INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 17 Control of Energy Conversion and Storage

CHAPTER 6

A Parabolic-Trough Simulator for Choosing a Setpoint Rather Than Defending One

Pages 134–162

Edited by John D. Hedengren

How to cite this chapter. "A Parabolic-Trough Simulator for Choosing a Setpoint Rather Than Defending One," in *Control of Energy Conversion and Storage*, John D. Hedengren, Ed., vol. 17 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 134–162.

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

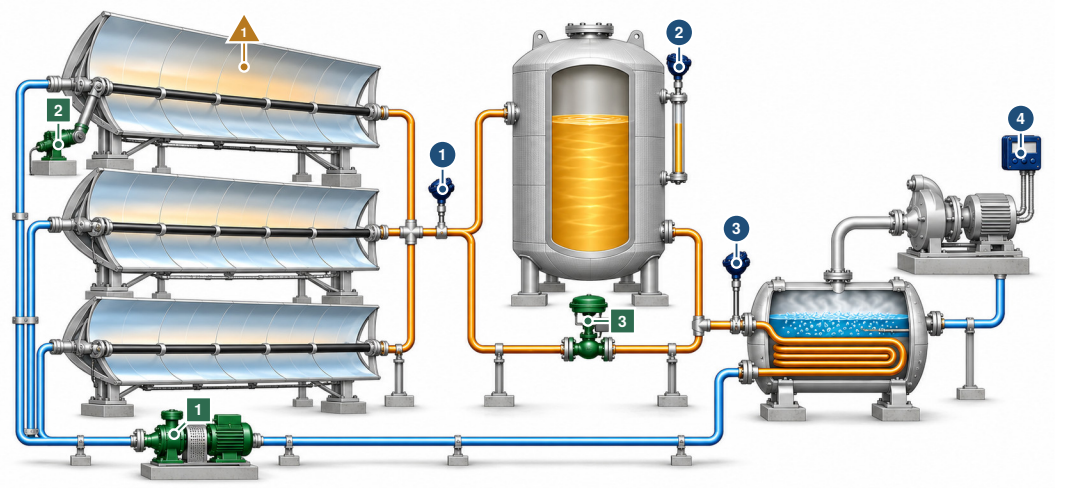
The app for this chapter
https://apmonitor.com/apps/parabolic-trough/

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

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

CHAPTER 6

A Parabolic-Trough Simulator for Choosing a Setpoint Rather Than Defending One



The controller moves 1 field pump speed 2 mirror drive that turns the troughs off the sun 3 bypass valve around the hot tank

The controller reads 1 oil temperature leaving the mirrors 2 hot tank level 3 oil temperature entering the steam generator 4 electric power out

What disturbs it 1 clouds that change the sunlight

Notice that the oil can reach the steam generator through the hot tank or around it, while the mirrors heat it only as fast as the sky allows. Illustration, not to scale.

IN PLAIN TERMS A Best Temperature That Moves With the Sky

The main idea. Long curved mirrors focus sunlight on a pipe, and oil in the pipe carries the heat to a storage tank and a power plant. A controller changes the pump speed to hold the hot oil at a chosen temperature.

The challenge. The pipe loses heat to the sky all the time, and more when it runs hotter, sunny or not. Under clouds, holding the usual high temperature wastes nearly a third of the heat, so the target must change with the sky.

Why it matters. On a broken day, a well chosen temperature makes about a third more electricity.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the history boxes of Section 6.2, the fact box of Section 6.3, the try-it boxes and the fact box of Section 6.7, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 6.7, Section 6.8, Section 6.9 and its open challenges

The control question A trough receiver loses heat as the fourth power of its temperature whether or not the sun is out, while the power block converts hotter oil more efficiently. When is the outlet temperature setpoint itself the decision worth automating, and when is the schedule of the power block worth more than either the setpoint or the tuning?

Learning objectives

- Compute the receiver loss, the field transit time and the outlet loop gain from the model, and explain how each moves with flow and irradiance.
- Explain why the optimal outlet temperature falls under cloud and why a fixed setpoint can be held perfectly and still be wrong.
- Describe how the level override, the high-temperature protection and the bypass interlock act on top of the four regulatory loops.
- Compare fixed and variable outlet temperature on a clear and a broken day, and the predictive controller against the outlet loop, from scenario runs.
- Separate what the tuning recovers from what the schedule earns in the graded challenge.

Prerequisites Sections 1.4.5 (a moving optimum), 1.4.6 (threshold and water value), 1.6.2 (velocity form), 1.6.3 (MPC option) and 1.6.4 (scenarios and scorecard).

The app https://apmonitor.com/apps/parabolic-trough/

See also Volume 9, Chapter 8 and Volume 9, Chapter 10 for supervisory setpoints that move with the weather or the load, and Volume 9, Chapter 6 for a delay set by the manipulated flow. Volume 5, Chapter 2 and Volume 5, Chapter 7 for control through a transport delay and a residence time.

6.1 Introduction

Ask a control engineer what to do about a collector outlet temperature and the answer arrives before the question finishes. Choose the design temperature, put a proportional-integral controller on the field flow, and defend it. The loop will do exactly that. On a cloudy morning it will defend 393 °C with complete success and throw away nearly a third of the hour doing it, and nothing about that failure is a tuning failure. The setpoint was wrong, and the plant has no mechanism for saying so.

That is the misconception this chapter corrects. A control course teaches regulation, and regulation assumes somebody upstream chose the right number to regulate to. In a solar field nobody did, because the right number moves. The receiver radiates as the fourth power of its metal temperature and does not care whether the sun is out, while the heat available to collect changes by a factor of five across a broken day. At full irradiance the conversion gain from running hot outweighs the radiation and the design temperature is roughly right. At a quarter of full irradiance, reaching the same temperature means running the field at a sixth of its flow, so the fluid spends six times as long under a receiver that is still losing 3.7 MW, and nearly a third of what is absorbed goes back to the sky.

The class of problem is steady-state optimization of an operating point that moves with a measurable disturbance, implemented as a supervisory layer that writes a setpoint into an otherwise conventional regulatory scheme. A fired heater with a variable draft, a solid oxide fuel cell trading polarization against thermal loss, a refrigeration plant choosing a condensing pressure against ambient, and a distillation column choosing a reflux ratio against feed composition all have this shape. In each, a loss term and a gain term depend on the same variable with opposite sign, the balance between them depends on the load, and holding the operating point fixed is a decision that somebody made once and nobody revisits. The formal treatment is real-time optimization above regulatory control, and the solar field is the clearest instance in this volume because the disturbance is visible in the sky. A data hall whose chilled water setpoint follows the outdoor wet bulb, in Volume 9, Chapter 8, has the same structure.

The chapter replicates on one plant a result already in the literature for this configuration [1, 2]. It reports the clear day and the broken day at the same length and finds that the variable temperature

strategy is worth 0.3 % of net generation on the first and 32 % on the second. The plant is a thirty-state distributed receiver model with a supervisory optimizer that runs twenty algebraic evaluations every three minutes, so it executes at interactive speed in a browser. Twelve scenarios separate what the regulatory layer can do from what only the supervisory layer can do, and the most valuable decision they expose is a setpoint rather than a gain.

Section 6.2 traces trough control from the SEGS plants to forecast-based dispatch. Section 6.3 gives the model, its parameters and its assumptions. Section 6.4 gives the open-loop character and the optimum outlet temperature against sky clarity. Section 6.5 gives the four regulatory loops, the override and the out-of-service logic, and Section 6.5.4 the optional predictive outlet controller. Section 6.6 gives the supervisory optimizer and the schedule. Section 6.7 reports the scenarios, the clear and broken days, the schedule, the override, the graded challenge and, in Section 6.7.6, the predictive controller against the outlet loop. Section 6.8 describes the hardware of a real field, Section 6.9 states the limitations and Section 6.10 concludes.

6.2 From SEGS to a Weather Forecast

The first commercial plants, told in the history box below, set the pattern the industry still follows: the collector loop outlet temperature is the controlled variable, the field circulation pump is the manipulated variable, and a fixed setpoint a few degrees below the stability limit of a synthetic aromatic fluid, about 400 °C, is defended through whatever the weather does [3, 4].

> ⧖ HISTORY **Nine Plants in the Mojave (1984 to 1991)**
>
> Nine commercial parabolic trough plants were built in the California Mojave Desert between 1984 and 1991, ranging from 14 to 80 MW of net output [5]. Falling energy prices and shrinking incentives stopped the building, and the developer filed for bankruptcy in 1991 when it could not finance a tenth plant [5]. The plants kept running. Many technical problems of the early units were found and corrected in the later ones, and together they gave more than 30 years of operating experience on the lifetime and performance of trough technology [5]. A loop in a modern field still takes oil in at about 292 °C and returns it at about 393 °C [5], the outlet temperature this chapter's fixed setpoint defends.

Control research on distributed collector fields is older than the commercial plants and was driven mostly by the ACUREX field at the Plataforma Solar de Almeria. Two decades of that work were collected in a survey whose first part covers the modeling and the classical schemes and whose second part covers adaptive, predictive and nonlinear approaches [6, 7], and in a book covering the same ground in depth [8]. The recurring difficulty in all of it is that the field is a distributed system with a transport delay that changes with the manipulated variable itself, so a controller tuned at design flow is wrong at a sixth of it. Feedforward from the irradiance, gain scheduling on the flow and series compensation of the delay are all standard answers, and all of them still take the setpoint as given.

Storage changed the question. Two-tank indirect storage decouples collection from generation, so a plant with a solar multiple above one can move its generation into the hours that pay for it [9, 10]. Once generation is schedulable the plant has an economic optimization to solve, not just a regulation problem, and the operating temperature becomes a decision variable rather than a constraint. Molten salt was proposed as the field fluid to raise that ceiling, and the engineering assessment of what it would cost in freeze protection is the clearest statement of why the synthetic aromatic survived [11].

> ⧖ HISTORY **Storage Comes to the Trough (2007)**
>
> Two-tank molten-salt storage was first demonstrated in a tower pilot plant in the 1990s. In the early 2000s a joint United States and European team adapted it to troughs by adding an intermediate heat exchanger between the field oil and the salt, and the concept was first used commercially in 2007 [5]. By 2020 about 31 trough plants around the world had between 3 and 10 hours of storage [5]. The chapter's two tanks hold the field oil itself, with no exchanger, and store 3.6 h at the power block rating.

The dynamic optimization of this exact plant configuration was published in a sequence of papers that the present simulator follows. A dynamic model of a trough plant with two-tank storage and a control scheme for it [12] was followed by optimization over a twenty-four hour horizon using weather forecasts [1] and then by the hybrid solar and fossil case [2]. The receiver loss correlations that make the optimum move are from the NREL receiver heat transfer model [13] and the heat-loss tests that calibrated it [14].

6.3 Process Model

The simulated plant has 73 800 m^2 of mirror aperture feeding 2815 m^2 of receiver tube, two tanks holding 1.8×10^6 kg of synthetic aromatic heat transfer fluid, and a power block of 27.8 MW_{th} rating at 36 % design efficiency. At clear noon the mirrors intercept 70.1 MW and the receivers absorb 50.8 MW, which is a solar multiple of 1.83 on the power block. Filled to its upper working level the hot tank holds 106.3 MWh above the cold tank, and 99.5 MWh of that lies between the two level limits, which is 3.6 h at the power block thermal rating. Figure 6.1 shows the arrangement and every symbol used below.

> ✦ DID YOU KNOW? **Two Million Square Meters of Mirrors**
>
> The nine first commercial trough plants held 354 MW of generating capacity and more than 2 000 000 m^2 of collectors. By the end of 2001 they had accumulated 127 years of operating experience [3]. The field of this chapter has 73 800 m^2 of mirror aperture.

The clock is in hours, temperatures in degrees Celsius, flows in kilograms per second, power in megawatts and tank level as a fraction of the total fluid inventory. Thirty states are integrated. Ten are receiver fluid node temperatures and ten are receiver metal node temperatures, which together carry the transport delay that dominates the outlet loop. The rest are the hot tank mass and temperature, the cold tank temperature, a steam generator lag, five actuator lags and a filtered sky state. The outlet and power loops read their measurements through first-order instrument lags of 2.2 s and 1.8 s, and the other two loops read the state directly.

6.3.1 Variables

Table 6.1 classifies what the application exposes. Four controlled variables pair with four manipulated variables, and one of those pairings is an override rather than a regulator.

6.3.2 The Sun on the Mirrors

Clear-sky direct normal irradiance and the optical efficiency are both written against the fraction of the solar day $x = (h - h_{up})/(h_{dn} - h_{up})$, zero at sunrise and one at sunset,

$$\mathrm{DNI} = \sigma\, \mathrm{DNI}_{max} \sin^{0.6}(\pi x), \qquad 0 < x < 1 \tag{6.1}$$

$$\eta_{opt} = \eta_{opt,0}\, \kappa \sin^{n_i}(\pi x)\, \Pi\!\left(\frac{x - 0.02}{0.09}\right) \Pi\!\left(\frac{0.98 - x}{0.09}\right) \tag{6.2}$$

where $\Pi(z) = \min(\max(z, 0), 1)$ ramps the row-end losses that make the first and last hour of the day nearly worthless. The absorbed power per receiver node follows,

$$q_a = \frac{A_{ap}}{10}\, \mathrm{DNI}\, \eta_{opt}\, (1 - f_d^a). \tag{6.3}$$

With the default parameters and 96 % mirror cleanliness, clear noon gives 950 W/m^2, 70.1 MW intercepted and 50.8 MW absorbed. The electricity price is a gaussian evening peak on a base with a midday depression, and the peak is at 19:30, after sunset. That offset is the entire commercial argument for thermal storage.

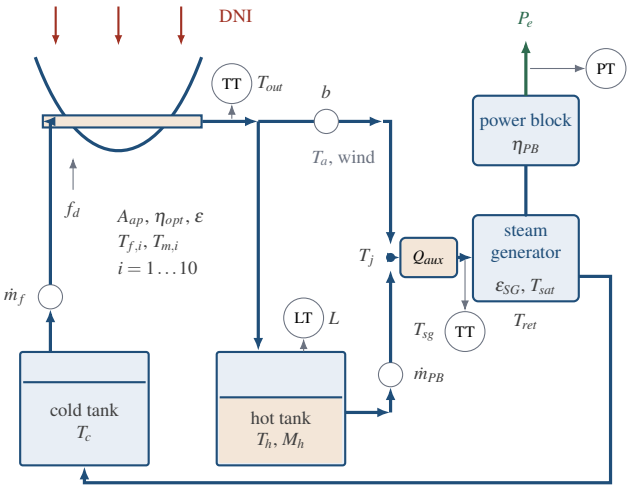


Figure 6.1 Plant arrangement and the symbols used in the equations. A single collector row stands for 73 800 m^2 of aperture, modeled as ten fluid nodes and ten metal nodes in series. Field fluid leaving the receiver either fills the hot tank or passes directly to the junction through the bypass, and it returns to the cold tank (not drawn) while it is less than 12 °C hotter than that tank. TT on the collector outlet supplies the controlled variable for the field flow loop, LT the hot tank level that drives the defocus override, TT after the heater the steam generator inlet T_{sg} that the bypass loop holds, and PT the net electrical output.

6.3.3 The Receiver

Each of the ten receiver segments carries a fluid node and a metal node. The metal node absorbs the sunlight, loses heat to the sky and conducts to the fluid, and the fluid node carries that heat downstream,

$$m_f c_p \frac{\mathrm{d}T_{f,i}}{\mathrm{d}t} = 3600\left[\dot{m}_f c_p (T_{f,i-1} - T_{f,i}) + hA(T_{m,i} - T_{f,i})\right] \tag{6.4}$$

$$m_m c_{p,m} \frac{\mathrm{d}T_{m,i}}{\mathrm{d}t} = 3600\left[q_a - hA(T_{m,i} - T_{f,i}) - q_{loss,i}\right] \tag{6.5}$$

with $m_f = M_f/10$, $m_m = M_m/10$ and $T_{f,0} = T_c$, the cold tank temperature. The loss from one node is radiation to the sky plus a convective envelope term that grows with wind,

$$q_{loss,i} = \varepsilon\sigma_{SB}\frac{A_{rec}}{10}\left[(T_{m,i}+273.15)^4 - (T_a+273.15)^4\right] + h_e\frac{A_{rec}}{10}(T_{m,i} - T_a) \tag{6.6}$$

$$h_e = h_{conv} + h_{wind}\, v_w^{0.6}. \tag{6.7}$$

Equation (6.6) is the origin of everything the application teaches. Summed over the field at an ambient of 25 °C and a 3 m/s wind, the loss is 0.12 MW at a metal temperature of 50 °C, 2.56 MW at 300 °C, 4.48 MW at 393 °C and 5.48 MW at 430 °C. None of those numbers depends on whether the sun is out.

Fluid properties are temperature dependent. Heat capacity is linear over the range the plant occupies, $c_p = c_{p,0} + c_{p,1}T$ in kJ/(kg °C) (entered in J/(kg °C) in the balances below, where it meets

Table 6.1 Variable classification

Symbol	Meaning	Unit	Range
Controlled variables			
T_{out}	collector outlet temperature	°C	312 to 400
P_e	net electrical output	MW	0 to 12
T_{sg}	steam generator inlet	°C	310 to 398
L	hot tank level (override)	%	0 to 100
Manipulated variables			
$\dot{m}_f$	field flow	kg/s	8 to 520
$\dot{m}_{PB}$	power block flow	kg/s	0 to 190
b	hot tank bypass	%	0 to 100
f_d	mirror defocus	%	0 to 100
Q_{aux}	supplemental heater	MW	0 to 12
Disturbances			
σ_∞	sky clarity	× clear	0.05 to 1.05
T_a	ambient temperature	°C	−10 to 48
v_w	wind speed	m/s	0 to 18
κ	mirror cleanliness	×	0.70 to 1.00
ε_{fc}	forecast error scale	×	0 to 1.5
σ_T	thermocouple noise, 1σ	°C	0 to 4
States			
$T_{f,i}$	receiver fluid node, $i = 1\ldots10$	°C	−30 to 520
$T_{m,i}$	receiver metal node, $i = 1\ldots10$	°C	−30 to 600
M_h	hot tank mass	kg	$0.001M$ to $0.998M$
T_h	hot tank temperature	°C	−30 to 500
T_c	cold tank temperature	°C	−30 to 500
T_{sg}	steam generator inlet lag	°C	−30 to 500
$\dot{m}_f^a$	applied field flow	kg/s	8 to 520
$\dot{m}_{PB}^a$	applied block flow	kg/s	0 to 190
b^a	applied bypass	—	0 to 1
f_d^a	applied defocus	—	0 to 1
Q_{aux}^a	applied heater duty	MW	0 to 12
σ	filtered sky clarity	—	0 to 1.05

conductances in watts), density is quadratic, and viscosity follows a Vogel form,

$$\mu(T) = \exp\left(\frac{544.149}{T + 114.43} - 2.59578\right) \quad [\text{mPa s}]. \tag{6.8}$$

Viscosity is in the model only because of the pump. At the design mean temperature of 346 °C the fluid is 0.243 mPa s and at 30 °C it is 3.23 mPa s. At design flow a field at 30 °C therefore needs 32.7 bar rather than 19.5 bar, although the pump power rises only 21 % because the cold oil is denser and occupies less volume. The model never limits the flow on that account.

6.3.4 Tanks, Junction and Power Block

Field fluid leaving the receiver splits three ways. A fraction b^a goes straight to the junction as bypass, as much of that as the power block will take is used directly, and everything else goes to the tanks,

$$\begin{aligned} \dot{m}_{byp} &= b^a \dot{m}_f, & \dot{m}_{dir} &= \min(\dot{m}_{byp}, \dot{m}_{PB}), \\ \dot{m}_{in} &= \dot{m}_f - \dot{m}_{dir}. \end{aligned} \tag{6.9}$$

What reaches the tanks is not automatically stored. Sending 200 °C oil into a 390 °C tank destroys the only thing the tank was holding, so the field return is recirculated to the cold tank until it is worth

more than a margin above it. The changeover is smoothed over a few degrees so that the right-hand side stays continuous for a Runge-Kutta step,

$$w_h = \Pi\left(\frac{T_{f,10} - (T_c + \Delta T_{rec})}{W_{rec}}\right), \quad \dot{m}_{hot} = w_h \dot{m}_{in}. \tag{6.10}$$

The tank balances follow, with $\dot{m}_{th} = \dot{m}_{PB} - \dot{m}_{dir}$ the draw from the hot tank and $M_c = M - M_h$. The hot pumps lose suction near the bottom of the tank. Below $M_h = 0.002M$ the part of their draw above $\dot{m}_{hot}$ fades linearly to zero at $0.001M$, and $\dot{m}_{PB}$ here and in the block equations is the flow they deliver, so an empty tank passes on only what flows into it and no state is clamped afterward,

$$\frac{\mathrm{d}M_h}{\mathrm{d}t} = 3600(\dot{m}_{hot} - \dot{m}_{th}) \tag{6.11}$$

$$M_h \frac{\mathrm{d}T_h}{\mathrm{d}t} = 3600\left[\dot{m}_{hot}(T_{f,10} - T_h) - \frac{UA_t(T_h - T_a)}{c_p}\right] \tag{6.12}$$

$$M_c \frac{\mathrm{d}T_c}{\mathrm{d}t} = 3600\left[\dot{m}_{PB}(T_{ret} - T_c) + (1 - w_h)\dot{m}_{in}(T_{f,10} - T_c) - \frac{UA_t(T_c - T_a)}{c_p}\right]. \tag{6.13}$$

The junction mixes bypass fluid with hot tank fluid, the supplemental heater raises it further, and a first-order lag of 9 s represents the steam generator metal,

$$T_j = \frac{\dot{m}_{dir} T_{f,10} + \dot{m}_{th} T_h}{\dot{m}_{PB}}, \qquad \tau_{SG} \frac{\mathrm{d}T_{sg}}{\mathrm{d}t} = T_j + \frac{10^6 Q^a_{aux}}{\dot{m}_{PB} c_p} - T_{sg}. \tag{6.14}$$

The power block is a counterflow steam generator of fixed effectiveness feeding a turbine with a part-load efficiency curve and a linear gain with inlet temperature,

$$\Delta T = \varepsilon_{SG}(T_{sg} - T_{sat}), \qquad Q = \frac{\dot{m}_{PB} c_p \Delta T}{10^6} \tag{6.15}$$

$$\eta = \eta_{PB}\, g(Q/Q_{PB})\,[1 + k_T(T_{sg} - 393)] \tag{6.16}$$

$$g(u) = \Pi^{1.05}_{0.2}\left(0.53 + 0.62u - 0.15u^2\right) \tag{6.17}$$

with the block producing nothing at all when $T_{sg} < T_{j,min}$ or when the load fraction falls below x_{min}, which is a genuine discontinuity and is what a real turbine does. At an inlet of 393 °C and 119 kg/s the block takes 26.4 $\mathrm{MW_{th}}$ and produces 9.33 MW gross at 35.4 %. At 340 °C the same flow gives 13.1 $\mathrm{MW_{th}}$ and 3.45 MW at 26.4 %, and at 310 °C the block is off.

Field pumping is quadratic in flow on a fixed hydraulic system with a viscosity correction,

$$\Delta p = \Delta p_0 \left(\frac{\mu(\bar{T})}{\mu_{ref}}\right)^{0.2} \left(\frac{\dot{m}_f}{\dot{m}_0}\right)^2, \quad P_{pump} = \frac{\Delta p\, \dot{m}_f}{\rho\, \eta_p\, 10^6}. \tag{6.18}$$

At the design flow of 222 kg/s the pressure drop is 19.5 bar and the pump takes 0.76 MW. At 350 kg/s it takes 2.97 MW and at the 520 kg/s limit it takes 9.75 MW, which is most of the plant output. Net electrical output subtracts pumping, heat tracing and a fixed auxiliary load from the gross, where P_{pump} is the field pump of Equation (6.18) plus 0.35 times the same law applied to the power block flow,

$$P_e = \max\left(0,\ \eta Q - P_{pump} - P_{fix} - P_{tr}\right). \tag{6.19}$$

6.3.5 Parameters

Table 6.2 gives every parameter with its default. The V column marks A for adjustable and shown, H for hidden from the beginner view, and F for a fixed physical constant.

6.3.6 Initialization and Numerics

Each scenario supplies a starting hot tank temperature, cold tank temperature and level. Before sunrise the field is not hot, so the receiver fluid nodes are initialized to the cold tank temperature with no profile along the row, and the metal nodes 2 °C above it. Giving a field at half past five in the morning a full temperature profile would credit the plant with a loop of collected heat it never collected, and that shows as a spurious couple of megawatt-hours in the harvest metric. When the clear-sky irradiance already exceeds 60 W/m^2, the nodes are initialized on a linear profile from the cold tank to the hot tank temperature, with the metal nodes 12 °C above the fluid.

Integration is classical fourth-order Runge-Kutta on a fixed step of 0.001 h, which is 3.6 s and about one time constant of the fastest actuator. The regulatory loops execute every 0.002 h, which is 7.2 s, the supervisory layer every 0.05 h, which is three minutes. The historian logs one sample every ninth physics step, 32.4 s. A post-step clip holds every temperature state, the tank mass, the two fractional actuators and the sky state inside their physical ranges.

6.3.7 Assumptions

1. **One collector row stands for the whole field.** All loops see the same irradiance, the same flow and the same inlet temperature. A real field has hundreds of parallel loops with a flow distribution, and the loop that runs hottest is the one that sets the defocus decision. The assumption understates the risk of a local overtemperature.

2. **The optical model is a smooth analytic surrogate.** The incidence angle modifier in Equation (6.2) is a power of the solar-day sine rather than a computed sun position, and the row-end term is a linear ramp. Neither is traceable to a collector test [14], and the shape is right while the numbers are representative rather than measured.

3. **Receiver loss is a single emissivity at all temperatures.** The NREL receiver model makes emissivity a function of the metal temperature [13], which steepens the loss curve above 350 °C. Using a constant understates the loss at the top of the range, so the value of lowering the setpoint is understated rather than exaggerated.

4. **Each tank is perfectly mixed and fully insulated to one *UA*.** Thermocline behavior, tank stratification and the wall thermal mass are all absent, so the tanks respond faster than real ones and never stratify on a partial charge [10].

5. **The power block is a steady-state map.** Equation (6.16) gives efficiency at once from inlet temperature and duty, with only a 9 s lag on the inlet. A real steam cycle has drum level, superheater metal and turbine thermal stress dynamics measured in tens of minutes, and the start and stop transients that carry most of the cost are missing.

6. **Fluid degradation is not modeled.** The stability limit at 400 °C is an alarm and a defocus trigger, not a mechanism. Exceeding it in the simulator costs nothing, where on a real plant it produces hydrogen that permeates into the receiver annulus and destroys the vacuum permanently.

7. **The price curve and the weather are synthetic.** The irradiance, the ambient and the price are smooth analytic profiles with a seeded forecast error. They are teaching data and are never labeled as measured.

Table 6.2 Model parameters and default values

Symbol	Meaning	Value	Unit	V
A_{ap}	mirror aperture area	73 800	m^2	A
A_{rec}	receiver tube surface area	2815	m^2	A
$\eta_{opt,0}$	peak optical efficiency	0.755	—	A
ε	receiver emissivity	0.12	—	A
h_{conv}	envelope loss, still air	0.35	W/(m^2 °C)	A
h_{wind}	wind term on envelope loss	0.25	W/(m^2 °C)	A
hA	metal to fluid conductance, per node	2.0×10^5	W/°C	A
M_f	fluid held in the field	60 000	kg	A
M_m	receiver tube mass	64 000	kg	A
$c_{p,m}$	tube heat capacity	490	J/(kg °C)	A
n_i	incidence-angle exponent	0.30	—	H
$c_{p,0}$	fluid c_p, 0 °C	1.60	kJ/(kg °C)	A
$c_{p,1}$	fluid c_p slope	0.0024	kJ/kg $°C^2$	A
ρ_0	fluid density, 0 °C	1074	kg/m^3	H
ρ_1, ρ_2	density slope, curvature	0.6367, 7.81×10^{-4}	kg m^{-3} $°C^{-1}$, $^{-2}$	H
T_{max}	fluid stability limit	400	°C	A
T_{freeze}	fluid freezing point	12	°C	A
$T_{f,min}$	freeze-protection setpoint	60	°C	A
M	total fluid inventory	1.8×10^6	kg	A
UA_t	tank heat loss, each	1500	W/°C	A
ΔT_{rec}	recirculation margin	12	°C	A
W_{rec}	recirculation changeover width	10	°C	H
L_{min}, L_{max}	working level limits	0.06, 0.94	—	A
Q_{PB}	power block thermal rating	27.8	MW	A
η_{PB}	block efficiency at rating	0.36	—	A
T_{sat}	evaporator saturation	285	°C	A
ε_{SG}	steam generator effectiveness	0.85	—	A
$T_{j,min}$	minimum inlet temperature	305	°C	A
x_{min}	minimum stable load	0.25	—	A
k_T	efficiency gain per degree	0.0013	1/°C	A
P_{fix}	fixed auxiliary load	0.32	MW	A
τ_{SG}	steam generator lag	9	s	A
Δp_0	field drop at design flow	2.0×10^6	Pa	A
$\dot{m}_0$	design field flow	222	kg/s	A
η_p	pump and drive efficiency	0.75	—	A
μ_{ref}	viscosity, design mean	0.277	mPa s	H
DNI_{max}	clear-sky noon irradiance	950	W/m^2	A
h_{up}, h_{dn}	sunrise, sunset	6.2, 19.3	h	A
τ_{sky}	cloud edge crossing time	0.04	h	A
t_0	clock time at the start of a run	4.0 (per case)	h	A
p_{base}	off-peak price	42	$/MWh	A
p_{peak}	evening price adder	165	$/MWh	A
c_{aux}	cost of supplemental heat	58	$/MWh	A
c_{st}	cost of a block start (not charged)	900	$	A
$T_{m,max}$	receiver metal limit	430	°C	A
$Q_{aux,max}$	heater rating	12	MW	A
ΔP_b	power tolerance band	0.4	MW	A
$\tau_f, \tau_{pb}, \tau_b, \tau_d$	actuator lags	3, 5, 5, 40	s	A
τ_{Tm}, τ_{Pm}	thermocouple, power transducer lags	2.2, 1.8	s	A
σ_{SB}	Stefan-Boltzmann constant	5.67×10^{-8}	W $m^{-2}K^{-4}$	F

Table 6.3 Optimal outlet temperature at solar noon against sky clarity

Sky × clear	DNI (W/m^2)	Q_{abs} (MW)	T_{opt} (°C)	Loss (MW)	$\dot{m}_f$ (kg/s)	Net gain over 393
1.00	950	50.8	393	3.98	207	0.0 %
0.70	665	35.6	393	3.81	140	0.0 %
0.50	475	25.4	340	3.15	238	8.8 %
0.35	333	17.8	327	2.96	232	21.4 %
0.25	238	12.7	325	2.89	167	36.0 %
0.15	143	7.6	325	2.84	81	–

6.4 Open-Loop Character

The field holds 60 000 kg of fluid and the design flow is 222 kg/s, so the transit from the cold tank inlet to the collector outlet takes 4.5 min. That is the dominant dynamic in the outlet loop and it is not a fixed dead time, because the flow is the manipulated variable. Cutting the flow to a sixth of design to defend a temperature under cloud stretches the transit to 27 min. A loop tuned at design flow is a different loop at a sixth of it, which is the standard difficulty in distributed collector fields [6]. The same delay set by the manipulated flow governs the buried network of Volume 9, Chapter 6.

Two thermal masses set the rest. The receiver metal is 31.4 MJ/°C and the fluid in the field is 146 MJ/°C, so the cloud of the third scenario, which removes 37 MW of absorbed power, moves the metal quickly and the fluid behind it. The tanks are three orders of magnitude larger, holding 99.5 MWh between the level limits, which is 3.6 h at the power block thermal rating.

Manual mode is where the inventory and the transit become visible. Starting at half past five in the morning with the whole field at the 295 °C of the cold tank, the operator sets a field flow and watches the outlet climb node by node. Setting it too high keeps the field cold and delays the first useful heat. Leaving it at the 40 kg/s the scenario starts with carries the outlet past the 400 °C stability limit at 07:42 and the receiver metal past its 430 °C limit at 07:50, while the inlet node is still near 310 °C. Raising the flow to 200 kg/s as the outlet passes 393 °C moves the outlet reading by 2 °C within 0.8 min, because a flow change acts on every node at once, but a fall of 20 °C takes 3.1 min. There is no setpoint anywhere in that run, and the reader has already met the central trade-off.

The trade-off is worth stating as a table. Table 6.3 reports the outlet temperature that maximizes net electrical output at solar noon for a range of sky clarity, computed by the application's own golden-section search over the supervisor's interval (325 to 393 °C for a 300 °C cold tank), against what the design setpoint of 393 °C would deliver at the same conditions.

Above half the clear-sky irradiance the optimum is the top of that interval, which is the design value, because the field is collecting enough that the flow needed to reach 393 °C is comfortable and the conversion gain from running hot outweighs the extra radiation. Below that the optimum falls sharply and levels out at 325 °C, which is the bottom of the interval, 25 °C above the cold tank, rather than a property of the field. With the interval opened down to 312 °C the quarter-sky optimum is 319 °C. At the bottom row the comparison stops being a percentage. Holding 393 °C at 143 W/m^2 produces exactly zero net electrical output, because the 3.51 MW of receiver loss leaves 4.1 MW to be carried at 18 kg/s, which is below the power block's minimum stable load, while the optimizer still extracts 0.89 MW.

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.5 Regulatory Control

6.5.1 The Discrete Algorithm

All four loops use the velocity-form PID of Section 1.6.2, Equations (1.10) to (1.13). The chapter writes the action a as γ, $+1$ for a reverse-acting loop and -1 for a direct-acting one, as in Chapter 1.

Table 6.4 Regulatory loops as implemented

CV	MV	K_c	τ_I (h)	Character
T_{out}	$\dot{m}_f$	2.5	0.05	direct, delay dominant
P_e	$\dot{m}_{PB}$	10	0.01	reverse, fast
T_{sg}	b	1.2	0.03	reverse, interlocked
L	f_d	8	0.30	direct, override

The tracking time is $T_t = \tau_I$ on all four loops and no output has a rate limit. In manual mode, and while a loop is out of service, the controller holds its integral so that the output equals the value applied, and in automatic the manual stations follow the applied commands at every scan, so both transfers are bumpless. Every loop in this application runs with $\tau_D = 0$, so there is no derivative action anywhere and no derivative filter to choose. The regulatory sample period is $\Delta t_c = 7.2\,\mathrm{s}$.

Two limits act after a controller output and are invisible to its integral. The power block command is cut back toward the field flow that reaches the block or the hot tank, none of it while the return is recirculated, as the hot tank approaches its lower level limit, and the level loop's defocus demand passes the high selector of Equation (6.20). In both cases the integral is back-calculated against the controller's own clipped output rather than the value applied. At 15:05 on the broken day of Figure 6.3, with the tank at its lowest working level, the power controller sits at its 190 kg/s limit while the plant receives 140 kg/s.

6.5.2 Four Loops, One of Them an Override

Table 6.4 gives the pairings and default tunings. Two loops are direct acting. More field flow makes the outlet colder, and more defocus makes the tank level fall, so both process gains are negative.

The outlet loop has a process gain of about 0.38 °C/(kg s) at the design point and about four and a half minutes of transport delay, and the gain rises as the irradiance falls. The conventional rule for a plant whose gain varies with the load is to tune it at the weak condition. The default tuning is not set that way. With it the second scenario's step from 370 back to 393 °C settles within 1 °C in 12 min without overshoot under a clear sky. With the sky set to 0.25 before the run, the same gains are still swinging between 367 and 374 °C in the 18 min before that step, and the step back takes 69 min to settle, its 4.3 °C overshoot clipped by the high-temperature protection described below. Tripling the outlet loop gain to 7.5 settles the same step in 26.5 min, and no tested gain or hour of the day makes the loop unstable.

The bypass loop is the interesting structural case. It holds the steam generator inlet T_{sg}, read after the heater and the steam generator lag and shown as the loop's trend, by sending field fluid straight past the hot tank, and its process gain is proportional to how much hotter the field is than the tank. When the field is colder, opening the bypass makes the junction colder rather than hotter, so the gain changes sign. No fixed-action controller survives a sign change, so the loop is interlocked out whenever $T_{f,10} < T_h + 3\,^{\circ}\mathrm{C}$ or $\dot{m}_f < 15\,\mathrm{kg/s}$. While it is out the valve closes and the integral tracks the closed position, so a field cooling after sunset does not go on feeding the junction. A real plant interlocks it for exactly the same reason.

The level loop is an override rather than a regulator. It sits against zero for most of the day, because the tank level is below its setpoint and the direct-acting controller therefore demands negative defocus, which the output limit refuses. It wakes when the tank fills. A second demand comes from a high-temperature protection that ramps defocus linearly once the outlet passes $T_{max} - 4\,^{\circ}\mathrm{C}$, and the two are combined by a high selector,

$$f_d = \max\left(f_{d,L},\ \Pi_0^{100}\left[40\,(T_o - (T_{max} - 4))\right]\right), \tag{6.20}$$

where T_o is the larger of $T_{f,10}$ and the thermocouple reading. The protection threshold sits deliberately above any setpoint the supervisor is allowed to choose. Placing it lower makes the two fight, because

the optimizer asks for the top of its range, the protection reads that as an excursion and defocuses part of the field, and the plant then collects less than it should while every loop reports success.

Pause. Try Exercise 6.2 before reading on.

6.5.3 Taking a Loop Out of Service

When the absorbed power falls below 2.5 MW the outlet temperature controller has no authority and no purpose. The application detects that condition, parks the setpoint on the measurement, and walks the field flow down to its floor at a rate limit rather than stepping it. Parking the setpoint matters for a reason that is about instrumentation rather than control. A setpoint left at 393 °C over a field that is at 180 °C after sunset puts a two hundred degree error on the trend and in the scorecard, and no action exists that would reduce it. The operator's own setpoint is kept behind the parked one, so a fixed 393 °C returns when the sun does. The same reasoning applies to the steam generator inlet setpoint, which is parked whenever the power block is not running.

Freeze protection is the other out-of-service mode and it runs in the opposite direction. When the coldest fluid node or the cold tank falls below 60 °C, electric heat tracing is energized until it is 8 °C above that, and the field flow floor is raised to 12 % of design so that the fluid keeps moving. The hot tank has its own tracing on the same rule, whose heat goes into whatever is left at its bottom. Both cost energy. The fluid freezes at 12 °C and a frozen header is an excavation rather than a restart, which is the assessment that kept synthetic aromatic in trough fields after molten salt was evaluated [11].

6.5.4 Model Predictive Control

A second outlet controller can take the field flow in place of the outlet loop of Table 6.4, within the general option of Section 1.6.3. It is a receding-horizon optimization [15] of the family developed for distributed collector fields at Almeria [7, 16], and it is an option rather than the default, since every scenario except the eleventh loads the proportional-integral loop. Nothing else changes. The supervisory setting still writes the outlet setpoint T_{SP} by the rules of Section 6.6, the division of labor proposed for trough plants in which an operating-point optimization sits above a predictive field controller [17], and the power, bypass and level loops, the power derate near an empty tank, the bypass interlock and the high selector of Equation (6.20) all stay in force over the plan just as over the loop. While the predictive controller is in charge the outlet loop tracks the applied flow at every scan with no preload, which makes switching between the two bumpless both ways.

Controller model. The plan predicts with the receiver lumped into five cells, each a fluid node and a metal node, written from Equations (6.4) to (6.7) with two plant nodes per cell, the declared defaults of Table 6.2 and the design weather (25 °C, 3 m/s, mirrors at 0.96), never the parameters of the running plant,

$$\frac{M_f}{5} c_p \frac{\mathrm{d}\hat{T}_{f,j}}{\mathrm{d}t} = 3600\big[\dot{m}_f c_p(\hat{T}_{f,j-1} - \hat{T}_{f,j}) + 2hA(\hat{T}_{m,j} - \hat{T}_{f,j})\big] \tag{6.21}$$

$$\frac{M_m}{5} c_{p,m} \frac{\mathrm{d}\hat{T}_{m,j}}{\mathrm{d}t} = 3600\Big[\frac{\hat{Q}_a}{5} - 2hA(\hat{T}_{m,j} - \hat{T}_{f,j}) - 2q_{loss}(\hat{T}_{m,j})\Big] \tag{6.22}$$

for $j = 1 \ldots 5$, with $\hat{T}_{f,0} = T_c$ from the inlet thermocouple and $\hat{Q}_a = A_{ap}\,\mathrm{DNI}\,\eta_{opt}\,(1 - \hat{f}_d)$. The irradiance is the pyrheliometer reading that the out-of-service test and the supervisor already use, and $\hat{f}_d$ is the defocus the high selector last commanded, passed through the declared 40 s drive lag. A cell holds 12 t of oil and 12.8 t of tube, so the transport delay is the declared inventory over the flow, $M_f/\dot{m}_f$, which is 4.5 min at 222 kg/s and 18 min at 55 kg/s. The delay therefore stretches with the manipulated variable by construction rather than by a schedule. Flow steps of $\pm 20\,\mathrm{kg/s}$ about 200 kg/s at clear noon give a gain of -0.436 on the plant against -0.423 on the cells, and steps

of $\pm 6\,\mathrm{kg/s}$ about 50 kg/s under a 285 W/m^2 noon give -1.644 against -1.589, all in °C per kg/s, so the cells read 2.9 and 3.3 % low. Plant and cells pass half of the response after about 37 000 kg of oil at both flows (37 440 and 36 720 kg on the plant), the signature of a delay that lives in fluid inventory rather than in seconds. The two step responses differ by 0.23 and 0.36 °C RMS at the two points, and over a sequence of flow steps and a cloud in manual mode the cells alone follow the outlet thermocouple to 1.35 °C RMS (4.30 °C at worst).

At each decision the cells are linearized numerically, by central differences, at the present estimate and at the flow at which the algebraic field of Equation (6.26) would hold $T_{SP} - b$ at the present irradiance, not about the applied flow, because the field's gain scales as $1/\dot{m}_f^2$ and its time scale as $1/\dot{m}_f$, so a model taken at a flow the plan is about to leave mispredicts the move it plans. The linearization is discretized with an exact zero-order hold, and its states are the ten cell temperatures and the previous planned flow.

Offset. An output bias removes steady error [18],

$$b_k = b_{k-1} + \left(1 - e^{-\Delta t_c/\tau_b}\right)\left(T_{out,k} - \hat{T}_{f,5} - b_{k-1}\right), \tag{6.23}$$

with $\tau_b = 72\,\mathrm{s}$, updated at every scan whichever controller acts, so a transfer to the predictive controller starts from a warm estimate, and held constant over the horizon. It absorbs what the cells omit, namely the true mirror cleanliness, the wind, the ambient temperature and the dispersion of ten nodes, and with it the open-loop gap quoted above falls to 0.28 °C RMS (1.01 °C at worst). The offset-free controller of Volume 9, Chapter 4 carries its offset in an estimated fuel bias instead.

Horizon and objective. A decision is taken every $T_s = 36\,\mathrm{s}$, five regulatory scans, over $N = 33$ decisions (20 min) grouped into twelve move blocks, four of one decision, three of two, two of three and one each of four, five and eight. The first move alone reaches the pump, and it stays there until the next decision. The plan minimizes

$$J = \sum_{k=1}^{N} \left(\hat{T}_k + b - T_{SP}\right)^2 + P\left(\hat{T}_{ss} + b - T_{SP}\right)^2 + \lambda K_u^2 \sum_{j=1}^{12} \Delta\dot{m}_j^2 + \rho s^2 + \sigma s \tag{6.24}$$

where $\hat{T}_k$ is the predicted outlet cell temperature, $\hat{T}_{ss}$ the outlet at which the linear model settles with the last planned flow held, K_u that model's steady gain from flow to outlet and $\Delta\dot{m}_j$ the change of flow at block j. A move is thus priced at λ times the square of the steady outlet shift it would cause, against one decision of squared error. The weights are $P = 10$, $\lambda = 0.02$, $\rho = 10^4$ per °C^2 and $\sigma = 10^3$ per °C. They were set once on the eleventh scenario and frozen: λ is the smallest value on a grid from 0 to 0.2 at which 1 °C of thermocouple noise raises the flow travel by no more than half (1.49 times at 0.02 against 1.67 at 0.01), and horizons of 10 to 40 min give 46.5 to 47.0 °C min of outlet error on that scenario, so the move price, not the horizon, limits the plan.

Constraints. The outlet is held under the threshold at which the protection of Equation (6.20) begins to defocus,

$$\hat{T}_k + b \le T_{max} - 4 + s, \qquad s \ge 0, \tag{6.25}$$

a soft constraint with one slack s for the whole horizon, on the run's own T_{max} as the protection reads it. The flow is held hard between the clamps of the outlet loop, 8 kg/s, or 26.6 kg/s (12 % of design) while freeze tracing is energized, and 520 kg/s. The defocus the level loop demands is held over the horizon, while the protection's share is not, because it releases as soon as the outlet returns under the threshold, which is what the plan is for. While the outlet reads less than the inlet plus ΔT_{rec} the field is still recirculating to the cold tank, and the flow is held on its floor, which is also where the loop leaves it. After that hold, until the outlet is within 2 °C of T_{SP} and one declared field inventory of 60 t

has been pumped, every planned move is also capped at the pump ramp the application declares for taking the field out of service, 600 kg/s per hour or 6 kg/s a decision. Uncapped, the first plan on every case that starts before sunrise drove the pump from 8 to 520 kg/s and back within a minute, pushing the oil warmed near the inlet past the oil the night had cooled at the outlet end. Capped, the flow stays at or below 110 kg/s in the first ten minutes after the hold on all seven such runs. Applied to every plan, the same cap raises the outlet error on the eleventh scenario from 46.8 to 1224.1 °C min, so it is confined to the start-up. Below 2.5 MW absorbed the plan stops and the field is taken out of service as described above.

Preview. The irradiance over the horizon enters the prediction as a known input. With the preview on it is the sky fraction the pyrheliometer reads now, persisted over the clear-sky profile of Equation (6.1) for the next 20 min, which is the rule the Forecast setting applies to its first hour. With the preview off the present irradiance is held. Neither carries a cloud edge, so every cloud in every scenario arrives unannounced, which is all a measurement at the field can offer.

Acceptance and fallback. The plan takes effect only if the solver returns an optimal status, every planned flow lies inside its bounds and inside the start-up cap where it applies, the first planned flow is the one applied, and its cost, slack and predicted outlets are all finite numbers. Otherwise the historian is marked, the outlet loop continues from its tracked state, and the next valid plan takes the flow back. Eight forced failures at 10:30 on the eleventh scenario, with the noise off, each changed the flow at the hand-over by 0.001 kg/s. The first plan is solved at $t = 0$, so a run started under the predictive controller is under it from the first logged row. In Chrome a plan costs a mean of 0.37 to 0.39 ms over the 350 decisions of the eleventh scenario (4.7 to 5.2 ms for the first, before the browser compiles it), which is 3.7 to 3.9 % of one core at one simulated hour per second and 18 to 19 % at five.

Three models kept apart. The plant is the thirty-state model of Section 6.3 at the run's parameters and weather. The identified model is a first-order-plus-dead-time fit to flow steps on that plant at 200 kg/s, a gain of −0.437 °C per kg/s, a dead time of 0.61 min and a lag of 3.19 min with a residual of 2.0 % of the response, which fixes only the retuned loop compared in Section 6.7.6. The controller's model is Equations (6.21) and (6.22) at the declared defaults, relinearized at every decision. A learner's change to any parameter is therefore a mismatch the bias must absorb, not knowledge the plan receives.

6.6 The Supervisory Layer

Three supervisory settings are available and the application exists for the difference between them. The first holds the operator setpoint. The second chooses the outlet temperature that maximizes net electrical output at the irradiance measured now. The third chooses it from the mean of the sky measured now and the sky forecast ten and twenty minutes ahead, and adds a schedule for the power block.

The objective is written once as algebra so that it can be evaluated without integrating the field. For a candidate outlet temperature T at inlet T_c, the absorbed power per node is known from Equation (6.3), the receiver metal sits above the fluid mean by whatever the conductance requires, and the flow follows from the energy that the ten nodes have to deliver,

$$\bar{T} = \tfrac{1}{2}(T_c + T), \quad T_m = \bar{T} + \frac{q_a}{hA}, \quad \dot{m} = \frac{10\,[q_a - q_{loss}(T_m)]}{c_p (T - T_c)}, \tag{6.26}$$

with q_{loss} the node loss of Equation (6.6). That flow is then passed through the power block map and the pump, and the objective is the net electrical output,

$$J(T) = P_e^{pb}\big(T, \min(\dot{m}, \dot{m}_{duty})\big) - P_{pump}(\dot{m}). \tag{6.27}$$

A golden-section search maximizes J over an interval in eighteen iterations, with no convergence test and no iteration count that depends on the data. The total cost of the supervisory layer is therefore

twenty evaluations of Equation (6.26) every three minutes, two to start the search and one per iteration. Any candidate that needs more than 2.2 times the design flow scores −99 MW, which marks it infeasible.

Two details of that search matter. The interval runs from 25 °C above the cold tank, or 312 °C if that is higher, to 393 °C, seven degrees below the fluid limit. A fixed lower end of 312 °C would ask for a rise of only 12 °C over a 300 °C cold tank, which at clear noon needs 1700 kg/s, more than three times the pump limit. Such probes score as infeasible and the search moves away from them, so in the present model a fixed lower end gives the same or a better answer at every condition tested (cold tank 295 to 310 °C, sky 0.25 to 1.0), and the margin binds only below about a third of clear sky, as Table 6.3 shows. The other detail is that the objective must be continuous. The plant has a genuine minimum stable load below which the power block produces nothing, and a cliff inside a golden section is not a search problem but a posing problem, because the objective stops being unimodal and the answer depends on where the first two probes landed. The optimizer is therefore handed a smoothed version of the same curve, gated over a band around the minimum load, while the plant keeps the discontinuity it actually has.

The chosen setpoint is rate limited before it reaches the loop, at a default of 500 °C/h. A supervisory setpoint that steps is a disturbance to the loop underneath it, and the loop underneath this one has four and a half minutes of transport delay.

The third setting adds a schedule. Twenty-eight half-hour blocks cover the remaining day, the net field power in each block is estimated from Equation (6.26) at the supervisor's current outlet temperature and the forecast sky (the sky measured now for the first hour), and the price is known. A merit order alone is not enough, because the tank holds 106.3 MWh and a clear day brings 370 to 380 MWh, so a schedule that saves everything for the evening fills the tank by early afternoon and then defocuses the rest off the mirrors. The schedule is therefore a causal forward pass that charges the tank block by block, runs the power block wherever the price clears a trial threshold, and counts whatever overflows,

$$\begin{aligned} \tilde{E}_k &= E_k + q_k \Delta \\ D_k &= \mathbb{1}[\pi_k \geq \theta]\, \min\!\left(Q_{PB}\Delta,\ \max(0, \tilde{E}_k)\right) \\ E_{k+1} &= \min\!\left(E_{cap},\ \tilde{E}_k - D_k\right) \end{aligned} \tag{6.28}$$

with the spill counted whenever the cap binds. Spill rises monotonically with the threshold θ, so fourteen bisections find the highest threshold that spills no more than 0.1 MWh, which is the dearest set of hours the tank can actually hold the energy for. The threshold plays the part of the water value of Section 1.4.6, found by forward simulation rather than by merit order. Nobody has to choose a threshold. The block runs whenever the price clears it, and also while a must-run rule holds, which latches when the hot tank comes within eight points of the level override setpoint and releases twelve points below it, so the turbine takes heat before the mirrors have to shed it. Measured on the evening scenario the threshold rises from $28.7 per megawatt-hour at noon to $205 at six in the evening, holding the block off in the cheap early afternoon, and then falls once the evening peak is being served.

6.7 Scenarios and Measured Results

Table 6.5 lists the twelve scenarios in selector order and gives the scorecard at the end of each window with the settings each case loads. Credit, Reset and reloading work as Section 1.6.4 describes, and the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

6.7.1 A Cloud Against a Fixed Setpoint

The third scenario runs at eleven in the morning and drops the irradiance to a quarter at 11:48. The temperature loop does what it is told. It cuts the field flow from 199 kg/s to a mean of 33 kg/s between

Table 6.5 Scorecard at the end of each scenario window, as loaded

#	Scenario	Start	Window (h)	Generated (MWh)	Harvest (MWh)	Loss (MWh)	Defocus (MWh)	Margin ($)
1	Sunrise by hand	05:30	5	0.0	62.3	22.7	0.0	0
2	Outlet loop and its lag	11:00	3	29.1	136.6	11.9	0.0	847
3	Cloud on a fixed setpoint	11:00	4	40.0	135.3	15.9	13.8	1 205
4	Clear day	06:00	14	119.3	370.6	51.4	3.4	7 869
5	Broken day	06:00	14	25.4	111.0	48.2	9.4	1 168
6	Bypass and mixing	09:00	6	59.5	246.2	24.1	5.0	1 885
7	Full tank and defocus	09:00	8	32.0	135.9	30.1	184.9	1 138
8	Evening peak	06:00	18	80.6	261.6	53.0	0.1	6 413
9	Forecast error	08:00	14	61.6	186.2	43.5	3.5	5 680
10	Cold night	18:00	14	8.3	7.7	31.8	3.1	1 357
11	PI against MPC, passing clouds	09:30	3.5	35.0	111.3	13.8	0.0	1 086
12	Challenge	06:00	16	3.2	49.5	55.9	188.2	210

Table 6.6 Fixed against variable outlet temperature

Day	Strategy	Gen. (MWh)	Harvest (MWh)	Loss (MWh)	Pump (MWh)	Collect (%)
clear, 16 h	fixed	120.8	370.6	55.1	4.01	87.9
clear, 16 h	variable	121.1	377.8	51.8	6.25	88.9
broken, 14 h	fixed	25.4	111.0	48.2	0.39	70.5
broken, 14 h	variable	33.5	129.9	39.3	2.71	78.1

The strategy gives the same run whether it is selected before Run or carried through Reset.

12:00 and 12:36, resting on its 8 kg/s floor for nine minutes, and after a 25 °C dip at the cloud edge the outlet comes back within 4 °C of 393 °C 18 min into the cloud and stays within 4.3 °C of it until the clearing. Meanwhile the receiver loses 3.75 MW of the 11.9 MW it absorbs, nearly a third. The clearing at 12:36 is harder than the cloud, because the integral has to carry the flow from about 30 back toward 200 kg/s. The outlet peaks at 399.95 °C and rides between 395 and 400 °C until 13:26 while the high-temperature protection cycles the mirrors in and out. That protection accounts for all 13.8 MWh defocused, because the tank never passes 58 %. Over the four hours the scorecard shows 15.9 MWh of receiver loss against 135.3 MWh harvested, and the collection efficiency, which is heat into the fluid over heat absorbed by the receivers, at 89.1 %. The instructive part is where to look. The temperature trend says the controller succeeded and the energy trend says the plant did not.

> **TRY IT IN THE APP** **3 · What a Cloud Does to a Plant Defending a Temperature**
>
> Load scenario 3 and press Run in Auto on Fixed T. At 11:48 the sky falls to a quarter. Watch Field flow on the outlet strip fall from about 199 kg/s and rest on its 8 kg/s floor, while the outlet returns to within 4 °C of 393 °C. Then compare receiver loss with absorbed on the Where the sunlight goes strip: about 3.75 of 11.9 MW.

6.7.2 Clear Day Against Broken Day

The comparison the application is built around runs the identical day twice, once holding 393 °C and once letting the supervisor choose. Table 6.6 gives both days, and Figure 6.2 draws the same numbers as bars.

On a clear day the supervisor is worth nothing. It harvests 1.9 % more heat and loses 5.9 % less from the receivers, and it spends almost the entire difference on pumping, because the lower

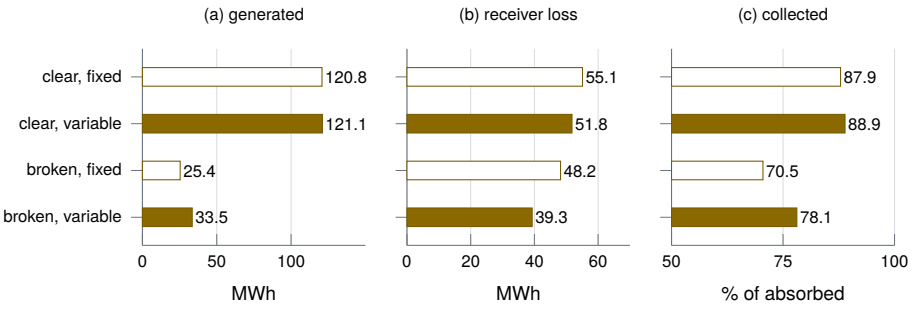


Figure 6.2 Table 6.6 as a chart: (a) net generation, (b) receiver loss and (c) the share of absorbed heat that reaches the fluid, for the fixed setpoint (white bars) and the variable setpoint (solid bars) on the clear day of 16 h and the broken day of 14 h. The axis of (c) starts at 50 %, so its bars show differences, not ratios. Notice that the variable setpoint cuts the loss on both days but raises generation only on the broken day.

outlet temperature it selects before 10:00 and after 16:00 (as low as 321 °C, while it holds 393 °C in between) requires more flow and pumping power rises as the cube of the flow. Net generation over sixteen hours is 121.07 MWh against 120.77 MWh, a difference of 0.3 % that is inside any reasonable uncertainty on the model. That is the correct result and it is worth reporting. A supervisory optimizer that reproduces the operator setpoint when the operator setpoint is right is a working optimizer, and a study that reported only the cloudy-day number would be reporting half a result.

> **DID YOU KNOW? More Heat, Almost No More Power**
>
> From the simulator: on the clear day the variable setpoint harvested 377.8 MWh against 370.6 MWh for the fixed one, but pumping took 6.25 MWh against 4.01 MWh, and net generation came to 121.1 MWh against 120.8 MWh (Table 6.6).

On a broken day the same optimizer is worth 32 %. It drops the outlet setpoint to about 318 °C during the cloudy stretches, the bottom of its interval, which keeps the flow near 134 kg/s rather than 30 kg/s and the receiver cooler, and the receiver loss falls from 48.2 MWh to 39.3 MWh at exactly the time the field can least afford it. Heat into the fluid rises by 17.0 % and net generation by 32.1 %, because the generation gain compounds the harvest gain through the part-load efficiency curve. The fraction of absorbed heat reaching the fluid goes from 70.5 % to 78.1 %. Published studies of this configuration report the same shape of result for the same reason [1, 2].

Two observations about the mechanism are worth making explicit. The optimizer is not clever, being a golden section on one scalar over an algebraic model. And it never touches an actuator. The ordinary temperature controller does all the work underneath, and it would have done exactly the same work for a wrong setpoint.

6.7.3 The Schedule Is Worth More Than the Setpoint

The eighth scenario runs a hazy day (sky 0.72) for eighteen hours from six in the morning, with the evening price peak at 19:30, after sunset. Under a fixed setpoint and a constant power demand the plant generates 78.1 MWh and books $2 875. Under the variable-temperature supervisor it generates 79.6 MWh and books $3 081. Under the schedule it generates 80.6 MWh, about the same, and books $6 413, which is 2.2 times as much. The schedule holds the power block off from 11:58 to 14:31, runs at rating through the late afternoon under the must-run rule (stopping only from 15:44 to 15:55) while the tank fills to 85 % by 17:14, lets the rule release at 18:21, and then runs into the evening peak from 19:06, when the price is $200 per megawatt-hour, until 21:11. Of its generation, 39.2 MWh

comes after 16:00, against 8.7 MWh under the fixed setpoint. On a clear sky the same schedule books only 1 % more than constant power ($8 208 against $8 112), because a field with a solar multiple of 1.83 and 3.6 h of storage has to run the block most of the day or overflow.

That ordering is the practical lesson of the application. Choosing the temperature is worth a few percent of energy. Choosing when to convert the energy is worth most of the revenue, and the two decisions are independent enough that a plant can do either without the other.

6.7.4 Defocus as an Override

The seventh scenario starts at nine in the morning with the tank nearly full and the power block asked for only 4 MW. Over eight hours the plant throws away 184.9 MWh, which is more than the 135.9 MWh it harvests. Every controller is holding its setpoint. There is simply nowhere to put the heat, so the level override takes the mirrors off the sun, and defocusing at noon on a plant sized for a whole day is normal operation rather than a fault. The setpoint that decides how much gets thrown away is the level override setpoint rather than any tuning constant. Raising it from 88 % to 96 % reduces the defocused energy from 189.5 MWh to 179.7 MWh over ten hours by giving the tank more room, and it leaves the drive less room to act before the tank is genuinely full. The power setpoint matters more. Asking the block for 10 MW instead of 4 MW over the same eight hours cuts the defocused energy to 69.4 MWh and raises generation from 32.0 to 80.0 MWh.

> **TRY IT IN THE APP** **7 · A Full Hot Tank, and Defocus as an Override**
>
> Load scenario 7 and run its eight hours in Auto with the Electrical output setpoint at 4 MW, then read Defocused: 184.9 MWh. Reset, set the setpoint to 10 MW before pressing Run, and run again. Defocused falls to 69.4 MWh and Generated rises from 32.0 to 80.0 MWh.

6.7.5 The Graded Challenge

The twelfth scenario opens sixteen hours from six in the morning with weak tuning, a fixed setpoint, mirrors ten percent dirty, the bypass in manual, the level override setpoint at 96 % and the power block asked for 10 MW from the start. To pass, one 16 h Auto run with the case's parameters and disturbances must generate more than 60 MWh, deliver more than 78 % of the absorbed heat into the fluid, defocus less than 20 MWh, book a margin above $4 500, and keep T_{sg} within 60 °C of its setpoint over the last third of the day, from 16:34. The default configuration generates 3.2 MWh, collects 49.8 %, defocuses 188.2 MWh, books $210 and holds T_{sg} within 5 °C of its setpoint, failing four of the five parts. The suggested answer is two changes, the schedule and the tuned gains, and with them the day generates 72.8 MWh, collects 84.1 %, defocuses 5.1 MWh, books $6 780 and holds T_{sg} within 29 °C. It passes every part, the narrowest being collection at 6.1 points above its threshold, and Reveal followed by Reset gives the same run. Together they are worth 69.6 MWh and $6 570, but they are worth different things. The tuned gains alone, inside the default structure, generate 68.7 MWh with 81.9 % collection and 11.2 MWh defocused, yet book only $2 663. The schedule with the weak gains books $4 751 but generates only 41.1 MWh and defocuses 75.0 MWh. The gains recover the energy and the schedule earns the money. Under the schedule the power loop tuning matters as much as the outlet loop tuning, since retuning the outlet loop alone gives 50.3 MWh and fails. The bypass in automatic, a 90 % level setpoint and a 388 °C inlet setpoint do not raise the score, and the lower level setpoint costs about $300, because the must-run rule then sells more of the day's heat in the afternoon (33.9 MWh against 31.5 MWh) and less after 18:00 (22.0 MWh against 24.7 MWh).

6.7.6 Predictive Control Against the Outlet Loop

Table 6.7 runs nine cases under three outlet controllers, each from $t = 0$ with the case's own settings, seed and duration, the controller chosen before the run and the supervisory setting the case loads (scenarios 4 and 5 also on Variable T, as their second runs ask). The shipped loop keeps its shipped tuning, and the challenge is run both as shipped and with its suggested settings. The predictive

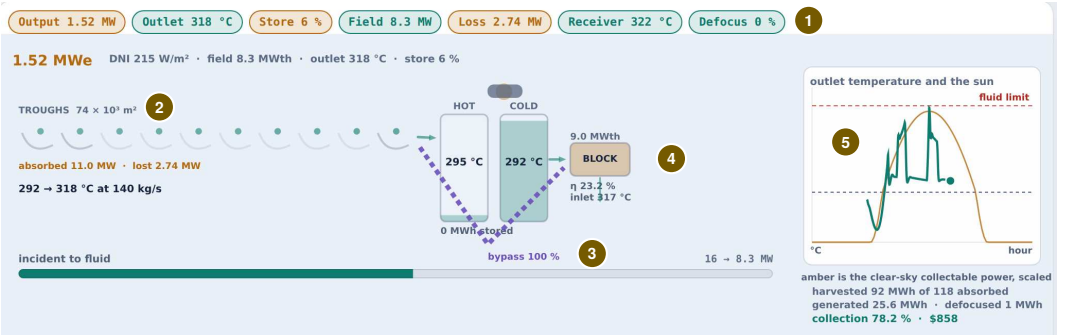


Figure 6.3 The application at 15:05 on the fifth scenario, nine hours into the broken day, with Variable T selected before the run. At 215 W/m^2 of direct normal irradiance the optimizer has selected 318 °C rather than 393 °C, and the field flow is holding at 140 kg/s rather than collapsing to defend a temperature the afternoon cannot support. Receiver loss is 2.74 MW against 11.0 MW absorbed, and the run has delivered 92 MWh of the 118 MWh absorbed into the fluid, a collection efficiency of 78.2 %. The hot tank is held at its lowest working level of 6 %, the consequence of an afternoon at a quarter of clear-sky irradiance. The view is cropped to the plant. Callouts: (1) the badge row with output, outlet, store, receiver loss and defocus, (2) the trough row with the absorbed power and the receiver loss, (3) the hot and cold tanks with the bypass carrying the whole block flow, (4) the power block with its efficiency and inlet temperature, and (5) the outlet temperature against the clear-sky collectable power, with the fluid limit marked.

controller runs with the eleventh scenario's settings on every case, with no retuning. The third row is the stronger of two classical structures given the same information, whichever has the lower outlet error on the case. One is the shipped loop retuned once by the SIMC rule on the identified model ($K_c = 6.0$, $\tau_I = 0.053$ h). The other is the shipped loop trimming a static feedforward, the flow that Equation (6.26) says holds the setpoint at the present pyrheliometer reading with the level loop's defocus held, started on the feedforward's own flow as the predictive controller starts on its own plan. Neither is in the application. Outlet error is the scorecard's IAE, the excursion columns count time and integrated degrees above the protection threshold $T_{max} - 4$ °C, and flow travel sums the absolute changes of the flow command. Defocused energy is reported rather than a count of defocus episodes, because the protection's command chatters across its first percent and a count depends on the counter.

The eleventh scenario is the case built for the comparison. It opens at 09:30 under a clear sky on Fixed T at 393 °C, three degrees under the protection threshold, and two clouds cross the field, a quarter of the sun from 10:00 to 10:48 and 0.45 of it from 11:36 to 12:00. The shipped loop answers the first cloud only as its error grows, so the outlet falls to 371.8 °C, and the integral then carries the flow down to its 8 kg/s floor after the outlet has already turned. The outlet climbs back to 396.8 °C inside the cloud, and when the cloud clears the flow is far too low, as shown in Figure 6.4. At 10:51 the loop's outlet peaks at 399.6 °C on 52.8 kg/s, while the predictive controller holds 394.9 °C on 142.7 kg/s, and 36 s later the protection has defocused 88 % of the loop's mirrors while the other run's mirrors are fully focused. Over the morning the loop spends 62.7 min above the threshold and defocuses 17.85 MWh, while the predictive controller keeps the outlet between 391.3 and 395.1 °C from the first cloud on. Net electricity is the same in both runs (34.99 against 34.98 MWh), because the block runs at its setpoint from the tank, and the difference is heat kept, with 37.8 MWh stored at 13:00 against 21.0 MWh.

Against the shipped loop the predictive controller lowers the outlet error on every case, by 11 % on the clear day at a fixed setpoint (4140.8 against 4666.4 °C min) and by 95 % on the eleventh scenario, and it spends no minute above the threshold on any case but the broken day on Variable T, where

Table 6.7 Outlet controllers compared on nine cases

Case (window)	Outlet controller	IAE (°C min)	Above 396 (min)	Above 396 (°C min)	Defocused (MWh)	Net (MWh)	Flow travel (kg/s)	Plans, rejected
11, passing clouds, Fixed T (3.5 h)	shipped PI	882.8	62.7	60.9	17.85	34.99	865	—
	MPC	46.8	0	0	0	34.98	622	350, 0
	MPC, preview off	47.1	0	0	0	34.98	622	350, 0
	PI with feedforward	**45.7**	0	0	0	34.98	574	—
3, one cloud, Fixed T (4 h)	shipped PI	624.8	44.0	49.1	13.84	39.99	616	—
	MPC	33.8	0	0	0	39.99	417	400, 0
	PI with feedforward	**31.0**	0	0	0	39.99	387	—
5, broken day, Fixed T (14 h)	shipped PI	8312.3	86.3	58.9	9.41	25.39	1050	—
	MPC	**7222.5**	0	0	0	**28.74**	999	1073, 0
	PI with feedforward	7461.9	11.4	11.6	0.89	28.51	905	—
5, broken day, Variable T (14 h)	shipped PI	5864.7	1.8	3.2	0.58	33.52	1870	—
	MPC	**5244.5**	1.4	1.0	0.19	32.40	6153	1073, 0
	PI retuned (SIMC)	5370.5	0	0	0	**33.61**	3277	—
4, clear day, Fixed T (14 h)	shipped PI	4666.4	18.2	29.2	3.43	119.34	625	—
	MPC	**4140.8**	0	0	0	**120.27**	737	1144, 0
	PI with feedforward	4304.9	14.9	25.9	2.87	119.57	654	—
4, clear day, Variable T (14 h)	shipped PI	3268.9	1.4	0.4	0.09	118.28	1252	—
	MPC	2805.9	0	0	0	**118.65**	4186	1144, 0
	PI with feedforward	**2701.7**	1.7	1.3	0.30	118.29	2027	—
9, forecast error, Forecast (14 h)	shipped PI	1826.6	31.9	18.9	3.55	61.58	1297	—
	MPC	**1020.0**	0	0	0	**63.13**	3935	1062, 0
	PI retuned (SIMC)	1260.9	10.0	5.1	0.96	62.85	2069	—
12, challenge as shipped (16 h)	shipped PI, weak tuning	7417.4	628	995.4	188.22	3.21	302	—
	MPC	**4326.7**	0	0	0	**72.78**	991	1136, 0
	PI retuned (SIMC)	4726.7	23.7	27.5	3.70	71.49	1130	—
12, challenge, suggested (16 h)	shipped PI, tuned	3593.3	28.6	19.0	5.11	72.85	1676	—
	MPC	**2839.7**	0	0	0.42	**74.63**	4439	1136, 0
	PI with feedforward	2896.1	10.7	9.2	2.77	73.55	2675	—

Bold marks the lowest outlet error and the highest net electricity of each case where the rows differ. The controller is selected before Run, and selecting it later and pressing Reset reproduces the same run. The defocused energy of the predictive controller on the suggested challenge is the level loop's, with no minute above the threshold.

it spends 1.4 min there against 1.8 min for the shipped loop. None of the 8 868 plans in the table was rejected. The comparison with the classical row says where that lead comes from. It is not the preview, because switching the preview off changes the outlet error by at most 0.64 % (on the eleventh scenario) and net electricity by at most 0.04 MWh, and it is not the receding horizon as such, because the loop trimming a static feedforward from the same pyrheliometer, started on its own flow, beats the predictive controller over the whole run on scenarios 11 and 3 (45.7 against 46.8 and 31.0 against

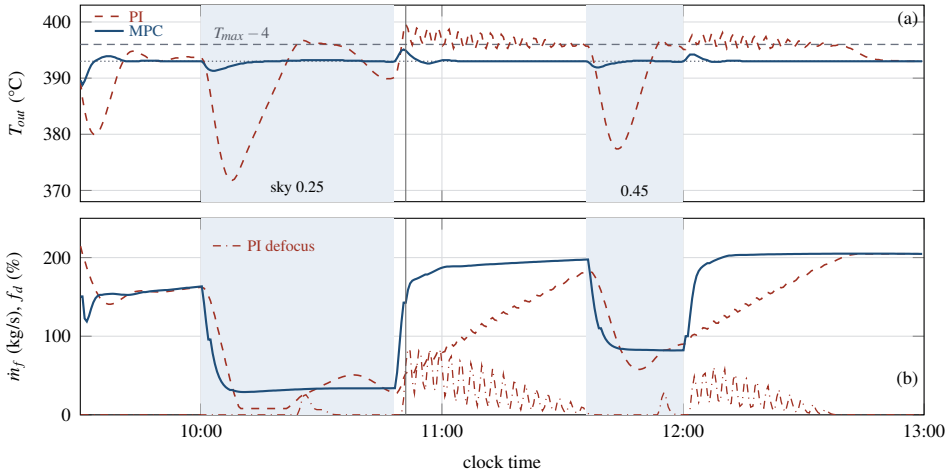


Figure 6.4 The eleventh scenario from 09:30 to 13:00 under the shipped outlet loop (red, dashed) and under the predictive controller (blue, solid), with the two clouds shaded and the instant 10:51 marked. (a) Outlet temperature against the 393 °C setpoint (gray, dotted) and the 396 °C protection threshold (gray, densely dashed). (b) Field flow under both controllers and the defocus the protection applies under the loop (red, dash-dotted), which reaches 88 % just after the first cloud clears.

33.8 °C min) and inside the clouds (22.3 against 34.9 and 15.0 against 22.8 °C min). The measured sun and a model of the field carry most of the result, whichever structure uses them. The predictive controller wins in two places. The first is the end of the morning approach, where on the clear day it begins raising the flow while the outlet is still about 80 °C below the setpoint and arrives without crossing the threshold, while the loop waits on its floor for the setpoint and overshoots to 400.0 °C, the fluid limit. The second is the cloud windows of scenarios 5, 9 and the challenge, where the outlet runs close to the threshold and the feedforward with its trim overshoots it for 10.1 to 11.4 min on the broken day at a fixed setpoint, on scenario 9 and on the suggested challenge, while the predictive controller holds the ceiling of Equation (6.25). Inside the clouds its error is 99.7 against 137.7 on the broken day, 919.2 against 1173.1 for the retuned loop on scenario 9, 859.7 against 1083.1 on the suggested challenge and 98.1 against 375.9 °C min for the retuned loop on the challenge as shipped.

🕮 **Pause.** Try Exercise 6.3 before reading on.

The predictive controller loses three cases on outlet error to the loop with feedforward, scenarios 11 and 3 as above and the clear day on Variable T (2805.9 against 2701.7 °C min), where the start-up cap slows the approach to the low early setpoints. It loses net electricity on the broken day on Variable T, 32.40 MWh against 33.52 MWh for the shipped loop and 33.61 MWh for the retuned one, with the same heat into the fluid (129.95 against 129.88 MWh). Of the 1.12 MWh difference, 1.05 MWh is lower gross output and 0.06 MWh is extra pumping (0.08 MWh more in the field and 0.02 MWh less at the block). The gross loss is thermal duty through the block (132.92 against 135.91 MWh at a conversion efficiency of 0.287 against 0.288), and the duty went into the hot tank. From the morning approach on, heat collects in the tank, and a tank kept warm leaves the field within 3 °C of it, so the interlock holds the bypass out, while under the loop the tank stays cooler and trips it less often. The bypass carries 40.1 % of the block's flow under the predictive controller against 88.4 % under the loop, 52.7 MWh enters the hot tank against 10.1 MWh, 0.30 MWh more leaks through its wall, and 2.79 MWh of the 2.99 MWh is still in the tanks when the run ends. Wherever the setpoint moves the predictive controller also moves the flow 2.6 to 3.3 times as far as the shipped loop. On the clear day on Variable T the loop with its feedforward read one transport delay ahead from the same preview makes 118.55 MWh, less than the predictive controller's 118.65 MWh.

plant control system: the four loops every 7.2 s, the supervisor every 3 min. Field pump lag 3 s, block pump and bypass 5 s, mirror drive 40 s. Thermocouple 2.2 s, power transducer 1.8 s.

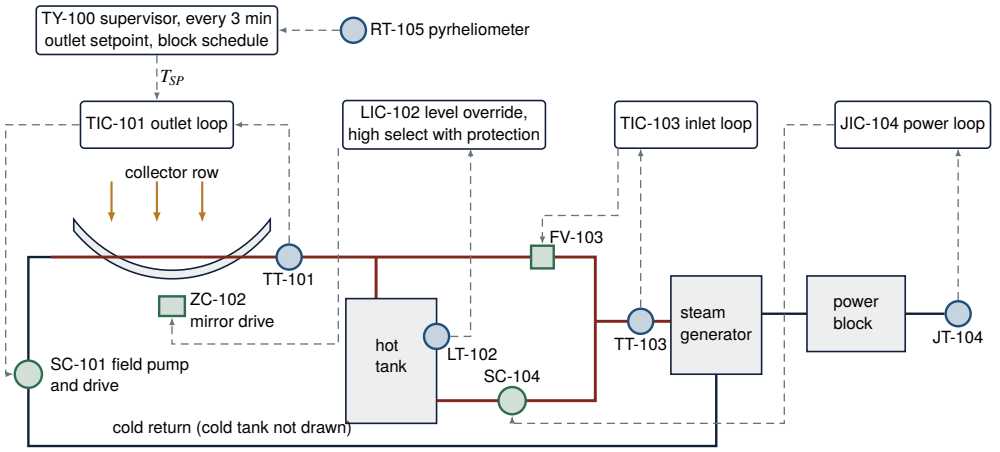


Figure 6.5 The four loops and the supervisor drawn with their hardware. The outlet thermocouple TT-101 feeds the outlet loop TIC-101, which sets the speed of the field pump SC-101, and the supervisor TY-100 writes that loop's setpoint from the pyrheliometer RT-105. The level LT-102 drives the mirror drive ZC-102 through the override LIC-102, high-selected with the outlet protection. TT-103 holds the steam generator inlet through the bypass valve FV-103, and the power transducer JT-104 sets the block pump SC-104. Notice that the supervisor moves no actuator: it writes one number into an ordinary loop.

The challenge's twentyfold difference, 72.78 against 3.21 MWh, belongs to the challenge's weak tuning, not to the predictive controller. With the weak gains the protection pins the outlet at 397 to 399 °C all day, the loop sees an error of only 4 to 6 °C and raises the flow by about 2 kg/s an hour, and 188.22 MWh is defocused. The shipped loop with only its outlet gains retuned makes 68.64 MWh, and the SIMC retune makes 71.49 MWh.

6.8 Technology in Practice

The loops of Section 6.5 read a temperature, a level and a power and move a pump, a valve and the mirrors, and the supervisor of Section 6.6 reads the sun. Figure 6.5 draws them with that hardware in place of the schematic's symbols.

What Measures It

The collector outlet temperature is the variable the whole chapter is about. A modern loop takes oil in at about 292 °C and returns it at about 393 °C, and some collectors in the loop are defocused when the outlet would otherwise pass about 395 °C [5]. The outlet is read by a temperature element in a thermowell. A pyrheliometer measures the direct normal irradiance, and a plant should have at least two or three instruments measuring it, checked often for alignment and cleanliness [5]. The chapter's thermocouple is a 2.2 s lag on the last fluid node, and its pyrheliometer reads the true irradiance with no error.

> **TECHNOLOGY** **Sensor: A Thermocouple at the Loop Outlet**
>
> A thermocouple is a junction of two dissimilar metal wires that produces a voltage approximately proportional to its temperature [19]. The common type K is listed for 32 to 2300 degrees Fahrenheit. In a pipe the element sits in a thermowell, which protects it but adds a first-order lag, and a lag severe enough can compromise the stability of feedback control [19]. The most common failure is to fail open, or burn out. An open input

can pick up electrical noise that reads as a wildly varying temperature, so the transmitter is set to read high or low when the element burns out [19], and the trend jumps to that end of its scale. The chapter's outlet thermocouple is a 2.2 s first-order lag with a noise setting of 0 to 4 °C.

What Moves It

The pump speed sets the field flow, which is also hard to measure on a hot oil loop. Plants have estimated it from the drive speed, the motor current or the pressure drop when their flow meters proved unreliable, at an accuracy typically not nearly as good as a working meter [5]. The mirrors are the second final element. The collector drive turns a trough partly or fully off the sun, and partial defocus of every collector in a loop by the same amount has worked well for outlet temperature control at a number of plants [5]. The chapter's defocus drive follows its command through a 40 s lag, and its field flow is known exactly.

TECHNOLOGY **Final Element: The Field Pump and Its Drive**

The field flow is set by the speed of the oil pumps. A variable-frequency drive sets the speed of an alternating-current motor by varying the frequency of the power that feeds it [19]. A study of operating plants recommends this capability on every oil pump, after one plant without it damaged its flow-control valves and shook its piping at each pump start [5]. Pump makers set limits on minimum flow, maximum flow and minimum speed, very low flow brings vibration, bearing wear and flow oscillations, and the failure reported most often is the pump seal [5]. The chapter's pump follows its command through a 3 s lag between 8 and 520 kg/s.

Where the Controller Runs

A trough plant is run from a distributed control system [5]. In that architecture a supervisory controller sends setpoints to the field controllers and requests data from them [20], and supervisory and lower-level loops run with cycle times from milliseconds to minutes [20]. The chapter's four loops execute every 7.2 s, and its supervisor every three minutes. The supervisor holds the chapter's lesson. It writes the outlet setpoint into an ordinary outlet loop, which does the same work for a right setpoint as for a wrong one. Tuning such a system takes time. Automation of oil flow control deserves higher priority, and tuning time should extend beyond commissioning to capture the plant's starts, stops and transients [5].

IN PRACTICE **Keep the Flow Up Under Cloud**

Operators report that loops at low flow under concentrated sun are where receivers break. Partial defocus used to hold the outlet temperature may put concentrated light on a receiver's glass-to-metal seals, and on a partly cloudy day a field whose flow is cut to drive up the outlet temperature may run some loops at low flow in full sun [5]. The recommended practice is a higher minimum flow, well into the turbulent region, also on partly cloudy days. One plant that holds about 50 % of design flow lost only 0.3 % of its receivers in its first ten years [5]. The chapter's outlet loop rests on its 8 kg/s floor in the cloud of scenario 3.

What Is Modernizing It

Three developments touch this chapter. Flow estimated from pump speed, current and pressure drop is a soft sensor that some plants already use in place of a failing meter, and two plants with newer automatic flow control schemes were pleased with them [5]. Performance models are asked to move from hourly to sub-hourly irradiance data, because hourly data do not capture the short cloud transients that a field feels [5], the transients the supervisor of this chapter answers. And as Ethernet, Internet Protocol and wireless devices replace older proprietary networks, the likelihood of cybersecurity vulnerabilities and incidents rises [20], which matters for a supervisor that writes setpoints into the field.

6.9 Limitations

The optical and receiver models are analytic surrogates rather than fitted correlations. The shape of the loss curve in Equation (6.6) drives every conclusion in this chapter, and a constant emissivity understates the loss at the top of the temperature range relative to the NREL receiver model [13]. The direction of that error is favorable to the fixed setpoint, so the reported advantage of variable-temperature operation is conservative, but no claim is made that the percentages would survive recalibration against a measured collector.

> **OPEN CHALLENGE** **Hydrogen in the Receiver Vacuum**
>
> A receiver tube sits inside an evacuated glass envelope. The oil slowly breaks down and gives off hydrogen, which permeates the steel into the vacuum. At a partial pressure of about 1 mbar the receiver's heat loss rises by a factor of 3 to 4, and the damage starts at the hottest receivers and works back along the loop over years [5]. Scrubbing hydrogen from the oil, injecting argon into the receivers and lower maximum operating temperatures are being considered, and the hydrogen in the oil is very difficult to measure [5]. The chapter holds one emissivity at all temperatures. Progress would be a measured receiver loss that the supervisor could weigh against the gain of running hot.

One collector row stands for the entire field. A real field has hundreds of parallel loops, a flow distribution across them, and a hottest loop that sets the defocus decision for all of them. Nothing in this model can pose the loop balancing problem, which is a substantial part of real trough plant operation.

The power block is a steady-state map with a nine-second inlet lag. Start and stop transients, drum level, superheater metal and turbine thermal stress are all absent, so the model understates the cost of the cycling that the schedule in Section 6.6 produces. The application counts a power block start in its exported data only after at least fifteen minutes off, and does not charge for starts in the margin. The eighth scenario records 3 starts in eighteen hours, while the block crosses its minimum stable load 22 times as the power loop works near it.

The schedule optimizes a threshold rather than a trajectory. It answers whether to run the block at rating in a given block, not at what load, so a plant that would be better served by running at part load for longer is not represented. The scheduler enforces tank capacity but makes no terminal-inventory guarantee. The application has no reserve setting, because holding energy back in trial draws does not impose a floor on actual discharge. A guaranteed reserve would require an operating constraint and a different scheduling objective.

The weather and the price are synthetic. The improvement numbers should be read as the shape of a result rather than as its magnitude, and any statement about what dynamic optimization is worth on a particular site requires that site's irradiance record and price curve.

The predictive outlet controller predicts with five cells at the declared design values, where the plant has ten nodes and whatever weather and parameters the run carries. It cannot see the true mirror cleanliness, the wind, the ambient temperature or a parameter the learner changes, it disperses the transit more than ten nodes do, and the output bias of Equation (6.23) carries all of that as one slowly varying offset. Its start-up cap is a rule rather than a model of the pump, borrowed from the ramp at which the field is taken out of service, and on the clear day at a variable setpoint it costs enough outlet error to lose that case to a feedback loop. Finally, the shipped loop is not the honest competitor. A proportional-integral loop trimming a static feedforward from the same pyrheliometer matches or beats the predictive controller through passing clouds, so a plant weighing the two should compare against that loop rather than against plain feedback.

OPEN CHALLENGE A Plant That Runs Itself

Trough plants are not yet automated end to end. A 2020 study found that their operating modes had not been refined to the point where the control system can run a plant automatically on a routine basis, so operators often write procedures as they go when equipment misbehaves [5]. Field flow is hard to automate because flow meters on the hot oil have been unreliable, and estimates from pump speed, motor current and pressure drop are less accurate than a working meter [5]. The chapter's supervisor trusts every measurement it reads. Progress would be a supervisory layer that keeps choosing a sound setpoint when a flow or temperature signal is lost.

6.10 Conclusion

A parabolic-trough field is a plant whose best operating point moves with a disturbance that anybody can see. Receiver loss depends on the fourth power of a temperature the operator chooses and not at all on the sunshine, while the heat available to collect changes by a factor of five across a broken day. The consequence is that the outlet temperature setpoint is a decision that has to be made repeatedly, and a control system with no mechanism for making it will defend a wrong number perfectly.

The measured comparison has two halves and both matter. Holding 393 °C on a clear day costs 0.3 % of net generation against an optimizer, because the extra harvest is spent on pumping. Holding it on a broken day costs 32 %, because the receiver goes on losing 3.7 MW while the field absorbs a fifth of what it does at clear noon. Above both sits a scheduling decision worth more than either. On a hazy day the schedule generates about what constant power does (80.6 against 78.1 MWh) and books 2.2 times the margin, by moving 39 MWh of generation past four in the afternoon. The regulatory loops underneath are ordinary, but their tuning is not a detail, because in the graded day the gains recover most of the energy while the schedule earns the money. An optional predictive outlet controller on a five-cell receiver model removes all but 1.4 min above the protection threshold on the nine measured cases, while a loop given a feedforward from the same pyrheliometer does as well through passing clouds, so the gain belongs to the measured sun and the model more than to the receding horizon.

KEY POINTS

- The best temperature for the oil leaving a solar field moves with the sky, so a controller can hold a fixed target perfectly and still be wrong.
- Holding 393 °C costs 0.3 % of net generation on a clear day and 32 % on a broken day.
- On a hazy day the schedule books 2.2 times the margin of constant power for about the same energy.
- A loop with feedforward from the same pyrheliometer does as well as the predictive controller through passing clouds.

Where else this applies. A supervisory setpoint that follows the weather also sets the chilled water of Volume 9, Chapter 8 and the aeration of Volume 9, Chapter 10.

References

[1] K. M. Powell, J. D. Hedengren, and T. F. Edgar, "Dynamic optimization of a solar thermal energy storage system over a 24 hour period using weather forecasts," in *2013 American Control Conference*, 2013, pp. 2946–2951.

[2] ——, "Dynamic optimization of a hybrid solar thermal and fossil fuel system," *Solar Energy*, vol. 108, pp. 210–218, 2014.

[3] H. Price, E. Lüpfert, D. Kearney, E. Zarza, G. Cohen, R. Gee, and R. Mahoney, "Advances in parabolic trough solar power technology," *Journal of Solar Energy Engineering*, vol. 124, no. 2, pp. 109–125, 2002.

[4] A. Fernández-García, E. Zarza, L. Valenzuela, and M. Pérez, "Parabolic-trough solar collectors and their applications," *Renewable and Sustainable Energy Reviews*, vol. 14, no. 7, pp. 1695–1721, 2010.

[5] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

[6] E. F. Camacho, F. R. Rubio, M. Berenguel, and L. Valenzuela, "A survey on control schemes for distributed solar collector fields. Part I: Modeling and basic control approaches," *Solar Energy*, vol. 81, no. 10, pp. 1240–1251, 2007.

[7] ——, "A survey on control schemes for distributed solar collector fields. Part II: Advanced control approaches," *Solar Energy*, vol. 81, no. 10, pp. 1252–1272, 2007.

[8] E. F. Camacho, M. Berenguel, and F. R. Rubio, *Advanced Control of Solar Plants*, ser. Advances in Industrial Control. London: Springer London, 1997.

[9] D. Kearney, U. Herrmann, P. Nava, B. Kelly, R. Mahoney, J. Pacheco, R. Cable, N. Potrovitza, D. Blake, and H. Price, "Assessment of a molten salt heat transfer fluid in a parabolic trough solar field," *Journal of Solar Energy Engineering*, vol. 125, no. 2, pp. 170–176, 2003.

[10] D. Cocco and F. Serra, "Performance comparison of two-tank direct and thermocline thermal energy storage systems for 1 MWe class concentrating solar power plants," *Energy*, vol. 81, pp. 526–536, 2015.

[11] D. Kearney, B. Kelly, U. Herrmann, R. Cable, J. Pacheco, R. Mahoney, H. Price, D. Blake, P. Nava, and N. Potrovitza, "Engineering aspects of a molten salt heat transfer fluid in a trough solar field," *Energy*, vol. 29, no. 5-6, pp. 861–870, 2004.

[12] K. M. Powell and T. F. Edgar, "Modeling and control of a solar thermal power plant with thermal energy storage," *Chemical Engineering Science*, vol. 71, pp. 138–145, 2012.

[13] R. Forristall, "Heat transfer analysis and modeling of a parabolic trough solar receiver implemented in engineering equation solver," National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep., 2003.

[14] F. Burkholder and C. Kutscher, "Heat loss testing of Schott's 2008 PTR70 parabolic trough receiver," National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep., 2009.

[15] E. F. Camacho and C. Bordons, *Model Predictive Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2007.

[16] A. J. Gallego and E. F. Camacho, "Adaptative state-space model predictive control of a parabolic-trough field," *Control Engineering Practice*, vol. 20, no. 9, pp. 904–911, 2012.

[17] E. F. Camacho and A. J. Gallego, "Optimal operation in solar trough plants: A case study," *Solar Energy*, vol. 95, pp. 106–117, 2013.

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

[19] 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.

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

6.11 Exercises

6.1 ○ Concept check · 2 min

At solar noon with a 300 °C cold tank, Table 6.3 puts the outlet temperature that maximizes net electrical output at 393 °C under a clear sky and at 325 °C at a quarter of it. Which explanation is right?

(a) Pumping power grows as the cube of the flow, so under cloud the optimizer runs cold to save pumping

(b) Receiver loss depends on the metal temperature and not on the sun, so under cloud it is a larger share of what is absorbed

(c) The fluid stability limit falls with the irradiance, so under cloud the optimizer must stay further below it

(d) The protection defocuses more of the field under cloud, so the optimizer moves away from its threshold

6.2 ○ Concept check · 1 min

For most of a clear day in Auto, the level loop's defocus output sits at zero. Why?

(a) The loop is interlocked out until the field is hotter than the hot tank

(b) Its integral time of 0.30 h is too long for it to act within a day

(c) The level is below its setpoint, so the direct-acting loop asks for negative defocus, which the limit refuses

(d) The high selector passes the protection's demand instead, which is larger all day

6.3 Concept check · 2 min

On scenario 11 the shipped outlet loop has an outlet error of 882.8 °C min. The predictive controller has 46.8 with its preview and 47.1 without it, and a PI loop trimming a static feedforward from the pyrheliometer, started on its own flow, has 45.7. What carries most of the predictive controller's lead over the shipped loop?

(a) The preview of the irradiance over the next twenty minutes, which shows each cloud edge in advance

(b) The receding horizon, because a plan over twenty minutes looks past the transport delay

(c) The soft ceiling at $T_{max} - 4\,^{\circ}\mathrm{C}$, which keeps the outlet out of the protection

(d) The measured irradiance and a model of the field, which any structure given both can use

6.4 Theory and calculation · 6 min

Use Equations (6.6) and (6.7) summed over the whole field, with $A_{rec} = 2815\,\mathrm{m}^2$, $\varepsilon = 0.12$, $h_{conv} = 0.35\,\mathrm{W/(m^2\,K)}$, $h_{wind} = 0.25\,\mathrm{W/(m^2\,K)}$, $\sigma_{SB} = 5.67 \times 10^{-8}\,\mathrm{W/(m^2\,K^4)}$, an ambient of 25 °C and a wind of 3 m/s. Take every metal node at the same temperature.

(a) Find h_e and the field loss with the metal at 393 °C, split into radiation and convection. *Checkpoint:* $h_e = 0.833\,\mathrm{W/(m^2\,K)}$, 3.62 MW plus 0.86 MW, 4.48 MW in all.

(b) Repeat at 300 °C, and give the ratio of the two radiative parts. *Checkpoint:* 2.56 MW, ratio 1.89.

(c) Express the loss at 393 °C as a share of the 50.8 MW absorbed at clear noon and of the 12.7 MW absorbed at a quarter of the sun. *Checkpoint:* 8.8 % and 35.3 %.

6.5 Theory and calculation · 6 min

The field holds $M_f = 60\,000\,\mathrm{kg}$ of oil with $c_p = 1.60 + 0.0024\,T$ in kJ/(kg K) (T in °C), and the receiver tube is $M_m = 64\,000\,\mathrm{kg}$ of steel with $c_{p,m} = 490\,\mathrm{J/(kg\,K)}$.

(a) Find the transit time of the oil through the field at the design flow of 222 kg/s and at a sixth of it. *Checkpoint:* 4.5 min and 27 min.

(b) Find the heat capacities of the tube metal and of the oil in the field, with c_p at the design mean of 346 °C. *Checkpoint:* 31.4 MJ/K and 146 MJ/K.

(c) The steady outlet rise over the inlet is $Q_{net}/(\dot{m}_f c_p)$. Find $\partial T_{out}/\partial \dot{m}_f$ at 222 kg/s with $Q_{net} = 46\,\mathrm{MW}$, and the factor by which it changes when both the heat and the flow fall to a quarter. *Checkpoint:* –0.384 °C per kg/s, four times larger.

6.6 Theory and calculation · 7 min

Use Equations (6.15) to (6.17) with $\varepsilon_{SG} = 0.85$, $T_{sat} = 285\,^{\circ}\mathrm{C}$, $Q_{PB} = 27.8\,\mathrm{MW}$, $\eta_{PB} = 0.36$, $k_T = 0.0013/^{\circ}\mathrm{C}$, a minimum stable load $x_{min} = 0.25$ and $c_p = 1.60 + 0.0024\,T$ in kJ/(kg K) evaluated at the mean of T_{sg} and T_{sat}. The block flow is 119 kg/s.

(a) At $T_{sg} = 393\,^{\circ}\mathrm{C}$, find the thermal duty, the efficiency and the gross output. *Checkpoint:* 26.4 $\mathrm{MW_{th}}$, 35.4 %, 9.33 MW.

(b) Repeat at $T_{sg} = 340\,^{\circ}\mathrm{C}$. *Checkpoint:* 13.1 $\mathrm{MW_{th}}$, 26.4 %, 3.45 MW.

(c) Show that the block is off at $T_{sg} = 310\,^{\circ}\mathrm{C}$. *Checkpoint:* load fraction 0.21, below 0.25.

6.7 App challenge · 8 min

Load scenario 3, *What a cloud does to a plant defending a temperature*, and change no setting. Run it at speed 0.5 with the default seed and pause between 15:00 and 15:10, four hours after the 11:00 start. Export the CSV file.

(a) Read the Defocused tile, and the highest value of the column `level_pct` in the file. *Checkpoint:* 13.8 MWh, within 0.1 MWh, and the level stays below 58 %.

(b) Average the columns `Qloss_MW` and `Qabs_MW` over the rows with `clock_h` from 12.00 to 12.60, inside the cloud, and give the loss as a share of the heat absorbed. *Checkpoint:* 3.75 MW of 11.9 MW, 31.5 % within one point.

(c) Say what the defocused energy was spent on, and which trend shows the controller succeeding while the plant fails. *Checkpoint:* The high-temperature protection after the clearing. The outlet strip against the "Where the sunlight goes" strip.

6.8 App challenge · 8 min

Load scenario 5, *The same comparison on an intermittent day*, which opens on Fixed T. Run it at speed 2 with the default seed and pause between 20:00 and 21:00. Then press Reset, set Supervisory strategy to Variable T before Run, and run again to the same window.

(a) Read Generated and Collection for the Fixed T run. *Checkpoint:* 25.4 MWh and 70.5 %, within one unit of the last digit.

(b) Read the same two tiles for the Variable T run. *Checkpoint:* 33.5 MWh and 78.1 %.

(c) Give the gain in net generation as a percentage of the Fixed T value, and explain why it is larger than the gain in collection. *Checkpoint:* About 32 %.

6.9 App challenge · 8 min

Load scenario 11, *PI against MPC on the outlet loop, through passing clouds*, which opens with Outlet loop on MPC. Use speed 0.5 and the default seed, and pause each run between 13:00 and 13:15. Run it as loaded. Then press Reset, set Outlet loop to PI and run again. Then set Outlet loop to MPC, set Preview to Off under Advanced, press Reset and run a third time.

(a) For each run read the IAE tile of the collector outlet temperature loop, in °C h, and convert it to °C min. *Checkpoint:* 46.8, 883 and 47.1 °C min, within 2 °C min for the PI and 0.2 °C min for the others.

(b) Read the Defocused tile for each run. *Checkpoint:* 0.0, 17.9 and 0.0 MWh.

(c) Give the reduction in outlet error from the PI to the MPC and the change from switching the preview off, both as percentages. *Checkpoint:* 95 % lower, about 0.6 % higher.

6.10 Application area · 3 min

The outlet setpoint of a trough field sits at 393 °C and the protection starts at $T_{max} - 4\,^{\circ}\mathrm{C}$. Say what sets T_{max}, what crossing it costs on a real plant and in this simulator, and why the industry kept the fluid that imposes it.

6.11 Application area · 4 min

A plant engineer proposes to use this simulator's minutes above 396 °C and its defocused energy to size the defocus margin of a real field with hundreds of parallel collector loops. What does the model hide that matters for that decision, in which direction does it bias the two numbers, and what plant data would the engineer need?

6.12 Review the agent · 5 min

An AI assistant reviews the graded challenge, scenario 12, and writes: “The chapter says the schedule earns the money. Switch Supervisory strategy to Forecast, leave everything else as shipped, and all five targets pass, since the weak gains only cost a little outlet error.” Find the flaw, and say how you would check it in the app.