INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 5

Keeping an Organic Rankine Expander Dry When Its Waste Heat Falls

Pages 94–124

Edited by John D. Hedengren

How to cite this chapter. "Keeping an Organic Rankine Expander Dry When Its Waste Heat Falls," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 5, pp. 94–124.

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

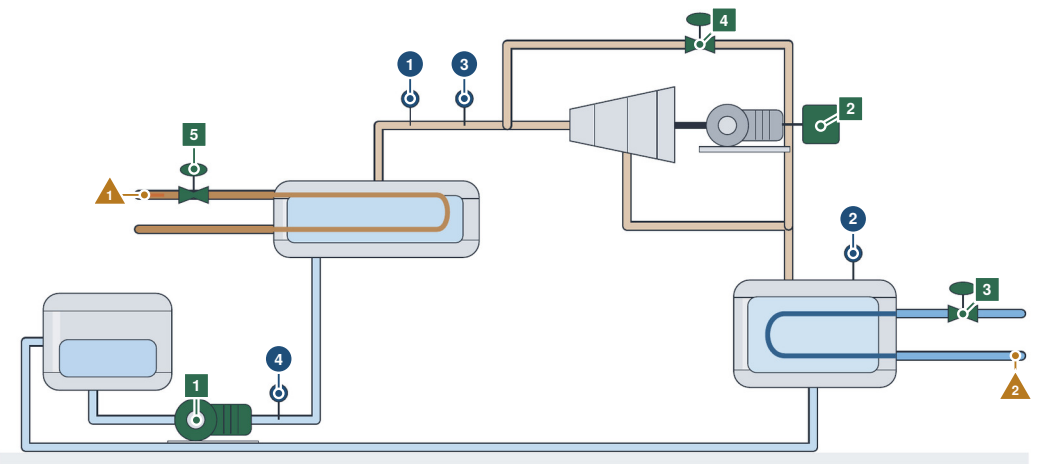
The app for this chapter
https://apmonitor.com/apps/organic-rankine/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v03-feedforward-cascade/#organic-rankine

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

CHAPTER 5

Keeping an Organic Rankine Expander Dry When Its Waste Heat Falls



The controller moves 1 feed pump speed 2 expander speed 3 cooling-water flow 4 vapor bypass around the expander 5 source bypass

The controller reads 1 evaporating pressure 2 condensing pressure 3 expander inlet temperature, for the superheat 4 feed flow

What disturbs it 1 hot source water: its flow falls or its temperature changes 2 cooling water warms

Notice that the vapor leaving the evaporator must stay above its boiling point on its way to the expander, and that the heat arriving in the hot water can fall faster than the thermometer reports it. Illustration, not to scale.

IN PLAIN TERMS **Keep the Vapor Dry When the Heat Drops**

The main idea. Hot water from an engine boils a fluid that boils more easily than water. The vapor spins a machine that makes electricity. The vapor must stay a little hotter than its boiling point, or liquid drops can harm the machine.

The challenge. When the hot water suddenly slows, a thermometer in a metal sleeve reports the cooling vapor late. A control loop that waits for it acts too late, but measuring the heat as it arrives lets the pump slow down first.

Why it matters. Any plant whose heat supply can drop without warning faces the same race between a slow sensor and a fast loss.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history of organic-fluid engines, the fact on engine heat and the feedforward on a truck engine (Section 5.2), the try-it segments of scenarios 4 and 5 and the fact on fouling (Section 5.5) and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 5.5), Technology in Practice (Section 5.6), Limitations (Section 5.7) and the open challenge

The control question When the waste heat behind an organic Rankine cycle falls by a third within seconds, can a well-tuned superheat loop keep the expander inlet dry, and which limits of the condenser and the heat source does no control structure remove?

Learning objectives

- Compute the expander's swept flow, fluid work, generator power and the cycle's net output from its model, and compare the efficiency with the Carnot bound.
- Derive the source-heat feedforward from a steady energy balance and convert a flow request into a pump speed.
- Explain why feedback through a thermowell, alone or with an override, loses the expander on a source loss that the feedforward rides through.
- Recognize a loop out of authority, such as the condensing loop on a hot day, and trade its setpoint against net power.
- Separate what a predictive controller gains from a forecast and what it gains from moving two handles together.

Prerequisites Sections 1.4.3 (static feedforward), 1.4.5 (inner loops), 1.4.2 (entry points), 1.6.2 (velocity form) and 1.6.3 (the MPC option). Chapter 2 for the structures on a simpler plant.

The app https://apmonitor.com/apps/organic-rankine/

See also Volume 4, Chapter 2 for a protective override that rests idle until its constraint nears, Volume 4, Chapter 9 for a power loop low-selected against an exhaust limit, and Volume 4, Chapter 10 for a closed power cycle with a predictive layer over its inner loops. Volume 17, Chapter 2 for the engines whose exhaust heat feeds a cycle like this one.



5.1 Introduction

An organic Rankine cycle (ORC) boils a refrigerant-like working fluid with heat that is too cool to raise useful steam and expands the vapor through a machine that drives a generator. In the simulated plant an intermediate water loop brings 959.95 kW from an on-site engine's exhaust at 150 °C to an R1233zd(E) evaporator at 12.90 bar (pressures are absolute throughout), a variable-speed positive-displacement expander delivers 91.74 kW at its generator, and after the feed pump, the cooling-water pumps and fans and the other auxiliaries the cycle nets 77.55 kW. That is 8.08 % of the accepted heat, against a Carnot bound of 29.54 % between the source and cooling-water supply temperatures, and the condenser rejects 863.96 kW, 9.4 times the generator output. The expander is specified for dry vapor, and its 10 K of superheat at the inlet is the quantity that the rest of the control system serves.

The simulator targets a specific misconception, the idea that a well-tuned superheat loop is enough to protect the expander. When the source-loop flow falls by 30 % in case 4, a cascade whose superheat controller runs at 0.08 kg/s per kelvin with a 20 s integral time lets the true superheat fall to 1.09 K, and the protection isolates the expander 17.3 s after the loss and keeps it out for 83.4 s. The controller sees the superheat through a 3 s thermowell after most of the evaporator's response has already happened. The same cascade with a feedforward computed from the source meters keeps the superheat between 9.39 and 11.08 K with the expander admitted throughout. A second belief, that a colder condenser always pays, fails on a hot day. With 35 °C cooling water, holding 1.95 bar drives the fans to their 30.0 kW limit and nets 8.69 kWh over eleven minutes, while letting the pressure rise to 2.6 bar nets 12.24 kWh.

The class of problem is general. An exchanger whose outlet state is held by the flow through it, while a second actuator sets the pressure that defines that state and an unmanaged source changes the heat, describes refrigeration evaporators and once-through boilers as well as every waste-heat ORC. A survey of the technology notes that both the pump speed and the expander speed must be controlled to hold the expander inlet conditions when the source varies [1], and the control of these cycles has been treated with multivariable designs [2] and with model-based feedforward on vehicle engines [3]. In process terms the superheat loop is a quality loop behind a measurement lag,

the inner flow loop is the secondary of a cascade (Section 1.4.5), the source-heat calculation is a measured-disturbance feedforward [4] (Section 1.4.3), and the overrides are selectors that hand an actuator to a constraint [5].

However, the published dynamic models of these cycles serve design and validation, with exchanger formulations compared against measured transients [6] and control strategies optimized on models [7] or validated on one engine [3]. A student who wants to watch the superheat fall through a thermowell, and then see a feedforward hold it, needs a plant whose loops, limits and protection are settings with consequences in kelvin, bar and kilowatts. This chapter separates what each element of a superheat control structure contributes to keeping the expander inlet dry on a waste-heat ORC, and locates the limits that the condenser and the heat source impose and no loop removes, on the simulator framework of Section 1.6.

The proportional-integral (PI) loops with back-calculation, the cascade, the feedforward and the minimum selector are textbook structures, and the chapter measures what each contributes in matched comparisons made on one plant with one random seed. Section 5.2 follows the domain from low-grade heat to the superheat loop. Section 5.3 builds a finite-volume cycle with one pressure per side, a two-zone receiver and audited property tables, whose working-fluid inventory closes to 3×10^{-12} kg and whose energy ledger closes to 8×10^{-5} of the accepted heat. Section 5.4 gives the loops, the selector, the protection layer and a predictive controller. Section 5.5 reports ten graded cases and the comparisons. Among them, the feedforward cuts the superheat error integrated over the source loss from 276.7 to 29.0 K s at the same gains, and doubling every exchanger's cell count moves the second figure to 64.8 K s without changing the ranking (Section 5.5.6). Section 5.6 describes the real instruments, drives and control platform, Section 5.7 states what the model leaves out, and Section 5.8 sums up.

5.2 From Low-Grade Heat to a Superheat Loop

Engines that run a Rankine cycle on an organic fluid have a long history, which a chapter-length account of the technology traces [8] and a review follows from the concept to current technology [9]. That review estimates the cumulative capacity of ORC systems that went into operation after 1995 at about 2000 MW of electricity, names geothermal reservoirs, waste heat from engines and industrial processes, biomass combustion and concentrated solar radiation as their sources, and notes that the same systems suit the cogeneration of heating or cooling [9]. Reviews presented the cycle for low-grade waste heat in 1997 [10] and catalogued its applications later [11], and manufacturers have sold ORC modules since the early 1980s [1].

> ⧖ HISTORY **From an Ether Engine to a Geothermal Plant, 1826 to 1952**
>
> An engine that runs a Rankine cycle on a fluid other than water is an old idea. In 1826 a patent described an engine with ether as its working fluid. Naphtha engines for launches followed, with the naphtha serving as fuel, working fluid and lubricant, and their low boiler pressure spared the owner the specialized operator a steam engine needed. In 1907 a flat solar collector boiled ether near 120 °C to drive a 2.6 kW engine. L. D'Amelio (1893 to 1967) worked out a solar plant on an organic fluid for irrigation in 1936, and in 1939 a 2.6 kW prototype for low-grade geothermal heat ran in a university laboratory. The first commercially operated geothermal plant of this kind ran for a short period in 1952. It used water at 91 °C and supplied 200 kW of electricity to a mining company [9].

Engine exhaust became a target early. A Rankine cycle on an internal combustion engine was first developed after the 1970 energy crisis, when a truck-engine prototype reported a 12.5 % reduction in fuel use over a 450 km road test, and present systems differ from that one mainly in their expansion machines and in the wider choice of working fluids [1]. The engines themselves are the plant of Volume 17, Chapter 2, where a module's power settles within seconds of a fuel step while its recovered heat needs minutes. The same survey treats the heat-rejection capacity as a constraint of its own,

because cooling fans can sharply reduce the net output [1], which is the trade of case 5. Data centers offer a further low-grade source, whose cooling technology and heat recovery opportunities have been reviewed [12], and case 8 measures how little a cycle makes from it.

> ✦ DID YOU KNOW? **At Least Half the Fuel Leaves as Heat**
>
> Even a modern engine cannot convert more than 50 % of its fuel energy into useful work. The rest is lost through the coolant and the exhaust gases, and the exhaust gases carry it at the highest temperatures [3].

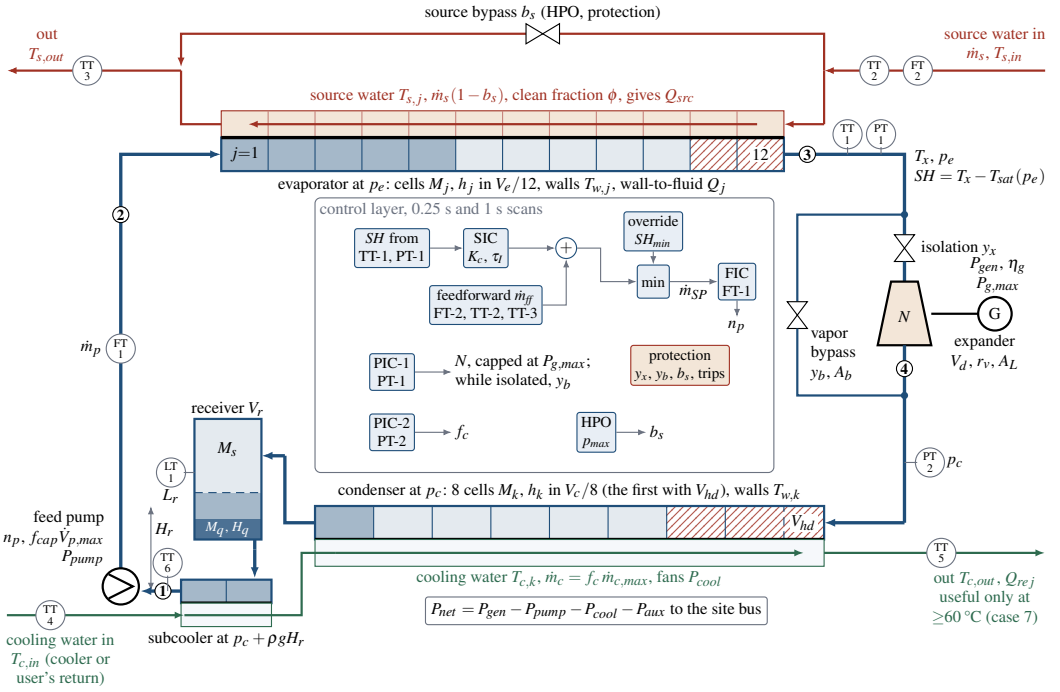
The working fluid and the machine decide where the superheat constraint comes from. A fluid whose saturated-vapor line has a positive slope stays superheated through an expansion, while a wet fluid such as water must be superheated at a turbine inlet to keep droplets from forming [1]. R1233zd(E) has been tested as a drop-in replacement for R245fa in an experimental ORC for low-temperature heat [13], and its reference equation of state, recommended as an international standard by the working group revising ISO 17584, holds from the triple point at 165.75 K to 450 K and 100 MPa [14]. Small units use positive-displacement expanders, often derived from compressors, which tolerate liquid during expansion and lose efficiency when the system's volume ratio departs from the machine's built-in volume ratio [1]. A semi-empirical scroll-expander model divides the expansion into a supply pressure drop, heat transfer, an isentropic expansion to the pressure the built-in volume ratio imposes and an expansion at constant machine volume, with the leakage mixed at the exhaust, and the machine it describes reached a maximum overall isentropic effectiveness of 68 % [15]. The expander simulated here keeps that structure in reduced form. Its dry-inlet requirement belongs to the chosen machine and is not a property of positive-displacement expanders in general.

Control came with the dynamic models. Dynamic ORC models were written for waste-heat recovery [16], the finite-volume and moving-boundary forms are the most widely used for exchangers with phase change, and the effect of the finite-volume form's number of control volumes has been examined against measurements [6]. Optimal control strategies for waste-heat cycles were built on such models [7], and a survey reports that re-optimizing the operating point continuously for a variable source raised the cycle's second-law efficiency by about 10 % [1]. Multivariable control has been applied to ORCs in waste-heat processes [2]. On a heavy-duty diesel engine a model-based feedforward on the pump speed held the superheat far closer to its setpoint than feedback alone, as the history segment below tells, and the superheat had to stay positive to prevent droplets that could damage the turbine [3]. The simulator reproduces that ordering with a steady energy-balance feedforward in place of a model inversion.

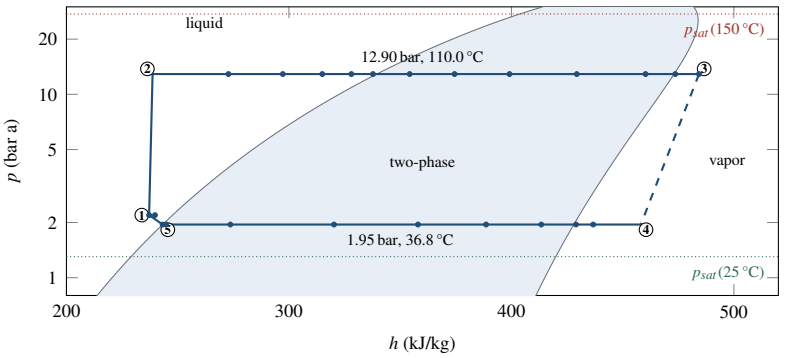
> ⧖ HISTORY **A Feedforward on a Truck Engine's Cycle, 2017**
>
> A vehicle's exhaust is a harder source than a geothermal well, because it changes strongly with the driving. A study published in 2017 fitted an organic Rankine cycle to a heavy-duty diesel engine in a test cell and held the superheat at the evaporator outlet with the pump speed. A gain-scheduled PID alone was compared with the same PID plus a dynamic feedforward, computed by inverting a reduced model of the evaporator. Under the same swings of engine torque and speed, the largest superheat error was about 1.9 K with the feedforward and about 10 K without it. A supervisor stepped the cycle through stop, cold idle, hot idle and power production, and over a road cycle of more than 50 min the superheat stayed tightly controlled [3].

5.3 A Closed R1233zd(E) Circuit Between Two Water Loops

The plant in Figure 5.1 was designed for teaching and copies no installation, and every size, efficiency and threshold in it is a synthetic value. Pressurized water from an engine's exhaust-heat exchanger arrives at 140 to 170 °C and passes a twelve-cell counterflow evaporator that preheats, boils and superheats the R1233zd(E). The vapor passes an isolation valve into the expander, or a vapor bypass around it, then an exhaust header and an eight-cell condenser cooled by water that has first passed a two-cell subcooler. The receiver between the condenser and the subcooler holds whatever inventory



(a) Process and control schematic



(b) Design point on the pressure–enthalpy plane

Figure 5.1 The simulated cycle as drawn on the application's Cycle view. (a) The source water ($\dot{m}_s$, $T_{s,j}$) gives Q_{src} to the evaporator's twelve cells at p_e, the expander at speed N behind the isolation valve y_x drives the generator, the vapor bypass y_b throttles around it, and the condenser's eight cells at p_c reject Q_{rej} to the cooling water at $\dot{m}_c$. The receiver holds a saturated zone M_s over a subcooled layer (M_q, H_q), and the subcooler and pump suction sit the static head H_r below it. Shading marks the design point's phases (dark liquid, light two-phase, hatched vapor). Circles are the instruments the controllers read, and the block diagram's outputs carry the actuator symbols. (b) The design point on the pressure–enthalpy plane from the application's property table, with the numbered states and every cell. Process model only, not to scale.

the rest of the circuit does not, and a positive-displacement feed pump returns the liquid to the evaporator. No other application runs, so the source is a declared boundary condition. Case 7 replaces the cooling water with a warm-water user's 50 °C return, and the condenser heat counts as useful only while it leaves at or above the user's 60 °C.

Table 5.1 sorts the variables by role. The integrated state has 84 components. Seventy-seven are physical, the two side pressures p_e and p_c, twelve evaporator cell masses, eleven low-side masses (eight condenser cells, the receiver and two subcooler cells), the subcooled layer's mass and enthalpy, 22 wall temperatures, 22 secondary-water temperatures and six actuator positions. Seven ledgers integrate the boundary fluxes so that conservation can be checked against independent states. Inside the model the units are SI, and the interface shows bar, °C and kW.

5.3.1 Cells at One Pressure per Side

Every working-fluid cell has a fixed volume V_j and shares its side's pressure p, so its balances read as shown in Equation (5.1).

$$\begin{aligned} \dot{M}_j &= \dot{m}_{j-1} - \dot{m}_j \\ M_j \dot{h}_j - V_j \dot{p} &= \dot{m}_{j-1}(h_{f,j-1} - h_j) - \dot{m}_j(h_{f,j} - h_j) + Q_j \end{aligned} \tag{5.1}$$

where $\dot{m}_j$ is the flow through the face that leaves cell j, $h_{f,j}$ the enthalpy that face carries and Q_j the heat from the wall. The second line is the energy balance $\dot{U}_j = \dot{m}_{j-1} h_{f,j-1} - \dot{m}_j h_{f,j} + Q_j$ for $U_j = M_j h_j - p V_j$ with h_j times the mass balance subtracted. The states are p and M_j. A cell's enthalpy, temperature and phase follow from p and $\rho_j = M_j / V_j$ through the property table, so its mass balance is $V_j[(\partial\rho/\partial p)_h \dot{p} + (\partial\rho/\partial h)_p \dot{h}_j] = \dot{m}_{j-1} - \dot{m}_j$. Eliminating $\dot{h}_j$ leaves one linear relation per cell between its two face flows and the pressure rate, as shown in Equation (5.2).

$$\begin{gathered} A_j \dot{m}_{j-1} - B_j \dot{m}_j = C_j \dot{p} + D_j \\ A_j = 1 - q_j(h_{f,j-1} - h_j), \quad B_j = 1 - q_j(h_{f,j} - h_j) \\ C_j = V_j\left[(\partial\rho/\partial p)_h + q_j\right], \quad D_j = q_j Q_j \end{gathered} \tag{5.2}$$

where $q_j = (\partial\rho/\partial h)_p / \rho_j$. Marching from the side's inlet flow makes every face flow affine in the pressure rate, $\dot{m}_j = \gamma_j + \delta_j \dot{p}$, and the outlet flow fixes it, $\dot{p} = (\dot{m}_{out} - \gamma_n)/\delta_n$. The outlet flow is the expander plus the bypass on the high side and the pump on the low side, each evaluated at the cell states of the moment, so the pressure is not relaxed toward anything and no cell carries a pressure of its own. A forward face carries the upstream cell's enthalpy, and the receiver drains its saturated liquid while it holds two phases. A reversed face carries the downstream cell's saturated vapor when that cell is two-phase, since vapor is the phase that can move back, and otherwise its enthalpy, except that a dense back-flow into a lighter cell with $B_j < 0.2$ carries the upstream enthalpy instead. The directions are iterated, at most $n+1$ times, until every flow agrees with the direction assumed for it. Each flux leaves one cell and enters the next, so the internal faces cancel exactly in both inventories.

This formulation departs from the per-cell (M, U) states with a pressure recovered in each cell that the development plan specified. That form was built first and, according to the build record, diverged within 1 s at the 0.05 s step from an exact steady start, growing about sixfold per step from the last condenser cell, which sits on the saturated-liquid line at the design point. A liquid cell's pressure is set by a change of about one part in a million in its density, and Runge-Kutta stages on either side of the dome saw compressibilities differing by orders of magnitude. That code is no longer in the source tree, so the divergence is quoted from the record and was not re-run for this chapter. With p and M_j as states, the inversion $h(p, \rho)$ is one-dimensional and monotone in every phase, because at fixed pressure the density falls steadily with enthalpy through liquid, dome and vapor. The (ρ, u) inversion remains implemented, and a check of the application confirms that every cell's density and internal energy return its side's pressure within 1×10^{-5} of it.

The price is a set of assumptions. There is no frictional pressure drop or momentum balance, pressure equalizes instantly along each side, and cells are homogeneous and in equilibrium. Where a subcooled stream meets a nearly liquid two-phase volume with no compressible vapor downstream, δ_n can pass through zero and Equation (5.2) lose its solution. The review stage found two such places, a receiver fed with subcooled condensate and an evaporator flooded by a pump held against a

Table 5.1 Variables Sorted by Role, Valued at the Design Point (Case 1, $t = 0$)

Variable	Role	Design	Range
SH	CV, expander-inlet superheat	10.0 K	SP 3–25
p_e	CV, evaporating pressure	12.90 bar	SP 6–18, trip 21
p_c	CV, condensing pressure	1.95 bar	SP 1.2–6
$\dot{m}_p$	CV (inner loop), feed flow	3.909 kg/s	0–6.25
n_p	MV, pump speed	62.71 %	0–100
N	MV, expander speed	2146 rpm	600–3600
f_c	MV, cooling-water flow	55.95 %	10–100
b_s	MV, source bypass (override)	0 %	0–100
y_x, y_b	protection, isolation, bypass	100, 0 %	0–100
$\dot{m}_s$	DV, source-loop flow	100 % of 5 kg/s	0–130
$T_{s,in}$	DV, source supply	150 °C	40–170
$T_{c,in}$	DV, cooling-water supply	25 °C	5–40, or 50
ϕ, f_{cap}	DV, clean source side, pump delivery	100, 100 %	30–100
σ	DV, instrument noise multiple	1	0–4
M_j	states, cell masses	Fig. 5.1b	charge 160 kg
M_r	state, receiver mass	93.72 kg	level 50.2 %
M_q, H_q	states, subcooled layer	0, 0	—
T_w, T_s, T_c	states, walls, water	per cell	—

lost source, resolved by the two-zone receiver below and by a flood switch that trips the plant first. Equation (5.1) holds to round-off while a cell keeps its phase, and the only energy residual comes from cells that cross the saturated-liquid line within a step, where the density partials jump between Runge-Kutta stages.

5.3.2 The Receiver in Two Zones and the Static Head

The receiver is a saturated zone, homogeneous at p_c with mass $M_s = M_r - M_q$ in the volume $V_r - V_q$, over a layer of subcooled condensate with its own mass M_q and enthalpy content H_q. Condensate more than 0.2 kJ/kg below saturation is routed to the layer, the layer supplies the outflow first and warms through an interface conductance, and a layer within 0.5 kJ/kg of saturation merges into the saturated zone, as shown in Equation (5.3).

$$\begin{aligned} \dot{M}_q &= \alpha \dot{m}_{in} - \beta \dot{m}_{out} - \mu \\ \dot{H}_q &= \alpha \dot{m}_{in} h_{in} - \beta h_q \dot{m}_{out} + Q_i - \mu h_q + V_q \dot{p}_c \\ \alpha &= \mathrm{clip}_{[0,1]}(h_l - h_{in} - 0.2), \quad \beta = \frac{M_q}{\max(M_q, 0.5)} \\ Q_i &= UA_i \max(0, T_{sat} - T_q) \min(1, M_q/0.5) \\ \mu &= k M_q \end{aligned} \tag{5.3}$$

where enthalpies are in kJ/kg and masses in kg inside the clips, $h_q = H_q/M_q$, α is the share of the arriving condensate routed to the layer, β the share of the outflow it supplies, $UA_i = 0.2\,\mathrm{kW/K}$, and $k = 1\,\mathrm{s}^{-1}$ times $\mathrm{clip}_{[0,1]}\big((h_q - h_l + 0.5)/0.5\big)$. The receiver's row of Equation (5.2) becomes the sum of the two zones' volume constraints, the layer's weighted by $\omega = \rho_s/\rho_q$, so that $A = (1-\alpha)A_s + \omega\alpha A_q$, $B = (1-\beta)B_s + \omega\beta B_q$ and $C = C_s + \omega C_q$, with the merge flow and the interface heat in D. The saturated zone must keep 3 % of the receiver's volume, or the run leaves the model. With saturated condensate the layer stays empty and the receiver is the homogeneous cell exactly. A check of the application confirms that manual pump steps to 40, 35, 33 and 0 %, which once ended outside the model, now stay inside it.

The receiver's liquid stands $H_r = 2.0\,\mathrm{m}$ above the pump inlet, so the subcooler and the pump suction sit at $p_S = p_c + \rho_{nom} g H_r$, a constant 0.245 bar above the condensing pressure with $\rho_{nom} = 1250\,\mathrm{kg/m^3}$. The potential energy, about 20 J/kg, is left out of the energy balance. The head is what lets a stopped pump restart. At standstill the circuit settles to the cooling-water temperature with the condensing

pressure at 1.303 bar, and the suction still has 5.11 K of subcooling, against 8.72 K at the design point.

5.3.3 Walls, Secondary Water and Conductances

Each heat-transfer cell has a wall node and a secondary-water node in counterflow, as shown in Equation (5.4).

$$\begin{aligned} C_w \dot{T}_{w,j} &= UA_{s,j}(T_{s,j} - T_{w,j}) - UA_{f,j}(T_{w,j} - T_j) \\ C_s \dot{T}_{s,j} &= \dot{m}_s(1-b_s)c_{p,s}(T_{s,j+1} - T_{s,j}) \\ &\quad - UA_{s,j}(T_{s,j} - T_{w,j}) \\ UA_{f,j} &= U(x_{h,j})\left(\dot{m}/\dot{m}_{ref}\right)^{0.6} \\ UA_{s,j} &= \frac{UA_s}{12}\,\phi\left(\frac{\dot{m}_s(1-b_s)}{\dot{m}_{s,0}}\right)^{0.8} \end{aligned} \tag{5.4}$$

where $Q_j = UA_{f,j}(T_{w,j} - T_j)$, $T_{s,13} = T_{s,in}$ and the source leaves at $T_{s,out} = T_{s,1}$. $U(x_h)$ blends the liquid, boiling and vapor conductances on the enthalpy quality x_h, liquid at $x_h \le 0$, a ramp to boiling up to 0.1, boiling to 0.85 and a dry-out ramp to vapor at 1, so it is continuous as a cell changes phase. The working-fluid flow $\dot{m}$ is the pump flow in the evaporator and subcooler and the expander plus bypass flow in the condenser, each flow ratio is floored at 0.05, and $\dot{m}_{ref} = 4.0\,\mathrm{kg/s}$. The low side is written the same way with the cooling water $\dot{m}_c = f_c \dot{m}_{c,max}$ entering the last subcooler cell and leaving the first condenser cell. The accepted heat is $Q_{src} = \dot{m}_s(1-b_s)c_{p,s}(T_{s,in} - T_{s,out})$ and the rejected heat $Q_{rej} = \dot{m}_c c_{p,c}(T_{c,out} - T_{c,in})$. Heat may cross any wall in either direction, because no check-valve or bypass regime prevents it. In case 7 the switch to the user's return moves 0.54 kWh from the water into the condenser metal and 0.12 kWh from the metal into the fluid, and in case 4 and the suggested challenge no heat flows against its normal direction.

5.3.4 Expander, Bypass, Pump and Auxiliaries

The expander swallows a fixed volume per revolution and expands along the isentrope to its built-in volume, after which the pressure jumps to the exhaust at constant volume, as shown in Equation (5.5).

$$\begin{aligned} \dot{m}_{sw} &= y_x \rho_{in} V_d N \\ \dot{m}_L &= y_x A_L \sqrt{\rho_{in}(p_e - p_c)} \\ w &= \eta_e \eta_{su}(h_{in} - h_{ad}) + (p_{ad} - p_c)/\rho_{ad} \\ \rho_{ad} &= \rho_{in}/r_v, \quad s_{ad} = s_{in} \\ \eta_{su} &= \max\left(0,\, 1 - c_{su}(N/N_{50})^2\right), \quad W_f = \dot{m}_{sw} w \\ P_{gen} &= \eta_g W_{sh}\ (W_{sh} \ge 0), \quad W_{sh}/\eta_g\ (W_{sh} < 0) \end{aligned} \tag{5.5}$$

where $W_{sh} = W_f - 2\pi\tau_f N$ is the shaft power, N is in revolutions per second, $N_{50} = 50\,\mathrm{s}^{-1}$ (3000 rpm), and the leakage does no work and mixes with the expanded flow, so $h_{out} = h_{in} - W_f/(\dot{m}_{sw} + \dot{m}_L)$. Under-expansion ($p_{ad} > p_c$) adds the blow-down term and over-expansion subtracts it. At the design point the swept flow is 3.567 kg/s with 0.341 kg/s of leakage (8.7 %), $w = 28.62\,\mathrm{kJ/kg}$ with $p_{ad} = 4.53$ bar, the fluid work is 102.09 kW, friction takes 4.49 kW, the shaft isentropic efficiency is 0.655 and the torque 434 N m. Rotor inertia is neglected, so the drive follows its command through a 0.5 s lag and a 300 rpm/s slew. The vapor bypass is an isenthalpic throttle, $\dot{m}_b = y_b A_b \sqrt{\rho_{in} \min(p_e - p_c, 0.5 p_e)}$, critical above a pressure ratio of two.

The feed pump is a positive-displacement machine drawing from the last subcooler cell, as shown in Equation (5.6), together with the remaining parasitics and the net output.

$$\begin{aligned} \dot{m}_p &= n_p \dot{V}_{p,max} f_{cap} \rho_1, \qquad h_2 = h_1 + \frac{p_e - p_S}{\rho_1 \eta_p} \\ P_{pump} &= \frac{\dot{m}_p (h_2 - h_1)}{\eta_m}, \qquad P_{cool} = P_{c,max} f_c^3 \\ P_{aux} &= P_{stby} + (P_a - P_{stby})\, r_a \\ r_a &= \mathrm{clip}_{[0,1]}\big(\max(n_p, y_x)/0.05\big) \\ P_{net} &= P_{gen} - P_{pump} - P_{cool} - P_{aux} \end{aligned} \tag{5.6}$$

At design the pump head is 10.705 bar, 6.10 kW enters the fluid and the motor draws 6.94 kW, the cooling water costs 5.26 kW and the other auxiliaries 2.00 kW. The actuator positions n_p, f_c, b_s, y_b and y_x respond to their commands as first-order lags with time constants of 1, 2, 3, 1 and 1 s. The electrical boundary is a site bus with nonnegative terms, generation plus import equal to the pump, cooling, auxiliary and generator-motoring demand plus export, with import and export never both nonzero and the residual logged at every step. No data-center load is modeled, so no service can be unmet or curtailed.

5.3.5 Property Tables for R1233zd(E)

The browser never evaluates an equation of state. A generator script evaluates the reference equation of state for R1233zd(E) [14] through the HEOS backend of CoolProp 8.0.0 [17], pinned in its own virtual environment, on four grids. The saturation table has 161 nodes uniform in $\zeta = \ln p$ from 0.2 to 30 bar, with both phases' temperature, densities, internal energies and entropies and their exact slopes along the saturation line, and a 57-node saturated-liquid table is uniform in temperature from 255.15 K to the saturation temperature at 30 bar. The vapor surface has 57×33 nodes in (ζ, ξ), where $T = T_{sat}(p) + \kappa(\xi)(T_{max} - T_{sat})$, $\kappa(\xi) = \xi(\xi + 0.25)/1.25$ and $T_{max} = 450$ K, the equation's limit. The liquid surface has 57×17 nodes in (T, ψ) with $p = p_{sat}(T) + \kappa(\psi)(32\,\mathrm{bar} - p_{sat})$, on a temperature axis because a liquid's density follows its temperature far more than its pressure. The vapor surface stores $\ln\rho$, u and s and the liquid surface ρ and u, each as a deviation from the saturated phase at the same pressure or temperature, so the dome edge is exact, and both are read by bicubic Hermite interpolation with fourth-order finite-difference slopes. The reference state is that of the International Institute of Refrigeration. The SHA-256 checksum of the stored arrays, recomputed by a check of the application from the arrays it reads, matches the one the generator wrote. R245fa was rejected because its equation of state in the same library ends at 440 K, below the 170 °C top of the source range.

The cell closure solves $\rho(p,h) = M_j/V_j$ at fixed p by Newton's method from the cell's previous accepted state, in one variable whatever the phase, and returns $(\partial\rho/\partial p)_h$ and $(\partial\rho/\partial h)_p$ from the same interpolants, so the coefficients of Equation (5.2) are the exact derivatives of the surface being inverted. Against 1500 CoolProp states drawn at random off the grids below 25 bar, the worst errors are 4.9×10^{-7} K in saturation temperature, 8.6×10^{-6} relative in vapor density, 4.2×10^{-8} in liquid density, and 0.68, 0.008 and 0.012 J/kg in the enthalpy the closure returns for vapor, liquid and two-phase cells. The inversions put every state in its phase, and the table's own audit records a Newton residual of 9.7×10^{-13}.

5.3.6 Instruments, Protection and Events

The controllers, the feedforward and the protection's interventions read only a measurement packet. PT-1 and PT-2 read the two pressures through 0.3 s lags with 0.005 and 0.003 bar of noise. TT-1 reads the expander-inlet temperature through a thermowell lag $\tau_{TT} = 3$ s with 0.05 K of noise, and the superheat transmitter computes $SH = T_x - T_{sat}(p_e)$ from the same table. FT-1 reads the feed flow

(0.5 s, 0.3 %), FT-2 the source flow (1 s, 0.5 %), TT-2 to TT-5 the water temperatures (5 s), TT-6 the pump inlet (3 s), giving the suction subcooling against $T_{sat}(p_c + \rho_{nom} g H_r)$, and LT-1 the receiver level (2 s, 0.5 %). Noise is added after the lag, scaled by the noise multiple, and every scan draws the same six deviates from a seeded generator. The CSV export marks each truth column with the prefix `truth_`, so that no replay can use it as a measurement.

Protection sits outside every controller, acts in Manual and changes state only at a step boundary, after the integration. Its interventions act on the instruments. A measured superheat below 3 K, a feed flow below 0.1 kg/s for 2 s, or generator motoring beyond 6 kW for 5 s isolates the expander and opens the vapor bypass. Readmission needs 20 s above 8 K with the feed flowing, a pressure ratio of at least 1.5 and a measured superheat that has fallen by no more than 0.5 K over the last 5 s. An evaporating pressure above 19 bar forces the source bypass open until it has stayed below 17.5 bar for 20 s. Latched trips are a measured superheat below 1 K while the expander still admits, a pressure above 21 bar, vapor or an empty receiver at the pump suction, liquid over half the evaporator's outlet cell for 5 s, and a state outside the table. A trip stops the pump, isolates the expander and opens both bypasses, and Reset trip requires 30 s, the pressure below 17.5 bar, 15 % of receiver level and 2 K of suction subcooling. Warnings at 6 and 25 K of superheat and 17.5 bar carry no action. All of these thresholds are teaching values. Actual ones would be taken from the expander and pressure-equipment manufacturers' data and the applicable codes.

5.3.7 Initialization and Integration

Each case starts at the steady state of its own disturbances and setpoints, solved from the 160 kg charge. Both sides are marched against their secondary water from the end whose state is known, the high side from the expander inlet back to the pump outlet for the flow that closes the march and the low side from saturated liquid at the condenser outlet back to the inlet for the cooling-water flow that does, with the pump work and the expander exhaust coupling them in a fixed-point loop. The receiver takes what the other cells leave of the charge, so its level, 50.2 % at design, is a result. At that state the largest scaled derivative is 1.6×10^{-8} and the steady first law, accepted heat plus pump work minus expander work and rejected heat, closes to 2.6×10^{-8} W. The integrator is classical fourth-order Runge-Kutta, inside the engine of Section 1.6.1. The fastest retained mode, the vapor in the first condenser cells, lies between 35 and 112 s^{-1} over seven operating points, outside the method's stability limit at the 0.05 s physics step, so each step takes three substeps of $1/60$ s, which put $|\lambda h|$ at or below 1.9. The slowest modes have time constants of 17 to 47 s. The controllers scan every 0.25 s and the historian logs one row per simulated second.

5.3.8 Assumptions

Parameters are listed in Table 5.2. Beyond the uniform pressure and homogeneous cells of Section 5.3.1, each wall is one node per cell, and with twelve evaporator cells the superheated region spans one or two cells, so the outlet temperature is resolved coarsely and the exchangers' effectiveness is not converged in the cell count, a sensitivity that Section 5.5.6 quantifies. The expander map is a teaching structure with no fitted machine behind it. The source and cooling water arrive at their slider temperatures whatever the plant does, because the engine and the cooler lie outside the boundary. Cold startup and filling, lubricant, non-condensable gas, pump cavitation and potential energy are outside the model. None of the values describes a particular plant.

5.4 Loops, a Selector and a Protection Layer in Code

One framework scan of $\Delta t_c = 0.25$ s runs the instruments, the flow loop, the evaporating-pressure loop and the high-pressure override. A resettable counter inside it runs the superheat loop, the low-superheat override and the condensing-pressure loop every fourth scan, each handed the time that actually elapsed. Table 5.3 lists the loops with their defaults. Every counter, filter and controller state is created again on a reset or a case reload.

Table 5.2 Model Parameters. Class T Is a Synthetic Teaching Value, L Comes From the Cited Equation of State Through CoolProp 8.0.0, and C Is Fixed in Code. Every Parameter Not Marked H Is Shown and Adjustable in the Parameters Panel

Symbol	Meaning	Value	Unit	Cl.
Source loop and evaporator				
$\dot{m}_{s,0}$, $c_{p,s}$	design flow, heat capacity	5.0, 4300	kg/s, J/(kg K)	T
M_s, UA_s	water held, conductance	55, 108	kg, kW/K	T
V_e, $C_{w,e}$	fluid volume, metal	60, 300	L, kJ/K	T
$UA_{l,e}$, $UA_{b,e}$, $UA_{v,e}$	liquid, boiling, vapor	54, 126, 21.6	kW/K	T
τ_{bs}	source bypass response	3	s	T
Expander, generator and bypass				
V_d, r_v	swept volume, built-in ratio	1.5, 3.0	L, —	T
A_L, η_e	leakage area, internal efficiency	40, 0.85	mm^2, —	T
c_{su}, τ_f	supply loss at 3000 rpm, friction	0.10, 20	—, N m	T
η_g, $P_{g,max}$	generator efficiency, rating	0.94, 120	—, kW	T
N_{min}, N_{max}	speed range	600, 3600	rpm	T
$\dot{N}_{max}$, τ_N	drive slew, response	300, 0.5	rpm/s, s	T
A_b, τ_v	bypass area, valve stroke	650, 1	mm^2, s	T
Condenser, receiver and subcooler				
V_c, V_{hd}, $C_{w,c}$	fluid, header, metal	120, 40, 200	L, L, kJ/K	T
$UA_{v,c}$, $UA_{c,c}$, $UA_{l,c}$	vapor, condensing, liquid	60, 300, 120	kW/K	T
UA_{cw}, M_{cw}	water side, water held	480, 80	kW/K, kg	T
$\dot{m}_{c,max}$, $\dot{m}_{c,0}$	full flow, conductance basis	40, 26	kg/s	T, V/H
$c_{p,c}$, $f_{c,min}$	heat capacity, least flow	4180, 10	J/(kg K), %	T, H/V
$P_{c,max}$, τ_c	pumps, fans at full flow, lag	30, 2	kW, s	T
V_r, H_r, UA_i	receiver, head, layer conductance	150, 2.0, 0.2	L, m, kW/K	T
V_{sc}, $C_{w,sc}$, $M_{sc,w}$	subcooler, metal, water	10, 20, 10	L, kJ/K, kg	T
UA_{sc}, $UA_{sc,w}$	subcooler fluid, water side	4, 16	kW/K	T
Pump, instruments and plant				
$\dot{V}_{p,max}$, τ_p	displacement flow, response	5.0, 1	L/s, s	T
η_p, η_m	hydraulic, motor efficiency	0.55, 0.88	—	T
τ_{TT}	thermowell lag (8 s in case 2)	3	s	T
M_{tot}	charge	160	kg	T
P_a, P_{stby}	auxiliaries, standby	2.0, 0.3	kW	T
$T_{u,ret}$, $T_{u,min}$	user return, least supply	50, 60	°C	T
Fixed in code				
$\dot{m}_{ref}$, exponents	flow basis, film exponents	4.0, 0.6, 0.8	kg/s, —	C
—	layer routing, ramp, merge	0.2, 1.0, 0.5	kJ/kg	C
ρ, u, s, T_{sat}	R1233zd(E) properties	tables	SI	L
—	cells: evaporator, condenser, sub.	12, 8, 2	—	C
ρ_{nom}, g	head density, gravity	1250, 9.80665	kg/m^3, m/s^2	C
—	layer mass, merge rate, zone	0.5, 1, 3	kg, s^{-1}, %	C
Δt, h	physics step, RK4 substep	0.05, 1/60	s	C

Table 5.3 Loops, Scans, Default and Suggested Tunings and Actions. Action −1 Is Direct Acting, +1 Reverse

Loop	PV → MV	Scan	K_c default (suggested)	τ_I	Act.
SIC, superheat	$SH \rightarrow \dot{m}_{SP}$	1 s	0.04 (0.08) kg/s/K	30 (20) s	−1
FIC, feed flow	$\dot{m}_p \rightarrow n_p$	0.25 s	8 %/(kg/s)	1.5 s	+1
override, low SH	$SH \rightarrow \dot{m}_{SP}$	1 s	0.08 kg/s/K, SP 5 K	15 s	−1
PIC-1, evaporating	$p_e \rightarrow N$	0.25 s	250 (800) rpm/bar	8 (5) s	−1
PIC-2, condensing	$p_c \rightarrow f_c$	1 s	40 %/bar	15 s	−1
HPO, high pressure	$p_e \rightarrow b_s$	0.25 s	30 %/bar, SP 17 bar	10 s	−1

5.4.1 The Shared PID Law

Every loop is an instance of the velocity-form PID of Section 1.6.2, written in this chapter's symbols in Equation (5.7).

$$\begin{aligned} e_k &= a(SP_k - PV_k), \quad P_k = K_c e_k \\ D_k &= -K_c \tau_D a \dot{PV}_{f,k} \\ z_k &= z_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ \hat{u}_k &= u_b + z_k + I_k, \qquad u_k = \mathrm{clip}_{[u_{min},u_{max}]}(\hat{u}_k) \\ I_{k+1} &= I_k + \Delta t_k \left[(K_c/\tau_I) e_k + (u_k - \hat{u}_k)/T_t \right] \end{aligned} \tag{5.7}$$

where $a = \pm 1$ is the action, u_b a bias, z_k the running proportional-plus-derivative total (S_k of Equation (1.10)), $\hat{u}_k$ the output before limiting (v_k of Equation (1.11)) and $T_t = \tau_I$ the back-calculation tracking time. An integral time of 3000 s or more removes integral action. All loops here run with $\tau_D = 0$ and without an output slew. A loop in Manual, or a controller that does not hold its actuator, tracks the applied value through its integral as Section 1.6.2 describes, and a PID is never both updated and tracked in the same scan.

5.4.2 Independent Loops and the Cascade

In the independent architecture the superheat controller's output, a flow in kg/s, becomes a pump speed through the pump's nominal capacity, $n_p = 100\,\dot{m}_{SP}/(\dot{V}_{p,max}\rho_{nom})$, so its gain means the same thing in both architectures. In the cascade it becomes the setpoint of the inner flow loop, which moves the pump until FT-1 reads it and so corrects a worn pump or a density change within seconds. The superheat controller's bias is the feedforward of Equation (5.8), the static compensator of Section 1.4.3 written as a steady energy balance.

$$\begin{aligned} \dot{m}_{ff} &= \frac{\hat{\dot{m}}_s c_{p,s} (\hat{T}_{s,in} - \hat{T}_{s,out})}{h(\hat{p}_e, T_{sat}(\hat{p}_e) + SH_{SP}) - \hat{h}_2} \\ \hat{h}_2 &= h\big(\hat{p}_S, \min(\hat{T}_1, T_{sat}(\hat{p}_S) - 0.05)\big) + \frac{\hat{p}_e - \hat{p}_S}{\rho_{nom}\eta_p} \end{aligned} \tag{5.8}$$

where hats are measurements from FT-2, TT-2, TT-3, PT-1, PT-2 and TT-6, and $\hat{p}_S = \hat{p}_c + \rho_{nom} g H_r$. It is zero when the enthalpy rise is below 10 kJ/kg, it passes a 2 s first-order filter, and switching it in or out moves its step into the controller's integral so the output does not jump. While the generator is at its rating the feedforward is capped at the flow the speed-capped expander can pass, because extra feed beyond it only floods the evaporator. It is a measured-heat feedforward, built from the source flow and both source temperatures, rather than a feedforward of the flow alone.

The low-superheat override joins through a minimum selector with external reset. Each supervisory scan both controllers' outputs are previewed from their states without changing them, $\hat{u} = u_b + K_c a(SP - PV) + I$, and the lower is selected, as shown in Equation (5.9).

$$\begin{aligned} I_o &\leftarrow \hat{u}_{SIC} - u_{b,o} \quad \text{while idle, so} \\ \hat{u}_o &= \hat{u}_{SIC} + K_{c,o}(SH - SH_{min}) \end{aligned} \tag{5.9}$$

The idle override therefore sits exactly its proportional action above the main request and takes the flow setpoint as soon as the superheat falls below $SH_{min} = 5$ K. While it is selected the main controller tracks its output, and when the pump runs at full speed the main controller's ceiling falls to the measured flow, so neither winds up while the pump cannot deliver more. A protective override that waits idle until its constraint nears is also the subject of Volume 4, Chapter 2.

5.4.3 Pressure Loops, the Overrides and Isolation

PIC-1 moves the expander speed between 600 and 3600 rpm. When the measured generator power exceeds 98 % of its 120 kW rating the drive caps the speed at $N(0.98P_{g,max})/P_{gen}$, and the loop's

upper limit follows that cap. While the protection has the expander isolated, PIC-1 keeps holding the evaporating pressure through the vapor bypass, whose opening is set to pass what the expander would swallow at the speed the loop asks, $y_b = (\hat{\rho} V_d N + A_L \sqrt{\hat{\rho}\Delta\hat{p}})/(A_b \sqrt{\hat{\rho}\,\min(\Delta\hat{p}, 0.5\hat{p}_e)})$ with $\hat{\rho}$ the vapor density at the measured inlet state, and the drive idles at that speed so that readmission hands the same flow back (a flying start). On a trip or an unload the bypass opens fully and the drive stops. PIC-2 moves the cooling-water flow between 10 and 100 %. The high-pressure override, available in the cascade, opens the source bypass when the evaporating pressure passes 17 bar and so removes heat at the source rather than fighting it with the expander.

The application keeps three models apart. The plant is everything in Section 5.3. No model is identified for the loop structures, and the gain a learner reads from case 1's step, about −24 K per kg/s at fixed expander speed, is the learner's. The loops' only internal models are the steady energy balance of Equation (5.8) and the bypass characteristic, both on the plant's own property table and parameters, so they meet mismatch only when a case changes the plant, as fouling does in case 6. The predictive controller of Section 5.4.4 is the exception, because its model is identified once at the design point and then held fixed.

5.4.4 Model Predictive Control

A third architecture replaces the superheat controller and PIC-1 with one model predictive controller (MPC), the option of Section 1.6.3. Every $T_s = 1\,\mathrm{s}$ it plans the flow setpoint $\dot{m}_{SP}$ of the inner flow loop and the expander speed N over $N_p = 120$ steps in 12 move blocks, whose lengths grow from one step for each of the first four to 75 steps for the last, and it holds the first move until the next plan. Its flow request takes the superheat controller's place in the minimum selector of Equation (5.9), so the overrides, the drive's power cap, PIC-2, the inner loop and the protection act exactly as under the cascade. Industrial predictive controllers are usually arranged in this way, writing setpoints to regulatory loops [18], and Volume 4, Chapter 10 plans the bypass and both inner-loop setpoints of a closed carbon dioxide power cycle in the same way.

The model was identified by step tests on the simulator and is fixed in the source, so it reads neither the plant nor a run's parameters. From case 1 lined out with the declared parameters and no noise, with the inner loop in Auto under a constant setpoint, PIC-1 in Manual and PIC-2 in Auto, each input was stepped up and down by 0.06 kg/s of flow setpoint, 60 rpm, 5 % of source flow or 3 K of supply, and the mean response of the measured superheat and PT-1 over 300 s was fitted by linear least squares [19], as Equation (5.10) states.

$$\begin{aligned} \hat{y}(t) = \sum_{j=1}^{4}\sum_{i=1}^{3} \Big[& \alpha_{ij}\big(1 - \mathrm{e}^{-t/\tau_i}\big) \\ & + \beta_{ij}\Big(1 - \big(1 + \tfrac{t}{\tau_i}\big)\mathrm{e}^{-t/\tau_i}\Big)\Big]\Delta v_j \\ v = \big(\ln\dot{m}_{SP},\ \ln N,\ \ln\dot{m}_s,\ T_{s,in}\big), \quad & \tau_i = 2,\ 8,\ 30\,\mathrm{s} \end{aligned} \tag{5.10}$$

Here $\hat{y} = (SH, p_e)$, each α_{ij} and β_{ij} holds one coefficient per output, and the thermowell, the inner loop, the drive and PIC-2 lie inside the fit, whose RMS residuals are 0.09 to 0.45 % of each response. The steady gains include −25.43 K and 1.298 bar per kg/s of flow setpoint and 0.022 06 K and −0.003 538 bar per rpm. Identified again at 70 % source flow, the gains per physical unit change by factors of 1.07 to 1.74 but those per unit of the logarithm only by 0.87 to 1.22, because the superheat follows the heat that each kilogram receives. Each coefficient is therefore scaled by its input's design value, 3.909 kg/s, 2145 rpm or 5 kg/s. The 24 lag states and two random-walk disturbances on $\ln\dot{m}_{SP}$ and $\ln N$ are discretized with an exact hold at T_s. Against five unfitted steps from lined-out points at 100, 70 and 60 % source flow, the open-loop superheat error is 0.23, 0.60 and 1.17 K RMS.

Each plan minimizes Equation (5.11), the objective of Equation (1.14) with move blocks and soft output limits, over the block values of v_1 and v_2, taken as deviations from the commands at the

estimator's anchor.

$$\begin{aligned}\mathscr{J} &= \sum_{k=1}^{N_p}\Big[q_{SH}\big(\widehat{SH}_k - SH_{SP}\big)^2 + q_p\big(\hat{p}_{e,k} - p_{e,SP}\big)^2\Big] \\ &\quad + \lambda\sum_{j=1}^{12}\Big[\Big(\frac{\Delta v_{1,j}}{s_1}\Big)^2 + \Big(\frac{\Delta v_{2,j}}{s_2}\Big)^2\Big] + \sum_{o=1}^{2}\big(\rho\,\varepsilon_o^2 + \sigma\varepsilon_o\big) \qquad (5.11)\\ \text{s.t. } \widehat{SH}_k &\geq \underline{SH}_k - \varepsilon_1, \quad \hat{p}_{e,k} \leq \overline{p}_k + \varepsilon_2, \quad \varepsilon_o \geq 0 \\ \dot{m}_{min} &\leq \dot{m}_{SP,j} \leq \dot{m}_{max}, \quad N_{min} \leq N_j \leq N_{lim}\end{aligned}$$

The weights $q_{SH} = 1\,\mathrm{K}^{-2}$ and $q_p = 100\,\mathrm{bar}^{-2}$ are the inverse squares of the loops' 2 K and 0.2 bar trend bands and were fixed before tuning. The move scales $s_1 = 0.010$ and $s_2 = 0.021$, about 0.04 kg/s and 45 rpm, each shift the superheat by about 1 K, and $\lambda = 0.03$. The bounds $\underline{SH}_k$ and $\overline{p}_k$ are the setpoints of the two overrides, $SH_{min} = 5\,\mathrm{K}$ and $p_{max} = 17\,\mathrm{bar}$, read as the overrides read them, and both are soft, with one slack per output priced at $\rho = 10^4$ per unit squared and $\sigma = 10^3$ per unit. A limit that the estimate already violates is approached along a funnel, $\underline{SH}_k = SH_{min} + (\widehat{SH}_0 - SH_{min})\mathrm{e}^{-kT_s/\tau_f}$ with $\tau_f = 20\,\mathrm{s}$, and the ceiling mirrors it. The flow setpoint is bounded hard by $\dot{m}_{min} = 0.25\,\mathrm{kg/s}$ and the superheat controller's own ceiling, and the speed by $N_{min} = 600\,\mathrm{rpm}$ and $N_{lim} = \min(3600\,\mathrm{rpm}, N_{cap}, 0.98NP_{g,max}/P_{gen})$, the drive's power cap of Section 5.4.3 and its formula applied before the cap engages. Each decision is one receding-horizon quadratic program [20], which a dual active-set solver [21] completes in a mean of 1.6 ms in a headless engine.

A steady-state Kalman filter for the 26-state model makes the plan offset-free [22], with measurement variances of $(0.055\,\mathrm{K})^2$ and $(0.005\,\mathrm{bar})^2$ and random-walk intensities $\gamma r/G^2$, where G is the gain of the flow on the superheat or of the speed on the pressure, r the variance of that output, and $\gamma = 0.1$. It runs under every architecture on the mean commands of the last second and on FT-2 and TT-2, stops while the expander is out and anchors again on its return. The plan holds the supply at TT-2 and the source flow at FT-2, the feedforward's own signal, unless the preview adds the case's announced source-flow changes at the declared 5 kg/s per 100 %. With the preview on, the MPC is therefore told in advance what the cascade can only measure, and no shipped option gives the cascade the schedule. In case 4 and the challenge the scheduled change is the source loss itself.

The values of λ and γ came from one pass on scenario 9 over 11 values of λ and seven of γ. Among settings with no rejected plan, trip, intervention or extra travel on either handle, it scored the mean of the superheat and pressure IAEs relative to the cascade with feedforward and overrides, and it took the largest λ within 10 % of the best score, then the smallest γ. The rated-speed term of N_{lim} was added after this pass as a correction, and repeating the pass with it selects the same values. No case is retuned.

The MPC acts only with the superheat and pressure loops in Auto and the expander in service. Otherwise, or when a plan fails to reach its optimum, holds a value that is not finite or leaves its bounds, the superheat controller and PIC-1, which track the applied commands every scan, take over, the historian records an MPC FALLBACK mark, and control returns with the first valid plan. With the noise off, three forced failures at 104 s of scenario 9 move the flow setpoint by at most 0.000 65 kg/s and the speed by 0.29 rpm per scan, against 0.003 71 kg/s and 1.30 rpm just before. Every protective action remains shipped code. After an unannounced loss of source flow to 20 % at 130 s in case 4, both structures isolate the expander on the first measured superheat below 3 K, at 183.30 and 192.30 s, and the MPC solves no plan until it acts again 0.5 s after the readmission.

5.5 What Structure Recovers and What the Plant Withholds

Table 5.4 lists the ten cases with the requirement that credits each one, in the scenario form of Section 1.6.4. Every result below was measured by a script that runs the shipped engine headless at the default seed. It loads a case as the case menu does, applies a named setting ahead of the

Table 5.4 The Ten Cases, Their Events and the Requirement That Completes Each

#	Case, controller and events	Lesson	Required to complete
1	Manual, every loop held; learner steps the pump	which signal moves first	After 10 s raise the pump to at least 64.2 %, hold it for 150 s, and confirm when the pressure, superheat and net power moved.
2	Independent PI 0.02, 90 s; SP 10 to 7 to 10 to 7 K at 60, 360, 780 s; noise ×3 and τ_{TT} 3 to 8 s at 660 s	tuning for the measurement	Set $K_c \geq 0.05$, hold the true superheat within 0.5 K of setpoint for 30 s between 420 and 540 s, run past the events, and confirm the comparison of IAE and travel.
3	Cascade, slow inner loop, PIC-1 in Manual; p_e SP 12.9 to 13.4 bar at 300 s and back at 600 s	loop interaction	Put PIC-1 in Auto before 300 s, see the superheat dip below 8 K, raise the inner gain to at least 5, and confirm the coupling.
4	Cascade, feedforward off; source flow to 70 % at 120 s, back at 480 s	feedback against feedforward	See the expander isolated with the feedforward off, switch it on and Reset, and keep the expander admitted with no intervention past 300 s.
5	Suggested cascade; cooling water 25 to 35 °C over 60 to 180 s	condensing capacity	See PIC-2 LIMITED after 150 s, raise its setpoint to at least 2.59 bar, hold net power above 60 kW for 60 s, and confirm the trade.
6	Cascade with feedforward, 0.04, 30 s; source side fouls to 60 % over 60 to 360 s	changed plant or bad tuning	Run past the fouling, see accepted heat below 870 kW with the superheat within 3 K for 30 s after 420 s, and confirm the reading of the temperatures.
7	Suggested cascade; condenser water becomes the user's 50 °C return at 60 s	useful heat costs work	After the switch raise the condensing setpoint to at least 4.19 bar, deliver useful heat for 60 s, and confirm the penalty.
8	55 °C server water at 20 kg/s, p_e 2.7 bar, SH 7 K; cooling water 15 to 22 °C over 120 to 300 s	when parasitics win	See net power below zero for 30 s after 120 s, unload the cycle and hold net above −1 kW for 30 s, and confirm the Carnot comparison.
9	MPC (Section 5.4.4) over the suggested cascade, SH_{SP} 7 K, schedule preview; source flow to 60 % at 120 s, back at 480 s, both announced	predictive against feedback control	Run to 600 s under the MPC, Reset, switch to the cascade before 120 s, run to 600 s again, and confirm the comparison of lowest superheat, time below 5 K and net energy.
C	Challenge, independent PI 0.03, 60 s: source flow to 70 % at 120, cooling water 25 to 32 °C over 480 to 540 s, pump delivery 70 % at 840 s	the whole problem	Run 1200 s with no trip, the true superheat within 5 to 20 K from 60 to 1200 s, back within ±2 K of setpoint no later than 120 s after each event, $p_e \leq 19$ bar, and the receiver at 15 to 85 % at the end.

first integration step, which makes it an initial condition, applies timed actions at their instants, and advances the plant at its own 0.05 s step. Each run starts from a private copy of the case. Integrated absolute error (IAE) is taken on the true superheat at every physics step over the stated window, and recovery is the last instant in the window at which the true superheat lies outside the stated band about its setpoint, after which it stays inside to the window's end. Pump travel is the total variation of the commanded pump speed, and stored inventory is the receiver's true liquid level at the window's start and end.

5.5.1 The Dynamic Signature

Case 1 shows the coupling before any loop closes. Raising the pump by two points adds 0.122 kg/s, the measured evaporating pressure moves 0.80 s later because the expander passes a fixed volume per turn, the measured superheat starts to fall after 2.55 s, and the true superheat covers 90 % of its change in 48.5 s and settles at 7.05 K, a gain of −24.1 K per kg/s. Warming the cooling water from 25 to 30 °C then raises the condensing pressure from 1.969 to 2.290 bar and lowers net power from 78.86 to 74.04 kW, and a five-point pump step reaches the 3 K intervention after 34.05 s. In case 3 the evaporating-pressure setpoint rises by half a bar, which lifts the saturation temperature by 1.85 K at once, so with PIC-1 in Auto the true superheat falls to 6.44 K although the pressure loop never touches the pump. The IAE after the step is 549 K s with a slow inner flow loop and 113 K s with the inner loop tuned first, against 7 K s with PIC-1 left in Manual. Small steps confirmed every loop's sign, the pump's at 13 of 13 lined-out points and the expander's at all five points below the generator's rating.

🕮 **Pause.** Try Exercise 5.2 before reading on.

Table 5.5 Matched Comparisons on the Simulator's Truth. Recovery Is the Last Exit From the Band About the Setpoint Within the Window. Limited Is the Time the Expander Was Isolated (A) or PIC-2 Was LIMITED (B). Energies Are Over the Window, the Receiver Level at Its Start and End

Setting	IAE (K s)	Max dev. (K)	Recovery (s)	*SH* range (K)	Limited (s)	Pump travel (%)	Net (kWh)	Accepted (kWh)	Receiver (%)
A. Case 4, 120–480 s, ±2 K									
independent PI, 0.08, 20 s	276.7	8.17	50.0	1.83–10.25	54.9	42.9	4.755	72.25	50.2 → 50.0
cascade, no feedforward	437.2	8.91	76.0	1.09–10.14	83.4	126.7	4.187	72.47	50.2 → 50.0
cascade with feedforward	29.0	1.08	0	9.39–11.08	0	129.4	5.819	71.83	50.3 → 50.0
MPC, preview off	27.9	0.75	0	9.25–10.24	0	126.6	5.812	71.82	50.3 → 50.0
MPC, loss announced	24.6	0.68	0	9.49–10.68	0	129.0	5.812	71.81	50.0 → 53.2
A. Case 4, 480–900 s, ±2 K									
independent PI	261.4	9.07	54.3	9.87–19.07	0	53.2	8.939	111.92	50.0 → 50.2
cascade, no feedforward	254.8	9.30	48.1	9.91–19.30	0	181.6	8.934	111.93	50.0 → 50.2
cascade with feedforward	38.0	1.08	0	8.92–10.61	0	186.4	8.938	112.58	50.0 → 50.3
MPC, preview off	36.9	1.36	0	9.70–11.36	0	185.2	8.941	112.64	50.0 → 50.2
MPC, loss announced	37.0	1.31	0	9.63–11.31	0	184.2	8.944	112.64	53.2 → 50.2
B. Case 5, 240–900 s, ±2 K									
p_c SP held at 1.95 bar	22.9	0.16	0	9.89–10.16	660	280.9	8.689	171.04	50.7 → 50.7
p_c SP 2.6 bar from 240 s	22.7	0.16	0	9.89–10.16	0	281.4	12.235	169.99	50.7 → 51.1
p_c SP 2.8 bar from 240 s	223.0	0.79	0	9.46–10.79	0	285.3	12.596	169.19	50.7 → 51.6
C. Case 2, 360–660 s, ±0.5 K, before the sensor change									
0.02, 90 s (shipped)	286.0	2.02	>300	7.98–9.40	0	6.4	6.512	—	—
0.04, 30 s	125.5	3.17	85.0	6.83–10.06	0	12.9	6.465	—	—
0.08, 20 s	50.8	3.34	31.5	6.66–10.17	0	26.2	6.444	—	—
C. Case 2, 780–1080 s, after it									
0.02, 90 s (shipped)	416.0	2.54	>300	—	0	19.2	—	—	—
0.04, 30 s	116.7	3.00	71.0	—	0	38.1	—	—	—
0.08, 20 s	64.9	3.04	47.1	—	0	76.2	—	—	—
D. Case 6, 360–900 s									
0.08, 20 s (clean-plant tuning)	844.0	2.92	—	7.86–12.92	0	247.9	10.193	125.83	—
0.04, 30 s (shipped)	68.6	0.75	—	9.74–10.75	0	192.0	10.213	125.23	—

A uses the suggested gains in every row (SIC 0.08 kg/s per K and 20 s, PIC-1 800 rpm/bar and 5 s), and the second row is case 4 as shipped. The MPC rows (Section 5.4.4) run over the same cascade with scenario 9's settings, and with the preview on the case's scheduled loss is announced to the plan. B and D run the cascade with the feedforward. C runs the independent architecture, and its windows are the 300 s after each 3 K setpoint step. A dash marks a quantity not computed for that window, and >300 a run still outside ±0.5 K at the window's end. Halving the physics step to 0.025 s gives IAEs of 276.66 and 29.01 K s in A's first window and 261.48 and 38.02 K s in its second.

5.5.2 Three Matched Comparisons

Each group in Table 5.5 shares the initial state, the seed, the instruments, the limits and the events, and differs only in the controller settings named. No load is served by this plant, so unmet and curtailed service are not applicable, and the accepted heat takes the place of fuel.

Comparison A tests nominal disturbance rejection on case 4's loss of 30 % of the source flow. The feedforward reads the lost heat from FT-2, TT-2 and TT-3 and lowers the flow setpoint within seconds, while either feedback structure waits for the thermowell, as shown in Figure 5.2a. Independent PI at the suggested gains lets the true superheat fall to 1.83 K, and the protection isolates the expander from 138.05 to 193.00 s. The case's own cascade without the feedforward does worse, 437.2 against 276.7 K s, mostly because it is readmitted later, at 220.75 s. Its low-superheat override shares the lagging measurement, takes the flow setpoint about 15 s after the loss, holds it for 41 s, and cannot prevent the intervention. The cascade with the feedforward stays within 1.08 K of setpoint with no intervention and makes 5.819 kWh against 4.755 and 4.187 kWh over the same six minutes. After the restoration at 480 s feedback lets the superheat reach 19.07 K and the feedforward 10.61 K, while all three runs keep the expander and their energies agree to 0.005 kWh. The feedforward costs actuator travel, 129.4 against 42.9 % of pump speed, mostly because the inner flow loop answers FT-1's noise every 0.25 s, as the first 120 s show with no event at all (48.4 against 9.0 %). On the pressure–enthalpy

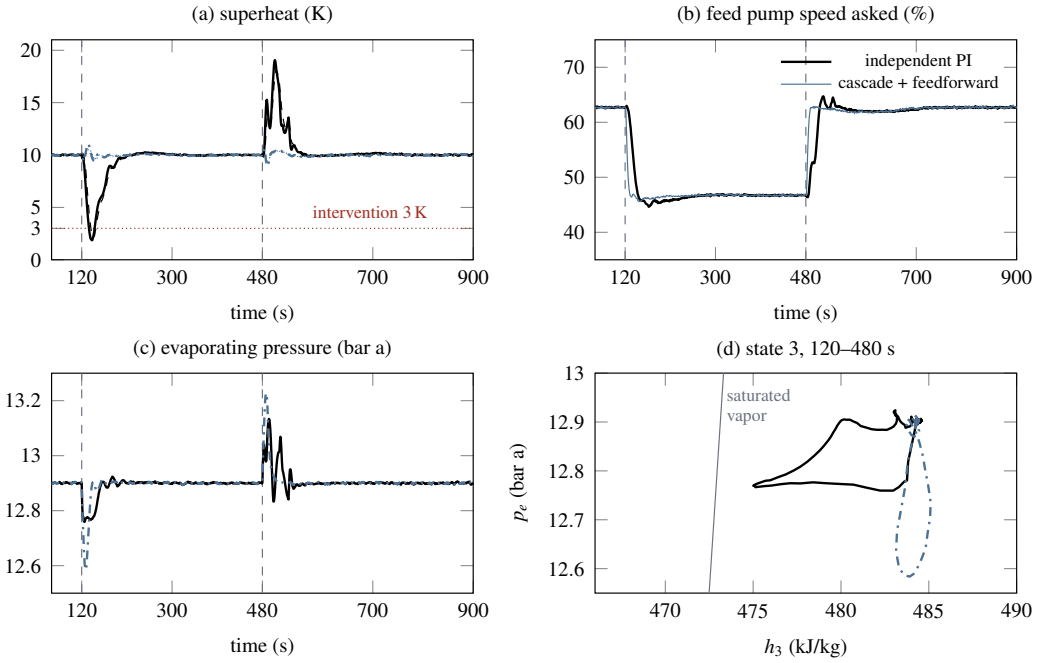


Figure 5.2 Comparison A, case 4, from the exported records: independent PI (black, solid) and the cascade with the source-heat feedforward (blue, dash-dotted) at the same gains, the source-loop flow falling to 70 % at 120 s and restored at 480 s (gray dashed verticals). (a) True superheat (thick) and the TT-1 and PT-1 measurement (thin, dashed for independent PI and dotted for the feedforward) against the 3 K intervention (dotted horizontal). (b) Commanded pump speed, independent PI thick and the feedforward thin. (c) True evaporating pressure. (d) The expander inlet, state 3, on the pressure–enthalpy plane from 120 to 480 s beside the saturated-vapor line (thin gray).

plane of Figure 5.2d the feedback run's inlet state slides to 475.0 kJ/kg, 1.9 kJ/kg from saturated vapor, and loops back at a lower pressure while the bypass holds the evaporator, and the feedforward run's stays within 1.2 kJ/kg of the design enthalpy. Figure 5.3 draws the first window of this comparison as bars.

> **TRY IT IN THE APP** **4 · Lose Source Flow: Feedback or Feedforward**
>
> Load the scenario and press Run. At 120 s the source flow falls to 70 %. Watch the superheat and the protection, which isolates the expander 17.3 s after the loss and keeps it out for 83.4 s. Then set Source-heat feedforward (cascade) to On, press Reset and run again. The true superheat stays between 9.39 and 11.08 K, and the expander never leaves the line.

Pause. Try Exercise 5.1 before reading on.

The second comparison is a capacity limit. In case 5 the cooling water warms from 25 to 35 °C over two minutes, PIC-2 reaches 100 % flow, and the condensing pressure floats to 2.383 bar, where the fans draw 30.0 kW and the cycle nets 47.47 kW, as shown in Figure 5.4a and b. No tuning of PIC-2 removes that limit, and the loop reads LIMITED for all 660 s of the window. Raising its setpoint to 2.6 bar at 240 s gives up generator energy, 15.303 against 15.807 kWh, to save cooling energy, 1.467 against 5.500 kWh, and nets 12.235 against 8.689 kWh. Lined out, the net power is 47.47 kW at the limit, 67.55 kW at 2.6 bar and 68.76 kW at 2.8 bar. The higher setpoint adds only 0.361 kWh over the window, and its larger move disturbs the superheat, an IAE of 223.0 K s against 22.7, because the

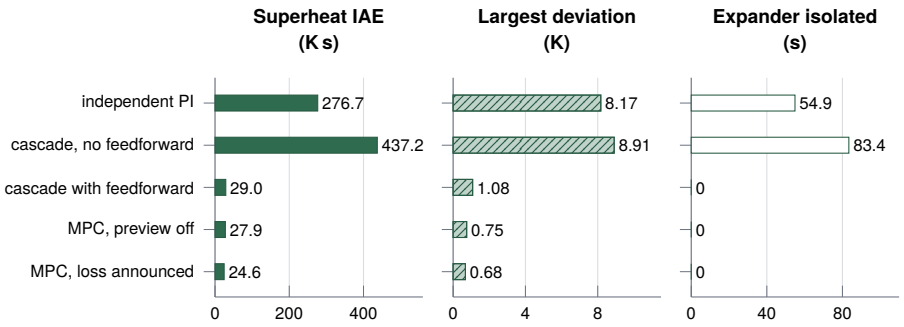


Figure 5.3 Comparison A of Table 5.5 redrawn: case 4 from 120 to 480 s, the source flow at 70 %, the PI and cascade rows at the suggested gains and the MPC rows with scenario 9's settings. Bars give the integrated absolute error of the true superheat, its largest deviation from setpoint and the time the protection kept the expander isolated. Notice that the two structures that wait for the thermowell lose the expander, while every row that acts on the source meters keeps it, with an error below 30 K s.

cooling power left to save shrinks with the cube of the flow. The superheat barely notices the sink in either of the first two runs (0.16 K at most), which is why the challenge's warm-water event costs power but not band time.

> **TRY IT IN THE APP** **5 · A Hot Day at the Condenser**
>
> Run the scenario as loaded. The cooling water warms from 25 to 35 °C, PIC-2 reads LIMITED with the cooling water at full flow, and the net power lines out near 47.47 kW. Run again and, at about 240 s, raise the Condensing pressure setpoint to 2.6 bar. PIC-2 leaves LIMITED, the cooling power falls, and the net power lines out near 67.55 kW.

Pause. Try Exercise 5.3 before reading on.

The third comparison is a sensor mismatch. Case 2 slows the thermowell from 3 to 8 s and triples the noise at 660 s, and the three tunings are compared over the 300 s after equal 3 K setpoint steps before and after the change. The fast tuning, 0.08 kg/s per K with 20 s, goes from 50.8 to 64.9 K s and from 26.2 to 76.2 % of pump travel, the moderate one from 125.5 to 116.7 K s with travel rising from 12.9 to 38.1 %, and the cautious shipped tuning never settles within 0.5 K in either window, as shown in Figure 5.4e and f. Over the longer window from 780 to 1200 s that the case's own texts use, the three IAEs are 493.5, 122.3 and 71.9 K s. The faster loop still tracks best, but its margin over the moderate one shrinks from 75 K s to 52 K s while its travel nearly triples. The likely cause is the phase margin that an 8 s lag removes from a loop whose plant answers in tens of seconds, while the tripled noise reaches the pump directly. A model mismatch gives the same lesson in case 6. When the source side fouls to 60 % of clean, accepted heat falls from 960.5 to 835.2 kW and the source leaves at 111.16 °C instead of 105.33 °C. The clean-plant tuning of 0.08 kg/s per K with 20 s then cycles between 7.86 and 12.92 K with a period of about 34 s, an IAE of 844.0 K s, while the shipped 0.04 kg/s per K with 30 s holds 9.74 to 10.75 K at 68.6 K s, and both net the same energy within 0.02 kWh. The fouling is a changed plant that the temperatures reveal.

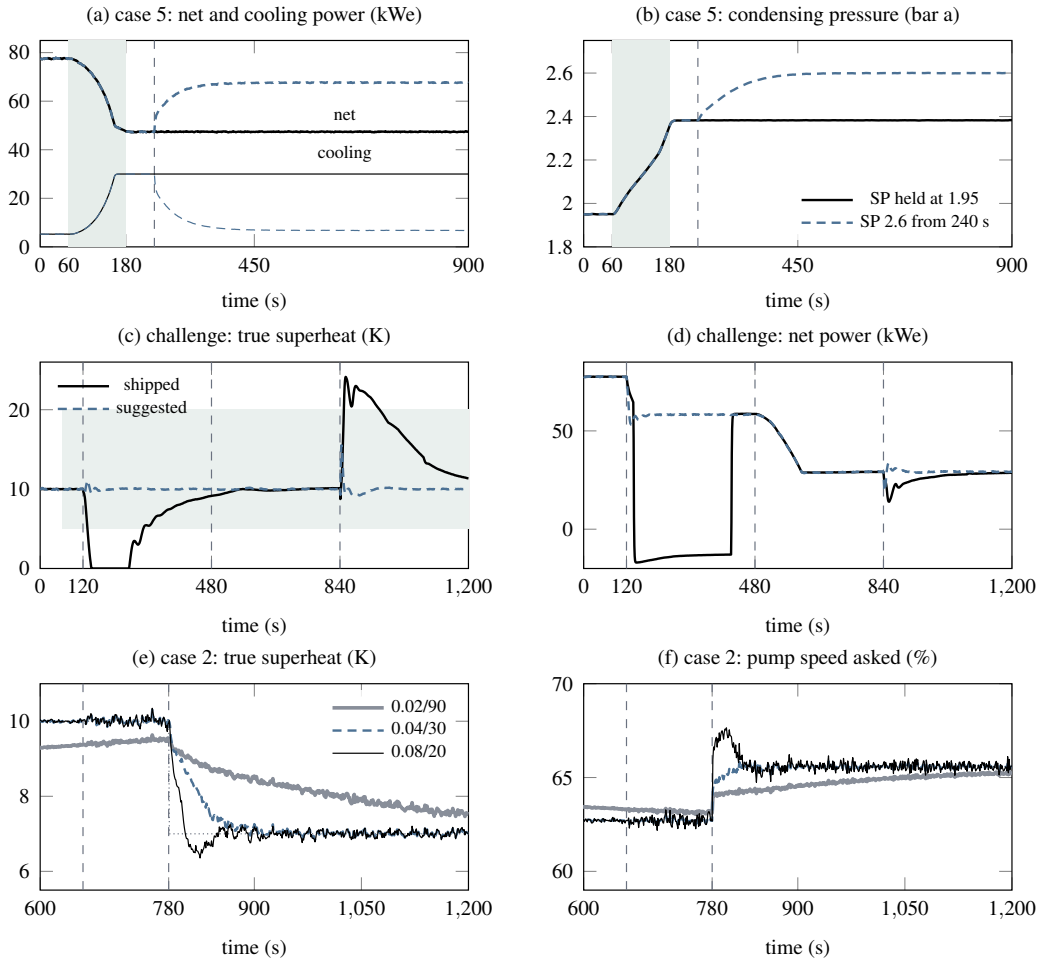


Figure 5.4 Comparisons B and C and the challenge, from the exported records. (a) Case 5, net power (thick) and cooling power (thin) with the condensing setpoint held at 1.95 bar (black, solid) and raised to 2.6 bar at 240 s (blue, dashed), the shaded band the cooling-water ramp from 25 to 35 °C. (b) The same runs' true condensing pressure. (c) The challenge's true superheat as shipped (black, solid) and with the suggested settings (blue, dashed) against the shaded 5 to 20 K band, the events at 120, 480 and 840 s as gray dashed verticals. (d) The challenge's net power, the same line styles. (e) Case 2, true superheat for the three tunings after the thermowell and noise change at 660 s and the setpoint step at 780 s: 0.02 and 90 s very thick gray, 0.04 and 30 s blue dashed, 0.08 and 20 s thin black, the setpoint dotted. (f) Their commanded pump speed.

> **+ DID YOU KNOW? The Calmer Tuning Wins on a Fouled Plant**
>
> From the simulator: with the source side fouled to 60 % of clean, the clean-plant tuning of 0.08 kg/s per K and 20 s cycles between 7.86 and 12.92 K, an integrated error of 844.0 K s. The calmer shipped tuning holds 9.74 to 10.75 K at 68.6 K s (Table 5.5).

5.5.3 Where Added Control Cannot Help

Case 8 runs the same machine on 55 °C server water at 20 kg/s with an evaporating pressure of 2.7 bar. At 100 s the cycle accepts 128.5 kW and its generator makes 3.29 kW, 2.56 % of the heat, at a

Table 5.6 The Challenge Over 1200 s. No Configuration Trips, Peaks Above 13.43 bar or Ends Outside 50.0 to 50.3 % of Receiver Level

Configuration	Out (s)	*SH* range (K)	Recovery (s)	Isol. (s)	IAE (K s)	Net (kWh)	Verdict
shipped, PI 0.03/60	256.7	0.00–24.14	266/0/305	274.2	4311	9.16	fails
PI 0.08/20	26.6	3.64–22.70	55/0/58	0	640	14.89	fails
PI 0.08/20, PIC-1 800/5	20.7	1.83–18.61	50/0/52	54.9	629	13.79	fails
cascade, no feedforward	39.6	1.09–15.57	76/0/9	83.4	521	13.27	fails
cascade + feedforward, shipped gains	0	6.77–18.89	42/0/16	0	395	14.92	passes
suggested	0	7.91–15.56	0/0/9	0	134	14.90	passes
MPC, loss announced	0	6.64–16.01	0/0/30	0	160	14.89	passes
MPC, preview off	0	6.64–16.01	0/0/30	0	161	14.89	passes

Out is the time outside 5 to 20 K from 60 to 1200 s. Recovery is the last exit from ±2 K after the events at 120, 480 and 840 s, allowed 120 s, and the shipped run is late twice. IAE is from 60 to 1200 s. The cascade rows use the suggested gains except the fifth. The MPC rows switch only the architecture of the challenge as loaded, and over the suggested settings the first is unchanged.

pressure ratio of 2.45, which leaves 1.02 kW after 2.28 kW of parasitics. As the cooling water warms from 15 to 22 °C, the condensing loop asks for more flow, net power turns negative at 204.3 s, and with the fans at full flow from 389 s it reaches −30.6 kW. The Carnot bound between the server water and the cooling water falls from 12.19 to 10.06 % over the ramp. A sweep of 192 steady settings at 20 °C cooling water, the evaporating pressure from 1.8 to 3.2 bar, the superheat from 3 to 12 K and the condensing pressure from 0.9 to 1.6 bar, solves every setting and finds no positive net output, the best being −0.59 kW. Unloading the cycle leaves −0.33 kW of standby load, so the correct operation is to stop, because the temperature head cannot pay for the pumps and fans. Routine server heat would not be sent to a power cycle, but the stylized case shows where parasitic power overtakes a vanishing temperature difference. The cooling of a data hall is itself the subject of Volume 9, Chapter 8.

The generator's rating sets a second limit at the other end of the source range. A step of the supply from 150 to 170 °C with the suggested settings caps the drive for 511 s, the generator ends at 117.7 kW, the evaporating pressure peaks at 17.12 bar, and the high-pressure override opens the source bypass by up to 4.4 %, with the superheat no lower than 3.57 K and no intervention. Without the feedforward the pressure peaks at 17.32 bar and the bypass opens to 9.5 %. Control keeps the plant inside its limits in both runs, but the heat the generator cannot convert is sent around the evaporator.

5.5.4 The Challenge

The challenge starts at the design point with the shipped independent PI at 0.03 kg/s per K and 60 s. It loses 30 % of the source flow at 120 s, warms the cooling water from 25 to 32 °C over 480 to 540 s, and wears the pump to 70 % of its rated delivery at 840 s, and Table 5.6 lists eight configurations. As shipped the expander is isolated from 139.55 to 413.75 s. The true superheat reaches 0 K at 144 s, while the closing isolation valve still admits 1.4 % and the bypass carries the wet vapor, and the wear later lifts it to 24.14 K. Each of the two scored events needs its own structural element. Faster tuning of the independent loops leaves 20.7 to 26.6 s outside the band, the cascade without the feedforward fixes the wear but not the source loss, and the cascade with the feedforward passes even at the shipped gains. The suggested settings keep the superheat between 7.91 and 15.56 K and recover within 8.7 s of the wear, as shown in Figure 5.4c and d. The warm water costs every configuration the same power, with net output at the end between 28.69 and 29.23 kW and the cooling water at its full flow. Over five further seeds the shipped run spends 256.0 to 256.8 s outside the band and fails each time, while the suggested run stays within 7.90 to 15.60 K and passes each time. Figure 5.5 shows the suggested run at 1200 s.

5.5.5 Predictive Control Through a Scheduled Turndown

Scenario 9 runs the superheat at 7 K, 2 K above the floor of the low-superheat override, and the on-site engine's announced schedule takes the source flow to 60 % at 120 s and back at 480 s. The thinner margin makes a little more power, and the turndown is deeper than case 4's. The case loads under the

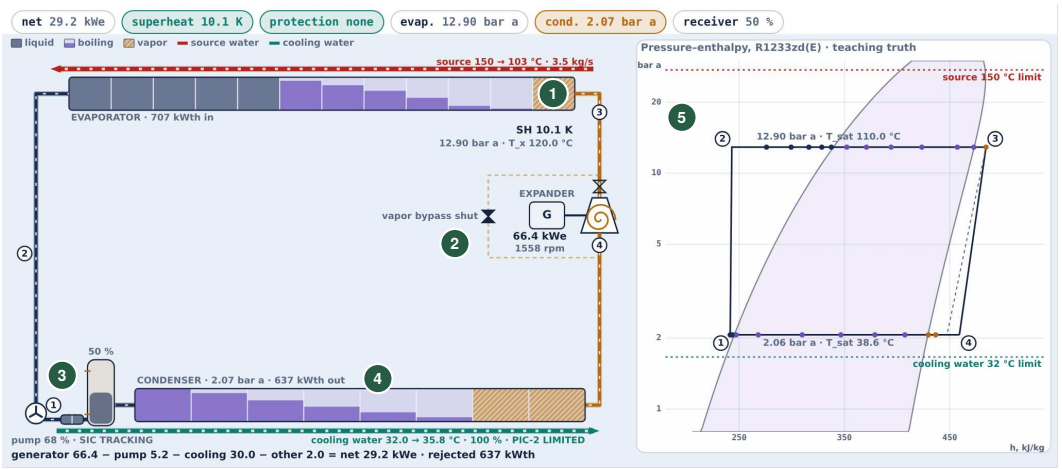


Figure 5.5 End of the suggested challenge run, $t = 1200$ s, on the Cycle view, cropped to the plant. (1) The hatched vapor cells at the evaporator outlet, where the superheat reads 10.1 K at 12.90 bar with the expander inlet at 120.0 °C. The evaporator accepts 707 kW from source water cooled from 150 to 103 °C at 3.5 kg/s. (2) The expander, its generator at 66.4 kW and the vapor bypass shut. (3) The receiver and the worn pump, running at 68 % under the flow loop. The flow setpoint of 2.93 kg/s that it delivers is read on the control panel, outside this crop. (4) The condenser, whose cooling water at 32.0 °C is at full flow, so PIC-2 is LIMITED and the condensing pressure reads 2.07 bar. (5) The cycle on the pressure–enthalpy plane. The bottom line gives the generator's 66.4 kW less 5.2 kW for the pump, 30.0 kW for cooling and 2.0 kW of auxiliaries, 29.2 kW net. The app's lead tile, also outside this crop, reports 0 s of 1140 s outside 5 to 20 K, the zero of Table 5.6.

Table 5.7 Scenario 9 Over 900 s, Default Seed. Upper Block as Scheduled, Lower Block With the Plant's Turndown Moved While the Announced Schedule Stays

	Cascade 7 K	Cascade 8.5 K	MPC, preview off	MPC, schedule
Lowest true *SH* (K)	2.43	5.46	5.11	5.80
at (s)	149	148	133	134
Time below 5 K (s)	21.0	0	0	0
SH IAE (K s)	275.0	174.2	135.9	122.2
p_e IAE (bar s)	16.15	14.17	6.35	5.88
Net energy (kWh)	16.848	16.727	16.857	16.857
Lowest true SH (K) / time below 5 K (s)				
10 s early	2.44/18.55	—	5.26/0	4.74/3.20
30 s early	2.45/20.90	—	5.10/0	4.91/4.00
30 s late	2.39/17.25	—	5.16/0	5.24/0
to 50 %, not 60	0.00/86.10	—	4.64/8.65	5.37/0
never comes	6.80/0	—	6.75/0	4.88/1.65

Both cascades run the suggested gains with feedforward and overrides, and 8.5 K is the lowest setpoint in steps of 0.5 K that keeps the floor. The cascade at 0.00 is isolated once.

MPC of Section 5.4.4, and Table 5.7 compares it over 900 s at the default seed with the strongest shipped structure, the cascade with feedforward and overrides at the suggested gains. Calmer cascade tunings fall further, to 3.69 K at 0.04 kg/s per K and 30 s and to 4.00 K at 0.02 and 60 s, and without the feedforward the expander is isolated and the plant trips. Each MPC run solves 900 plans and rejects none.

The MPC with the schedule keeps the true superheat at or above 5.80 K, reached at 134 s, while the cascade falls to 2.43 K at 149 s and spends 21.0 s below the floor, as shown in Figure 5.6a. The

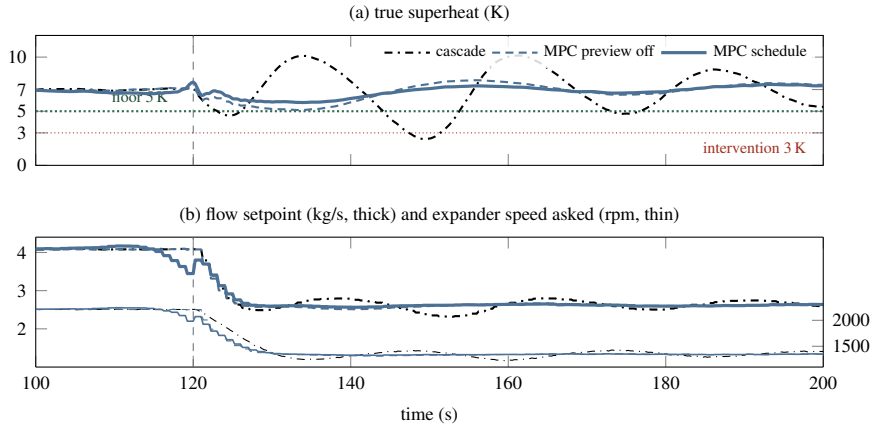


Figure 5.6 Scenario 9 from 100 to 200 s, default seed, the scheduled turndown at 120 s gray dashed: the cascade with feedforward and overrides (black, dash-dotted), the MPC with the schedule (blue, solid) and with the preview off (blue, dashed). (a) True superheat against the 5 K floor and the 3 K intervention (dotted horizontals). (b) Flow setpoint (thick, left axis) and expander speed asked (thin, right axis).

cascade's measured superheat stops at 3.20 K, 0.20 K above the intervention. Its first trough, 4.58 K at 124.5 s, follows the feedforward's filtered reading of FT-2, and the deeper one is its superheat loop swinging back while PIC-1 restores the pressure. The superheat IAE is 122.2 against 275.0 K s and the pressure IAE 5.88 against 16.15 bar s. Over ten noise seeds the lowest values span 5.73 to 5.86 K under the MPC and 2.23 to 2.43 K under the cascade, and the MPC's 5.80 K is unchanged when the run is extended to 3600 s. Part of the margin is information. With the preview off, given only the source flow that the cascade's feedforward also reads, the MPC still holds 5.11 K at an IAE of 135.9 K s, and that is the like-for-like comparison. The schedule is worth 0.69 K of the low point. At the same setpoint the two structures make the same energy, 16.857 against 16.848 kWh. The cascade keeps the floor from a setpoint of 8.5 K, where its lowest value is 5.46 K, and then nets 16.727 kWh, 0.130 kWh or 0.8 % less than the MPC at 7 K over the 15 min. Raised in steps of 0.5 K from 7 to 8.5 K, its setpoint lifts the cascade's lowest superheat through 2.43, 3.43, 4.44 and 5.46 K. Moving two handles together is therefore worth the gap from 2.43 K to the MPC's 5.11 K without a preview, and being told the future is worth the further 0.69 K.

What holds the superheat up is the early action and the two handles moving together, not the floor's price. Told of the turndown, the plan cuts the flow setpoint from 4.16 to 3.45 kg/s and the speed from 2239 to 1982 rpm between 110 and 120 s, before FT-2 moves, and by 130 s it asks for 2.60 kg/s and 1376 rpm, as shown in Figure 5.6b. Because the speed falls with the flow, the evaporating pressure stays within 12.78 to 13.09 bar after 100 s, against 12.45 to 13.39 bar under the cascade. The plan rests on the floor in 127 of its 900 plans, before each scheduled change. With the floor priced at zero, or moved to 3 K, the true low point is still 5.80 K, although the unpriced plan predicts −1.21 K. With the speed held by a move price a million times larger, or the flow held in the same way, the expander is isolated and the true superheat reaches 0 K. The result also rests on the output weights fixed before tuning, because with the pressure weight cut to a hundredth the low point falls to 0.79 K and the expander is isolated.

A wrong schedule misleads the plan, as the lower block of Table 5.7 shows for a plant whose turndown departs from the one announced. A turndown 10 s early takes the true superheat to 4.74 K for 3.20 s, 30 s early to 4.91 K, and a turndown announced but never made to 4.88 K at 482 s, after the plan has acted for a return that does not come. A late or deeper turndown stays above the floor, at 5.24 and 5.37 K. The floor constrains the measured superheat that the plan predicts, which is why in

the 10 s early run the measurement never falls below 5.36 K while the true value reaches 4.74 K.

The other cases load the MPC with scenario 9's settings and no retuning. In case 4 the preview announces the loss, and the MPC holds 9.49 to 10.68 K over it at an IAE of 24.6 K s, against 29.0 K s for the cascade with feedforward and 27.9 K s with the preview off, at the same energy within 0.007 kWh, as Table 5.5 shows, and it halves the pressure IAE over 900 s, 4.19 against 9.29 bar s. Its IAE is 151.6 against 1078.6 K s for case 2's cautious loop and 14.5 against 30.0 K s in case 8. It loses narrowly in case 5, 40.8 against 34.8 K s at equal energy, and in case 7, 163.1 against 147.3 K s. It loses clearly in two places. In case 6 it never settles, swinging between 8.0 and 13.1 K with a period of about 50 s to 3600 s, an IAE of 5583 against 335 K s. At 60 % fouling a small flow step takes 30 s to make 63 % of its superheat change, against 16 s in the model, with a gain per logarithmic unit 15 % higher, so a plan with $\lambda = 0.03$ overreaches, and with $\lambda = 3$, as a diagnostic, the superheat stays within 9.20 to 10.46 K. With case 4's events and the supply raised to 170 °C at 10 s, the superheat settles at 22.0 to 22.5 K from 600 to 1800 s against its 10 K setpoint, an IAE of 14 870 against 965 K s. The generator's rating holds the speed, and with a pressure weight of 100 per bar^2 against 1 per K^2 the plan cuts the flow to hold the pressure, which peaks at 15.35 bar. A weight of 10, as a diagnostic, lowers the IAE to 600 s from 2509 to 806 K s. That run nets 46.86 against 43.89 kWh with no intervention, where the cascade is isolated once, but its true superheat touches 3.34 K while the measurement stays above 5.43 K. The cascade is the better choice for fouling and for the hottest supply.

The challenge is met by the MPC alone, with only the architecture switched and with the preview on or off, as Table 5.6 shows. The superheat stays within 6.64 to 16.01 K and recovers 30.1 s after the pump wear, at an IAE of 160 against 134 K s for the suggested settings and 14.89 kWh of net energy. The challenge's rules are unchanged.

5.5.6 Verification

The numbers in this subsection check the implementation, not control performance. The working-fluid inventory stays at its charge within 3×10^{-12} kg in every run reported here, and the stored energy of fluid, walls and water matches the integrated boundary fluxes within 8×10^{-5} of the accepted heat, as does a separate ledger written from the face fluxes (2.8×10^{-5} in case 4, 5.8×10^{-5} in the suggested challenge). Nearly all of that residual sits in the last condenser cell, which crosses the saturated-liquid line 58 to 95 times per run, and in its neighbor, which crosses it up to four times, while every other cell closes within 0.022 kJ. Halving the physics step to 0.025 s changes comparison A's IAEs by at most 0.06 K s, and the application's own check holds the challenge's scored superheat and pressure within 1 % of their spans. The cell count is the larger numerical sensitivity. Doubling it to 24, 16 and 4 cells raises the design point's accepted heat by 6.60 %, from 959.95 to 1023.28 kW, and its net power by 6.84 %, from 77.55 to 82.86 kW. On case 4 the doubled mesh moves the IAE over the loss from 276.7 to 301.1 K s for independent PI, 437.2 to 445.5 K s for the cascade and 29.0 to 64.8 K s with the feedforward, deepens the feedback runs' dips to 0.41 and 0.00 K, and leaves every intervention count. In the challenge the shipped run's time outside the band becomes 303.3 s instead of 256.7 and the suggested run's range 6.78 to 14.60 K, so no verdict and no ranking changes. A simulated hour costs 23 s in a headless engine, and all 61 of the application's checks pass.

Agreement with literature is a narrower claim. The tables reproduce CoolProp's implementation of the equation of state to the errors of Section 5.3.5, a statement about interpolation only. The shaft isentropic efficiency of 0.655 lies below the 68 % maximum overall effectiveness measured on one scroll expander [15], and the feedforward's reduction of the superheat deviation, from 8.17 to 1.08 K, has the same direction as the 1.9 against 10 K measured on a vehicle cycle [3], but nothing in the model has been checked against data from an operating ORC. No learner study has been made, so this chapter claims nothing about learning gains.

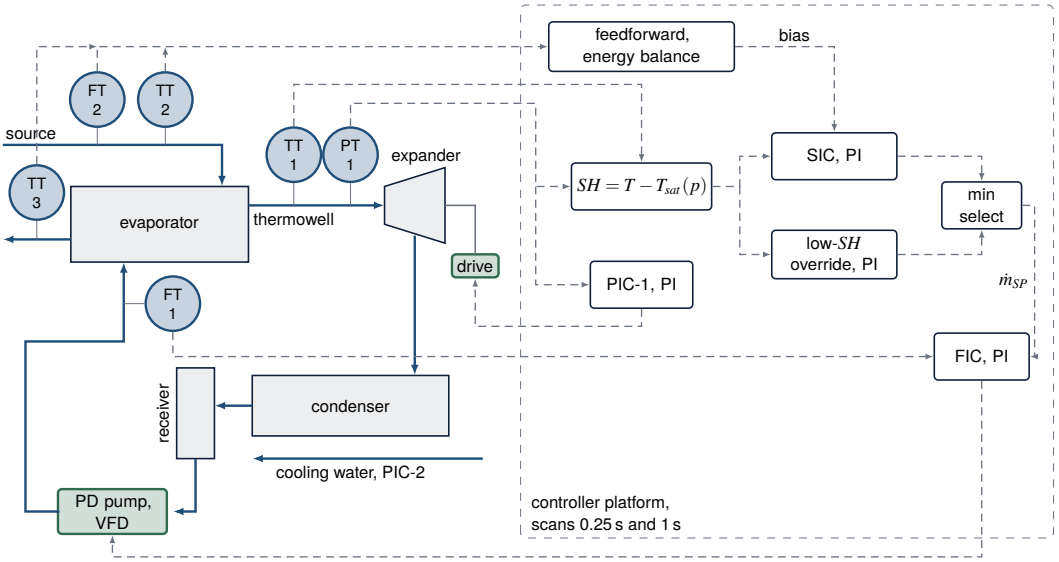


Figure 5.7 The superheat cascade on real hardware. TT-1 reads the expander inlet through its thermowell and PT-1 the evaporating pressure, and the platform computes the superheat from the two. SIC, biased by the feedforward from the source meters FT-2, TT-2 and TT-3, and the low-superheat override meet in the minimum selector, whose output is the flow setpoint of FIC. FIC reads FT-1 and sets the speed of the positive-displacement (PD) feed pump on its variable-frequency drive (VFD), and PIC-1 holds the pressure with the expander's drive. Solid arrows are process streams and dashed gray arrows are signals. Notice that the feedforward reaches the pump without passing the thermowell.

5.6 Technology in Practice

Figure 5.7 draws the superheat cascade of Section 5.4 and its feedforward with their hardware.

What Measures It

The controlled variable is computed rather than read. TT-1 measures the temperature at the expander inlet, PT-1 the evaporating pressure, and the transmitter subtracts the saturation temperature at that pressure. The 3 s lag of TT-1 stands for the thermowell and the 0.3 s lag of PT-1 for the pressure transmitter. A truck engine's cycle had the same data, pressure and temperature pairs between its components, and no sensor inside its heat exchangers [3]. The feedforward needs three instruments more, on the source loop: the flow meter FT-2 and the supply and return thermometers TT-2 and TT-3.

TECHNOLOGY **Sensor: A Resistance Thermometer in a Thermowell, Read With a Pressure**

Superheat is not measured directly. It is computed from a temperature and a pressure through the fluid's saturation properties, as the pressure and temperature pairs between the components of a vehicle cycle are used to estimate its superheat and subcooling [3]. The temperature element is commonly a resistance temperature detector, a fine platinum or copper wire whose resistance rises with temperature and is very linear. It sits in a thermowell, a pressure-tight sheath that lets the element be replaced without depressurizing the line. The sheath's metal adds a first-order lag, and an air gap left at the bottom of the well adds a second lag that can push a loop toward oscillation [23]. On a trend, a slow well shows as a measured superheat that trails every real change.

What Moves It

The superheat moves through the feed pump, whose speed sets the flow because the pump is of the positive-displacement kind. The model gives its drive a 1 s lag. The expander's speed is the second handle, held by PIC-1 through a drive that the model gives a 0.5 s lag and a 300 rpm/s slew, the kind of acceleration limit a motor drive carries [23]. The cooling flow is set by pumps and fans whose power the model takes as the cube of the flow, and the source bypass is the valve of the high-pressure override.

TECHNOLOGY **Final Element: A Positive-Displacement Pump on a Variable-Frequency Drive**

A positive-displacement pump discharges a fixed volume with every revolution of its shaft, so its speed sets the flow [3, 23]. Its motor is run by a variable-frequency drive, which sets an alternating-current motor's speed by the frequency of the power it supplies and its torque by the voltage and current. Such a drive needs no speed feedback to know roughly how fast the motor turns, although precise control at low speed may need a speed sensor. Drives limit current, torque, acceleration and deceleration, and a control system can change those limits over a network while the motor runs [23]. A worn pump delivers less per turn, which the chapter's inner flow loop answers by raising the speed.

Where the Controller Runs

Control loops in a plant are often nested or cascaded, with cycle times from milliseconds to minutes [24]. In the app the flow loop and PIC-1 scan every 0.25 s and the superheat loop every 1 s, against a thermowell lag of 3 s and slowest plant modes of 17 to 47 s. On a fieldbus platform a feedforward needs one more analog input block for the load transmitter, wired to the feedforward input of the PID block [23]. The energy balance of Equation (5.8) would sit in calculation blocks between the source meters and that input. The minimum selector is a control selector block whose back-calculation output tells a de-selected controller to halt its integral action [23], the external reset of Equation (5.9). On the truck engine the model-based feedforward was computed robustly enough for an embedded implementation, under a supervisor that moved the cycle through its stop, idle and production modes [3].

IN PRACTICE **Size the Feedforward, Then Time It**

A feedforward that keeps a mass or energy balance takes its size from that balance. It is applied to the dominant loads only, because compensating every load seldom repays the cost. Its timing is then tested. The engineer measures the lag with which the load moves the process variable and the lag with which the final element moves it, and compares the two. A quicker final element gets a delay or lag on the feedforward signal, so that it does not act too soon, and a slower one gets a lead. Even an imperfect feedforward is often far better than none [23]. This chapter's feedforward is static apart from a 2 s first-order filter.

What Is Modernizing It

Three developments bear on this loop. The first is a supervisory layer, a real-time optimization that hands optimal setpoints to the low-level superheat and pressure controllers. The truck study built

its controllers to serve such a layer [3], and the predictive controller of Section 5.4.4 is one form of it. The second is estimation. The same study added an extended Kalman filter that estimates the evaporator wall temperatures no sensor reads, so that a second controller could track a pressure setpoint with the evaporator bypass while the superheat stayed regulated, shown in simulation [3]. The third is virtual prototyping. A review expects simulation of the whole system, detailed enough for design, to overhaul the traditional sequential and iterative design process, with dynamic simulation bringing transient operation and control into design from its early stages [9].

5.7 Limitations

Every size, conductance, efficiency and threshold is a synthetic teaching value chosen so that the design point closes, not fitted to a plant, and passing a simulated case qualifies no equipment and estimates no savings. Uniform pressure per side omits the evaporator's frictional drop and every pressure wave. The exchangers are not converged in their cell count, so the operating point moves by about 7 % when the cells double, although no ranking does. The expander has the structure of a semi-empirical model without its fitted parameters and with no rotor inertia, and its dry-inlet specification is stricter than the liquid tolerance of many positive-displacement machines. The source loss of case 4 and the challenge falls by exactly 30 % in one step, which would not occur on a real engine loop, but the step isolates the thermowell's lag from every other cause. The protection acts on the measured superheat, which in the shipped challenge still reads 3 K when the true value is about 1.4 K. The engine behind the source loop and the cooler behind the cooling water do not respond to the cycle, and the loops' internal models share the plant's parameters, so mismatch enters only where a case changes the plant.

The predictive controller's model is linear in logarithms about one design point. At 60 % fouling the plant's flow response is about twice as slow and 15 % stronger than the model's, and the MPC never settles in case 6. It holds no model of the overrides, the protection or the condensing side, and its trades are set by output weights fixed before tuning, which at a 170 °C supply leave the superheat at 22 K. It takes the announced schedule as true, and its floor bounds the measured superheat rather than the true one, so a turndown 10 s early takes the true value to 4.74 K. Its gains come from the simulator, not from tests on an operating cycle.

> **OPEN CHALLENGE** **Staying in Production Through Transients**
>
> Recovery cycles on vehicle engines have reached prototypes and a few demonstration vehicles, yet it is still unclear how viable and efficient they are in real use. They often struggle to perform over a broad range of transient conditions, and a cycle that must shut down too often, or takes too long to reach power production, loses the gains it was built for [3]. A review counts automotive heat recovery, together with the conversion of concentrated sunlight, among the applications whose control aspects may be critical [9]. Progress would be a cycle that stays in production through a whole driving cycle, its superheat never below what the expander tolerates.

5.8 Conclusion

This chapter used a waste-heat ORC simulator with one pressure per side, a two-zone receiver and audited R1233zd(E) tables to weigh each part of a superheat control structure. On a 30 % loss of source flow the superheat loop alone, independent or in a cascade, lets the expander inlet fall to 1.83 or 1.09 K and loses the expander to protection, while a feedforward from the source meters holds it within 1.08 K of setpoint and cuts the IAE from 276.7 to 29.0 K s. On a hot day the condenser's fan limit, not the loop, sets the condensing pressure, and raising the setpoint to 2.6 bar nets 12.24 against 8.69 kWh. A slower, noisier thermowell erodes the advantage of fast tuning and nearly triples its pump travel. Server water at 55 °C cannot pay for the pumps and fans at any of 192 settings. A predictive controller told of a turndown holds the superheat at 5.80 K where the cascade falls to 2.43 K, or 5.11 K without the forecast, but its design-point model loses to the cascade under fouling.

The figures carry a 7 % cell-count sensitivity and rest on synthetic parameters. A fuller model would add frictional pressure drop, a converged mesh, a fitted expander map and a learner study. The simulator shows that the superheat a dry-inlet expander needs is protected by measuring the heat before it arrives.

KEY POINTS

- The expander needs dry vapor, 10 K above its boiling point, and the thermowell reports it after most of the evaporator's response.
- On a 30 % loss of source flow, the superheat loop alone loses the expander to protection, while a feedforward from the source meters holds the superheat within 1.08 K of setpoint.
- On a hot day the condenser's fan limit, not the loop, sets the condensing pressure, and raising the setpoint to 2.6 bar nets 12.24 against 8.69 kWh.
- A predictive controller told of a turndown holds 5.80 K where the cascade falls to 2.43 K, but its design-point model loses to the cascade under fouling.

Where else this applies. Volume 18, Chapter 4 feeds the steam flow forward to a drum level cascade, Chapter 6 feeds the raw water forward to a coagulant dose, and Chapter 4 biases a temperature loop with the measured process flow.

References

[1] S. Quoilin, M. Van Den Broek, S. Declaye, P. Dewallef, and V. Lemort, "Techno-economic survey of organic Rankine cycle (ORC) systems," *Renewable and Sustainable Energy Reviews*, vol. 22, pp. 168–186, 2013.

[2] J. Zhang, W. Zhang, G. Hou, and F. Fang, "Dynamic modeling and multivariable control of organic Rankine cycles in waste heat utilizing processes," *Computers & Mathematics with Applications*, vol. 64, no. 5, pp. 908–921, 2012.

[3] J. Peralez, M. Nadri, P. Dufour, P. Tona, and A. Sciarretta, "Organic Rankine cycle for vehicles: Control design and experimental results," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 3, pp. 952–965, 2017.

[4] 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.

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

[6] A. Desideri, B. Dechesne, J. Wronski, M. Van den Broek, S. Gusev, V. Lemort, and S. Quoilin, "Comparison of moving boundary and finite-volume heat exchanger models in the Modelica language," *Energies*, vol. 9, no. 5, p. 339, 2016.

[7] S. Quoilin, R. Aumann, A. Grill, A. Schuster, V. Lemort, and H. Spliethoff, "Dynamic modeling and optimal control strategy of waste heat recovery organic Rankine cycles," *Applied Energy*, vol. 88, no. 6, pp. 2183–2190, 2011.

[8] L. Y. Bronicki, "History of organic Rankine cycle systems," in *Organic Rankine Cycle (ORC) Power Systems*, E. Macchi and M. Astolfi, Eds. Elsevier, 2017, pp. 25–66.

[9] P. Colonna, E. Casati, C. Trapp, T. Mathijssen, J. Larjola, T. Turunen-Saaresti, and A. Uusitalo, "Organic Rankine cycle power systems: From the concept to current technology, applications, and an outlook to the future," *Journal of Engineering for Gas Turbines and Power*, vol. 137, no. 10, p. 100801, 2015.

[10] T. C. Hung, T. Y. Shai, and S. K. Wang, "A review of organic Rankine cycles (ORCs) for the recovery of low-grade waste heat," *Energy*, vol. 22, no. 7, pp. 661–667, 1997.

[11] B. F. Tchanche, G. Lambrinos, A. Frangoudakis, and G. Papadakis, "Low-grade heat conversion into power using organic Rankine cycles – a review of various applications," *Renewable and Sustainable Energy Reviews*, vol. 15, no. 8, pp. 3963–3979, 2011.

[12] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renewable and Sustainable Energy Reviews*, vol. 31, pp. 622–638, 2014.

[13] S. Eyerer, C. Wieland, A. Vandersickel, and H. Spliethoff, "Experimental study of an ORC (Organic Rankine Cycle) and analysis of R1233zd-E as a drop-in replacement for R245fa for low temperature heat utilization," *Energy*, vol. 103, pp. 660–671, 2016.

[14] R. Akasaka and E. W. Lemmon, "An international standard formulation for *trans*-1-chloro-3,3,3-trifluoroprop-1-ene [R1233zd(E)] covering temperatures from the triple-point temperature to 450 K and pressures up to 100 MPa," *Journal of Physical and Chemical Reference Data*, vol. 51, no. 2, p. 023101, 2022.

[15] V. Lemort, S. Quoilin, C. Cuevas, and J. Lebrun, "Testing and modeling a scroll expander integrated into an organic Rankine cycle," *Applied Thermal Engineering*, vol. 29, no. 14–15, pp. 3094–3102, 2009.

[16] D. Wei, X. Lu, Z. Lu, and J. Gu, "Dynamic modeling and simulation of an organic Rankine cycle (ORC) system for waste heat recovery," *Applied Thermal Engineering*, vol. 28, no. 10, pp. 1216–1224, 2008.

[17] I. H. Bell, J. Wronski, S. Quoilin, and V. Lemort, "Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library CoolProp," *Industrial & Engineering Chemistry Research*, vol. 53, no. 6, pp. 2498–2508, 2014.

[18] 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.

[19] L. Ljung, "Perspectives on system identification," *Annual Reviews in Control*, vol. 34, no. 1, pp. 1–12, 2010.

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

[21] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

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

[23] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[24] 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.

5.9 Exercises

5.1 ⌕ Concept check · 2 min

In case 4 the cascade without the feedforward loses the expander to the protection, although its low-superheat override takes the flow setpoint once the superheat falls below 5 K. Why does the override not prevent the intervention?

(a) It reads the same lagging superheat measurement, so it acts after most of the evaporator's response

(b) Its gain of 0.08 kg/s per K is too small to move the flow setpoint far enough

(c) The minimum selector keeps the main controller's request while the superheat is falling

(d) It opens the source bypass, which removes heat from the evaporator instead of feed

5.2 ⌕ Concept check · 1 min

In case 3 the evaporating-pressure setpoint rises from 12.9 to 13.4 bar with PIC-1 in Auto, and the true superheat falls to 6.44 K, although the pressure loop never touches the feed pump. What lowers the superheat first?

(a) PIC-1 raises the pump speed to lift the pressure and floods the evaporator

(b) The saturation temperature rises by about 1.85 K at once, and the vapor follows only slowly

(c) The thermowell lag grows with the pressure, so the transmitter reads low

(d) The expander speeds up, draws more vapor and cools the evaporator outlet

5.3 ⌕ Concept check · 2 min

In case 5 the cooling water warms to 35 °C, and PIC-2 reports LIMITED with the cooling water at 100 %. Which action recovers the most net power?

(a) Raise the condensing-pressure gain so that PIC-2 reaches its setpoint sooner

(b) Lower the condensing-pressure setpoint so that the loop asks for still more cooling

(c) Raise the condensing-pressure setpoint to about 2.6 bar, so that the flow and fan power fall

(d) Switch the superheat loop to the cascade with feedforward, so that the expander makes more work

5.4 Theory and calculation · 7 min

At the design point the expander of Equation (5.5) runs at 2146 rpm with a swept volume $V_d = 1.5\,\mathrm{L}$, friction torque $\tau_f = 20\,\mathrm{N\,m}$ and generator efficiency $\eta_g = 0.94$. It swallows 3.567 kg/s, leaks 0.341 kg/s, and the swept flow does $w = 28.62\,\mathrm{kJ/kg}$ of work.

(a) Compute the vapor density at the expander inlet and the leakage as a share of the total flow through the machine. *Checkpoint:* 66.5 kg/m^3, 8.7 %

(b) Compute the fluid work, the friction loss, the shaft power and the generator output. *Checkpoint:* 102.09, 4.49, 97.60 and 91.74 kW

(c) Compute the shaft torque, and say why the leakage adds to the flow the expander passes but not to its work. *Checkpoint:* 434 N m

5.5 Theory and calculation · 7 min

The feedforward of Equation (5.8) divides the measured source heat by the enthalpy rise the working fluid needs. At the design point the evaporator accepts 959.95 kW and the pump delivers 3.909 kg/s. The pump's displacement flow is 5.0 L/s at full speed and its nominal density 1250 kg/m^3.

(a) Compute the enthalpy rise from the pump outlet to the expander inlet at the design point. *Checkpoint:* 245.6 kJ/kg

(b) In case 4 the source flow falls to 70 % of 5 kg/s. Lined out, TT-2 reads 150 °C and TT-3 102.4 °C, with $c_{p,s} = 4300\,\mathrm{J/(kg\,K)}$. Take the enthalpy rise of part (a) as unchanged and compute the feedforward flow. *Checkpoint:* 716.4 kW, 2.917 kg/s

(c) Convert the flows of parts (a) and (b) into pump speeds through the nominal capacity, as the independent architecture does, and convert the suggested superheat gain of 0.08 kg/s per K into percent of pump speed per kelvin. *Checkpoint:* 62.5 %, 46.7 %, 1.28 % per K

5.6 Theory and calculation · 7 min

Use the design values of Section 5.3: generator 91.74 kW, pump motor 6.94 kW, cooling 5.26 kW at 55.95 % flow, other auxiliaries 2.00 kW, accepted heat 959.95 kW, source supply 150 °C and cooling water 25 °C.

(a) Compute the net power, the net efficiency and the Carnot bound between the two supply temperatures. *Checkpoint:* 77.54 kW, 8.08 %, 29.54 %

(b) The pump raises 3.909 kg/s of liquid at about 1247 kg/m^3 through 10.705 bar with a hydraulic efficiency of 0.55 and a motor efficiency of 0.88. Compute the enthalpy rise across the pump, the power that enters the fluid and the motor power. *Checkpoint:* 1.561 kJ/kg, 6.10 kW, 6.93 kW

(c) The cooling pumps and fans draw $P_{c,max} f_c^3$ with $P_{c,max} = 30\,\mathrm{kW}$. Compute the cooling power at the design flow, at full flow and at 61 % flow, the case 5 flow at a 2.6 bar setpoint. *Checkpoint:* 5.25, 30.0 and 6.81 kW

5.7 App challenge · 8 min

Load scenario 4, Lose source flow: feedback or feedforward, at the default seed. Run each case to 480 s, export the CSV and read the column `truth_sh_K` from 120 to 480 s, the column `expander_isolated`, and the Trips · interventions and Net energy tiles.

(a) Run the scenario as shipped, with Source-heat feedforward (cascade) Off. Give the lowest true superheat and when it occurs, and how long the expander is isolated. *Checkpoint:* 1.10 K at 141 s, isolated from 137.3 to 220.75 s, about 83 s

(b) Set Source-heat feedforward (cascade) to On, press Reset and run to 480 s again. Give the true superheat range and the two tiles. *Checkpoint:* 9.41 to 11.07 K, 0 · 0

(c) Compare the Net energy tiles of the two runs and explain where the difference comes from. *Checkpoint:* 6.77 against 8.40 kWh

Accept the superheat values within 0.05 K and the energies within 0.02 kWh.

5.8 App challenge · 7 min

Load scenario 5, A hot day at the condenser, at the default seed. The cooling water warms from 25 to 35 °C between 60 and 180 s.

(a) Run to 900 s as shipped. Read the status of PIC-2, the cooling-water flow, the true condensing pressure, the cooling power and the net power from the last row of the CSV, and the Net energy tile. *Checkpoint:* LIMITED, 100 %, 2.38 bar, 30.0 kW, 47.3 kW, 12.94 kWh

(b) Reload the scenario, run to 240 s, raise the Condensing pressure setpoint to 2.60 bar and run to 900 s. Read the same quantities. *Checkpoint:* 61 %, 2.60 bar, 6.8 kW, 67.6 kW, 16.49 kWh

(c) Explain why no tuning of PIC-2 recovers this power, and why the generator makes less in run (b) although the net power is higher.

Accept net power within 0.4 kW, which covers the noise on the last row, and energies within 0.02 kWh.

5.9 App challenge · 8 min

Load scenario 9, MPC or cascade through a scheduled turndown, at the default seed. It ships with Architecture MPC, MPC preview (load schedule) on Schedule and a superheat setpoint of 7 K. Run each case to 600 s and read the column `truth_sh_K` of the exported CSV.

(a) Run as shipped. Give the lowest true superheat and its time. *Checkpoint:* 5.80 K at 134 s

(b) Press Reset, set Architecture to Cascade before pressing Run, and run again. Give the lowest true superheat, its time and the time spent below the 5 K floor. *Checkpoint:* 2.45 K at 149 s, about 21 s

(c) Press Reset, set Architecture back to MPC and MPC preview (load schedule) to Off, and run again. Give the lowest true superheat, then split the gap between parts (a) and (b) into the part the forecast is worth and the part that moving two handles together is worth. *Checkpoint:* 5.11 K, 0.69 K and 2.66 K

Accept the superheat values within 0.05 K and the time below the floor within 2 s.

5.10 Application area · 4 min

In the shipped challenge the protection still reads 3 K of superheat when the true value is about 1.4 K.

(a) TT-1 is a first-order lag with $\tau_{TT} = 3$ s. For a superheat falling at a steady rate, a first-order lag reads low by the rate times its time constant. Estimate the rate at which the true superheat was falling. *Checkpoint:* about 0.53 K/s

(b) The chapter calls every threshold of its protection a teaching value. Where would the thresholds of an operating unit come from, and name two changes to the protection that the gap of part (a) suggests.

5.11 Application area · 3 min

A data-center operator asks for an ORC on 55 °C server return water with cooling water at 15 to 22 °C. Use case 8 to answer.

(a) Compute the Carnot bound at both cooling-water temperatures, and the share of the accepted 128.5 kW that the generator converts at 3.29 kW. *Checkpoint:* 12.19 and 10.06 %, 2.56 %

(b) The parasitics take 2.28 kW. What does the chapter's sweep of 192 settings say about operating the cycle, and what is the advice?

5.12 Review the agent · 5 min

An AI assistant reviews scenario 9 and writes: "The MPC holds the superheat at 5.80 K where the cascade falls to 2.43 K. Predictive control is therefore worth 3.4 K of superheat margin over the cascade on any loss of source flow." Find the flaw and say how you would check it.