass can only subtract. What the fast path does change is the approach. Without it the temperature loop overshoots its own setpoint and takes the metal to 624.4 °C, spending 12.7 min above the 620 °C limit, while with it the peak is 618.3 °C and the limit is never reached. The second row is the same structure with three percent of the turbine work bypassed and nothing to close it, and it cannot hold even the 930 kW it started from.

The third layer supplies what the others lack, which is headroom. Raising the discharge pressure setpoint takes the compressor from 85.0 to 92.0 % speed and the discharge from 20.58 to 22.89 MPa, and (10.4) converts that into mass flow the turbine can expand. The plant reaches 1050.0 kW, enters the band in 6.8 min, peaks at 617.0 °C without reaching the metal limit, and settles at 41.76 % efficiency against 41.46 % for the temperature-only structure and 40.59 % for the bypassed one. The integral error is smaller by a factor of 37 than the two-layer structure and 50 than the bypassed one. The handover is slow to finish, because the turbine inlet temperature is still returning to its 593 °C reference five hours after the step while the power has been on setpoint since the first ten minutes.

10.6.2 The Same Comparison Downward

Reducing dispatch is a different problem, because the fast actuator now points the right way.

Table 10.8 is the ordinary mid-ranging result and then something better. The bypass fast path is twenty-two times faster into the band than the firing rate alone and it pays 242.9 kWh for the privilege, which is the trade the textbook describes. Adding the third layer is faster still, at 1.3 min, and it discards 107.0 kWh, which is less than half, because backing the compressor speed down to 80.5 % removes the power at the source and lets the bypass reclose to 2.0 %. The three-deep structure is simultaneously the fastest and the most efficient of the four, at 38.77 % against 38.70 % for the structure that discards nothing at all, which is unusual enough to be worth the extra controller. Figure 10.2 sets the two comparisons side by side.

> ✦ **DID YOU KNOW? The Fastest Structure Is Also the Most Efficient**
>
> From the simulator: on a dispatch cut from 930 to 800 kW, the three-deep structure is the fastest into the band, at 1.3 min, and also the most efficient, at 38.77 % against 38.70 % for the firing rate alone with the bypass shut. It discards 107.0 kWh, less than half the 242.9 kWh of the two-layer structure (Table 10.8).

10.6.3 Load Rejection

Scenario 9 opens the breaker at 10 min and demands 450 kW immediately. With the fast path in service the plant is within 50 kW of the new dispatch 3.4 min later and inside the ±25 kW band at

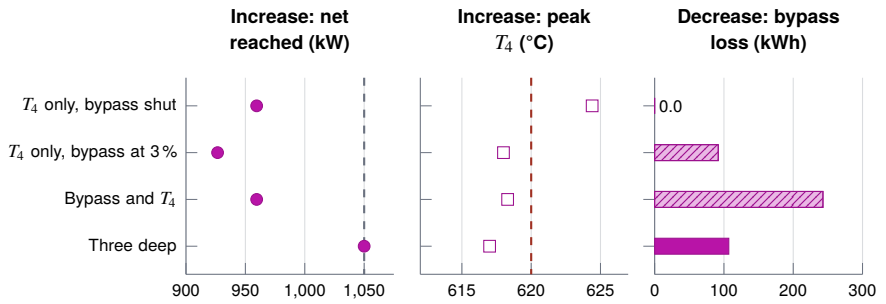


Figure 10.2 Tables 10.7 and 10.8 redrawn. Left: the net power each structure settles at after the dispatch increase, against the 1050 kW request (dashed). Middle: the peak turbine inlet temperature on the way, against the 620 °C metal limit (dashed). Right: the energy discarded through the bypass on the dispatch reduction, with the three-deep structure solid. Notice that only the three-deep structure reaches the request, that the temperature-only structure with the bypass shut is the only one to cross the metal limit, and that the three-deep structure discards less than half of what the two-layer structure discards.

4.9 min, for an integral error of 607 kW min over the following eighty minutes. With the fast path switched off before Run, the firing rate is the only handle, the plant is still at 751.4 kW eighty minutes later, and the integral error is 26 440 kW min, a factor of 44. This is the scenario in which the bypass, which the previous two subsections treat as an expense to be minimized, is the only thing standing between the plant and an overspeed trip.

10.6.4 Where the Third Layer Stops

Scenario 8 asks for 1080 kW at 15 min and then 1270 kW at 220 min. The first request is answered by the third layer as before, and at 180 min the plant holds 1080 kW with the bypass back at 4.3 %, the discharge at 23.3 MPa and the turbine inlet at 607.8 °C. The second request is not answered. All three setpoints reach their stops at the scan that receives it, and Figure 10.3 catches the plant fifteen minutes later. The discharge pressure setpoint is at its 23.8 MPa clamp against a 24.0 MPa rating, the turbine inlet setpoint is at its 618 °C clamp with the temperature at 617.4 °C and still rising, the bypass is fully shut, and the plant makes 1188 kW of a requested 1270 kW with no actuator left to move. The turbine inlet then overshoots its clamped setpoint to 620.8 °C and spends 7.2 min above the metal limit, and the plant settles at 1189.9 kW with the inlet at 618.0 °C.

That is the correct behavior and it is worth staging deliberately. A cascade of valve position controllers does not remove constraints, it reorders which one binds. The purpose of the ladder is to spend the cheap authority before the expensive authority and the expensive authority before the unavailable authority, and when all three are gone the answer is that the plant cannot deliver the dispatch.

> **TRY IT IN THE APP** **8 · Three Deep, and the Rating That Ends It**
>
> Load the scenario, press Run and let it pass 220 min, when the dispatch rises to 1270 kW. Watch the bypass, turbine inlet and discharge pressure strips. All three handles reach their stops, the turbine inlet overshoots to 620.8 °C and spends 7.2 min above the metal limit, and the plant settles at 1189.9 kW. The ladder reorders which constraint binds and creates no capacity.

10.6.5 A Loop Outside the Hierarchy

Scenario 3 warms the cooling tower from 15 to 29 °C with the compressor inlet controller in service. The controller does what it can, driving the coolant valve to 100 %, and the inlet temperature still

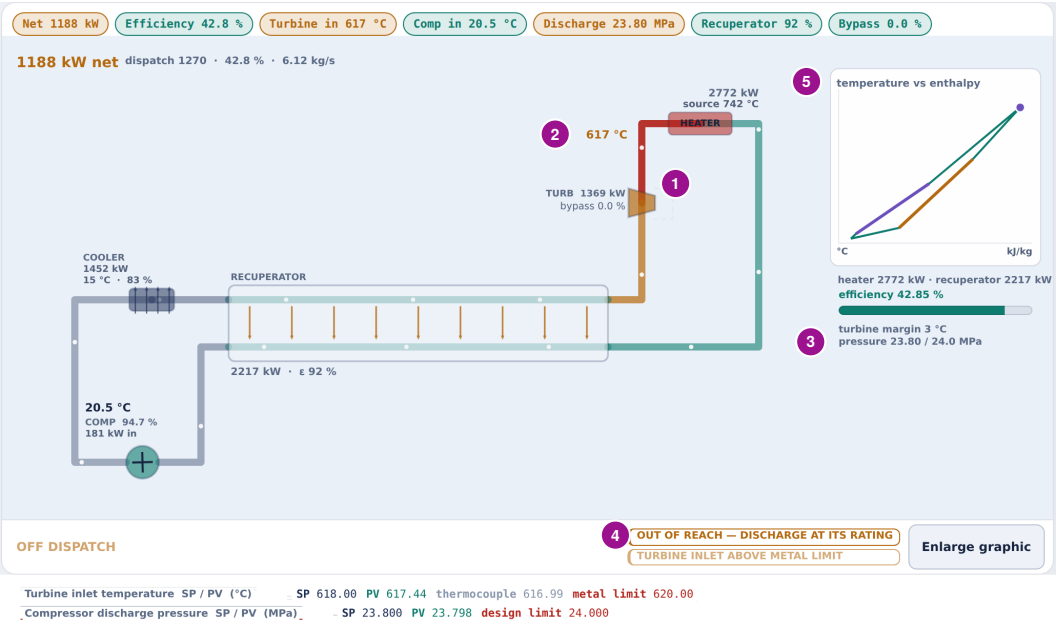


Figure 10.3 Scenario 8 at 3.92 h (235.2 min), fifteen minutes after a dispatch request of 1270 kW. The image is a composite of the application's window: the plant view, and below it the names and readings of two trend strips. All three layers of the hierarchy have run out at once. The bypass (1) is shut at 0.0 %, the turbine inlet setpoint is at its 618 °C clamp with the temperature (2) at 617.4 °C, 2.6 °C below the metal limit, and the compressor discharge (3) is 23.80 MPa against a 24.0 MPa rating. The plant makes 1188 kW, and the banner (4) reports the dispatch out of reach with the discharge at its rating. The inset (5) is the cycle drawn on temperature and enthalpy axes with the recuperator as the shared segment.

settles at 29.80 °C. Net power falls from 928.9 kW to 597.4 kW and thermal efficiency from 40.94 % to 35.33 %, and every controlled variable except the compressor inlet is exactly where it was.

> **TRY IT IN THE APP** **3 · The Cheapest Loop on the Plant**
>
> Load the scenario and press Run. At 20 min the coolant supply warms from 15 to 29 °C. Watch the compressor inlet strip. The coolant demand runs to 100 %, the inlet still settles at 29.80 °C, and net power falls from 928.9 to 597.4 kW, while every other controlled variable stays where it was.

The comparison worth making is against Section 10.6.4. The entire three-layer hierarchy, designed and tuned carefully, raises the reachable output from the 959.5 kW at which the turbine inlet alone stops to 1189.9 kW, which is about 230 kW of additional dispatch range. An afternoon costs 332 kW. On a real plant that argument is the case for evaporative or hybrid cooling, and it is an argument the control engineer is usually not in the room for.

10.6.6 Fouling, and a Controller That Hides It

Scenario 10 fouls the recuperator in two steps with the full hierarchy in service and the dispatch unchanged.

Table 10.9 is a well-behaved control system in front of a plant losing an eighth of its efficiency. Net power stays between 927.2 and 931.3 kW for four hours while the heater duty rises 14.7 % and the thermal efficiency falls 5.2 points, which is 12.8 % of the efficiency the plant started with. No

Table 10.9 Recuperator fouling at constant 930 kW dispatch. The controlled variable stays within 3 kW of its setpoint.

t (min)	ε	Recuperated	Heater (kW)	Net (kW)	η (%)
19	0.925	47.9 %	2298	930.4	40.48
60	0.869	44.8 %	2418	930.1	38.46
180	0.773	39.2 %	2638	930.0	35.25
240	0.773	39.1 %	2635	930.0	35.29

Table 10.10 The graded seven hour dispatch day, read from the application's own judgment at 420 min. The recommended configuration is the suggested setting described in the text, and each row after the second changes one element of it before Run.

Configuration	Out of band (min)	Over 620 °C (min)	Peak P_2 (MPa)	Bypass (kWh)
As loaded	420.0	0.0	20.58	0
Recommended	61.7	12.6	23.80	280
no third layer	339.1	40.2	20.58	127
$T_4^{\mathrm{ref}} = 593\,^{\circ}\mathrm{C}$	120.3	30.3	23.80	233
$b^{\mathrm{ref}} = 0\,\%$	116.3	9.7	23.62	132
$b^{\mathrm{ref}} = 8\,\%$	**17.6**	13.3	23.80	**509**

process alarm fires, because no measured variable has moved. The only warning is a calculated one. The application flags the recuperator as degraded when its computed effectiveness falls below 0.80, which happens for 0.8 min just after the first step, before the wall temperature adjusts to the lower conductance, and for good from the second step at 120 min. The first step therefore costs two points of efficiency with no lasting indication at all. The signal is in the ratio of heater duty to net power, which is a calculation nobody is required to look at, and the application draws it for that reason.

10.6.7 The Graded Dispatch Day

The challenge is seven hours of dispatch against a warming afternoon, a turbine efficiency loss and an evening peak, on a plant whose recuperator is already fifteen percent down. The target requires the plant to stay within ±25 kW for all but 80 min, spend no more than 20 min above the 620 °C metal limit, keep the discharge below 24 MPa and discard less than 400 kWh through the bypass. The suggested settings, which the recommended configuration of Table 10.10 uses, are the Cascade with the bypass and speed paths, the tunings of Table 10.6, a 3 % resting bypass and a 580 °C turbine inlet reference, with the compressor inlet setpoint and the bypass station left at the case's 20.5 °C and 0 %.

The recommended configuration passes all four clauses, inside the four limits by 18.3 min, 7.4 min, 0.20 MPa and 120 kWh, and the four ablations fail for four different reasons. As loaded, with the inner loops holding their own setpoints, weak tunings and nobody watching the power, the plant is out of band for the whole day, although the other three clauses hold. Removing the third layer fails two clauses at once, because without a route to discharge pressure the plant answers every dispatch increase with turbine inlet temperature, spends 40.2 min above the metal limit and is still out of band for 339.1 min. Leaving the turbine inlet reference at its design 593 °C rather than lowering it to 580 °C fails the same two clauses more narrowly, because the headroom the evening peak needs was spent in the morning.

The last two rows are the same change in opposite directions and they are the point of the scorecard. The resting bypass position is the amount of throughput held permanently in reserve so that the fast path has authority in both directions. Setting it to zero makes the plant maximally efficient and removes its ability to answer a load increase quickly, which costs 116.3 min out of band. Setting it to eight percent gives the best tracking on the sheet, 17.6 min, and discards 509 kWh, which fails the energy clause. Three percent passes both. The resting position is not a tuning parameter but

Table 10.11 The MPC against the cascade with the bypass and speed paths, shipped seed, noise off, full length of each case

	Power IAE (kW min)				Out of band (min)				Bypassed (kWh)				Over 620 °C (min)				MPC
Case	S	R	H	M	S	R	H	M	S	R	H	M	S	R	H	M	plans
5	943	86	78	90	11.9	0.6	0.5	1.1	147.1	0.8	156.0	0.6	0	0	0	0	600
6	42	35	7	4	0.0	0.0	0.0	0.0	79.4	0.3	83.1	0.1	0	0	0	0	300
7	180	80	30	40	1.2	0.1	0.1	0.4	152.5	0.7	146.9	0.6	0	0	0	0	600
8	24625	23075	22923	23310	309.1	281.8	281.2	281.8	115.5	1.4	144.6	0.1	15.4	11.7	11.4	0	1000
9[a]	15206	1711	592	542	136.2	23.9	4.0	3.1	116.3	16.2	133.2	12.9	15.1	0	0	0	480
10	87	41	12	46	0.0	0.0	0.0	0.0	224.9	0.4	222.3	0.1	0	0	0	0	800
11	8425	1050	633	781	101.1	9.6	8.1	10.3	143.7	8.2	152.4	2.5	12.0	3.4	0	0	510
12, suggested[b]	5098	1581	937	1482	61.7	8.6	6.4	14.7	280.2	15.2	240.3	2.5	12.6	0	0	0	840
12, as shipped[c]	66900	65709	26460	37429	420.0	375.1	284.6	309.8	47.0	37.5	228.0	72.2	0	0	0	0	840

S is the cascade at the case's own tunings, which scenarios 8 and 10 load, R the retuned cascade, H the hindsight setting and M the MPC. The MPC plans every 0.5 min, so its plans give each case's length, and none was rejected. [a]Scenario 9 loads the bypass path alone, which gives 5654 kW min, 82.3 min, 475.0 kWh and none over 620 °C. [b]The challenge with its suggested settings, the configuration recommended in Table 10.10. [c]The challenge with its own weak inner-loop tunings, which as loaded give 69 139 kW min and 420.0 min out of band.

a purchase of responsiveness with fuel, priced in this scenario at 3.8 kWh per minute of tracking recovered.

10.6.8 Model Predictive Control Against the Hierarchy

Table 10.11 sets the MPC beside three settings of the cascade with the bypass and speed paths on scenarios 5 to 11 and on the challenge, each run for its full length at the shipped seed with the noise off. All four read the same instruments and none is given a preview. The shipped setting uses each case's own tunings, the suggested ones on the challenge. The retuned setting rests the bypass at 0 % with fast valve position controllers ($K_c = 10\,°\text{C}/\%$ and $\tau_I = 10$ min for the first, 0.3 MPa/°C and 5 min for the second), and the hindsight setting is the least power error among 288 combinations of resting position, inlet reference and five tuning constants, tried on each case alone. Both stronger settings were therefore picked in hindsight, while the MPC keeps scenario 11's settings on every case. Integral error is the scorecard's, read from the logged rows.

Scenario 11 runs the tuned plant of scenario 8 through a dispatch day of 930, 820, 1050, 1180 and 930 kW, with steps at 15, 75, 135 and 195 min. Under the MPC it leaves 781 kW min of power error against 8425 for the shipped cascade, sends 2.5 kWh through the bypass against 143.7, and never passes 620 °C, where the shipped cascade spends 12.0 min above it. Most of that gain is not the optimizer. The retuned cascade, which moves the pressure within minutes and rests the bypass shut, leaves 1050 kW min and 8.2 kWh, with fewer minutes out of the band than the MPC (9.6 against 10.3). The third handle carries the plan's result, since with its pressure setpoint frozen the same controller leaves 21 050 kW min on this day, 112 220 on scenario 8 and 49 447 on the challenge.

What separates the plan from the retuned cascade is the turbine inlet (Figure 10.4). After the step to 1050 kW at 75 min both put the pressure setpoint on its 23.8 MPa clamp and are in the band for good at 77.1 min, the plan with its inlet setpoint raised only to 597.7 °C and the retuned cascade with that setpoint on its 618 °C clamp. The shipped cascade raises its pressure setpoint by 1.6 MPa in the first 20 min, overshoots to 625.7 °C at 95.7 min and stays out of the band until 113 min. After the step to 1180 kW at 135 min the plan and the retuned cascade both put the inlet setpoint on its clamp. By 142.5 min, with the inlet at 601.0 °C, the plan has begun to cut that setpoint, which stands at 600.8 °C by 145 min, ahead of the overshoot its second-order model predicts, and the inlet peaks at 613.9 °C at 150.5 min. The retuned cascade holds its clamp until 145 min, and its inlet reaches 620.0 °C at 152.7 min and 620.6 °C at 181 min after the setpoint has returned to the clamp, 3.4 min above the limit in all. The metal row is never active on this day, and a copy of the MPC whose inlet model is first order with the same 63 % time crosses 620 °C for 5.2 min, with a peak of 621.1 °C.

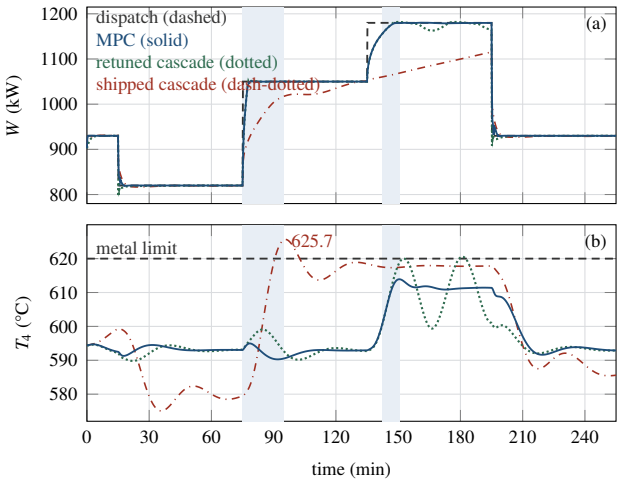


Figure 10.4 Scenario 11 run by the MPC (solid), by the shipped cascade with the bypass and speed paths (dash-dotted), and by the retuned cascade (dotted). (a) Net power against the dispatch (dashed). (b) Turbine inlet temperature against the 620 °C metal limit (dashed). The shaded windows, 75 to 95 and 142.5 to 151 min, follow the steps to 1050 and 1180 kW.

Scenario 8's request for 1270 kW is beyond every structure. The MPC spends no time above 620 °C (peak 618.3 °C), against 15.4, 11.7 and 11.4 min for the shipped, retuned and hindsight settings, at power errors within 8 % of each other. Either the clipped target or the metal row alone keeps the inlet under the limit, with peaks of 618.3 and 618.1 °C without the other, and without both the MPC spends the same 11.7 min above it as the retuned cascade.

In the load rejection of scenario 9 the plan and the shipped fast path both act on the scan the breaker opens. The MPC is in the band for good 0.86 min after the step to 450 kW, against 4.88 min for the bypass path the scenario loads, 3.42 for the shipped cascade with the speed path and 23.7 min for the retuned one, and its integral error is 542 kW min against 5654 for the bypass path. The price costs it the first scan, in which it sheds 24 kW less than the bypass path (729.0 kW at 10.1 min against 705.5), because it opens the bypass to 32.6 % rather than 43.7 and counts on the pressure setpoint it cuts to 14.8 MPa on the same scan. With $e^* = 0$ it reaches 708.0 kW on that scan, and at 10.2 min it already sheds more, at 624.4 kW against 637.9.

The MPC loses three comparisons, each with a mechanism. The hindsight setting has less power error on every case except scenarios 6 and 9 (scenario 11: 633 against 781 kW min), but each such setting rests the bypass at 3 % and throws 144.6 to 240.3 kWh away, against 0.1 to 2.5 kWh for the MPC on the same cases with tuned inner loops. The retuned cascade has less error on scenarios 5 and 10 (86 against 90 and 41 against 46 kW min) and fewer minutes out of the band on the challenge (8.6 against 14.7). When the coolant warms at 70 min, the plan, which holds the measured compressor inlet constant over its horizon, puts the pressure setpoint on its clamp by 73 min but raises the inlet setpoint only from 75 min, and it is out of the band for 6.5 min where the retuned cascade, moving both at once, is out for 1.6. The third loss is the first scan of the load rejection described above.

The plan is made on the inner loops of its identification and is only as good as that assumption. With the challenge's own weak inner-loop tunings it is out of the band for 309.8 min of 420, against 66.5 min once the pressure loop alone takes the identified tuning, 44.2 min with both inner loops and 7.5 min with the compressor inlet loop as well, and on scenario 11 doubling the inlet loop's gain takes the inlet over 620 °C for 7.2 min. Under those weak loops every structure is slow, and scenario 11 leaves 20 541, 22 279 and 20 459 kW min under the MPC and the shipped and retuned cascades.

On the challenge, the suggested settings with the supervisory cascade replaced by the MPC meet

the target, at 14.7 min out of the band, no time above 620 °C, a peak discharge of 23.800 MPa and 2.5 kWh bypassed, against 61.7, 12.6, 23.80 and 280 for the recommended cascade of Table 10.10 and 8.6 min out of the band for the retuned one. The run breaks none of the limits the rule leaves unscored. The turbine inlet peaks at 606.6 °C, the compressor inlet stays at least 7.11 °C under its alarm, and the recuperator effectiveness stays at or above 84.2 %, clear of the 0.80 at which the application flags it as degraded.

10.7 Technology in Practice

A real supercritical carbon dioxide loop reads and moves the quantities of Section 10.5 with hardware that has its own limits. Figure 10.5 draws the seven loops with plant devices in place of the schematic's symbols.

What Measures It

On the test loop of the history segments, resistance thermometers and thermocouples sat in the carbon dioxide at the inlet and exit of every component, pressure transducers at the inlet and outlet of every component, and an orifice flow meter after the compressor outlet [10]. The model keeps four of these readings as first-order lags of 0.05 to 0.35 min. The reading that matters most is the one at the compressor inlet.

> **TECHNOLOGY Sensor: A Resistance Thermometer at the Compressor Inlet**
>
> A resistance temperature detector (RTD) holds a sensing element whose electrical resistance changes with temperature. The test loop of the first history segment used RTDs with an accuracy of about 0.2 K beside thermocouples of about 1 K, and warned that heat conducted along an inserted probe can spoil the reading in small ducts. Its report judged it very difficult to know the real fluid temperature near the critical point to within a few tenths of a degree [10]. In a thermowell the response depends on heat transfer, which changes strongly with density there. An aging RTD drifts in calibration or answers more slowly, and a later study modeled the fault as a time constant growing from 2 s to 6 s [3]. The model's compressor inlet thermocouple lags 0.25 min.

What Moves It

In the reactor study of the second history segment, the settings of the slower handle were chosen so that the turbine bypass valve stays around 10 % open [16]. That is in effect the resting position of Section 10.6.7, which holds throughput in reserve so that the fast path has authority in both directions. On the cold side, the test loop moved its compressor inlet temperature with the speed of the cooling water and with a valve that recirculated warm water, against an evaporative cooler whose supply varied from about 278 to 298 K between winter and summer [10]. Pilot-plant studies keep the cooling water leaving the cooler below 50 °C to prevent scale and pump warm water back into the entering stream [19]. The model's coolant valve saturates between a supply of 19 and 22 °C (Table 10.5).

> **TECHNOLOGY Final Element: A High-Speed Motor Drive**
>
> On a test loop the compressor speed handle is the drive of a high-speed permanent-magnet motor-alternator whose stator sits in the carbon dioxide itself. The loop of the first history segment used a variable-speed controller with insulated-gate bipolar transistors and pulse-width modulation, built to motor at 50 kW and to generate up to 150 kW, and it ran the shaft to 65 000 rpm. A coil embedded in the stator replaced a speed proximity sensor that proved too noisy. Starting the motor, keeping the speed signal clean, deciding when to shift to speed control and handling internally sensed faults took considerable work, and the rotor had to accelerate quickly through its first two vibration modes at 17 000 and 32 000 rpm [10]. In the model, discharge pressure follows a speed step with a 63 percent time of 0.35 min.

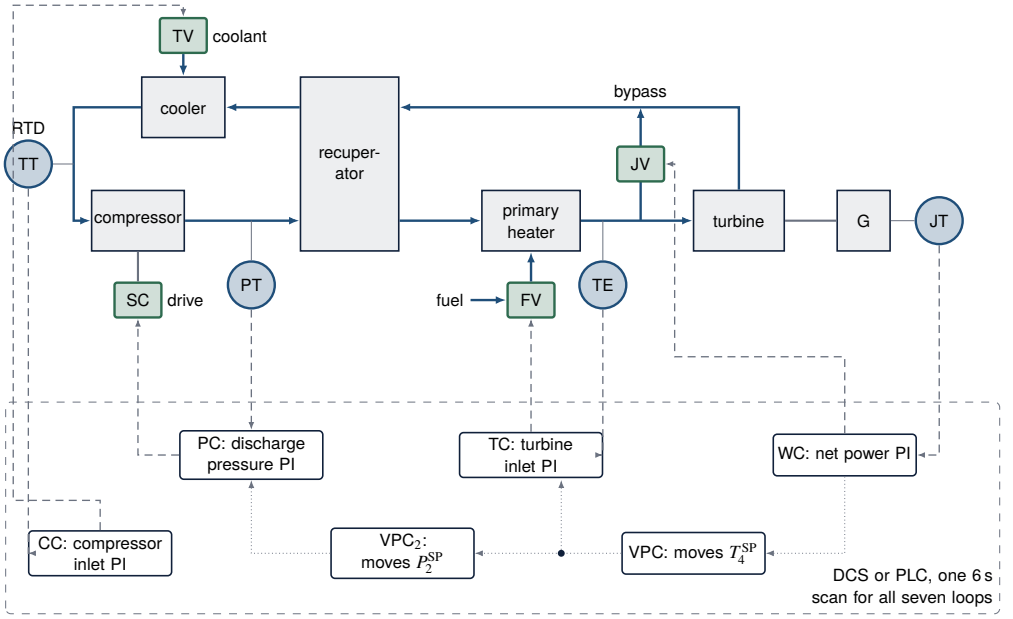


Figure 10.5 The loops of Figure 10.1 on plant hardware. A resistance thermometer (TT, RTD) at the compressor inlet feeds the inlet PI (CC), which moves the coolant valve (TV). The discharge pressure transmitter (PT) feeds the pressure PI (PC), which sets the motor drive (SC) of the compressor. The turbine inlet thermocouple (TE) feeds the firing PI (TC) and its fuel valve (FV), and the power transducer (JT) on the generator (G) feeds the power PI (WC), which moves the turbine bypass valve (JV). Dotted lines are the ladder: the valve position controller (VPC) reads the bypass loop's demand before the clamp and sets T_4^{SP}, and the second (VPC_2) reads that setpoint and sets P_2^{SP}. Notice that the coolant loop, the most valuable on the plant, has no link to the ladder.

Where the Controller Runs

Control loops are often nested, with the setpoint of one loop based on the variable of another, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [20]. The test loop ran on-off and PID loops for its charge, its cooling water flow and its warm-water recirculation from a data acquisition and control system that sampled about four times a second [10]. The application's seven controllers share a 6 s scan. Each valve position controller is an ordinary PID block whose measurement is the output of another controller, which is why the block must be able to read the demand before the clamp.

> **IN PRACTICE** **Showing the Operator Where the Critical Point Is**
>
> The test loop's operators watched more than trends. Its control screen plotted the measured compressor inlet, compressor outlet and chiller states on live temperature-entropy and temperature-density diagrams, so they could see where the loop ran relative to the saturation dome and the critical point, and another tab drew the inlet on the compressor map for the last 100 recorded time steps. The aim was an inlet within about 1 to 2 °C and 100 to 400 kPa of the critical point. A Coriolis meter at the compressor inlet, which reads density directly, proved a valuable control signal while carbon dioxide was pumped in or vented [10].

What Is Modernizing It

The first development is predictive control. A study for a reactor power converter found that one predictive controller over several inputs can coordinate the components of the cycle, where separate PID loops need individual tuning and other measures against coupling [3]. The second is the move

from test loops to pilot plants. A 2017 roadmap planned a pilot stage operated off grid by scientists, then an on-grid demonstration operated by industry, and it listed dry cooling among the motivations for small reactor, concentrating solar and natural gas applications [21]. Dry cooling makes the warm afternoon of Section 10.6.5 the usual case. The third is dynamic modeling for control design before a plant runs: pilot-plant studies built dynamic models of a 10 MW recompression cycle for load changes, startup, shutdown and trips, and for the control of its cooler [22].

10.8 Limitations

There is no recompression branch. A real high-efficiency sCO_2 cycle splits the flow before the cooler and recompresses part of it to balance the recuperator, and the split fraction is the main efficiency handle at part load [14]. Its absence removes an actuator that would change several of the conclusions above, and the load-following literature for recompression cycles treats it as central [13]. The application is a simple recuperated cycle and should be read as one.

Properties are constant specific heats per segment. Real carbon dioxide near the critical point has a specific heat that varies by an order of magnitude across the cooler and the low-temperature end of the recuperator, which is precisely where the constant-c_p assumption is worst. The near-critical behavior that matters for control is carried by (10.1) and (10.4) as fitted functions, so the shapes are right and the numbers are not transferable to a specific machine.

The heat exchangers are single lumped walls. A printed circuit recuperator has a distributed temperature profile, a pinch that can move along its length and transient behavior that a one-wall model cannot reproduce [12]. The application reports a pinch and that number should be read as an indicator rather than a design quantity.

Inventory control is a slider rather than a loop. Adding or removing carbon dioxide from the loop is a standard part-load strategy on real plants and is slower than everything in Table 10.3, so it would be a fourth and slowest layer of the same ladder. Turbomachinery is quasi-steady with no shaft dynamics, no surge line and no windage, so the application cannot represent the transients an overspeed event actually produces. The heat source is a first-order lag toward a temperature proportional to firing rate, which suits a fired heater and does not represent a nuclear or solar source with its own dynamics [11]. A solar field with a transit lag of its own is the plant of Volume 17, Chapter 6. The compressor inlet pressure is held constant by assumption, so the loop cannot show the pressure excursion a real inventory responds to when the loop temperature changes. The four-minute fired-heater lag and the inventory parameter are both deliberate simplifications, so that the mid-ranging argument is made with three actuators rather than five.

The predictive controller's model is the inner loops at one tuning and a power map relinearized at each decision. It holds the measured compressor inlet constant over the horizon, which costs it the warm afternoon of the challenge, and it carries no row for that inlet, which no compared run brought within 4.25 °C of its alarm. A detuned inner loop changes dynamics that the output biases cannot correct, the last block holds the commands for the final 7.5 min of a plan, and the shared solver allots a single slack to the whole horizon instead of one to each step.

OPEN CHALLENGE Holding the Compressor Inlet Above the Critical Point

Control studies for a 10 MW pilot plant state the problem plainly. Its controllers must hold the main compressor inlet at a design 33.5 °C with a lower constraint of 32 °C, keep the bypass compressor at least 10 % above its surge limit, and allow for a cooler outlet change of 1 °C that, left uncontrolled, can move the turbine inlet by 10 °C [22]. Proportional-integral (PI) controllers have fixed gains that cannot follow the steep property changes, and sensor delays add to the difficulty. A predictive controller kept a simulated pre-cooler stable where a PI controller with a degraded sensor did not [3]. Progress would be such a controller proven on a running plant, through startup, shutdown and trips.

10.9 Discussion

Three results generalize past this plant.

The first is that mid-ranging is a statement about cost as much as about speed. The textbook presentation pairs a fast small actuator with a slow large one and recovers the fast one's range. What Tables 10.7 and 10.8 add is that the ordering by expense need not agree with the ordering by bandwidth. The bypass is the fastest and the most expensive. The firing rate is the slowest and moderately expensive. Discharge pressure is intermediate in speed and free, which is why it belongs at the bottom of the ladder and why leaving it out costs both the tracking and the metal.

The second is that a saturating actuator must be handed upward by its unsaturated demand. A clamped output carries no information about how far past the clamp the controller wanted to go, and a supervisory layer reading the clamped value inherits a plant whose gain is zero on one side of the resting position. This is a small implementation detail with a large and asymmetric consequence, and it is not specific to valve position control, because the same requirement appears in override selectors, in split-range schemes and anywhere a signal is passed between controllers [6]. Section 1.4.4 treats the companion question for an override, which signal a controller left out by a selector should track.

The third is that a hierarchy reorders constraints rather than removing them, and a plant should be run until they all bind so that the operator knows which one arrives first. Section 10.6.4 reaches the pressure rating and the metal limit within a few minutes of each other, which is a well-designed plant and also a plant with no margin left. Knowing that is worth more than another point of tracking performance [8]. The corollary is that a variable held perfectly is not evidence of a healthy plant, because the fouling of Table 10.9 costs an eighth of the efficiency behind a power trend that stays within 3 kW.

10.10 Conclusion

A supercritical carbon dioxide power block gives the control engineer one controlled variable and three actuators whose open-loop response times span 0.14 min to 17.0 min and whose marginal costs run from zero to everything. A dispatch increase of 120 kW is not slow but unreachable for the two obvious architectures, because both are stopped by the 618 °C clamp on the turbine inlet setpoint at 959.5 kW. Adding a third mid-ranging layer that converts turbine inlet temperature back into compressor discharge pressure reaches 1050.0 kW in 6.8 min with an integral error smaller by a factor of 37, a higher efficiency and 3.0 °C of metal margin in hand, and it does so with one additional PID and one careful decision about which signal to hand upward. The same layer is the fastest and the cheapest on the way down, discarding 107.0 kWh against 242.9 kWh for the two-layer structure. Behind all of it sits a compressor inlet loop that is not part of the hierarchy, cannot be improved by better tuning, and whose saturation on a warm afternoon costs 332 kW, about one and a half times the 230 kW of range the whole hierarchy adds. Offered beside the hierarchy, a model predictive controller over the same three handles leaves 781 kW min of power error on the dispatch day of scenario 11 against 8425 for the shipped hierarchy and 1050 for a retuned one, and what it keeps for itself, a turbine inlet held under its limit by the model of the inner loop and a load rejection settled in 0.86 min, depends on inner loops tuned as that model assumes.

KEY POINTS

- One power output has three handles: a fast bypass that throws work away, a free compressor speed and a slow firing rate that buys fuel.
- On a 120 kW dispatch increase the two obvious architectures never arrive, because the 618 °C clamp on the turbine inlet setpoint stops them at 959.5 kW.
- A third layer that moves the discharge pressure reaches 1050.0 kW in 6.8 min, with 37 times less integral error.
- A saturating actuator must be handed upward by its unsaturated demand.
- A warm afternoon costs 332 kW through the compressor inlet loop, more than the 230 kW of range the whole hierarchy adds.

Where else this applies. Chapter 9 returns a fast electric chiller to its reserve with a valve-position controller on chilled water, and Volume 17, Chapter 4 hands one variable between two actuators on a wind turbine.

References

[1] Y. Ahn, S. J. Bae, M. Kim, S. K. Cho, S. Baik, J. I. Lee, and J. E. Cha, "Review of supercritical CO_2 power cycle technology and current status of research and development," *Nuclear Engineering and Technology*, vol. 47, no. 6, pp. 647–661, 2015.

[2] F. Crespi, G. Gavagnin, D. Sánchez, and G. S. Martínez, "Supercritical carbon dioxide cycles for power generation: A review," *Applied Energy*, vol. 195, pp. 152–183, 2017.

[3] R. Ponciroli and R. Vilim, "Analyses of model-based predictive control for a S-CO_2 Brayton cycle power converter," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep., 2016.

[4] T. Hägglund, "A feedforward approach to mid-ranging control," *Control Engineering Practice*, vol. 108, p. 104713, 2021.

[5] 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.

[6] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[7] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

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

[9] R. Span and W. Wagner, "A new equation of state for carbon dioxide covering the fluid region from the triple-point temperature to 1100 K at pressures up to 800 MPa," *Journal of Physical and Chemical Reference Data*, vol. 25, no. 6, pp. 1509–1596, 1996.

[10] S. A. Wright, R. F. Radel, M. E. Vernon, P. S. Pickard, and G. E. Rochau, "Operation and analysis of a supercritical CO_2 Brayton cycle," Sandia National Laboratories, Albuquerque, NM, USA, Tech. Rep. SAND2010-0171, 2010.

[11] G. S. Delsoto, F. G. Battisti, and A. K. da Silva, "Dynamic modeling and control of a solar-powered Brayton cycle using supercritical CO_2 and optimization of its thermal energy storage," *Renewable Energy*, vol. 206, pp. 336–356, 2023.

[12] M. Marchionni, L. Chai, G. Bianchi, and S. A. Tassou, "Numerical modelling and transient analysis of a printed circuit heat exchanger used as recuperator for supercritical CO_2 heat to power conversion systems," *Applied Thermal Engineering*, vol. 161, p. 114190, 2019.

[13] H. Wang, J. Chen, G. Yu, and M. Liu, "Research of wide-range load following control strategy for sCO_2 recompression Brayton cycle system," *Energy Conversion and Management*, vol. 349, p. 120870, 2026.

[14] F. Correa, R. Barraza, Y. C. Soo Too, R. Vasquez Padilla, and J. M. Cardemil, "Optimized operation of recompression sCO_2 Brayton cycle based on adjustable recompression fraction under variable conditions," *Energy*, vol. 227, p. 120334, 2021.

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

[16] A. Moisseytsev and J. J. Sienicki, "Development of the ANL plant dynamics code and control strategies for the supercritical carbon dioxide Brayton cycle and code validation with data from the Sandia small-scale supercritical carbon dioxide Brayton cycle test loop," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep. ANL-ARC-218, 2011.

[17] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[18] K. R. Muske and T. A. Badgwell, "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.

[19] J. Albright, E. Liese, S. Zitney, P. Mahapatra, and D. Bhattacharyya, "Load following for a supercritical CO_2 Brayton cycle," Presentation, Annual AIChE Meeting, Minneapolis, MN, National Energy Technology Laboratory, NETL-PUB-21179, OSTI 1409315, 2017.

[20] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[21] C. M. Mendez, G. Rochau, and M. Wilson, "Technology roadmap and systems engineering model for sCO_2 power cycle development," Poster, Sandia National Laboratories, SAND2017-12272C, OSTI 1482300, 2017.

[22] E. A. Liese, P. Mahapatra, J. T. Albright, and S. E. Zitney, "Regulatory control of a 10 MWe supercritical CO_2 recompression closed Brayton cycle," Presentation, 2019 Annual Review Meeting for Crosscutting Research, Pittsburgh, PA, National Energy Technology Laboratory, NETL-PUB-22283, OSTI 1623322, 2019.

10.11 Exercises

10.1 Concept check · 2 min

In Manual the application offers four handles on the cycle. Which of the following is *not* one of them, so that adding it would change several conclusions of this chapter?

(a) The recompression flow split fraction
(b) The compressor speed
(c) The coolant valve demand
(d) The turbine bypass position

10.2 Concept check · 2 min

The plant is 100 kW short of dispatch and the bypass is clamped shut at 0 %. The bypass position controller of Equation (10.20) reads b^{raw}, the bypass loop's demand before the clamp. What does that signal carry that the applied position does not?

(a) The resting position, which the clamp removes from the applied position
(b) The size of the shortfall, since the demand settles near $K_c(W - W^{\mathrm{SP}})$, about −9 %
(c) The integral of the shortfall, which grows without limit while the clamp holds
(d) The power thrown away through the bypass, which the shut valve reports as zero

10.3 Concept check · 1 min

Scenario 10 fouls the recuperator in two steps while every loop holds its setpoint. Which of these quantities does an instrument on the plant measure?

(a) The recuperator effectiveness
(b) The cycle thermal efficiency
(c) The net electrical power
(d) The carbon dioxide mass flow

10.4 Theory and calculation · 5 min

The compressor inlet is held at 20.5 °C, below the near-critical threshold T_c, and the speed is raised from 85 to 92 %. Use Equations (10.1) to (10.4) and Table 10.2.

(a) Compute the pressure ratio and the discharge pressure P_2 at 92 % speed. *Checkpoint:* 3.378 and 22.90 MPa

(b) Compute the specific compression work w_c at 85 and 92 %. *Checkpoint:* 20.84 kJ/kg and 24.34 kJ/kg

(c) Compute the mass flow at 92 % and the compressor power $\dot{m}w_c/\eta_m$ at both speeds. Why does net power still rise with speed? *Checkpoint:* 5.95 kg/s; 120.7 kW and 152.5 kW

10.5 Theory and calculation · 6 min

On a warm afternoon the coolant valve is saturated and the compressor inlet settles at 29.80 °C with the speed at 85 %, while the inner loops hold the turbine inlet and the discharge pressure at their design values (scenario 3). Use Equations (10.1), (10.2), (10.4) and (10.15) with Table 10.2.

(a) Compute the penalty u, the factor in square brackets of Equation (10.1) and the specific compression work. *Checkpoint:* $u = 0.955$, factor 1.782, $w_c = 37.13\,\mathrm{kJ/kg}$

(b) Compute the mass flow and the compressor power $\dot{m}w_c/\eta_m$, and compare it with the 120.7 kW it takes at design. *Checkpoint:* 4.40 kg/s, 172.0 kW

(c) The turbine specific work stays at 201.7 kJ/kg, and the exported file shows the coolant pump power rising from 15.5 to 82.7 kW. Split the fall in net power into its three parts and compare the sum with Table 10.5. *Checkpoint:* 213.0 kW, 51.3 kW and 67.2 kW, 331.5 kW in all

10.6 Theory and calculation · 7 min

The bypass loop has $K_c = 0.09\,\%$ per kW and the bypass position controller $K_c = 3.0\,^\circ\mathrm{C}/\%$, $\tau_I = 30\,\mathrm{min}$, action $a = +1$ and resting position $b^{\mathrm{ref}} = 3\,\%$ (Table 10.6). The bypass is held shut against a steady power shortfall.

(a) Using the steady state of Equation (10.18) with $T_t = \tau_I$, find the bypass demand c for shortfalls of 5, 100 and 500 kW. *Checkpoint:* –0.45 %, –9 % and –45 %

(b) Find the error of the position controller and the rate at which its integral action raises T_4^{SP} for each shortfall, first when it reads the demand and then when it reads the applied position. *Checkpoint:* 3.45, 12 and 48 %: 0.345, 1.2 and 4.8 °C/min; 3 % and 0.3 °C/min for all three

(c) At 100 kW short, how long does integral action alone take to carry T_4^{SP} from 593 to its 618 °C clamp in each case? *Checkpoint:* 20.8 min against 83.3 min

10.7 App challenge · 8 min

Load scenario 3, *The cheapest loop on the plant*, with its own settings and the default seed. Run to 140 min at speed 20 (twenty simulated minutes a second), then export the CSV.

(a) From the `cool_cmd_pct` column, find when the coolant demand first reaches 100 %. Just before 140 min, read `T1_C`, `power_kW` and the Efficiency tile. *Checkpoint:* 22.0 min (±0.2 min); 29.80 °C, 597.4 kW (±0.5 kW), 35.33 %

(b) At the same time read `mdot_kgs`, `w_comp`, `T4_C` and `P2_MPa`. Which controlled variables are still on setpoint, and why does no loop recover the lost power? *Checkpoint:* 4.40 kg/s, 37.13 kJ/kg, 593.0 °C, 20.58 MPa

10.8 App challenge · 8 min

Load scenario 8, *Three deep, and the rating that ends it*, with its own settings and the default seed. Run to 400 min at speed 20, then export the CSV.

(a) At 180 min, read the net power, the bypass position, the discharge pressure, the compressor speed and the turbine inlet temperature. Which actuator now carries the extra 150 kW? *Checkpoint:* 1080.0 kW, 4.3 %, 23.26 MPa, 93.1 %, 607.8 °C (±0.2 °C)

(b) From `metric_overtemp_min` at 219.9 and 400 min and `T4_C` after 220 min, find the minutes above 620 °C after the step to 1270 kW and the peak. *Checkpoint:* 7.2 min (±0.1 min), 620.8 °C

(c) Read the net power at 400 min and name the three stops that bind. *Checkpoint:* 1189.9 kW, 80.1 kW short

10.9 App challenge · 9 min

Load the challenge, *A dispatch day*, with the default seed and press Reveal suggested tuning. Run the full 420 min at speed 20 and export the CSV.

(a) Read the Off dispatch tile. In the last row of the file, read the minutes over 620 °C in `metric_overtemp_min` and the bypassed energy in the column `metric_bypass_loss_kWh`, then find the largest `P2_MPa`. Does the run pass, and by how much on each clause? *Checkpoint:* 61.7 min, 12.6 min, 280 kWh (±1 kWh), 23.80 MPa

(b) Reload, press Reveal suggested tuning again, set Bypass resting position (Advanced) to 8 % and repeat the run. Which clause fails? *Checkpoint:* 17.6 min, 13.3 min, 509 kWh (±1 kWh), 23.80 MPa

(c) Table 10.10 gives 116.3 min and 132 kWh for a 0 % resting position. What does a minute of tracking cost in bypassed energy between 0 and 8 %? *Checkpoint:* 3.8 kWh per minute

10.10 Application area · 3 min

The rating of a supercritical carbon dioxide plant is quoted with a cooling temperature attached.

(a) From Table 10.5, what share of the output does a coolant supply of 29 °C remove, and how much energy is lost over a six-hour afternoon at that temperature? *Checkpoint:* 35.7 %, about 1990 kWh

(b) Why can the control system not recover this loss, and what does the chapter name as the remedy?

10.11 Application area · 4 min

Table 10.9 shows the recuperator fouling while net power stays on setpoint.

(a) Compute the ratio of heater duty to net power at 19 and 240 min, and the extra heat the heater supplies. *Checkpoint:* 2.47 and 2.83; 337 kW

(b) If heat costs $0.02 per kWh and the plant runs all year (8760 h) with the fouled recuperator, what does the extra heat cost? What would you put on the operator's screen? *Checkpoint:* about $59 000 a year

10.12 Review the agent · 5 min

A student asks an AI assistant how to make the three-layer hierarchy more robust. The assistant replies: "The bypass position controller should read the applied bypass position b, not the bypass loop's internal demand. The demand can go negative, which has no physical meaning, while the applied position is what the valve really does. The change costs nothing on a load decrease and removes a source of nonsense values from the cascade." Find the flaw and say how you would check it.