INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 9 Model Predictive Control and Economic Optimization

CHAPTER 8

A Data Center Cooling Simulator Whose Optimum Moves With the Wet Bulb

Pages 172–201

Edited by John D. Hedengren

How to cite this chapter. "A Data Center Cooling Simulator Whose Optimum Moves With the Wet Bulb," in *Model Predictive Control and Economic Optimization*, John D. Hedengren, Ed., vol. 9 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 172–201.

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

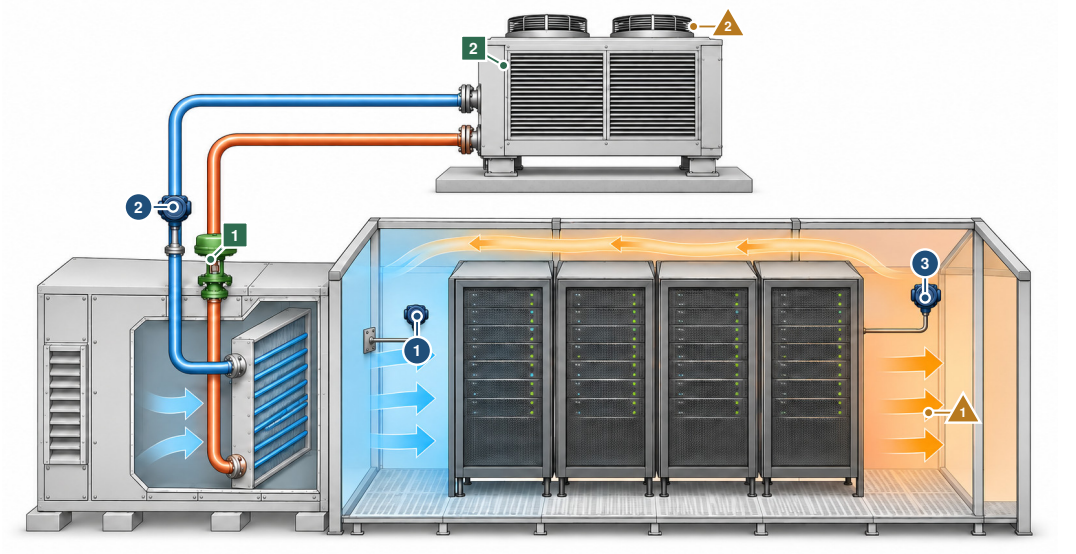
The app for this chapter
https://apmonitor.com/apps/data-center-cooling/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v09-mpc-optimization/#data-center-cooling

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

CHAPTER 8

A Data Center Cooling Simulator Whose Optimum Moves With the Wet Bulb



The controller moves 1 chilled water valve at the coil 2 the plant's chilled water setpoint and tower spray

The controller reads 1 cold-aisle (server inlet) temperature 2 chilled water supply temperature 3 hot-aisle temperature

What disturbs it 1 heat from the servers 2 outdoor dry and wet bulb

One valve holds the cold aisle, while the layer above it chooses how cold the plant makes the water and how wet it runs the tower, the choice that sets the bill. Illustration, not to scale.

IN PLAIN TERMS A Cool Server Room at the Lowest Cost

The main idea. A data center is a room of computers that turn power into heat, and a cooling plant carries the heat outdoors. A simple controller holds the air the computers take in at one temperature, and a planner above it picks how the plant runs.

The challenge. Colder water costs more power, and a wetter cooling tower saves power but uses up water. The cheapest mix that keeps every limit moves as the outdoor air turns damp or dry.

Why it matters. In this chapter's hall the planner cuts the cooling bill by more than a tenth while the room stays just as cool.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in the Introduction, the facts in Sections 8.2 and 8.6, the try-it segments in Section 8.6, and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the open challenge

The control question A loop that holds the server inlet at its setpoint says nothing about what the hall costs, because the chilled water setpoint and the tower's evaporative fraction beneath it decide the electricity and the water. Where does the cheapest operating point that meets every limit sit, what holds it there as the wet bulb moves, and does a predictive layer on the same cost find a better one?

Learning objectives

- Compute the steady water draw, cost, chiller COP and recirculation threshold of the hall from its model equations.
- Explain why the cost is not unimodal in the evaporative fraction and why the supervisory layer therefore searches a grid.
- Identify the binding constraint at an operating point and predict how weather, coil fouling, a chiller trip and a water price move it.
- Distinguish what tuning, feedforward and a cascade improve from what only the supervisory layer improves, using cold-aisle error, energy, water and cost.
- Compare the grid search with an economic predictive layer on the same cost, before and after a chiller trip.

Prerequisites Sections 1.6.2 (velocity form), 1.4.6 (soft constraints), 1.4.7 (steady-state optimization), 1.4.8 (economic MPC) and 1.6.4 (how an exercise names a run), and Volume 3, Chapter 2 for feedforward and cascade structure.

The app https://apmonitor.com/apps/data-center-cooling/

See also Volume 17, Chapter 6 for an optimum outlet temperature that moves with the irradiance, and Volume 17, Chapter 13 for a supervisory layer weighed against simpler heuristics. Volume 3, Chapter 6 for a feedforward from a measurement earlier in the causal chain.

8.1 Introduction

The cooling plant of a data hall is a control problem whose difficulty sits one level above its regulatory loops. A chilled water valve lowers the supply air temperature, the supply air lowers the cold-aisle temperature, and a single loop settles a setpoint step to within 0.1 °C in about three minutes, so a proportional-integral controller holds the server inlet temperature at 24.00 °C against every disturbance the coil and chiller have capacity for. The regulatory loop does not decide the chilled water supply temperature, the evaporative fraction of the cooling tower, or whether the water-side economizer is in service. Those decisions set the electricity and water consumption and they do not appear on the temperature trend. In the simulator of this chapter the shipped configuration holds 24.00 °C exactly while spending $98.00 per hour against a lowest-cost feasible point of $86.18, a gap of $11.81 per hour, or roughly $103,000 per year on one 3 MW hall.

Both resources are publicly accounted. A recalibration of global data center energy use placed 2018 consumption near 1 % of world electricity and found that a large rise in computing output had been absorbed almost entirely by efficiency gains rather than by more power [1]. A spatially resolved footprint of United States data centers put them at about 1.8 % of national electricity use, attributed about 0.5 % of national greenhouse gas emissions to them, and found that one fifth of the direct water footprint of their servers is drawn from moderately to highly water-stressed watersheds [2], with the water spent indirectly in generating that electricity larger again than the water evaporated in the towers [3]. Neither quantity is decided by the hardware once the plant is built, because both are decided by setpoints.

A trend holding its setpoint is evidence about the regulatory loop and not about the operating cost of the plant, and the simulator is built to separate the two. Tuning the regulatory loop over a whole day of weather and then adding a load feedforward improves the peak cold-aisle error by a factor of four and changes the day's operating cost by 18 cents. The distinction does not appear in a block diagram, because neither the objective function nor the binding constraint is part of the loop.

> ⧖ HISTORY **2007: A Report to Congress Counts the Servers**
>
> A report to the United States Congress in 2007 estimated that the nation's data centers had used about 61 billion kilowatt-hours in 2006, 1.5 percent of national electricity, for about $4.5 billion. That was more than double the estimate for 2000. A national laboratory's update of 2016 found the growth had slowed: about 70 billion kilowatt-hours in 2014, about 1.8 percent of national use, up about 4 percent from 2010 to 2014 against nearly 90 percent from 2000 to 2005. It traced the change to slower growth in server shipments, to virtualization, and to very large data centers run at high utilization, which need fewer servers for the same services [4]. The cooling plant of every hall sits inside those totals.

The class of problem is supervisory optimization over a fast regulatory layer, in which the objective encodes a chosen price ratio and the binding constraint changes with an uncontrolled input, the steady-state optimization of Section 1.4.7. Which variables to hold constant and which to reoptimize is the plantwide question of self-optimizing control [5], the same structure appears wherever a slow layer feeds setpoints to fast loops against a live objective [6, 7], and the building services version of it is a documented retrofit practice [8, 9]. The same structure appears in any plant where a boiler steam pressure, a compressor discharge pressure or a refrigeration suction pressure was fixed at commissioning and never reoptimized, and Volume 17, Chapter 6 meets it in a solar field whose best outlet temperature moves with the irradiance.

The physical content of the model is stated with its fundamentals. Each air and water node is an unsteady energy balance over a well-mixed volume, the two heat exchangers are effectiveness models, the chiller is a Carnot term in series with a loss ceiling, and the layers are separated in time, with the regulatory loop settling in three to eight minutes and the supervisory layer executing every three minutes against a steady-state model and moving its setpoints at a rate limit [10–12]. The scenarios teach process gain that changes with operating point, feedforward against a measured load, cascade structure, anti-windup at a saturated valve and constrained steady-state optimization above regulatory control. The behavior matters industrially because chilled water and condenser water setpoints fixed at commissioning are a documented source of avoidable energy use in central cooling plants [9, 13].

The data center literature divides in a way that leaves this shape uncovered. One half is thermal, concerned with rack densities, airflow fields and the arrangement of the room [14–16], and the other half optimizes the facility as a whole with the plant treated as a set of efficiency curves [17]. Neither provides an interactive plant whose feasible region moves while the search runs. This chapter sets out a browser-hosted data hall and its cooling plant in enough detail for a reader to rebuild both, and then establishes where its constrained optimum sits, what holds it there, and which of its two decision variables the weather moves. Two measured results carry the argument. The cost surface is not unimodal in the tower evaporative fraction, so an exhaustive search is a modeling decision rather than a convenience. Over a full day of weather the chilled water setpoint takes three of the twenty-two values the search grid offers while the evaporative fraction sweeps all fifteen, from the first to the last. The plant has ten states and one steady-state evaluator shared by the cost map, the optimizer and the initializer. Above it sit a constrained grid search transcribed from the executing code and an optional predictive layer on the same cost, measured against it in Section 8.6.8. The simulator reports electricity, water, inlet margin and cost separately rather than combined into one score. The chiller correlation, the tower approach model and the search itself are established methods.

Section 8.2 reviews how a raised floor delivering 13 °C air became a hall held at 24 °C with a supervisory layer above it. Section 8.3 sets out the plant model, its parameters and its assumptions. Section 8.4 chooses the two supervisory decisions by sweeping the steady-state model, and Section 8.5 describes the regulatory loops, the grid search and the predictive layer. Section 8.6 measures the scenarios, from the hall by hand to the graded day and the predictive layer against the grid search. Section 8.7 describes the sensors, valves and controllers of real halls, Section 8.8 states the limitations, Section 8.9 concludes, and the exercises in Section 8.10 run the same cases in the app.

8.2 Background

The first computer rooms were cooled to a narrow band of temperature and humidity specified by the mainframe vendor, held by computer room air conditioning units around the perimeter of a raised floor. Air was pushed into the plenum, rose through perforated tiles, crossed the equipment above them and returned to the units through the room. The controlled variable was the return air temperature at the unit, an average over the whole room, so the only safe way to protect the worst rack was to make the supply air cold enough for it. Halls delivering air at 13 °C were common, and the excess cooling was paid for twice, in chiller lift and in mixing losses.

Rack power densities ended that arrangement. When a cabinet went from one or two kilowatts to ten or more, the room stopped behaving like a well-mixed volume and thermal management became a subject in its own right [15]. The response that mattered for control was the proposal to provision cooling where the heat is generated and to treat the plant as a dynamically operated system rather than a statically sized one [14]. Cold-aisle and hot-aisle rack arrangement came first and physical containment of one aisle came after it, which is what makes the recirculation term of Section 8.3 small rather than dominant.

Permission to run a hall warmer came from industry guidelines rather than from the control literature. The ASHRAE thermal guidelines for data processing environments, first issued in 2004 and revised repeatedly since, separated a recommended envelope from wider allowable classes and moved the recommended upper limit up as manufacturers agreed to it [16]. The two inlet limits in this application, 27 °C recommended and 32 °C allowable, come from that lineage, and they enter the model as constraints rather than as setpoints because that is what they are. Raising the inlet is not free. An analysis of hall temperature against facility efficiency showed the gain is not monotone, because the chiller savings from a warmer hall are eventually overtaken by the server fan power that the warmer hall causes [18], which is why the server fan ramp of Equation (8.3) is in the model at all.

A wider envelope makes free cooling reachable. Air side economizers, indirect evaporative units and water side economizers of the kind modeled here all exploit the same fact, which is that a hall permitted to run warm can often be cooled by the weather instead of by a compressor [19]. Free cooling behaves as a threshold in the wet bulb rather than as a gradient, so site selection became largely an argument about climate and the cooling tower became the component that decides how close to the wet bulb a plant can get.

Two ratios were adopted to make the argument quantitative. Power usage effectiveness, the ratio of total facility power to IT power, was promoted by the Green Grid consortium and became the standard reported metric. Water use effectiveness, in liters of make-up per kilowatt-hour of IT energy, followed as its companion once it became clear that a plant could lower one ratio by raising the other, which is the trade this application is built around. Both are reported by the application, and at the design point of Section 8.3 the hall runs at a PUE of 1.291 and a water use effectiveness of 1.77 L/kWh.

> ✦ DID YOU KNOW? **About a Third for Cooling and Power**
>
> Facility infrastructure, in which cooling dominates, took 36 percent of United States data center electricity in 2018 and 31 percent in 2024, as the national average PUE fell from 1.55 to 1.45 [20]. This chapter's hall runs at 1.291 at its design point (Table 8.5).

The setpoint question remains. Chilled water is still commonly made at a fixed temperature chosen at commissioning, condenser water is still commonly held at a fixed approach, and resetting either against load or weather is sold as an energy services retrofit rather than treated as ordinary operation [9, 13]. Present interest in direct-to-chip and immersion liquid cooling changes the hall side of the problem substantially and the plant side very little, because the heat still has to leave the building through a tower or a dry cooler, and how cold to make the water and how much of it to evaporate is the same question in every case.

8.3 Process Model

The plant is a 3 MW data hall with hot-aisle containment. Server fans pull air out of the cold aisle, through the racks and into the hot aisle, and computer room air handlers (CRAH) take that return air across a chilled water coil and deliver it back to the cold aisle. A water-cooled centrifugal chiller makes the chilled water, a hybrid cooling tower with wet and dry cells rejects the heat outdoors, and a plate heat exchanger across the chiller provides free cooling of the chilled water return whenever the condenser water is cold enough. The arrangement and every symbol in the equations below are shown in Figure 8.1.

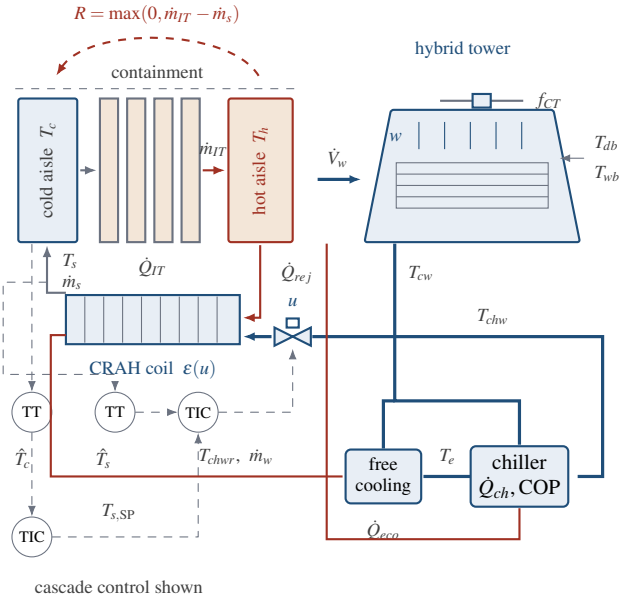


Figure 8.1 The hall, the chilled water loop and the hybrid tower. Server fans draw $\dot{m}_{IT}$ from the cold aisle at T_c through the racks into the hot aisle at T_h, and the CRAH delivers $\dot{m}_s$ back at T_s across a coil of effectiveness $\varepsilon(u)$. Whenever the racks pull more air than the floor delivers, the shortfall R is made up over the containment and lifts the inlet above the supply air. Valve u is the only manipulated variable. The free-cooling exchanger takes $\dot{Q}_{eco}$ out of the return T_{chwr} against condenser water at T_{cw} and the chiller finishes what it can, while the tower rejects $\dot{Q}_{rej}$ to a sink set by the evaporative fraction w between the dry bulb T_{db} and the wet bulb T_{wb}, spending make-up water $\dot{V}_w$ to do it.

The controlled variable is the cold-aisle temperature, which is the server inlet temperature, and its setpoint is 24 °C. The manipulated variable is the chilled water valve at the coils, and it is the same manipulated variable in every control architecture the application offers. Table 8.1 classifies everything the loop interacts with. Ten states are integrated, and the time base is minutes throughout, with thermal capacities carried in kW min/°C so that $\mathrm{d}T/\mathrm{d}t = \dot{Q}/C$ needs no conversion factor.

8.3.1 Air-Side Energy Balances and Recirculation

Hot-aisle containment is never perfect, and the airflow through the racks is set by the server fans rather than by the CRAH. An unsteady energy balance over each aisle as a well-mixed node gives Equations (8.1).

$$
\begin{aligned}
C_c \frac{\mathrm{d}T_c}{\mathrm{d}t} &= \dot{m}_s c_p (T_s - T_c) + R\, c_p (T_h - T_c) \\
C_h \frac{\mathrm{d}T_h}{\mathrm{d}t} &= \dot{m}_{IT} c_p (T_c - T_h) + \dot{Q}_{hall} + B\, c_p (T_c - T_h)
\end{aligned}
\tag{8.1}
$$

Table 8.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at 24 °C Dry Bulb and 17 °C Wet Bulb.

Symbol	Role	Nominal	Range
T_c	CV, measured	24.00 °C	18–27 °C SP
u	MV	26.7 %	0–100 %
$T_{chw,\mathrm{SP}}$	supervisory decision	7 °C	5–20 °C
w	supervisory decision	85 %	0–100 %
f_{CT}, f_{CRAH}, d_{eco}	operator controls	80, 85, 100 %	see Sec. 8.4
$\dot{Q}_{IT}$	DV, measured	3000 kW	20–125 %
T_{db}	DV, measured	24 °C	−10–45 °C
T_{wb}	DV, measured	17 °C	−12–30 °C
ϕ fouling	DV, unmeasured	0	0–40 %
a_{ch} chiller	DV, measured	100 %	40–110 %
a_p pump	DV, measured	100 %	45–110 %
b_c inlet bias	DV, unmeasured	0 °C	±3 °C
T_c, T_h, T_s	air states	24.00, 38.6, 24.00	°C
T_{chw}, T_{cw}	water states	7.0, 22.7	°C
$\dot{Q}_{ch}$	chiller duty	2923	kW, $\le \dot{Q}_{cap}$
a_v, a_w, a_f, a_d	actuators	0.267, 0.85, 0.80, 1.00	0–1

where C_c and C_h are the two air capacities in kW min/°C, c_p is the specific heat of air, $R = \max(0, \dot{m}_{IT} - \dot{m}_s)$ is air recirculated over the containment when the racks demand more air than the CRAH delivers, $B = \max(0, \dot{m}_s - \dot{m}_{IT})$ is cold air bypassing straight to the return, and $\dot{Q}_{hall}$ is every watt dissipated in the hall. Both sides carry units of kW. Solving (8.1) at steady state collapses the pair to Equation (8.2).

$$T_c = T_s + \frac{R\,\dot{Q}_{hall}}{\dot{m}_s M c_p}, \qquad M = \max(\dot{m}_{IT}, \dot{m}_s) \tag{8.2}$$

While the CRAH moves at least as much air as the racks do, R is zero and the cold aisle is exactly the supply air. The moment the racks pull more, the shortfall comes out of the hot aisle and the inlet rises with the supply air unchanged. At the design point the CRAH moves 221 kg/s against 210 kg/s of server airflow, so R is zero with 11 kg/s to spare.

The server fans consume that margin, because they hold a base speed until the inlet passes a knee temperature and then ramp, with fan power following the cube of speed, as shown in Equation (8.3).

$$\begin{aligned} s &= \mathrm{clamp}\big(1 + k_f(T_c - T_{knee}),\ 1,\ s_{max}\big) \\ \dot{m}_{IT} &= s\,\dot{m}_{IT,0}, \qquad P_f = s^3 P_{f,0} \end{aligned} \tag{8.3}$$

where s is fan speed as a multiple of base, k_f is the ramp slope and s_{max} is the upper limit. The airflow term in (8.3) is stabilizing, while the power term raises the cooling load and the operating cost. Server fan power is 180 kW at the knee, 281 kW at 27 °C and 494 kW at 30 °C, and all of that power is dissipated as heat in the same hall.

8.3.2 Coil Effectiveness and the Chilled Water Loop

The coil is written as an air-side effectiveness that saturates in water flow, as shown in Equation (8.4), because the last third of the valve stroke adds little duty while pump power rises with the cube of flow.

$$\begin{aligned} \varepsilon(u) &= \varepsilon_{max}(1 - \phi)\,\frac{1 - \mathrm{e}^{-k_{coil}u^{0.8}}}{1 - \mathrm{e}^{-k_{coil}}} \\ \tau_{coil}\frac{\mathrm{d}T_s}{\mathrm{d}t} &= T_h - \varepsilon(T_h - T_{chw}) - T_s \end{aligned} \tag{8.4}$$

where u is the valve fraction, which sets both the effectiveness and the water flow, ϕ is the fouling disturbance, k_{coil} is a shape factor and the exponent 0.8 is the usual turbulent water-side dependence.

The saturation in (8.4) is the reason the process gain the regulatory loop sees is a function of where the supervisory layer has put the chilled water setpoint, which Section 8.6 measures.

Chilled water flow follows the valve above a minimum-flow floor,

$$\dot{m}_w = \max(\dot{m}_{w,max}\, a_p u,\ 0.06\,\dot{m}_{w,max}),$$

because without a bypass the return temperature grows without bound as the valve closes. The chilled water loop and the chiller's own capacity control are Equations (8.5).

$$\begin{aligned} C_{chw}\frac{\mathrm{d}T_{chw}}{\mathrm{d}t} &= \dot{m}_w c_w (T_e - T_{chw}) - \dot{Q}_{ch} \\ \tau_{ch}\frac{\mathrm{d}\dot{Q}_{ch}}{\mathrm{d}t} &= \min\left(\dot{Q}_{req},\ \dot{Q}_{cap}\right) - \dot{Q}_{ch} \end{aligned} \tag{8.5}$$

where T_e is the water temperature leaving the free-cooling exchanger, $\dot{Q}_{req} = \dot{m}_w c_w (T_e - T_{chw,\mathrm{SP}})$ is the duty that would land the supply exactly on setpoint, $\dot{Q}_{cap}$ is the capacity available now, and τ_{ch} is the vane or drive lag. The min in (8.5) is the capacity constraint and nothing else in the model enforces it. Unsaturated it returns $T_{chw} = T_{chw,\mathrm{SP}}$ exactly, and saturated it lets the supply drift up until the warmer evaporator hands the machine enough capacity back to hold.

The economizer removes the first part of the return duty before the chiller sees it, as shown in Equation (8.6).

$$\begin{aligned} T_e &= T_{chwr} - \varepsilon_{eco} a_d \max(0,\ T_{chwr} - (T_{cw} + \Delta T_{HX})) \\ \dot{Q}_{eco} &= \dot{m}_w c_w (T_{chwr} - T_e) \end{aligned} \tag{8.6}$$

with $T_{chwr} = T_{chw} + \dot{Q}_{coil}/(\dot{m}_w c_w)$. Free cooling is a threshold rather than a gradient, because the max in (8.6) is zero until the condenser water plus the exchanger approach falls below the chilled water return. With 16 °C chilled water and a fully wet tower the chiller switches off entirely at any wet bulb below 9.3 °C. At 10 °C free cooling still carries 92 % of the load, at 12 °C it carries 70 %, and at 14 °C it carries 48 %. Two degrees of wet bulb near that threshold change the operating cost more than any control action available to the plant.

8.3.3 Chiller Performance and Capacity

Chiller performance is written as a fixed fraction of the Carnot value between evaporator and condenser, in series with a lift-independent ceiling, as shown in Equation (8.7). The two-resistance form is the reduced thermodynamic account of a vapor compression machine, in which a reversible term carries the lift dependence and a lumped internal loss carries what does not scale with lift [21].

$$\begin{aligned} \frac{1}{COP} &= \frac{T_{cond} - T_{chw}}{\eta_C (T_{chw} + 273.15)} + \frac{1}{COP_\infty} \\ T_{cond} &= T_{cw} + \Delta T_{cd} \end{aligned} \tag{8.7}$$

where η_C is the fraction of Carnot achieved, COP_∞ is the ceiling and ΔT_{cd} is the condenser approach. The correlation reproduces both of the usual rules of thumb, since one degree of warmer chilled water improves the COP by 3.8 % at the design condition and one degree of colder condenser water improves it by 3.6 %, with the lift dependence carried entirely by the Carnot term. The Carnot term alone extrapolates poorly at low lift, because at 16 °C chilled water it passes a COP of 40 at about 4.5 °C of lift and keeps climbing, which no machine does, since motor, drive, oil and control losses do not shrink as the lift shrinks. The ceiling term is quantitatively necessary. At the design point the Carnot term alone gives 8.52 against 6.06 for the combined model, and at the optimum 15.57 against 8.94, so the error would be largest at the operating points the supervisory layer selects.

Capacity moves with both ends of the machine, as shown in Equation (8.8).

$$\dot{Q}_{cap} = a_{ch}\, \dot{Q}_{ch,0}\, [1 + k_{ev}(T_{chw} - 7) - k_{cd}(T_{cw} - 30)] \tag{8.8}$$

where a_{ch} is the fraction of the machine in service and $\dot{Q}_{ch,0}$ is the nameplate at 7 °C chilled water and 30 °C condenser water. The rated 3600 kW becomes 3960 kW at 12 °C and 3330 kW against a 35 °C condenser.

The second correction concerns which temperature (8.7) and (8.8) are evaluated at. Both use T_{chw}, the temperature the machine is actually leaving water at, rather than $T_{chw,\mathrm{SP}}$, the setpoint it may be failing to reach, and that distinction decides whether a plant short of cooling recovers or diverges. Evaluated at the setpoint, a saturated chiller never recovers the capacity a warmer evaporator provides, so the shortfall grows, the hall warms, the server fans ramp and (8.3) carries it away. An early implementation evaluated at the setpoint ran the hall to 45 °C on a capacity fault instead of settling. Evaluated at the real leaving temperature the same fault settles, which Section 8.6 demonstrates on a chiller trip that stops at 19.1 °C.

8.3.4 Hybrid Cooling Tower and Make-Up Water

A hybrid tower has wet cells and dry cells, and the evaporative fraction w decides how much of the rejection is which. Fully wet cells chase the wet bulb and fully dry cells chase the dry bulb, as shown in Equation (8.9).

$$\begin{aligned} T_{sink} &= w\,T_{wb} + (1-w)\,T_{db} \\ a &= \mathrm{clamp}\left(a_0 \left(\frac{\dot{Q}_{rej}}{\dot{Q}_{ref}}\right)^{0.6} f_{CT}^{-0.8},\ a_{min},\ 30\right) \\ \tau_{CT}\frac{\mathrm{d}T_{cw}}{\mathrm{d}t} &= T_{sink} + a - T_{cw} \end{aligned} \tag{8.9}$$

where a is the approach above the sink, a_0 is the design approach, f_{CT} is fan speed and $\dot{Q}_{rej} = \dot{Q}_{ch} + P_{ch} + \dot{Q}_{eco}$ is everything the tower has to reject. The gap between the dry bulb and the wet bulb is the wet-bulb depression, and at the design condition of 24 °C dry bulb and 17 °C wet bulb it is 7 °C. The evaporative fraction determines how much of that 7 °C depression the tower reaches.

Only the evaporative fraction spends water, and blowdown follows evaporation through the cycles of concentration, as in Equation (8.10).

$$\dot{V}_w = 60\left(1 + \frac{1}{N-1}\right)\frac{f_{lat}\,w\,\dot{Q}_{rej}}{h_{fg}} \tag{8.10}$$

where N is the cycles of concentration, f_{lat} is the latent share of wet-cell rejection, h_{fg} is the latent heat and the factor 60 converts kg/s to L/min. Water chemistry is not modeled, because basin level and cycles as a controlled variable belong to Volume 2, Chapter 5, and here the tower is treated as a heat sink with a metered water consumption. At the design point it rejects 3736 kW, of which the economizer contributes 331 kW, and makes up 88.5 L/min, a water use effectiveness of 1.77 L/kWh of IT energy.

8.3.5 Objective Function and Constraints

The model reduces to one scalar objective, given in Equation (8.11), and five inequalities, given in Equation (8.12).

$$\begin{aligned} J &= c_e P_{cool} + 0.06\,p_w \dot{V}_w \\ P_{cool} &= P_f + P_{CRAH} + P_{ch} + P_{pump} + P_{cw} + P_{CT} \end{aligned} \tag{8.11}$$

where $P_{ch} = \dot{Q}_{ch}/COP$, $P_{pump} = P_{pump,0}(0.12 + 0.88u^3)$, $P_{CT} = P_{CT,0} f_{CT}^3 + P_{spray,0} w$, $P_{CRAH} = P_{CRAH,0} f_{CRAH}^3$, P_{cw} is the constant condenser water pump, c_e and p_w are the electricity and water prices in \$/kWh and \$/m^3, and the factor 0.06 converts L/min to m^3/h. IT power is excluded from (8.11) because it is the same number whatever the plant does. The five limits are Equation (8.12).

$$\begin{aligned} T_c &\le T_{c,max}, & \dot{Q}_{req} &\le \dot{Q}_{cap} \\ T_{cw} &\le T_{cw,max}, & \dot{V}_w &\le \dot{V}_{max} \end{aligned} \tag{8.12}$$

Table 8.2 Model Parameters. Apart From c_p, c_w and h_{fg}, All Constants Are Adjustable in the Application, and the Four Limits Decide Where the Optimizer Stops.

Symbol	Meaning	Value	Unit
Hall and air side			
$\dot{Q}_{IT,0}$	design IT load	3000	kW
$\dot{m}_{s,0}$	CRAH airflow at full fan	260	kg/s
$\dot{m}_{IT,0}$	server airflow at base speed	210	kg/s
C_c	cold-aisle capacity	60	kW min/°C
C_h	hot-aisle capacity	100	kW min/°C
$P_{f,0}$	server fan power at base speed	180	kW
T_{knee}	fan ramp knee	25	°C
k_f	fan speed rise per degree	0.08	1/°C
s_{max}	maximum server fan speed	2.2	—
$P_{CRAH,0}$	CRAH fan power at full speed	120	kW
c_p	specific heat of air	1.005	kJ/kg/°C
Coil and chilled water			
ε_{max}	coil effectiveness at full flow	0.72	—
k_{coil}	coil NTU shape factor	2.6	—
τ_{coil}	coil and duct lag	0.4	min
$\dot{m}_{w,max}$	chilled water flow at full valve	150	kg/s
C_{chw}	chilled water loop capacity	600	kW min/°C
$P_{pump,0}$	chilled water pump at full flow	140	kW
c_w	specific heat of water	4.186	kJ/kg/°C
Chiller			
$\dot{Q}_{ch,0}$	capacity at 7 °C/30 °C	3600	kW
η_C	fraction of Carnot achieved	0.63	—
COP_∞	lift-independent COP ceiling	21	—
ΔT_{cd}	condenser approach	5	°C
k_{ev}	capacity gain per degree of CHW	0.020	1/°C
k_{cd}	capacity loss per degree of CW	0.015	1/°C
τ_{ch}	capacity control lag	2.0	min

Symbol	Meaning	Value	Unit
Tower and economizer			
a_0	approach at design load, full fan	4.0	°C
a_{min}	minimum achievable approach	1.5	°C
$\dot{Q}_{ref}$	reference heat rejection	3900	kW
τ_{CT}	condenser loop time constant	3	min
$P_{CT,0}$	tower fan power at full speed	90	kW
$P_{spray,0}$	spray pump power fully wet	15	kW
P_{cw}	condenser water pump	60	kW
N	cycles of concentration	4	—
f_{lat}	latent share of wet rejection	0.85	—
h_{fg}	latent heat of vaporization	2440	kJ/kg
ΔT_{HX}	free-cooling exchanger approach	1.5	°C
ε_{eco}	economizer effectiveness	0.90	—
Limits and prices			
$T_{c,max}$	recommended inlet limit	27	°C
$T_{c,allow}$	allowable inlet limit	32	°C
$T_{cw,max}$	chiller high-head limit	36	°C
$\dot{V}_{max}$	water budget	95	L/min
c_e	electricity price	0.10	$/kWh
p_w	water price	2.00	$/m^3
Actuators, instruments, supervisor			
τ_v	valve stroke time	0.25	min
τ_w	spray staging time	1.0	min
τ_f	tower fan ramp time	0.5	min
τ_{eco}	economizer stroke time	0.75	min
τ_{Tc}	cold-aisle sensor filter	0.5	min
τ_{Ts}	supply air sensor filter	0.10	min
T_{opt}	optimizer execution interval	3	min
—	chilled water setpoint rate	0.2	°C/min
—	spray fraction rate	10	%/min

together with a fifth that is structural rather than numerical, namely that the coil must be able to hold the cold-aisle setpoint at all. Every limit in (8.12) came from outside the control system, since the inlet limit is an industry guideline, the capacity a nameplate, the condenser limit a protection setting and the water budget a supply agreement.

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

8.3.6 Parameters, Initialization and Integration

Constants are collected in Table 8.2. Apart from the three physical properties c_p, c_w and h_{fg}, every constant is displayed in the app and is adjustable, a deliberate difference from the identification apps elsewhere in this series. No constant is withheld, because the task here is not to recover a hidden parameter from a trend but to search a cost surface whose shape has been fully declared.

No scenario begins with a transient left over from its own set-up. The initializer evaluates the same steady-state function that the supervisory layer searches and that generates the displayed cost map, which is the only place the steady-state algebra is written down, so the cost map, the optimizer and the plant cannot disagree. A manual scenario is started with the operator's own valve position held fixed instead.

The engine of Section 1.6.1 carries the states forward on a 0.005 min (0.3 s) fixed fourth-order Runge-Kutta step, with the minute as the time base. Controllers execute once per 0.05 min (3 s) and ten samples per simulated minute reach the historian. Only scenario 10 carries sensor noise, drawn from its own fixed seed of 90210 (the application default is 20260829), so every run repeats exactly.

Manual operation and five automatic architectures are provided, and a finished run exports at the historian rate with the true and measured inlet, the valve command and position, the six power terms, the water flow and the lowest feasible cost with its binding constraint on every row (a zero stand-in and the flag *no legal point* while no grid point is feasible), so a run can be reanalyzed offline. Twenty-nine internal checks cover the steady state, the coil energy balance, the sign of each actuator, the recirculation term, the capacity relation (8.8), the free cooling threshold, the water arithmetic, the agreement between the supervisor and a brute-force scan of its own grid, the reporting of an undefined optimum and, in thirteen checks, the predictive layer of Section 8.5.3.

8.3.7 Assumptions

Each simplification below fails somewhere, and the failure mode is stated beside it.

Two well-mixed air nodes with one recirculation term. A real hall has a temperature field, and the number a control system uses is an average over a handful of rack-inlet sensors. Hot spots, rack-level maldistribution and the underfloor pressure profile cannot appear in (8.1), and they are the reason a real hall runs a lower average setpoint than its worst rack could tolerate.

Airflow is commanded as a mass flow. There are no fan curves, no plenum pressure and no filter pressure drop, so moving the CRAH fan here is easier and more linear than in a raised-floor hall with a fixed grille layout. The cube-law fan power in (8.3) is the one piece of that physics kept.

The coil is sensible only. Equation (8.4) has no latent cooling, no dehumidification and no condensate, so nothing stops the supply air setpoint going below the dew point except the lower output limit set on the outer loop. A real coil driven below the dew point spends duty on condensation and the hall must then be rehumidified.

The chiller is one machine with one lag. Real machines have a minimum load, surge limits at low lift, staging across several units and part-load curves that are not as smooth as (8.7). Chilled water flow is proportional to valve position above a minimum-flow floor, so pump staging and low delta-T syndrome are outside the model.

The tower is a lumped approach to a linear blend. Equation (8.9) has no fill performance curve, no plume, no freeze protection and no basin, and the linear blend between wet bulb and dry bulb is a modeling convenience rather than a psychrometric result. Cycles of concentration in (8.10) are a constant rather than a controlled variable. The weather driving it is two numbers with an optional sine, with no relative humidity, no barometric pressure and no solar gain on the envelope, and the wet bulb clamped half a degree below the dry bulb rather than derived from a psychrometric relation.

The economics are two constants. Real electricity is priced by time of use with demand charges, and real water is often a stepped tariff or a hard allocation. A time-of-use price would repaint the map hour by hour and would make (8.11) worth solving over a horizon rather than instantaneously. Limits annunciate and no equipment trips, so a transient remains observable past the point where equipment protection would have ended it.

Two disturbances are stylized. A cold-aisle sensor that acquires exactly 2 °C of bias at $t = 0$ is not the behavior of a real transmitter, but it isolates the one failure mode that optimization does not detect. The diurnal weather is a symmetric sine, which real weather is not, so that the optimum trajectory of Table 8.6 is traversed within one simulated day.

8.4 Selection of the Supervisory Decision Variables

The supervisory layer is given exactly two decisions, and the choice of which two is the most consequential modeling decision in the application. Both were selected by sweeping the steady-state model rather than by argument.

8.4.1 Chilled Water Setpoint Sweep

The chilled water setpoint trades electricity against inlet margin, capacity and pump power at the same time. Warmer water means less lift and a better *COP* in (8.7), but the coil in (8.4) is less effective, so the valve opens further, the pump works harder as the cube of its flow, and eventually the coil cannot hold the cold aisle at all. Measured at design weather with the tower 90 % wet, the valve goes 23.6, 26.7, 33.0, 39.1, 58.8, 73.9 and 92.0 % at chilled water setpoints of 5, 7, 10, 12, 15.6, 17 and 18 °C, and at 19 °C it is wide open with the hall at 24.70 °C and the setpoint lost. Cost over the same sweep runs \$100.38, \$97.17, \$92.50, \$89.61, \$86.18, \$87.28 and \$91.23 per hour. The cost is unimodal with an interior minimum, and each term in it belongs to a different piece of equipment.

8.4.2 Tower Fan Speed Sweep

Tower fan speed is the obvious candidate for the second decision, and the sweep rules it out. More fan means a colder condenser, a better *COP* and less heat to reject in (8.10), which means *less* water rather than more. Measured at 7 °C chilled water and 85 % spray at design weather, taking the fan from 20 % to 100 % moves the condenser water from 32.71 to 21.93 °C, the chiller from 721 to 450 kW, and the water ledger from 94.2 down to 87.7 L/min. Both objectives point the same way. The only opposing term is the fan's own cube-law power, which puts an interior minimum in cost at 80 % of \$98.00 per hour against \$100.83 at 60 % and \$99.11 at 100 %, as shown in the upper panel of Figure 8.2. The fan speed is therefore a single-objective tuning problem with no conflict, so it remains an operator control.

8.4.3 Evaporative Fraction Sweep

The evaporative fraction behaves differently. It is the only variable in the plant that spends water and also the only one that moves the tower sink in (8.9) toward the wet bulb, so water against electricity is a genuine trade set by the wet-bulb depression. Measured at 15.6 °C chilled water and the fan at 80 %, taking spray from 0 to 100 % moves the water ledger from 0 to 99.0 L/min and the cost through \$87.95, \$88.24, \$88.49, \$88.67, \$86.18 and \$83.58 per hour, as shown in the lower panel of Figure 8.2. Cost *rises* over the first three quarters of the range and then falls sharply, because the economizer only comes into reach at the wet end. Free cooling carries 0 % of the load at 75 % spray, 7 % at 90 % and 14 % at 100 %.

Two conclusions follow from that one sweep. The surface is not unimodal in the spray fraction, so a gradient method started at the dry end would stop at \$87.95 and report convergence while a \$86.18 point existed elsewhere in the range. And the global minimum of the sweep, \$83.58 at a fully wet tower, is not available at all, because 99.0 L/min exceeds the 95 L/min budget in (8.12). The optimum is decided by the constraint and not by the gradient, which is the justification for the exhaustive search of Section 8.5.

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

8.5 Controller Design

8.5.1 Regulatory Layer

Every loop runs the velocity-form algorithm of Section 1.6.2, with the derivative on the measurement, written here per scan as Equations (8.13) and (8.14).

$$e_k = \alpha\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{8.13}$$

$$\Delta u_k^P = K_c\,(e_k - e_{k-1})$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c$$

$$\Delta u_k^D = -\,K_c\,\tau_D\,\alpha\,\frac{\mathrm{PV}_k^f - 2\mathrm{PV}_{k-1}^f + \mathrm{PV}_{k-2}^f}{\Delta t_c} \tag{8.14}$$

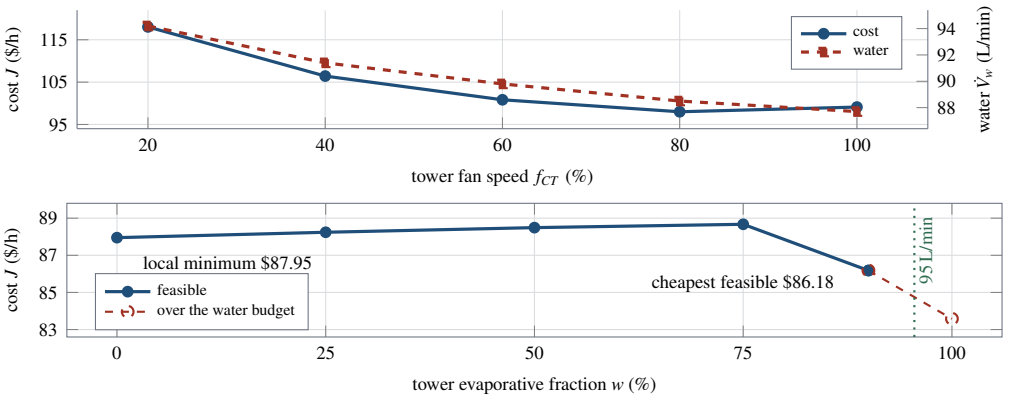


Figure 8.2 Steady-state sweeps of the tower fan speed (upper) and the evaporative fraction (lower) at design weather. In the upper panel cost is the solid line with filled circles and water the dashed line with squares. In the lower panel feasible points are filled circles on a solid line, the points over the water budget open circles on a dashed line, and the budget a dotted vertical line. The fan sweep at 7 °C chilled water and 85 % spray lowers water monotonically while cost passes an interior minimum at 80 %, so the two objectives do not conflict. The spray sweep at 15.6 °C raises cost to $88.67 before it falls to $86.18, so the surface is not unimodal, and the lowest point on the curve, $83.58 at a fully wet tower, is infeasible because 99.0 L/min exceeds the budget.

where $\alpha = \pm 1$ carries the sign convention. The increments are combined and the demand is clipped to the actuator limits, in Equation (8.15).

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= u_{bias} + s_k + I_{k-1} \\ u_k &= \mathrm{clamp}(c_k,\ u_{min},\ u_{max}) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{\tau_t}\,(u_k - c_k) \end{aligned} \tag{8.15}$$

In the symbols of Chapter 1, α is the action a, s_k the running sum S_k, c_k the output before limiting v_k and u_{bias} the output bias u_b. Anti-windup is back-calculation with tracking time $\tau_t = \tau_I$. At $\Delta t_c = 0.05$ min each scan removes $1/60$ of the clipped excess on the single loop ($\tau_I = 3$ min), $1/100$ on the outer cascade loop and $1/20$ on the inner one. The derivative reads a measurement smoothed by the recursion $\mathrm{PV}_k^f = \beta \mathrm{PV}_{k-1}^f + (1-\beta)\mathrm{PV}_k$, $\beta = \exp(-\Delta t_c/\tau_f)$, although every shipped tuning has $\tau_D = 0$. No controller output carries a rate limit. The valve is itself a state with a 0.25 min stroke, which is what separates the commanded position from the reached position.

Four architectures share that algorithm and the same manipulated variable. *Single* drives the valve from cold-aisle temperature with $K_c = 8\,\%/^\circ\mathrm{C}$ of valve, $\tau_I = 3$ min, $\tau_D = 0$, limits of 0 and 100 % and $\alpha = -1$, because opening the valve makes the hall colder. *Plus feedforward* adds the branch-circuit IT load measurement as the output bias $u_{bias} = K_{ff}(\dot{Q}_{IT}/\dot{Q}_{IT,0} - 1)$ in (8.15) with $K_{ff} = 0.3$, which opens the valve while the cold aisle is still on setpoint. *Cascade* makes the cold-aisle loop output a supply air setpoint with $K_c = 2.0\,^\circ\mathrm{C}/^\circ\mathrm{C}$, $\tau_I = 5$ min, limits of 12 and 26 °C and $\alpha = +1$, because a hall that is too warm needs a *lower* supply air setpoint, and a fast inner loop holds that setpoint with the valve at $K_c = 10\,\%/^\circ\mathrm{C}$ and $\tau_I = 1$ min. The feedforward is then added to the outer output, with its sign flipped and scaled by the 14 °C setpoint span, because more load asks for colder supply air rather than a larger number, and the sum is clipped again to 12 and 26 °C. That second clip lies outside the outer controller, whose back-calculation sees only its own limits. *Optimizer* leaves the cascade unchanged and puts the supervisory layer above it.

What the outer controller's output represents changes with the architecture, a valve fraction in two of them and a temperature in the other two, so carrying an integral term across that change would drive the valve closed or command a supply air setpoint of zero. The switch is handled as a transfer instead, in which every controller tracks the valve position the plant already has for one execution, which is the same seeding of I_k in (8.15) that makes the transfer from manual bumpless (Section 1.6.2).

8.5.2 Supervisory Layer

The layer is a transparent constrained grid search rather than a model predictive controller, and it runs the same steady-state evaluator that generates the cost map. It solves the steady-state optimization of Section 1.4.7 by enumeration. Every $T_{opt} = 3\,\text{min}$ it evaluates Equation (8.16) over a 22×15 grid of cell centers.

$$T_{chw,\text{SP}}^{(i)} = 5 + \frac{15(i+\frac{1}{2})}{22}, \qquad w^{(j)} = \frac{j+\frac{1}{2}}{15} \tag{8.16}$$

for $i = 0 \ldots 21$ and $j = 0 \ldots 14$, which is 330 points. The bounds in (8.16) are the declared operating ranges of the two supervisory decisions, so the search cannot command a value outside them. Every point violating (8.12) is discarded, the lowest-cost feasible point is selected, and the four grid neighbors of that point are then probed. A neighbor that is lower in cost but infeasible identifies the constraint that limits further reduction, and that constraint is reported as *binding*. A lower-cost direction that falls outside a decision variable's own range is reported as *range*, because an optimum pressed against the first or last cell of the grid is on a bound rather than interior. A neighbor that is feasible and costlier means the optimum is interior and nothing binds. The two setpoints move toward the selected point at 0.2 °C/min and 10 %/min, so the chiller is never commanded a step change in lift.

Three design choices matter. The search runs once at reset through the engine's priming call rather than waiting for its own schedule, so the reported lowest feasible cost is correct from the first logged sample. It also runs in manual mode and in every architecture, because the identity of the binding constraint is informative even when no controller acts on it. And the layer is deliberately slow, since the plant underneath settles a setpoint step in three minutes on the single loop and eight on the cascade, and a faster supervisor would respond to its own transients rather than to the weather. The economizer valve is opened rather than optimized, because free cooling costs nothing once the tower is already running.

Exhaustive enumeration is chosen over a gradient method because of the measurement in Figure 8.2 rather than for convenience. At 330 evaluations of a fixed-point solve every three simulated minutes the search is inexpensive enough to run in the browser at the fastest simulation speed, and it cannot be trapped by the local minimum at the dry end of the spray range.

8.5.3 Predictive Supervisory Layer

A fifth architecture, the default only of its own scenario, replaces the grid search with a model predictive controller [22, 23] above the same cascade. Every $T_s = 1\,\text{min}$ it plans the same two decisions, $u = (T_{chw,\text{SP}}, w)$, over 30 min in ten 3 min blocks against the grid search's own cost, which makes it an economic predictive controller [24] in the sense of Section 1.4.8. A first design that also planned the supply air setpoint failed the challenge's 0.5 °C band at 0.874 to 0.905 °C over its settings, through its own overshoot of 0.90 °C after the 25 % load drop, and reached 0.763 °C with the shipped feedforward added, so the cold aisle stays with the cascade.

The plan's dynamics are the lags of its own moves, identified at the design point from step responses with the cascade closed, which give two equal lags of $\tau_{chw} = 2.063\,\text{min}$ from the chilled water setpoint to the supply (residual 4.2 % of the response), and $\tau_a = 1.061\,\text{min}$ from the spray command to its position followed by $\tau_{ct} = 3.169\,\text{min}$ to the condenser water (0.45 %). They run open loop on the applied commands, because no shipped controller measures either water loop. The steady map is the grid search's evaluator itself at the applied setpoints and the present weather, load

and parameters, whose central differences of 0.25 °C and 0.05 of spray give the cost's slope g and diagonal curvature H and the slopes of the five constraint quantities in the lagged states. The plan solves Equation (8.17)

$$\min_{u,s} \frac{T_s}{60} \sum_{k=1}^{30} \left(g^\top \delta_k + \tfrac{1}{2} \delta_k^\top \bar{H} \delta_k \right) + \sum_j \Delta u_j^\top S \Delta u_j + \sum_i \left(\rho_i s_i^2 + \sigma_i s_i \right) \tag{8.17}$$

subject to the lag model $x_{k+1} = Ax_k + Bu_k$, the rows $Cx_k \le y_{max} + s$ with $s \ge 0$, the ranges of (8.16), and $|u_1 - u_0| \le rT_s$ and $|\Delta u_j| \le 3rT_s$. Here $\delta_k = u_k - u_0$ is the departure from the applied setpoints u_0, $\bar{H}$ is H with its diagonal floored at 2 and 80 \$/h per unit squared so that a flat stretch is still a bowl, $S = \mathrm{diag}(0.002, 0.02)$ prices each block boundary move Δu_j, x holds the four lags and a constant, and r is the supervisor's own rate limit of 0.2 °C/min and 10 %/min, so both layers walk the map at one speed.

The five rows are the constraints of (8.12), each with one slack over the horizon priced (ρ_i, σ_i) per unit squared and per unit, the two prices of Equation (1.7). The full-valve cold aisle below $T_{c,max}$ and the condenser water below $T_{cw,max}$, the two limits the application raises as alarms, carry $(10^4, 10^3)$ per °C, the make-up water carries $(10^3, 10^2)$ per 10 L/min held 1 L/min inside $\dot{V}_{max}$, the full-valve cold aisle 0.25 °C below the setpoint, which keeps the cascade's authority, carries $(10^2, 10)$, and the capacity margin $\dot{Q}_{cap} - \dot{Q}_{req} \ge 0$ carries $(1, 1)$ per 100 kW. A slack for every third step was not taken, because it cost four to eight times the solve time, and on the condenser and water rows, at the first condenser price, it traded 159.1 for 20.9 min over 36 °C on the 45/30 °C day of Section 8.6.8 against 1.5 for 35.6 min over the water budget in scenario 8.

The plan tracks no output and re-evaluates its rows on the plant's steady model every minute, so it carries no disturbance estimator, and it carries no forecast because the application has none, holding the weather and the load at their present values as the grid search does. It reads only the cold-aisle setpoint, the applied commands, the rate limits and the evaluator's output. A plan takes effect when its status is optimal, its moves lie in their boxes with the first within one decision's travel, and every number is finite. Otherwise the grid search, which keeps running, takes both setpoints from where they are without a step until the next valid plan, which no run reported here needed. In Chrome the 20 decision variables and five slacks solve in a mean of 540 µs over scenario 13 (largest 6.0 ms), 0.054 % of a core at real time and 3.24 % at the 60-fold top speed.

8.6 Results

The fourteen scenarios of Table 8.3 are ordered by concept rather than by equipment. Each fixes an opening mode, an architecture, the three tuning constants, the disturbances in force and the events scheduled during the run, and each poses a question to predict and provides a measured answer. Every number below was measured on the running application at the prices of Table 8.2 unless the scenario sets its own.

8.6.1 Open-Loop Behavior and the Regulatory Baseline

Scenario 1 runs the hall with the valve parked at 27 % and no controller, so it opens at 23.83 °C. The IT load steps 8 % at $t = 5$ min and the cold aisle settles by itself at 25.09 °C, because the coil in (8.4) removes duty in proportion to how far the return air is above the chilled water, so a warmer hall is cooled more strongly. The required valve position is precise, because near the steady position of 28.9 % the inlet moves about 0.5 °C per percent of valve, landing at 22.5, 21.6 and 20.9 °C for valve positions of 32, 34 and 36 %.

The self-regulation runs out at 111 % of design load. With the valve back at 27 % the inlet creeps upward, crosses 27 °C 20.3 min after the load change and 32 °C four minutes later, peaks at 47.2 °C and comes to rest at 47.0 °C only when the server fans reach their speed limit, while at 110 % it settles at 25.5 °C. Both terms in (8.3) are responsible. Fan power rises as the cube of speed and lands

Table 8.3 Scenario Sequence and the Mechanism Each One Exercises

#	Scenario	What it exercises
1	Manual operation	self-regulation and its load limit
2	Single-loop regulation	regulation against a priced setpoint
3	IT load step	load step, gain sweep, oscillation
4	Measured-load feedforward	feedforward and its optimum gain
5	7 °C to 15 °C chilled water	a disturbance the feedforward misses
6	Supervisory layer against a fixed setpoint	the four-objective trade under a budget
7	Hot humid afternoon	economizer shutoff at high wet bulb
8	Chiller capacity loss to 60 %	an infeasible search region
9	Coils fouled 25 %	equipment moves the setpoint axis
10	Cold-aisle sensor bias +2 °C	sensor bias propagates to every layer
11	Water at \$6.00 per m^3	the price ratio as a policy choice
12	Twenty-four hours of weather	the diurnal optimum trajectory
13	The same cost as a 30 min plan	grid search against a predictive layer
14	Graded day, three load changes	four performance measures at once

back in the hall as heat, and the ramping fans eventually pull more air than the CRAH delivers, at which point R in (8.2) becomes positive and the inlet rises with the supply air unchanged. The limit between 110 and 111 % is set by the airflow rather than by the chiller, so the containment airflow margin, not the chiller, sets the overload capacity of the hall.

Scenario 2 closes the loop and the cold aisle holds 24.00 °C exactly. The plant spends \$98.00 per hour where the cheapest feasible point costs \$86.18, so the configuration as shipped spends \$11.81 per hour more than necessary, about \$103,000 per year on one hall. The regulatory loop is not at fault. The chilled water setpoint beneath it is.

8.6.2 Load Steps and Setpoint Moves Across Architectures

Table 8.4 carries the 18 % IT load step of scenarios 3 and 4, read over the pair of steps at $t = 5$ and $t = 35$ min and a run to $t = 60$ min. Half a minute of sensor filter, four tenths of a minute of coil lag and two minutes of chiller capacity control separate a load change from its appearance at the measurement, so the single loop responds late at any gain.

Table 8.4 Cold-Aisle Response to an 18 % IT Load Increase at $t = 5$ min and Its Removal at $t = 35$ min, Read Over Both Events. Peak Is the Largest Inlet Deviation, IAE the Scorecard Value From 0 to 60 min

Configuration	Peak (°C)	IAE (°C min)	Note
$K_c = 4,\ \tau_I = 8$	1.05	18.18	detuned
$K_c = 8,\ \tau_I = 3$	0.78	5.39	as shipped
$K_c = 12$	0.66	3.88	best of the sweep
$K_c = 16$	0.59	8.89	rings
$K_c = 20$	7.64	100.98	unstable
$K_c = 8,\ K_{ff} = 0.15$	0.44	3.65	half the job
$K_c = 8,\ K_{ff} = 0.30$	0.27	2.59	as shipped
$K_c = 8,\ K_{ff} = 0.45$	0.56	3.33	sign reversed
$K_c = 8,\ K_{ff} = 0.60$	1.02	4.98	sign reversed

Raising the proportional gain improves the peak error up to a limit. At 16 %/°C the peak is smaller than at 12 %/°C, 0.59 against 0.66 °C, while the integrated error more than doubles from 3.88 to 8.89 °C min because the response rings and decays only slowly, and at 20 %/°C the loop is unstable at 7.64 °C and spends 7.1 min above the recommended inlet limit. The feedforward reaches 0.27 °C for one additional measurement and no additional actuator, a 65 % reduction from 0.78 °C, and its gain

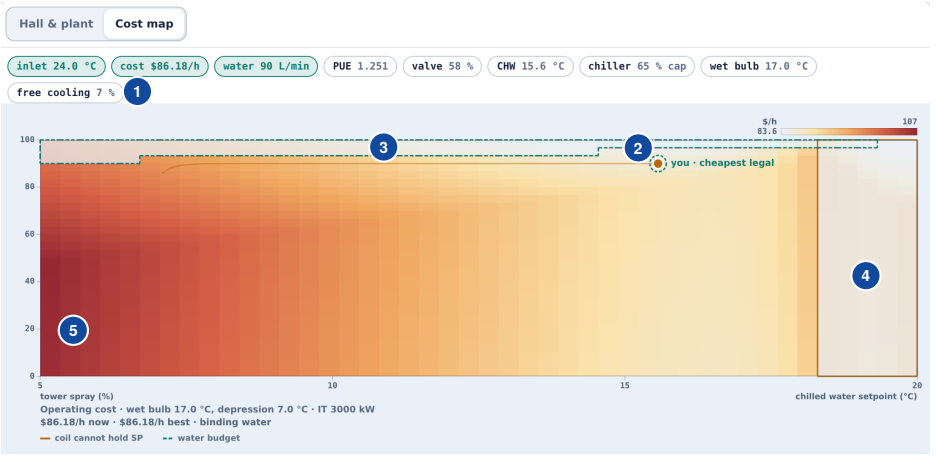


Figure 8.3 Scenario 6 after its two-hour run, showing the cost map over the plane of the two supervisory decisions, chilled water setpoint across and tower spray fraction up. Color is operating cost. (1) The readout chips: inlet, cost rate, make-up water, valve, chilled water and free cooling. (2) The operating point. The supervisory layer has moved the plant from the shipped 7 °C chilled water setpoint to 15.6 °C with the tower 90 % wet, reached at 43 min, which is the lowest-cost feasible point at design weather (Table 8.7). (3) The water budget. The infeasible region beyond it is lower in cost than the operating point and is shown dimmed rather than removed. (4) The coil capability limit, traced on the right. (5) The cold and dry corner, where the color is darkest. Over the run the cold aisle stays within 0.05 °C of 24 °C while the cost falls from $98.00 to $86.18 per hour, make-up water settles at 90 L/min against the 95 L/min budget (Table 8.2), and the run totals are 1531 kWh and 10.73 m^3 with no time above 27 °C (Table 8.5). The readout under the map reports water as the binding constraint.

has an interior optimum. Up to the optimum the hall warms when the job starts and cools when it ends, as it does with no feedforward. Above it the signs reverse, because the valve moves further than the load needs in both directions. At $K_{ff} = 0.6$ the hall is overcooled by 0.79 °C at the start of the job and runs 1.02 °C warm one minute after it ends, when the feedforward has closed the valve too far. A feedforward gain twice its optimum value gives a larger peak error than no feedforward at all.

Scenario 5 separates the two structures. The chilled water setpoint is raised from 7 to 15 °C at $t = 6$ min to reduce chiller power, and the IT load is held constant. Read over a run to 30 min, the single loop peaks at 1.18 °C with an integrated error of 11.50 °C min, and the feedforward gives exactly the same two values, because the disturbance does not appear in the measured IT load the feedforward acts on. The cascade gives 0.17 and 1.59 °C min, about seven times better on both, because its inner loop measures the supply air, which responds within seconds of the coil losing duty, and it moves the valve from 27 to 54 % before the cold aisle deviates. A disturbance that enters inside the inner loop is the one a cascade rejects, the structural argument of Volume 3, Chapter 2. The same move lowers the operating cost from $98.00 to $87.69 per hour as the *COP* rises from 6.1 to 8.5, so the cascade structure is what makes an economically motivated setpoint move safe to execute.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.6.3 Supervisory Optimization Against a Fixed Setpoint

Scenario 6 assigns the two setpoints to the supervisory layer, which moves the plant across the cost map until it reaches a constraint, as shown in Figure 8.3. The result against the identical hours at a fixed 7 °C setpoint is given in Table 8.5. Over two hours at design weather the optimizer spends 1531 kWh against 1748 kWh, a saving of 217 kWh, and 10 732 L of water against 10 619 L, which is 113 L more water.

Table 8.5 The Supervisory Layer Against a Fixed 7 °C Chilled Water Setpoint With the Tower 85 % Wet, Over Three Runs. PUE Is Total Facility Energy Divided by IT Energy Over the Run, Peak Error the Largest Cold-Aisle Deviation, and a Dash Indicates a Quantity Not Defined for That Run.

	Two hours, design weather			Twenty-four hours, diurnal			Graded day, three load changes		
	fixed	optimizer	change	fixed	optimizer	change	fixed	optimizer	change
Cooling energy (kWh)	1748	1531	−12.4 %	15 959	12 344	−22.7 %	19 824	16 574	−16.4 %
Water (L)	10 619	10 732	+1.1 %	120 311	90 728	−24.6 %	130 817	111 124	−15.1 %
Cost ($)	195.99	174.57	−10.9 %	1836.55	1415.88	−22.9 %	2244.07	1879.60	−16.2 %
Cost rate ($/h)	98.00	86.18	−$11.81	—	—	—	—	—	—
PUE	1.291	1.255	−0.036	1.222	1.171	−0.050	1.262	1.219	−0.043
Peak error (°C)	0.00	0.05	—	0.00	0.06	—	1.73	0.42	−76 %

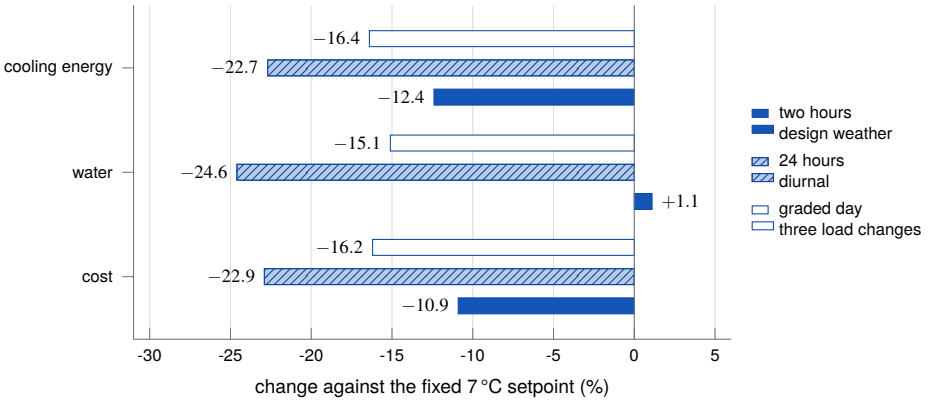


Figure 8.4 The supervisory layer against the fixed 7 °C chilled water setpoint, redrawn from the change columns of Table 8.5: dark bars for the two hours at design weather, hatched bars for the diurnal day and open bars for the graded day. Notice that the optimizer spends 1.1 % more water over two hours at design weather but 24.6 % less over a full day, from the same objective.

> **TRY IT IN THE APP** **6 · The Supervisory Layer Against a Fixed Setpoint**
>
> Load the scenario, open the Cost map and press Run. Watch the marker walk from the shipped 7 °C chilled water setpoint to 15.6 °C with the tower 90 % wet and stop on the water budget. The Binding tile then reads water, and the cost rate has fallen from $98.00 to $86.18 per hour (Table 8.5).

The additional water is a property of the objective rather than an artifact. At that weather the direction of decreasing cost is toward a wetter tower, and (8.11) spends more of the resource that is cheaper per unit of cooling delivered. The water budget limits how far that substitution goes, and the optimum sits on it.

Over a full day the sign of the water change reverses, and Table 8.5 reports both cases for that reason. The fixed baseline sprays the tower at 85 % through the night, when evaporative operation delivers no benefit, and the optimizer does not. Measured at 6 °C dry bulb and 3 °C wet bulb with 16 °C chilled water, a 3 % wet tower costs $41.01 per hour and 2.7 L/min while a fully wet one costs $53.02 per hour and 90.7 L/min, with free cooling carrying 100 % of the load in both cases. Night spraying is charged twice, in spray pump power and in water, and returns nothing. Both results come from the same objective function evaluated at different weather. Figure 8.4 draws the three change columns of Table 8.5 side by side.

The supervisory layer also changes the process the regulatory loop was tuned against, which is a

gain-scheduling problem stated by measurement. At 7 °C the valve sits at 26.7 %, on the steep part of (8.4), where a percent of valve delivers a large change in duty. At the 15.6 °C optimum it sits at 58.5 %, on the saturating part, where a percent of valve delivers much less. Tested with a 0.5 °C step in the cold-aisle setpoint after the plant has settled, the single loop settles at 15 %/°C and sustains an oscillation at 15.5 %/°C with 7 °C chilled water, while at 15.5 °C it settles at 53 %/°C and oscillates only from 54 %/°C, a usable gain about 3.5 times larger. The outer cascade loop moves its limit from between 2.5 and 2.6 °C/°C at 7 °C to between 5 and 6 °C/°C at 15.5 °C, about twice as large. The direction of the change is the operating hazard, because a loop commissioned at the warm end and later run at the cold end during a period of coil cleaning carries a tuning that was never validated there.

8.6.4 Diurnal Weather

Scenario 12 runs twenty-four hours from midnight with the dry bulb swinging between 9 and 27 °C and the wet bulb between 6 and 16 °C, peaking at three in the afternoon, and with the IT load held constant so that the weather is the only input that moves. The lowest-cost feasible point through that day is given in Table 8.6, evaluated on the same steady-state function the supervisor searches, which the plant then trails by a few minutes through its rate limits.

Table 8.6 The Lowest-Cost Feasible Point Through the Day of Scenario 12, Evaluated on the Same Steady-State Function the Supervisor Searches. The IT Load Is 3000 kW and the Cold-Aisle Setpoint 24 °C Throughout, and the Weather Profile Is Symmetric About Its 1500 Peak.

Time	T_{wb} (°C)	$T^{\star}_{chw,\mathrm{SP}}$ (°C)	$w^{\star}$ (%)	free cool. (%)	$J^{\star}$ ($/h)
03:00	6.0	15.6	3	99	40.76
06:00	7.5	16.9	17	92	46.13
09:00	11.0	16.9	97	89	56.98
12:00	14.5	16.9	97	41	71.75
15:00	16.0	16.3	97	21	79.51
span	10.0	1.3	94	78	38.75

The result of interest is which decision variable did not move. Across the whole day the chilled water setpoint ran between 15.6 and 16.9 °C, a span of 1.3 °C, while the evaporative fraction ran from 3 to 97 % and back. In terms of the search grid in (8.16), the setpoint took three of the twenty-two available values in twenty-four hours, all three interior, while the evaporative fraction swept all fifteen from the first to the last. The weather is carried almost entirely on the water axis, so monitoring only the chilled water setpoint would indicate that the supervisory layer had not acted at all.

The mechanism follows from (8.7) and (8.6). On the cold night the economizer carries the whole load and the chiller is off, so the lift the chilled water setpoint controls no longer contributes to cost and the only remaining term on that axis is pump power, which is minimized at the lowest setpoint that keeps the valve away from its upper limit. On the hot afternoon the chiller carries three quarters of the load, warmer water improves the *COP* substantially, and the setpoint rises. The two conditions give optimum setpoints within 1.3 °C of each other. The evaporative fraction has almost no effect on cost at night and a large effect in the afternoon, because it alone determines whether the free-cooling threshold in (8.6) is crossed.

Which constraint binds is a property of the operating condition rather than of the plant, and scenario 12 is a case where no process constraint binds. Sampled every ten minutes across the day, no process constraint is reported at any of the 144 samples, and the closest approach is at the 1500 peak, where the optimum draws 94.6 L/min against the 95 L/min budget. What is reported instead is *range* at 116 of those samples against *none* at 28, so for 81 % of the day the lowest-cost feasible point lies at an end of the spray axis, fully dry or fully wet, rather than on a plant constraint. The

chilled water optimum is interior at every sample. Over most of the day the limit on operating cost is therefore the travel of the evaporative fraction itself, between a fully dry and a fully wet tower, which no wider search can extend. The design condition of 24 °C dry bulb and 17 °C wet bulb is more humid than anything this profile reaches, and at that condition the water budget binds. Table 8.7 collects the conditions that move the optimum and reports the constraint that limits each one.

Table 8.7 The Lowest-Cost Feasible Operating Point Under Seven Conditions, With the Binding Constraint Reported. Range Marks an Optimum at the Fully Dry or Fully Wet End of the Spray Axis. The Regulatory Loop and Its Tuning Are Identical in Every Row.

Condition	$T^{\star}_{chw,\mathrm{SP}}$ (°C)	$w^{\star}$ (%)	$J^{\star}$ ($/h)	Binding
Design, 17 °C wet bulb	15.6	90	86.18	water
Hot humid, 26 °C wet bulb	16.3	90	105.87	water
Coils fouled 25 %	9.4	3	101.93	range
Chiller at 60 %	—	—	—	infeasible
and inlet SP 26 °C	18.3	97	81.11	range
Mild weather, water $2.00	16.9	97	60.26	range
Mild weather, water $6.00	15.6	3	75.84	range
Mild weather, budget 40 L/min	16.3	37	68.61	water

Coil fouling moves the setpoint axis further in one step than a whole day of weather does. A 25 % loss of coil effectiveness takes the optimum from 15.6 to 9.4 °C and the cost from $86.18 to $101.93 per hour, and it takes the tower from 90 % wet to 3 % and the water consumption from 90 to 4 L/min. The mechanism is the free-cooling threshold. Colder chilled water means a colder chilled water return, which puts the threshold in (8.6) out of reach, so the only remaining benefit of a wet tower is the *COP* improvement. At that operating point a 90 % wet tower saves $9.94 per hour of electricity and costs $10.82 per hour of water, a net penalty of $0.88 per hour, so the optimum is a dry tower. A change in coil condition therefore changes the optimal water policy, which is an argument against a fixed table of seasonal setpoints.

8.6.5 Chiller Capacity Loss

Scenario 8 removes chiller capacity and leaves 60 % in service. The direction of decreasing cost on this map is warmer chilled water, and (8.8) states that capacity increases in the same direction, so the two appear compatible. They are not. With 60 % of the machine and a 24 °C cold-aisle setpoint no feasible point exists anywhere on the grid, because the capacity constraint excludes the cold end, the coil cannot hold 24 °C at the warm end, and the two infeasible regions meet. The optimizer has walked the chilled water setpoint from 7 to 9 °C when the trip arrives at 10 min, finds no feasible point after it and stops moving the setpoint, and the chilled water itself drifts to 19.1 °C and settles there. The valve pins fully open at 34.7 min and is at its upper limit for 71 % of the first two hours, and the hall settles 0.78 °C above setpoint at 24.78 °C. The drift terminating is the demonstration that evaluating (8.7) and (8.8) at the actual leaving water temperature gives a stable response, as Section 8.3.3 states. Holding the setpoints while no grid point is feasible, 168 of the first 180 min, is deliberate policy, and the optimum is then reported as undefined. The gap and binding readouts show *no legal point*, and the export logs zero for the best cost and setpoints, the cost as the gap and that flag, where an earlier version logged a zero gap. The correction moved only those five columns, in this scenario and scenario 13.

The margin the plant requires is in a variable the supervisor is not permitted to move. Raising the cold-aisle setpoint to 26 °C when the valve pins restores feasibility at 18.3 °C chilled water and a 97 % wet tower, $81.11 per hour at the steady-state optimum and below the $98.00 the plant started at, after an excursion to 26.73 °C inside the recommended band. The plant reaches that point but does not settle on it. A 26 °C inlet ramps the server fans in (8.3) past the 221 kg/s the CRAH delivers once

the inlet exceeds 25.66 °C, so the recirculation R in (8.2) switches on and off around the setpoint, and the hall cycles between 25.5 and 26.5 °C with a period of 5.0 min while the cost averages $81.09 per hour. Capacity is recovered by spending inlet margin, at a rate given by k_{ev} in (8.8), where one degree of chilled water is worth 2 % of capacity.

8.6.6 Sensor Bias

Scenario 10 applies +2 °C of bias to the cold-aisle transmitter and changes nothing else. The controller closes the loop on the biased measurement, so it holds the hall two degrees colder than intended while the trend shows no error, with the measurement reading 23.90 °C at three hours while the true inlet is 21.98 °C. The additional cooling is not free, because colder supply air requires a more open valve and colder chilled water, and colder chilled water lowers the COP in (8.7). The cost is $103.89 per hour against $98.21 with an unbiased sensor, which is 2781 kWh over three hours against 2622 kWh, a 6 % electricity penalty for a two-degree instrument error.

Enabling the supervisory layer from the start does not recover the loss. The same bias costs $93.81 per hour against $86.40 at three hours. The search never reads the measurement, because it evaluates the steady-state model at the 24 °C setpoint, so it finds the same $86.18 optimum at 15.6 °C and 90 % spray as with an honest sensor while the loop beneath it holds the hall at 22 °C. The application reports the result as a $7.63 per hour gap between the plant and that optimum, against $0.22 with an honest sensor. The optimizer solves the correct problem for a hall that is not the one being cooled, and nothing in the control system acts on the persistent gap that is the only symptom of the bias. The countermeasure is a redundant measurement at a second rack or a calibration check with a portable instrument, which belongs to the commissioning plan rather than to the control configuration.

> ✦ DID YOU KNOW? **A Bias the Single Loop Never Sees**
>
> From the simulator: a supply air sensor that reads 3 °C high leaves the single loop with a worst true error of 0.00 °C, because the single loop never reads that sensor. The cascade, whose inner loop holds on it, reaches 2.23 °C, and the supervisory layers built on the cascade reach 1.46 to 2.10 °C (Table 8.9, case W9).

8.6.7 Water Price and Water Budget

Scenario 11 runs mild weather at 18 °C dry bulb and 12 °C wet bulb, and changes only the water price in (8.11). At $2.00 per cubic meter the optimizer runs 16.9 °C chilled water with the tower 97 % wet, drawing 496 kW and 89 L/min at $60.26 per hour. At $6.00 it runs 15.6 °C with the tower 3 % wet, drawing 746 kW and 3 L/min at $75.84 per hour. Over two hours that is 1052 kWh and 10 673 L against 1488 kWh and 795 L, so tripling the water price saves 9878 L at a cost of 436 kWh. No part of the control system changed. One constant in the objective changed.

A hard budget behaves differently from a price. Returning the price to $2.00 and capping water at 40 L/min places the optimum between the two cases, at 37 % spray and 643 kW at $68.61 per hour, drawing 35.6 L/min and reporting water as the binding constraint where the same weather at the same price had reported only the wet end of the spray range. A price shifts the optimum continuously, while a budget makes the optimum sit on a boundary.

The graded scenario reports the same distinction in one set of numbers. Twenty-four hours of weather and three load changes run first on a detuned single loop at 4 %/°C with eight minutes of reset, where the peak cold-aisle error is 1.73 °C and the day costs $2244.07. Tuning the loop to 8 %/°C takes the peak to 1.21 °C and adding the load feedforward takes it to 0.41 °C, at $2243.84 and $2243.89. Eighteen cents separate the worst regulatory performance in the set from the best. The supervisory layer takes the same day to $1879.60 at a peak error of 0.42 °C, within 0.01 °C of the best fixed-setpoint run, as recorded in Table 8.5. Against the challenge's four limits, a ±0.5 °C band, no time above 27 °C, $2000 and 120 m^3 of water, the day as delivered fails three, at 1.73 °C, $2244 and 130.8 m^3, although the inlet never exceeds 27 °C. The supervisory layer passes all four, at 0.42 °C, $1880 and 111.1 m^3, with margins of 0.08 °C, $120 and 8.9 m^3 and no time above 27 °C.

Regulatory performance and operating cost are therefore separate measures, and improving one does not improve the other.

8.6.8 Predictive Layer Against the Grid Search

Scenario 13 repeats the design weather of scenario 6 and trips one chiller to 60 % at 60 min. Both layers walk from 7 °C at the same rates, arriving at 42.6 and 46.0 min, and settle on the water budget, the grid at its nearest legal cell of 15.57 °C and 90 % (90.1 L/min) and the plan on its own row at 15.96 °C and 94.5 % (94.0 L/min), at $84.96 per hour against $86.18 over 70 to 120 min of scenario 6. After the trip no grid point is feasible for 120 of the 180 min, the grid search holds its setpoints, and the hall settles at 24.78 °C, as in scenario 8. The plan raises the spray to 95.8 % and the hall settles at 24.38 °C for $0.73 per hour more (Figure 8.5). Its chilled water setpoint of 18.53 °C changes nothing, since the chiller at capacity delivers 18.69 °C regardless, and with either 15.57 or 18.5 °C held the hall settles at 24.379 °C at 95.8 % spray and 24.778 °C at 90 %. Holding the setpoints fixed in this way assigns both differences to the spray, the $1.22 per hour the plan saves before the trip and the 0.40 °C it gains after it. Over 180 min the plan spends $265.65, 2322 kWh and 16.71 m^3 against $265.29, 2335 kWh and 15.88 m^3.

> **TRY IT IN THE APP** **13 · The Same Cost as a 30-Minute Plan (MPC)**
>
> Load the scenario and run three hours on MPC with the Cost map open, through the chiller trip at 60 min. The plan raises the spray to 95.8 % and the hall settles at 24.38 °C. Press Reset, set Control architecture to Optimizer and run again: the grid search holds 90 % spray and the hall settles at 24.78 °C (Figure 8.5).

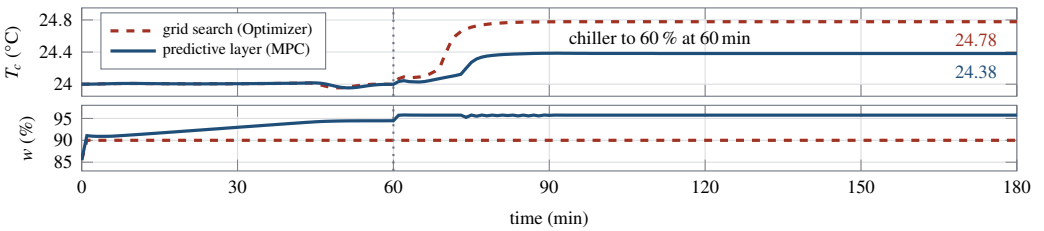


Figure 8.5 Scenario 13, true cold aisle (upper) and spray command (lower), under the grid search (dashed) and the predictive layer (solid). The dotted vertical line marks the chiller trip. After the chiller trip the grid search holds 90 % spray with no feasible point, and the plan's 95.8 % holds the hall 0.40 °C cooler.

Table 8.8 Cost ($) and Minutes Over the Water Budget of the Supervisory Layers. The Last Three Are Scripts, the Lowest Cost in Bold.

Case	Grid search	Grid, retuned	MPC	Grid 88×60	Cont.	Cont., margins
13, 180 min	265.29	**262.87**	265.65	265.38	264.91	265.13
min over	0	0.6	0	12.5	12.4	6.5
12, day	1415.88	1411.08	1409.40	1407.54	**1405.04**	1409.49
min over	3.5	3.7	0	89.9	180.4	16.1
14, graded	1879.60	1874.59	1873.74	1873.53	**1871.95**	1873.63
min over	21.7	24.6	19.0	82.1	305.5	33.3
3, job, 90 min	139.33	**135.91**	137.84	138.22	138.16	138.38
min over	0.8	0	3.9	24.8	26.2	11.4
13, T_c after trip	24.78	24.78	24.39	24.38	24.39	24.47

Table 8.8 sets the plan against the shipped layer and four stronger forms of it, with nothing retuned per case and no fallback in the plan's 90 to 1440 solves per case. The retuned column is the grid

Table 8.9 Nine Worst Cases Under Every Structure, 180 min Unless Noted. Peak T_c, T_{cw} (°C) and $\dot{V}_w$ (L/min) With Minutes Over 27 °C, 36 °C or the Budget in Parentheses, Worst True Error $|e|$ (°C), and Minutes at Capacity (cap) or Saturated (sat). W2, W4, W5, W7, W9 Start From Scenario 6, Events at 10 min (W6 at 30 min). Feedforward Equals Single Except W5 ($|e| = 0.372$).

Case: measures	Single	Cascade	Grid search	Grid, retuned	MPC	Grid 88×60	Cont., margins
W1 scenario 8: T_c, $\dot{V}_w$, sat	25.14, 88.5 (0), 127	25.14, 88.5 (0), 150	24.78, 92.4 (0), 145	24.78, 91.0 (0), 158	24.39, 95.3 (1.5), 123	24.38, 97.9 (6.2), 122	24.47, 96.5 (4.0), 132
W2 pump to 55 %: $\dot{V}_w$	116.1 (2.0)	116.1 (1.9)	117.4 (2.1)	101.6 (1.1)	118.7 (3.1)	123.6 (19.0)	122.2 (7.2)
W3 scenario 7: T_{cw}, $\dot{V}_w$	32.48, 95.7 (4.8)	32.48, 95.6 (4.5)	31.79, 96.1 (9.0)	31.79, 100.7 (7.4)	32.07, 94.8 (0)	31.88, 97.9 (9.6)	31.85, 96.5 (8.1)
W4 weather to 45/30 °C: T_{cw}, $\dot{V}_w$, $\|e\|$	37.14 (162.4), 97.5 (166.1), 0.136, 163 cap	37.14 (162.3), 97.4 (166.1), 0.032, 163 cap	36.35 (158.9), 101.4 (166.5), 0.051	36.27 (158.0), 102.6 (165.8), 0.196	36.10 (147.5), 100.0 (142.9), 0.236, 5.0 sat	35.46 (0), 107.9 (173.8), 0.057	35.67 (0), 106.3 (171.5), 0.055
W5 load to 125 %, 120 min: $\dot{V}_w$, $\|e\|$	114.3 (109.2), 1.079, 2.0 cap	114.9 (109.3), 0.460, 5.4 cap	104.9 (3.5), 0.393	91.0 (0), 1.004, 5.4 cap, 5.8 sat	95.7 (0.3), 0.394, 2.6 cap	111.3 (10.6), 0.398	109.8 (8.3), 0.397
W6 scenario 7, chiller to 60 %: T_c	50.80 (142.6)	50.80 (139.4)	50.26 (142.1)	50.26 (144.7)	49.18 (142.0)	50.35 (142.1)	50.32 (142.1)
W7 coils fouled 40 %: T_c	59.22 (169.5)	59.22 (169.4)	58.88 (169.4)	58.88 (169.5)	58.20 (169.4)	58.49 (169.4)	58.58 (169.4)
W8 scenario 10: $\|e\|$, $\dot{V}_w$	2.37, 93.2 (0)	2.39, 100.6 (0.8)	2.28, 104.3 (1.6)	2.37, 103.2 (2.6)	2.28, 105.0 (40.6)	2.28, 107.5 (146.3)	2.28, 106.9 (49.9)
W9 supply air bias +3 °C: $\|e\|$, $\dot{V}_w$	0.00, 88.5 (0)	2.23, 102.5 (0.4)	2.10, 100.4 (0.4)	1.46, 97.7 (3.1)	2.10, 101.9 (5.7)	2.09, 106.8 (140.6)	2.10, 105.3 (13.9)

search under the tuning rule fixed for the plan before any run, which chose execution every minute and rates of 0.8 °C/min and 40 %/min. The rule tied the plan's own 27 weightings at $265.65 to the cent, and the water peak broke the tie afterward. The last three columns are scripts of the same cost, the grid on 88×60 points, a pattern search from the grid's best cell, and that search held to the plan's 1 L/min and 0.25 °C margins.

The plan is the cheapest on none of the four cases. Its gain over the shipped layer, $6.48 on the ambient day and $5.86 on the graded day, comes from continuous decisions and not from the horizon, because its only dynamics are its own lags and the rate limits set the walk. A grid sixteen times denser matches or beats it on every case but the job step, and the continuous search held to its margins comes within $0.54 of it on every case with more minutes over the budget (6.5, 16.1, 33.3 and 11.4 against 0, 0, 19.0 and 3.9). Free of the margins it is the cheapest on both days but rides the budget for 180.4 and 305.5 min, so the margins and the lag model buy water discipline, not cost. The retuned grid search beats the plan on scenario 13 and the job step, because the walk from 7 °C is the largest cost of the morning, and loses the days by $1.68 and $0.85. On the job step the plan is a local method on a cost with two basins in the spray at 118 % load, staying near 77.9 % mean spray over the 30 min job where the grid search goes to 21.6 %, and pays $51.52 against $50.92 before recovering $2.09 over the next 55 min. Both layers pass the challenge, at worst errors of 0.411 and 0.416 °C.

In Table 8.9 only W6 and W7 pass 27 °C, under every structure, through the server fan feedback of (8.3). The plan is worse than the shipped layer on six levels. In W4 the valve saturates for 5.0 min against none and the worst error is 0.236 against 0.051 °C, because the lags were identified at the design valve of 27 %. Near a full valve the chilled water loop is second order with a damping ratio of 0.36 and trails a ramp by 1.0 min instead of 3.5, so the plan's lag state trailed the plant by 0.45 to 0.62 °C and the chilled water overshot to 18.43 °C when the plan reversed. At the first condenser price, judging the coil row on the commands, or every row on both, cut the saturation from 4.7 to 1.1 and 2.1 min but put scenario 8 17.7 and 24.3 min over the water budget, so the transient stands as a limitation. The other four levels are the water of W1, W2, W8 and W9. The W2 peak comes 0.1 min after the pump trip, before any decision, and scales with the spray in place, 91.2 against 90.0 %. In W8 and W9 the steady map assumes the cascade holds the true cold aisle on its setpoint, the biased sensor holds it 2 °C colder, and the plant's water sits 0.91 L/min above the plan's row (94.90 against 94.00), nearly the whole margin, so the premise that a re-evaluated steady model leaves nothing

to estimate fails under sensor bias. The grid search hides the same error (90.74 against 90.0) in its 5 L/min cell. A 2 L/min margin clears W8 to 1.6 min but costs $3.98 and $1.87 on the two days and raises the trip error of scenario 13 from 0.386 to 0.454 °C. The plan is better on the cold aisle of W1, the water of W3 and W5, the peaks of W6 and W7 (by 1.08 and 0.68 °C) and both broken limits of W4, the last only since a review priced the condenser row with the inlet row, as the application ranks its alarms. At the water row's price the plan had bought 5.4 L/min of water with 0.04 °C of head, peaking at 36.43 °C for 159.1 min, and the repricing moved no other case by a cent.

8.7 Technology in Practice

What Measures It

Performance monitoring of a data hall puts temperature and humidity sensors at the air inlets of the IT equipment, and wireless networks of sensors across the rack inlets are becoming more common. Supply air temperature and humidity are monitored at each room air handler, and the facility's energy monitoring and control system trends every measured value and archives it for at least a year. Linking the servers' own on-board temperature sensors to that system is a developing technology [25]. In this chapter the cold-aisle transmitter and the supply air transmitter of Figure 8.1 stand for those two measurements, each a first-order filter on the true temperature.

> **TECHNOLOGY** **Sensor: The Rack Inlet Thermometer**
>
> The temperature a data hall is held to is read where air enters the IT equipment, at heights set by the industry thermal guidelines, by wired or wireless sensors outside the rack or by the servers' own on-board sensors read over the network. Guidance asks that supply temperature and airflow be controlled from these intake readings and not from the return air [25]. An electrical sensor of this kind is a resistance detector, a metal element whose resistance rises almost linearly with temperature, or a thermistor, a semiconductor of high sensitivity over a narrow, strongly nonlinear range [26]. Accuracy must fit the use: a pair of sensors good to ± 1.5 °F says nothing about a chilled water difference that can fall to 5 °F [25]. The chapter's cold-aisle sensor has a 0.5 min filter and an optional bias b_c (Tables 8.1 and 8.2).

What Moves It

Variable-speed drives on the room air handlers vary the airflow with the load. On the plant side, variable-flow pumping and chillers with variable-speed compressors keep the plant efficient at low load. A water-side economizer is a heat exchanger between the chilled water and the cooling tower water. Placed upstream of the chillers, it cools the returning chilled water first: when it can reach the chilled water setpoint the chillers are bypassed, and otherwise they run at reduced load to finish the job. Larger towers with variable-speed fans approach the wet bulb more closely [25]. The chapter's plate exchanger in series (Equation (8.6)), its chiller with a 2 min capacity lag and its tower approach model are this arrangement in lumped form.

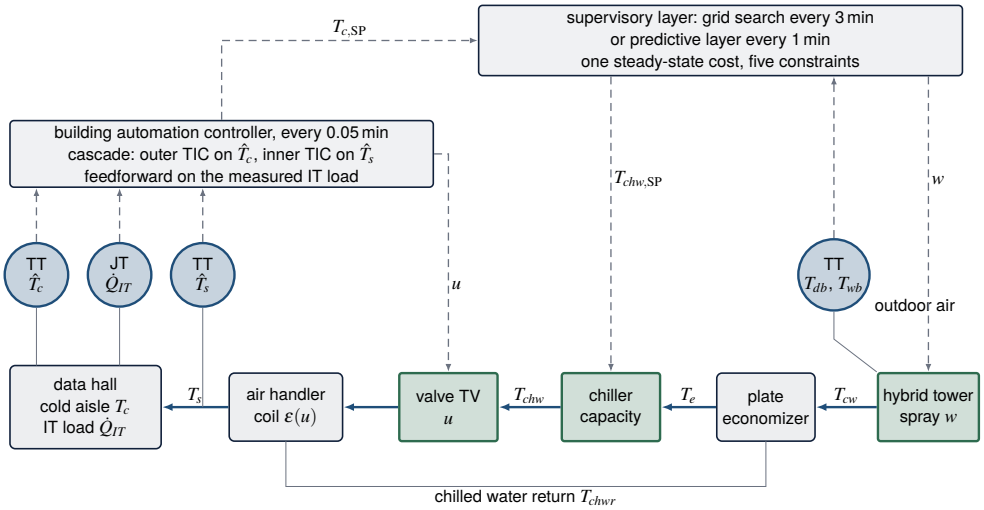


Figure 8.6 The chapter's loops on field hardware, drawn from Sections 8.3 and 8.5. Green boxes are the final elements (the chilled water valve TV, the chiller's capacity control and the tower spray), blue circles the sensors (cold aisle, supply air, IT load, outdoor dry and wet bulb), dashed arrows signals. Notice that the supervisory layer moves only the chiller and the tower, while the valve, the one manipulated variable of every architecture, stays with the regulatory loops.

TECHNOLOGY **Final Element: The Chilled Water Valve**

A control valve throttles the chilled water to each air handler's coil. In a globe valve a plug moves against a fixed ring seat, and the shape of the plug gives a quick-opening, linear or equal-percentage characteristic, the last two the best for throttling. Hysteresis makes the flow depend on the direction the valve last moved, and when it grows large it can hold a loop in oscillation around its setpoint [27]. Guidance for data centers asks for two-way valves rather than three-way, so that variable-speed pumps, not a bypass, set the chilled water flow [25]. The chapter's valve is a 0.25 min stroke with a minimum-flow floor and no stuck-valve fault (Section 8.8).

Where the Controller Runs

A building automation system controls heating, ventilation and air conditioning among a building's other systems, and its sensors reach building controllers over wired or wireless paths [28]. Supervisory loops and the loops below them run with cycle times from milliseconds to minutes [28]. Guidance asks that the control system be programmed to maximize the efficiency of the cooling under changing weather and IT load, with all air handlers communicating and working to the same temperature [25]. In the chapter the loops execute every 0.05 min, the grid search every 3 min and the predictive layer every 1 min, as Figure 8.6 draws them.

> **IN PRACTICE** **Raise the Setpoints, a Little at a Time**
>
> A federal best-practices guide asks operators to raise the supply air temperature until the most demanding equipment intake sits as close to 80 °F as possible and, leaving room for error, suggests a setpoint of 77 to 79 °F, reached in small steps and only after air management is in place. Supply air and chilled water temperatures are to be set as high as the needed cooling capacity allows. Chilled water at 55 °F or above helps the chiller and adds economizer hours, and the chilled water supply may be reset higher at low load [25]. The supervisory layer of this chapter makes that reset against a 24 °C inlet setpoint.

What Is Modernizing It

Three developments stand out. A hybrid of dry and wet heat rejection, priced as this chapter prices it, has run at one computing center. A dry cooler that moves heat in a passive refrigerant cycle was placed upstream of the cooling towers, and its fan controls watch the difference between the entering water and the dry bulb and the ratio of the water price to the electricity price, adding dry cooling as either rises. In its first year the site's water use effectiveness was 0.70 L/kWh against 1.27 L/kWh with towers alone, a saving of 4400 m^3 [29]. Direct liquid cooling moves heat into a fluid at or near the chip, and its warmer chilled water pairs well with a water-side economizer [25]. The data centers with the lowest PUEs are those that cool their artificial intelligence servers with liquid [20]. Supervisory optimization is the third: the chilled water reset that guidance gives as a rule, chosen by a search against prices and limits.

8.8 Limitations

The supervisory layer searches a fixed grid and moves at a rate limit. It has no prediction horizon, no model of the transient its own move causes, and no mechanism to accept a higher cost now for a lower cost later, and it does not account for chiller starts, staging or minimum run times. It also solves against exactly the same steady-state model the plant is built from, which no real optimizer has available. Plant and model mismatch, and the parameter estimation or bias updating that a real-time optimization layer requires to handle it, are the largest omission in the application.

Only two decisions are optimized. The cold-aisle setpoint, the tower fan and the CRAH fan remain operator controls, and scenario 8 demonstrates that the margin the plant requires is in a variable the supervisor may not move. The control scope is PID with a feedforward, a cascade, a grid-search supervisor and the linear predictive layer of Section 8.5.3, with no nonlinear model predictive control and no explicit gain scheduling on the chilled water setpoint even though Section 8.6 measures a change of about a factor of 3.5 in usable proportional gain across the range the supervisor moves through. Instruments are exact apart from the bias, the noise and the filters set by the user, so drift, a stuck valve and an inlet sensor mounted at an unrepresentative rack are absent, and the last of those is the most consequential real failure the application does not represent.

The predictive layer models only its own lags at the design valve, with no cascade transient, cross curvature, forecast, water measurement or estimator, so it corrects neither the 45/30 °C day's saturation nor the 0.91 L/min error under sensor bias. Moving horizon estimation on the same model could estimate the unmeasured coil fouling ϕ and the cold-aisle sensor bias b_c from the valve position, the supply air temperature and the chiller duty, which is the soft sensing whose absence scenarios 9 and 10 demonstrate.

Nothing here supports a quantitative claim about a real hall. With no airflow field, no psychrometrics, no latent cooling, a lumped tower and two constant prices, the model cannot predict a PUE, a water consumption or the setpoint a particular site should run. What the application does establish is the structure of the decision, which is that four objectives conflict, that the optimum lies on a constraint, and that the constraint moves with the uncontrolled inputs. Where each assumption fails is set out in Section 8.3.7.

OPEN CHALLENGE **Demand That Moves in Time**

The 2025 update of United States data center electricity use asks for research that characterizes how data center demand varies in time, so that the grid or on-site generation can serve it, and that explores ways of shifting the timing of loads, called demand flexibility. It also asks how the sector's growth may be constrained by water resources [20]. This chapter prices electricity and water as two constants and holds the weather and load at their present values, so its predictive layer gains from continuous decisions and not from its horizon. Progress would be a plant objective with time-varying prices and forecasts, over which a plan's horizon has something to buy.

8.9 Conclusion

This chapter has set out a data hall and its cooling plant as an interactive laboratory for supervisory optimization above a regulatory layer. The simulator integrates ten states covering two air nodes with a recirculation term, a coil saturating in water flow, a chilled water loop whose capacity clamp is the constraint, a condenser loop and four actuator positions, with a chiller written as a fraction of Carnot in series with a ceiling of 21 and evaluated at the temperature the machine is actually leaving water at. The regulatory loop holds 24.00 °C wherever the plant has the capacity, and at the shipped 7 °C chilled water setpoint it costs $98.00 per hour against $86.18 for the cheapest point that satisfies all five constraints. The supervisory layer is a 22×15 constrained grid search, chosen over a gradient method because an evaporative fraction sweep at 15.6 °C rises from $87.95 to $88.67 before falling to $86.18. Over two hours at design weather it saves 217 kWh and spends 113 L more water, and over a diurnal day it saves 22.7 % of the electricity and 24.6 % of the water, while the chilled water setpoint moves 1.3 °C and the evaporative fraction moves from 3 to 97 %. A cold-aisle sensor two degrees high produces a trend with no visible error, a hall two degrees colder than intended and a 6 % electricity penalty that the optimizer propagates rather than detects. A predictive layer on the same cost saves $6.48 and $5.86 on the two days through continuous decisions rather than its horizon. A fuller model would add latent cooling with a psychrometric model, a time-of-use electricity price, plant and model mismatch with a bias-updating layer, moving horizon estimation and nonlinear predictive control on the same model, and a gain schedule on the chilled water setpoint. In its present form the simulator separates regulatory performance from plant operating cost in a browser, with nothing to install.

KEY POINTS

- A loop that holds the server inlet on its setpoint says nothing about the cost, which the chilled water setpoint and the tower spray decide.
- The cost is not unimodal in the evaporative fraction, so the supervisory layer searches a grid rather than following a gradient.
- Over a day the chilled water optimum moves 1.3 °C while the evaporative fraction sweeps from 3 to 97 %.
- A predictive layer on the same cost saves $6.48 and $5.86 on the two days through continuous decisions rather than its horizon.

Where else this applies. An office building's supervisory layer (Chapter 9) also gains little from a plan's horizon: on the design day the 24-hour plan's bill is $557.61 against $552.74 for its supervisor retuned once (Table 9.5).

References

[1] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, 2020.

[2] M. A. B. Siddik, A. Shehabi, and L. Marston, "The environmental footprint of data centers in the United States," *Environmental Research Letters*, vol. 16, no. 6, p. 064017, 2021.

[3] B. Ristic, K. Madani, and Z. Makuch, "The water footprint of data centers," *Sustainability*, vol. 7, no. 8, pp. 11 260–11 284, 2015.

[4] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United states data center energy usage report," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-1005775, 2016.

[5] S. Skogestad, "Plantwide control: the search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[6] S. Engell, "Feedback control for optimal process operation," *Journal of Process Control*, vol. 17, no. 3, pp. 203–219, 2007.

[7] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

[8] S. Wang and Z. Ma, "Supervisory and optimal control of building HVAC systems: A review," *HVAC&R Research*, vol. 14, no. 1, pp. 3–32, 2008.

[9] Z. Ma and S. Wang, "Supervisory and optimal control of central chiller plants using simplified adaptive models and genetic algorithm," *Applied Energy*, vol. 88, no. 1, pp. 198–211, 2011.

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

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

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

[13] S. Huang, W. Zuo, and M. D. Sohn, "Improved cooling tower control of legacy chiller plants by optimizing the condenser water set point," *Building and Environment*, vol. 111, pp. 33–46, 2017.

[14] C. D. Patel, C. E. Bash, R. Sharma, M. Beitelmal, and R. Friedrich, "Smart cooling of data centers," in *2003 International Electronic Packaging Technical Conference and Exhibition, Volume 2*. ASME, 2003, pp. 129–137.

[15] R. R. Schmidt, E. E. Cruz, and M. Iyengar, "Challenges of data center thermal management," *IBM Journal of Research and Development*, vol. 49, no. 4.5, pp. 709–723, 2005.

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

[17] L. Parolini, B. Sinopoli, B. H. Krogh, and Z. Wang, "A cyber-physical systems approach to data center modeling and control for energy efficiency," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 254–268, 2012.

[18] M. K. Patterson, "The effect of data center temperature on energy efficiency," in *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*. IEEE, 2008, pp. 1167–1174.

[19] H. Zhang, S. Shao, H. Xu, H. Zou, and C. Tian, "Free cooling of data centers: A review," *Renewable and Sustainable Energy Reviews*, vol. 35, pp. 171–182, 2014.

[20] S. J. Smith, A. Hubbard, A. Newkirk, M. Ganeshalingam, B. Holecek, D. Sartor, M. Mills, and A. Shehabi, "United states data center energy usage report: 2025 update," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-2001758, 2026.

[21] J. M. Gordon and K. C. Ng, "Thermodynamic modeling of reciprocating chillers," *Journal of Applied Physics*, vol. 75, no. 6, pp. 2769–2774, 1994.

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

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

[24] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[25] O. Van Geet and D. Sickinger, "Best practices guide for energy-efficient data center design," National Renewable Energy Laboratory for the U.S. Department of Energy Federal Energy Management Program, Golden, CO, Tech. Rep. DOE/GO-102024-6283, 2024.

[26] A. Boghossian, J. Brown, and S. Zak, "3.2: Temperature sensors," in *Chemical Process Dynamics and Controls*, P. Woolf, Ed. University of Michigan open textbook, LibreTexts, 2023.

[27] D. Katzman, J. Moreno, J. Noelanders, and M. Winston-Galant, "Valves: Types, kinds, and selection," in *Chemical Process Dynamics and Controls*, P. Woolf, Ed. University of Michigan open textbook, LibreTexts, 2023, ch. 3.9.

[28] 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.

[29] D. Sickinger, O. Van Geet, S. Belmont, T. Carter, and D. Martinez, "Thermosyphon cooler hybrid system for water savings in an energy-efficient HPC data center: Results from 24 months and impact on water usage effectiveness," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-2C00-72196, 2018.

8.10 Exercises

8.1 Concept check · 2 min

The supervisory layer prices all 330 points of its grid every three minutes instead of following the slope of the cost downhill. Which measurement in Section 8.4 makes that a modeling decision rather than a convenience?

(a) Cost falls steadily toward a wetter tower, so a slope step would overshoot the water budget

(b) The chilled water sweep has two minima, one near 7 °C and one near 15.6 °C

(c) The tower fan sweep has an interior minimum at 80 %, which a slope method cannot find

(d) Cost rises over most of the spray range before it falls, so a search started dry stops early

8.2 Concept check · 1 min

In Scenario 5 the chilled water setpoint rises from 7 to 15 °C at 6 min while the IT load stays constant. The single loop and the feedforward both peak at 1.18 °C, the cascade at 0.17 °C. Why does the cascade do better?

(a) Its inner loop measures the supply air, which answers within seconds, and opens the valve early

(b) Its feedforward enters the outer output, so it sees the setpoint change before the cold aisle does

(c) It moves the CRAH fan as well as the valve, so it has a second handle on the supply air

(d) Its outer loop has a larger gain than the single loop, so it reacts harder to the same error

8.3 Concept check · 2 min

What does the supervisory layer minimize every three minutes?

(a) The deviation of the cold aisle from its 24 °C setpoint, weighted by the price of electricity

(b) The cost of cooling electricity and make-up water, subject to five limits that include holding the setpoint

(c) The power usage effectiveness of the hall, subject to the water budget and the chiller capacity

(d) The make-up water the tower draws, subject to the inlet limit and the condenser high-head limit

8.4 Theory and calculation · 6 min

At the design point the tower runs 85 % wet and rejects 3736 kW, the cooling equipment draws 874 kW and the IT load is 3000 kW. Use Equations (8.10) and (8.11) and Table 8.2.

(a) Compute the make-up water $\dot{V}_w$ in L/min. *Checkpoint:* 88.5 L/min.

(b) Compute the water use effectiveness, liters of make-up per kilowatt-hour of IT energy. *Checkpoint:* 1.77 L/kWh.

(c) Compute the cost rate J, split into electricity and water, and the PUE. *Checkpoint:* \$87.40 and \$10.62, \$98.02 per hour, PUE 1.291.

8.5 Theory and calculation · 6 min

At the design point the chiller makes 7 °C water against condenser water at 22.71 °C. Use Equation (8.7) with $\eta_C = 0.63$, $COP_\infty = 21$ and $\Delta T_{cd} = 5\,^\circ\mathrm{C}$.

(a) Compute the Carnot term alone, $\eta_C(T_{chw} + 273.15)/(T_{cond} - T_{chw})$. *Checkpoint:* 8.52.

(b) Compute the combined COP and the chiller's electric power for a duty of 2923 kW. *Checkpoint:* 6.06 and 482 kW.

(c) Raise the chilled water to 8 °C and find the percentage gain in COP. Explain why the error of the Carnot term alone grows at the operating points the supervisory layer selects. *Checkpoint:* 6.29, a gain of 3.8 %.

8.6 Theory and calculation · 7 min

The server fans follow Equation (8.3) with $T_{knee} = 25\,^\circ\mathrm{C}$, $k_f = 0.08/^\circ\mathrm{C}$, $\dot{m}_{IT,0} = 210\,\mathrm{kg/s}$ and $P_{f,0} = 180\,\mathrm{kW}$. The CRAH units deliver $\dot{m}_{s,0} f_{CRAH}$, with $\dot{m}_{s,0} = 260\,\mathrm{kg/s}$, and their fans draw $P_{CRAH,0} f_{CRAH}^3$ with $P_{CRAH,0} = 120\,\mathrm{kW}$.

(a) Find the fan speed, server airflow and server fan power at a 27 °C inlet. *Checkpoint:* 1.16, 243.6 kg/s, 281 kW.

(b) At the shipped $f_{CRAH} = 85\,\%$, find the inlet temperature above which R in (8.2) is positive. Why does a hall held at a 26 °C setpoint cycle? *Checkpoint:* 221 kg/s, 25.65 °C.

(c) Repeat (b) at $f_{CRAH} = 90\,\%$ and find what the extra CRAH fan power costs per hour at \$0.10 per kWh. *Checkpoint:* 234 kg/s, 26.43 °C, 13.8 kW or \$1.38 per hour.

8.7 App challenge · 8 min

Load Scenario 4, *The load is known before the temperature moves*, at the default seed and switch to the Advanced level. Make three runs to 60 min at a speed of 5 (five simulated minutes per second), each with *IT load feedforward gain* set before Run (Reset between runs keeps the new value): 0.3 as shipped, 0, and 0.6. At the end of each run, pausing between 60 and 65 min, read the IAE tile and the *Cost now* tile. Export the CSV (its first lines start with #, and the time column reads in hours) and find the largest positive and the most negative value of `inlet_error_C`, which is the cold aisle minus its setpoint.

(a) Give the IAE and the two extreme errors of the shipped gain and of gain 0, to 0.05 °C min and 0.02 °C. By what percentage does the feedforward cut the larger error? *Checkpoint:* 0.3: 2.59, $+0.27$ and -0.26. 0: 5.39, $+0.78$ and -0.77. About 65 %.

(b) Give the same three numbers for gain 0.6 and say when each extreme occurs. Explain the signs. *Checkpoint:* 4.98, $+1.02$ at about 36 min and -0.79 at about 6 min.

(c) Compare *Cost now* across the three runs and explain. *Checkpoint:* \$98.08 to \$98.10 per hour at 60 min and \$98.01 to \$98.02 at 65 min, the three runs within \$0.02 of one another at the same minute.

8.8 App challenge · 9 min

Load Scenario 8, *A chiller trips and capacity becomes the constraint*, at the default seed. Make three runs to 180 min at a speed of 20, with every setting made before Run and Reset between runs, which keeps them: (1) as loaded, (2) with *Cold-aisle setpoint* at 26 °C, and (3) with the Advanced level and *CRAH fan speed* also at 90 %. For each run read the *Binding*, *Cost now* and *Over 27 °C* tiles at the end, and from the exported CSV the lowest and highest `inlet_C` between 2.5 and 3.0 h.

(a) Give where run 1's hall settles and what the Binding tile reads. *Checkpoint:* 24.78 °C, no legal point.

(b) Give run 2's band of `inlet_C` and the range of `cost_usd_h` over the same window, to 0.05 °C and $0.10. *Checkpoint:* 25.46 to 26.54 °C, $79.38 to $84.08 per hour.

(c) Give run 3's cold aisle and *Cost now*, and explain the difference from run 2. *Checkpoint:* 26.00 °C, $78.63 per hour.

8.9 App challenge · 9 min

Load Scenario 13, *The same cost as a 30-minute plan (MPC)*, at the default seed and run it as loaded to 180 min at a speed of 20. Press Reset, set *Control architecture* to Optimizer before Run, and run the same 180 min. For each run export the CSV and read, in the row at 3.00 h, `energy_kWh`, `water_liters`, `cost_usd`, `inlet_C` and `spray_pct`.

(a) Give the energy, water and cost of each run, to 2 kWh, 20 L and $0.05. Which layer is cheaper over the three hours? *Checkpoint:* MPC 2322 kWh, 16 706 L, $265.65. Optimizer 2335 kWh, 15 878 L, $265.29.

(b) Give the cold aisle and the spray of each run after the trip, and the difference in cold aisle. *Checkpoint:* MPC 24.38 at 95.8 %, Optimizer 24.78 °C at 90 %, 0.40 °C.

(c) The plan is cheaper per hour before the trip. Why is it not cheaper over the run, and what does the Optimizer do after the trip that the plan does not?

8.10 Application area · 4 min

A site on a $2.00 per m^3 water tariff is asked to cut make-up water on mild days. Use the Scenario 11 results of Section 8.6.

(a) Compare pricing water at $6.00 per m^3 inside the objective with keeping $2.00 and capping water at 40 L/min, in cooling power and water.

(b) Name one fact from the chapter's Introduction that argues against judging the trade on the tower's water alone.

8.11 Application area · 3 min

A commissioning engineer sees a cold-aisle trend sitting on its 24 °C setpoint with no visible error, while the optimizer's *Above the optimum* tile holds at about $7.6 per hour. What does the chapter say this pattern indicates, what does it cost, and what check would find the cause?

8.12 Review the agent · 5 min

An AI assistant reviews the graded challenge and writes: *"The day fails because the single loop is detuned at 4 %/°C. Switch to + Feedforward with* $K_c = 8\,\%/^\circ\mathrm{C}$, $\tau_I = 3\,\mathrm{min}$ *and* $K_{ff} = 0.3$. *The worst error falls from 1.73 to 0.41 °C, and the tighter control will also bring the day under $2000 and 120* m^3 *of water."* Find the flaw and say how you would check it in the app.