Model resolution and the choice of cooling action in a data-center facility during a heatwave

Ricardo Marquez, Ph.D.
rmarquez123@gmail.com

September 2026

Abstract

A data center is cooled through a hierarchy — a plant, halls, rows of racks and the racks themselves — but the model behind an operator's dashboard resolves only part of that hierarchy, and the cooling action taken during a hot spell depends on how much of it the model resolves. This report asks, for one 96-hour heatwave at Santa Clara, CA, whether the resolution at which a modeled facility of 30 racks is read changes which cooling action is safe and what that action costs in electricity. The facility is simulated once as a transient heat-balance network of racks, cold and hot aisles, rooms and buildings on a supply-air temperature the outdoor weather sets, with each rack's power rising with its own temperature; the simulated temperatures are then read at four resolutions — facility rack-mean temperature, worst rack per hall reading, worst rack per row reading and every rack, reconciled against physics — and the action each resolution can justify is simulated in turn and priced with a staged plant of two code-minimum air-cooled chillers on the Gordon–Ng law. The mean over all racks peaks at 33.15°C, below the 36.85°C limit, while rack B3_r5 reaches 39.12°C, so the coarsest resolution recommends no action and misses a violation. The three finer resolutions each hold every rack under the limit, at 2.259, 1.783, 1.592 MWh of cooling electricity; the hall-wide action costs 41.9 % more than the rack-sized one. One reading, that attributed to rack B3_r4, stands 1.500 K above the temperature the simulation gives that rack, which is the offset written into it. The energy ordering of the three safe actions holds in all 192 combinations of the registered plant choices, with a smallest margin of 0.0611, but it rests entirely on the assumed penalty for depressing the supply setpoint: priced without that penalty the three actions differ by thousandths of a MWh and the ordering reverses. The study is a forward one — every input is stated and nothing is estimated from data — and costs five transient solves of 1,152 steps each.

1 Introduction

Data centers consume a few percent of global electricity, and a large part of that goes not into computing but into removing the heat computing produces; Masanet, Shehabi, Lei, Smith and Koomey (2020) recalibrate the global estimate and show how strongly it depends on how facilities are operated, and Barroso, Hölzle and Ranganathan (2019) describe the cooling overhead of warehouse-scale facilities. The heat is not spread evenly. Schmidt, Cruz and Iyengar (2005) survey the thermal management of data centers and document the local hot spots that arise from uneven rack loading and from air that leaves a hot aisle and returns to a cold one, and Sharma, Bash, Patel, Friedrich and Chase (2005) show that where heat is generated, and how the cooling responds, changes the energy a facility spends. An operator facing a forecast heat event must decide what cooling response to put in place, and the decision is made against a dashboard whose underlying model may resolve the facility to a single average, to a hall, to a row, or to every rack.

Work on this problem falls into a few groups. Detailed airflow simulation resolves the room to the scale of a rack face and is used to site equipment and to grade containment, at a computational cost that rules it out for an hour-by-hour decision. Reduced models of the kind used here write a heat balance on a coarse network of racks, aisles, rooms and plant and trade spatial detail for speed. Whatever the model, a decision made from sensor readings has to contend with the sensors: Katipamula and Brambley (2005) review fault detection and diagnosis for building systems and find sensor faults, and biases in particular, among the most common and the hardest to notice, because a biased reading looks exactly like a real excursion. The process industries answer this with data reconciliation, in which readings are corrected against the conservation laws the plant must obey and readings inconsistent with those laws are flagged; Crowe (1996) reviews the approach and its gross-error tests. Cooling cost, finally, is governed by the chiller, whose efficiency depends on its load and on the two water temperatures through the thermodynamic model of Gordon and Ng, which Chua, Ng and Gordon (1996) test against measurements, with the minimum efficiency of new machines fixed by ASHRAE Standard 90.1 (ASHRAE, 2019) at the rating conditions of AHRI Standard 550/590 (AHRI, 2023).

TERRA is a general-purpose engine for probabilistic inference on physical networks. It treats a network as a graph on which the physics is written locally: a conservation law at every node, a constitutive law on every edge, and the boundary and component laws that close them, attached to the same graph. Each of these is a relation the state of the network must satisfy, and the engine assembles them and solves them together. Domains enter through plugins that supply their laws, units and file formats; the solving and the bookkeeping are shared. This study uses its data-center plugin to write the heat balances of the racks, aisles, rooms and buildings of a two-hall facility whose supply-air temperature follows the outdoor weather, and its chiller-plant model to price the heat the plant removes. Section 3 states the framework and each of these constructs precisely.

This report makes no claim of novelty. It asks one question of one modeled facility during one real heatwave: does the resolution at which the facility is read change which cooling action is safe, and what does each coarser resolution cost? It answers by simulating the facility once with no action, reading the result at four resolutions, simulating and pricing the action each resolution can justify, and testing how far the resulting energy ordering depends on the choices made in modeling the plant. It is a forward study: every input is stated and nothing is estimated from data. Section 2 states the facility, the data and the model, Section 3 the methods, Section 4 the results, Section 5 their limits and Section 6 the answers; the appendices hold the reproduction recipe, the provenance of the inputs and the notation.

2 System and data

The facility. Figure 1 draws the modeled facility. It has two halls, each one room holding three rows of five racks, 30 racks in all. A row has a cold aisle, the air its racks take in, and a hot aisle, the air they exhaust; the row's recirculation ratio is the conductance of the path that carries hot-aisle air back to the cold aisle divided by that of the supply path from the room, and is zero for perfect containment. Every rack draws a base power of 4.29 kW plus 5 W for every kelvin its temperature exceeds 300 K, and the temperature no rack may exceed is 310 K (36.85°C). The halls are identical except in three stated respects: hall B moves 85 % of the nominal room airflow; its row B3 has a recirculation ratio of 0.45 against 0.08 in every other row and only 60 % of the nominal rack-face airflow, the pattern left by missing blanking panels; and rack B3_r5 in that row carries 1.7 times the load of its neighbors. The asymmetry is deliberately mild and deliberately placed, so that the coarsest reading of the facility is misleading and the question is whether a finer reading corrects

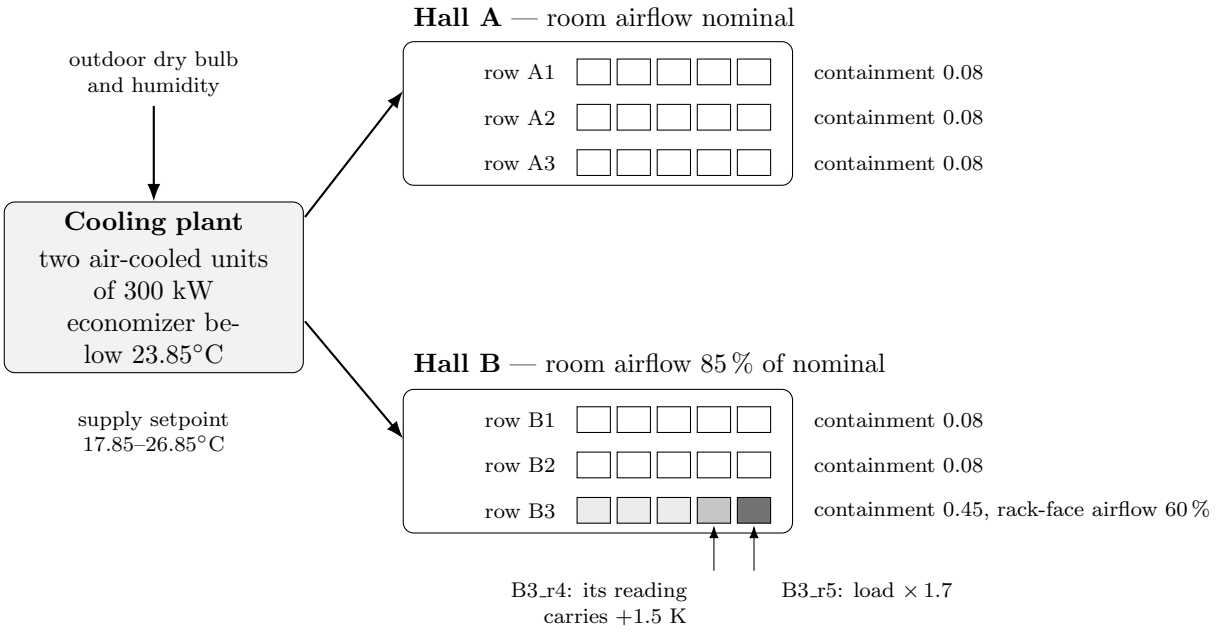


Figure 1: The modeled facility. One weather-driven cooling plant of two air-cooled chiller units of 300 kW serves two halls; each hall is one room holding three rows of five racks, 30 racks in all. The plant sets the supply-air temperature from the outdoor dry bulb and uses outdoor air alone while the dry bulb stays below 23.85°C. Small squares are racks. The halls are identical except for three stated differences: hall B moves 85 % of the nominal room airflow, its row B3 has poor containment (a recirculation ratio of 0.45 against 0.08 elsewhere, light gray) and 60 % of the nominal rack-face airflow, and rack B3_r5 in that row (dark) carries 1.7 times its neighbors' load. The reading attributed to rack B3_r4 (medium gray) is the simulated temperature of that rack plus a fixed offset of 1.5 K.

the decision.

The plant. One plant serves both halls. It sets the supply-air temperature from the outdoor dry bulb, following it with an approach of 3 K and holding the result between 17.85 and 26.85°C. Below an outdoor temperature of 23.85°C outdoor air alone holds that setpoint and the plant pays only for its fans and pumps, 3 % of the heat removed; above it the chillers take a share that grows linearly and reaches the whole load a band of 2 K higher. The chillers are two identical air-cooled units of 300 kW (85.3 tons of refrigeration) each, and each is the least efficient machine that meets both ASHRAE 90.1 minimums for its class and size. They condense against the outdoor dry bulb, subject to a head-pressure floor, and a second unit is staged on when the load passes 90 % of the running unit's capacity.

The data. The weather is real. One hourly record for Santa Clara, CA (37.37°N, 121.92°W) for 2023, 8,751 hours in all, is taken from the ERA5 reanalysis (Hersbach et al., 2020) through a public archive, and the study uses the 96-hour window beginning at row 6,631 of that file, from 2023-10-04 15:00:00+00:00 to 2023-10-08 14:00:00+00:00, which holds the year's highest temperature, 36.7°C. The rack load is not measured. Its magnitude is set by scaling the regional electricity demand of the same hours, which ranges over the window from 20,663 to 36,715 MW, linearly into a per-rack band of 2,600 to 4,400 W, and then holding the value at the hour of the outdoor peak, 4.29 kW, for

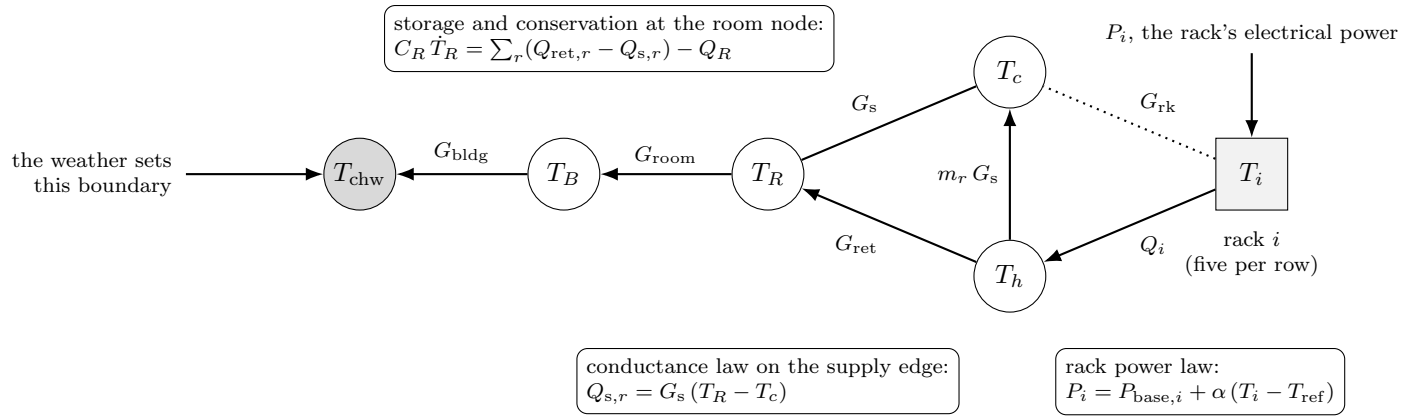


Figure 2: The model of one row with its room and building. Circles are unknown temperatures, the shaded circle the chilled-water boundary the weather sets, and the square a rack; arrows are heat flows on conductance edges. The dotted edge draws the rack's heat from the air it takes in, the cold aisle, and the solid edge below it delivers that heat to the hot aisle, from which one part returns to the room and another recirculates to the cold aisle through the mixing edge. The relations of the model are of four kinds, three of them written out here and the fourth marked by the arrow on the left: a stored-energy balance at every node, a conductance law on every edge, the temperature boundary the plant imposes, and the rack power law, which makes a rack draw more power as it warms.

the whole window. The load is therefore constant and the weather is the only time-varying driver. No rack, aisle, room or plant measurement enters the study, and no random numbers are drawn.

The readings. What the report calls a reading is a quantity constructed from the simulation, not a measurement. The simulation gives the true temperature of every rack at every step. The reading attributed to a rack is that temperature unchanged, except for rack B3_r4, to whose reading a fixed offset of 1.5 K is added throughout. That rack is the neighbor of the one truly at risk, and the offset is large enough to put its reading over the limit while the rack itself stays below. The purpose of the construction is to ask what each resolution does with a reading the physics of the facility does not support; it is not an estimate of a sensor bias from noisy data.

The model. Figure 2 shows the model of one row with its room and building. The unknowns are the temperature of every rack, aisle, room and building, the electrical power of every rack, and the heat flow on every edge; every rack, aisle, room and building node also carries a stored energy, so the model is transient and is advanced in time from a steady start. The plant is not part of the network: the supply-air temperature it delivers is a boundary the weather sets, and the electricity it draws is priced from the heat the halls reject after the network has been solved.

Assumptions.

- Units: temperatures in K in the equations and in degrees Celsius in the text and figures, powers and heat flows in W, conductances in W/K, thermal capacities in J/K, energies in MWh over the window.
- The rack load is constant at 4.29 kW per rack and identical in every scenario, so no action changes the computing work done.

- The supply-air temperature is a boundary imposed on both halls; it does not respond to the halls' state, and an action that depresses it is delivered exactly, subject only to a floor of 11.85°C.
- Rooms and rows within a hall are identical except where Section 2 states otherwise, and the racks of a row are identical except for the one dense rack.
- The plant condenses against the outdoor dry bulb with no approach beyond the head-pressure floor, and both halls' heat is priced on one combined duty.
- The cost of depressing the setpoint is charged as a fixed electrical penalty per kelvin of depression, stated in Section 3 and not derived from the plant model.
- The actions are the four listed in Section 3; no action is optimized, and the cost of making any of them possible is outside the model.

The question. Does the resolution at which the facility is read change which cooling action is safe, and how much electricity does each resolution's action cost? Two supporting questions follow: what does a reading the physics does not support cost a resolution that cannot check it, and how far does the resulting energy ordering depend on the choices made in modeling the plant and in charging the setpoint penalty?

3 Methods

3.1 The thermal network of a hall

A hall is written as a network of nodes carrying temperatures and edges carrying heat flows. The unknowns of one hall are, for every rack i, its temperature T_i, its electrical power P_i and the heat flow Q_i it passes to the air; for every row r, the cold-aisle and hot-aisle temperatures $T_{c,r}$ and $T_{h,r}$ and the supply, recirculation and return flows $Q_{\mathrm{s},r}$, $Q_{\mathrm{m},r}$ and $Q_{\mathrm{ret},r}$; and, for the hall, the room temperature T_R with the flow Q_R it passes to the building node, and the building temperature T_B with the flow Q_B it rejects to the supply air at T_{chw}. Each node stores energy in proportion to its temperature, so its balance is a rate equation; each edge carries a heat flow proportional to the temperature difference across it. Stored energy and conservation at the rack (1), the rack power law (2), the conductance of its heat path (3), the two aisles (4) and (5), the supply, recirculation and return paths (6)–(8), and the room and building (9)–(12) give

$$C_{\mathrm{rk}}\,\dot{T}_i = P_i - Q_i, \tag{1}$$

$$P_i = P_{\mathrm{base},i} + \alpha\,(T_i - T_{\mathrm{ref}}), \tag{2}$$

$$Q_i = G_{\mathrm{rk},r}\,(T_i - T_{c,r}), \tag{3}$$

$$C_{\mathrm{a}}\,\dot{T}_{c,r} = Q_{\mathrm{s},r} + Q_{\mathrm{m},r}, \tag{4}$$

$$C_{\mathrm{a}}\,\dot{T}_{h,r} = \textstyle\sum_{i\in r} Q_i - Q_{\mathrm{m},r} - Q_{\mathrm{ret},r}, \tag{5}$$

$$Q_{\mathrm{s},r} = G_{\mathrm{s}}\,(T_R - T_{c,r}), \tag{6}$$

$$Q_{\mathrm{m},r} = m_r\,G_{\mathrm{s}}\,(T_{h,r} - T_{c,r}), \tag{7}$$

$$Q_{\mathrm{ret},r} = G_{\mathrm{ret}}\,(T_{h,r} - T_R), \tag{8}$$

$$C_R\,\dot{T}_R = \textstyle\sum_r (Q_{\mathrm{ret},r} - Q_{\mathrm{s},r}) - Q_R, \tag{9}$$

$$Q_R = G_{\mathrm{room}}\,(T_R - T_B), \tag{10}$$

$$C_B\,\dot{T}_B = Q_R - Q_B, \tag{11}$$

$$Q_B = G_{\text{bldg}}\,(T_B - T_{\text{chw}}). \tag{12}$$

where a dot is a time derivative and the sums run over the racks of a row and the rows of a room. The supply flow $Q_{\text{s},r}$ is counted from room to cold aisle and is negative when the aisle is the warmer of the two; the recirculation path carries heat from the hot aisle back to the cold aisle with a conductance m_r times the supply's, which is how poor containment raises the air the racks take in; and the rack's heat flow is driven by its temperature above that same cold aisle, which is why the rack-face airflow of a row enters through $G_{\text{rk},r}$. Every constant of (1)–(12) is listed with its value in Table 5, and the departures from those nominal values are the three of Section 2: the room conductance of hall B is scaled by 0.85, row B3 takes $m_r = 0.45$ and has its rack-face conductance scaled by 0.6, and one rack of that row takes 1.7 times the base power of the rest. Each hall is solved on its own, since the two halls meet only at the plant and the plant does not move the supply temperature; the system is started from its steady state at the first hour's weather and advanced over the window with a second-order implicit method at a fixed step of 300 s, 1,152 steps, each step solved to a residual tolerance of 0.001 in at most 40 iterations.

3.2 The network model

The relations of the previous subsection are one instance of a general construction, in which the unknowns of a network and the relations they must satisfy are written on its graph and solved together; this subsection states it. A physical network is a graph, and the engine writes its physics on that graph. The state is a vector z of unknowns: one potential per node, such as a temperature or a head, one flow per edge, and whatever else the physics determines. Every relation the state must satisfy — a conservation law at a node, a constitutive law on an edge, a boundary condition, a component law — is written as a residual $r_k(z)$ that vanishes when the relation holds. With as many independent relations as unknowns and no readings to reconcile, the state is the unique solution of

$$r(z) = 0, \tag{13}$$

found by Newton iteration on the sparse Jacobian J of the residuals, $J_{ki} = \partial r_k / \partial z_i$, from a stated start until the residual norm falls below a tolerance ϵ:

$$z \leftarrow z - J^{-1} r(z) \quad \text{until } \|r(z)\| < \epsilon. \tag{14}$$

A transient model adds a storage term to the conservation relations and advances the state in time with an implicit integrator, solving a system of the same kind at every step. Nothing in this report is estimated from data: every input is stated, and the results are the solutions of these relations and quantities computed from them. Everything that follows is a choice of unknowns and relations on this graph.

3.3 The supply-air policy and the cost of cooling

The boundary T_{chw} of (12) follows the outdoor dry bulb by an approach a, bounded on both sides (15); an action may ask for a depression δ of that setpoint over a stated window, subject to a floor T_{fl} (16), and the depression actually delivered, which the floor can cut short, is Δ (17),

$$T^0(t) = \min\big(T_{\max},\ \max(T_{\min},\ T_{\text{db}}(t) + a)\big), \tag{15}$$

$$T_{\mathrm{chw}}(t) = \max\big(T_{\mathrm{fl}},\ T^{0}(t) - \delta\big), \tag{16}$$

$$\Delta(t) = T^{0}(t) - T_{\mathrm{chw}}(t). \tag{17}$$

with $a = 3$ K, $T_{\min} = 291$ K, $T_{\max} = 300$ K and $T_{\mathrm{fl}} = 285$ K, and $\delta = 0$ outside an action's window.

The heat both halls reject is removed by one plant. A chiller unit removing heat q from chilled water entering its evaporator at T_{chi}, with condenser air entering at T_{ci}, draws compressor power given by the Gordon–Ng law, derived from the first and second laws applied to the refrigeration cycle and tested against measurements by Chua et al. (1996). It carries three loss coefficients: b_1, the internal entropy generation, b_2, an equivalent heat leak, and b_3, the total heat-exchanger resistance. Solved for the power it gives (18); below a minimum stable load $q_{\min}$, a fraction 0.1 of the unit's capacity, the unit cycles and its time-averaged power is the linear (19); and its condenser cannot be driven below the head-pressure floor (20), a function F of the unit's load fraction x,

$$P_{\mathrm{GN}}(q) = \frac{T_{\mathrm{ci}}\, q + b_1 T_{\mathrm{chi}} T_{\mathrm{ci}} + b_2\,(T_{\mathrm{ci}} - T_{\mathrm{chi}})}{T_{\mathrm{chi}} - b_3\, q} - q, \tag{18}$$

$$P_{\mathrm{cyc}}(q) = \frac{q}{q_{\min}}\, P_{\mathrm{GN}}(q_{\min}), \tag{19}$$

$$T_{\mathrm{ci}}(t) = \max\big(T_{\mathrm{db}}(t),\ F(x(t))\big). \tag{20}$$

with $b_1 = 53.13$ W/K, $b_2 = 82{,}635$ W and $b_3 = 5.815 \times 10^{-5}$ K/W. Chilled water enters each unit at its rating temperature, 285.37 K, and the floor starts at 313.15 K at zero load and falls by 16.7 K per unit of load fraction. The three coefficients are those of a machine that just meets both ASHRAE 90.1 minimums for an air-cooled chiller under 150 tons, COP 2.96 at full load and COP 4.02 on the integrated part-load value; the part-load minimum is the binding one, and at the full-load rating point the unit reaches COP 4.316.

The plant is then priced step by step. The duty is the sum of the two halls' rejected heat, each clamped at zero so that a hall briefly giving heat back cannot cancel the other's load (21); the share the chillers carry grows linearly across the crossover band (22); units are staged by threshold and share their duty equally (23)–(24); the share carried by outdoor air costs a parasitic fraction $f_{\mathrm{par}} = 0.03$ of the heat removed (25); the whole is multiplied by a penalty for running the setpoint below where the weather would put it (26); and the electricity over the window is the integral (27), evaluated by the trapezoidal rule on the simulation's own time grid,

$$Q(t) = \textstyle\sum_h \max\big(0,\ Q_B^{h}(t)\big), \tag{21}$$

$$\varphi(t) = \min\Big(1,\ \max\Big(0,\ \tfrac{T_{\mathrm{db}}(t) - (T_{\max} - a)}{\Delta_{\mathrm{b}}}\Big)\Big), \tag{22}$$

$$n_{\mathrm{run}}(t) = \min\big(n,\ \lceil \varphi Q / (0.90\, C) \rceil\big), \tag{23}$$

$$P_{\mathrm{ch}}(t) = n_{\mathrm{run}}\, P_{\mathrm{unit}}\big(\varphi Q / n_{\mathrm{run}}\big), \tag{24}$$

$$P_{\mathrm{cool}}(t) = (1 - \varphi)\, f_{\mathrm{par}}\, Q + P_{\mathrm{ch}}, \tag{25}$$

$$P(t) = P_{\mathrm{cool}}(t)\,\big(1 + \beta\, \Delta(t)\, w(t)\big), \tag{26}$$

$$E = \int_0^{\tau} P(t)\, \mathrm{d}t. \tag{27}$$

with $\Delta_{\mathrm{b}} = 2$ K, $n = 2$, $C = 300$ kW and τ the length of the window. The penalty coefficient is $\beta = 0.12$ per kelvin. It stands for the extra work a chiller does to deliver colder air, which does not appear in the duty, since in steady state the heat removed is the computing load whatever the setpoint. The weight w says how much of the plant's duty the penalty is charged on: the primary choice charges it on the whole plant, $w = 1$, on the grounds that one loop serves both halls, and the alternative charges it only on the precooled hall's share of the duty. Both are stated choices, not derived ones, and Section 4 reports what each is worth.

3.4 The four resolutions and their actions

The facility is first simulated with no action, which gives the true temperature of every rack at every step. That one simulation is then read at four resolutions. The coarsest reads the mean of all 30 rack temperatures, a single number for the facility. The next reads the highest reading in each hall, the next the highest reading in each row, and the finest reads every rack and compares each reading with the temperature the physics gives that rack. A resolution detects a violation when the quantity it reads crosses the limit, and it can place a violation only as precisely as the quantity it reads: the coarsest cannot place it at all, the second places it in one of two halls, the third in one of 6 rows, and the finest in one rack.

The reading attributed to rack i and the residual of that reading against the physics are

$$y_i = \max_t T_i(t) + b_i, \tag{28}$$

$$\rho_i = y_i - \max_t T_i(t). \tag{29}$$

with b_i the offset written into rack i's reading, zero for every rack but one. By construction ρ_i returns b_i exactly. This is a consistency check between two quantities the study itself produces, and it says only that a reading inconsistent with the simulated state can be separated from a rack that is genuinely hot when the state of every rack is available; it is not the recovery of an unknown bias from noisy measurements, and no noise is present.

Each resolution is then given the action it can justify from what it reads, and that action is simulated over the whole facility and priced by (21)–(27). The coarsest reads nothing above the limit and takes no action. The hall resolution knows only that hall B holds a reading above the limit and depresses that hall's supply setpoint by $\delta = 3.5$ K for the whole window. The row resolution knows that the readings above the limit are both in row B3, restores that row's containment and rack-face airflow to nominal, and, because it cannot tell which of the two readings is real, adds a precautionary trim of $\delta = 1$ K to the hall setpoint. The rack resolution restores the same row and takes no setpoint action, because the reading it would have acted on is the one the residual (29) rejects. The actions are stated, not optimized; no search over depressions or airflow settings is performed.

3.5 Sensitivity of the energy ordering

The energy ordering of the safe actions rests on the plant model and on how the setpoint penalty is charged, and both involve choices a reader could make differently. Six are varied. The loss scale multiplies the anchored unit's three coefficients by 0.8 or 1.2, or leaves them as they are. The staging rule is the threshold rule (23) or least power, which runs the number of units that draws least among those that can carry the load. The unit count is n units of 300 kW or $2n$ units of half that, each re-anchored to the code minimums for its size. The head-pressure floor takes one of four registered forms: 40°C SCT target (313.15 K at zero load, falling by 16.7 K per unit of load fraction); no floor (the condenser follows the weather with no floor at all); pressure ratio 2.8 (311.75 K at zero load, falling by 16.7 K per unit of load fraction); load-independent 296.45 K (a fixed 296.45 K). The chilled-water model holds the water entering each unit at its fixed rating temperature or lets it fall with load under constant evaporator flow. The load share w of (26) charges the penalty on the whole plant or on the precooled hall's share of the duty. The product of the levels gives 192 combinations. In each, the solved thermal duties are held fixed and re-priced, and the ordering is scored by

$$c_1 = (E_2 - E_3)/E_4, \tag{30}$$

$$c_2 = (E_3 - E_4)/E_4, \tag{31}$$
$$g = \min(c_1, c_2). \tag{32}$$

with E_2, E_3 and E_4 the energies of the hall, row and rack actions. Both clauses are positive while the ordering holds, and the margin g turns negative as soon as either gap closes. The table reports the registered plant, each combination that differs from it in exactly one choice, and the combination with the smallest margin; ties among combinations would go to the one with the lowest level indices in the order the choices are listed here. This is a sensitivity analysis over modeling choices; it is not a statement about how uncertain any of these numbers is.

3.6 Metrics

The safety question is answered by the worst true rack temperature under an action, in degrees Celsius: the maximum over every rack and every step of the simulated rack temperature, read against the limit of 36.85°C, which is the number the action exists to respect. It is paired with two counts, both read against zero: missed violations, racks above the limit that the resolution does not act on, and false alarms, readings above the limit that the resolution acts on although the rack is below it. The cost question is answered by the cooling electricity E of (27), in MWh over the 96-hour window, which is what the meter records; its reference is the cost of doing nothing, 1.593 MWh, since an action that costs no more than that is free at the meter. The resolution question is answered by the peak each resolution reads, in degrees Celsius against the same limit, and by how finely it can place a violation, against the 30 racks of the facility. The reading check is answered by the residual ρ of (29), in kelvin, whose reference is the offset written into the reading. The robustness of the energy ordering is answered by the margin g of (32), dimensionless, read against zero, with its size against the gaps themselves saying how close the ordering is to turning.

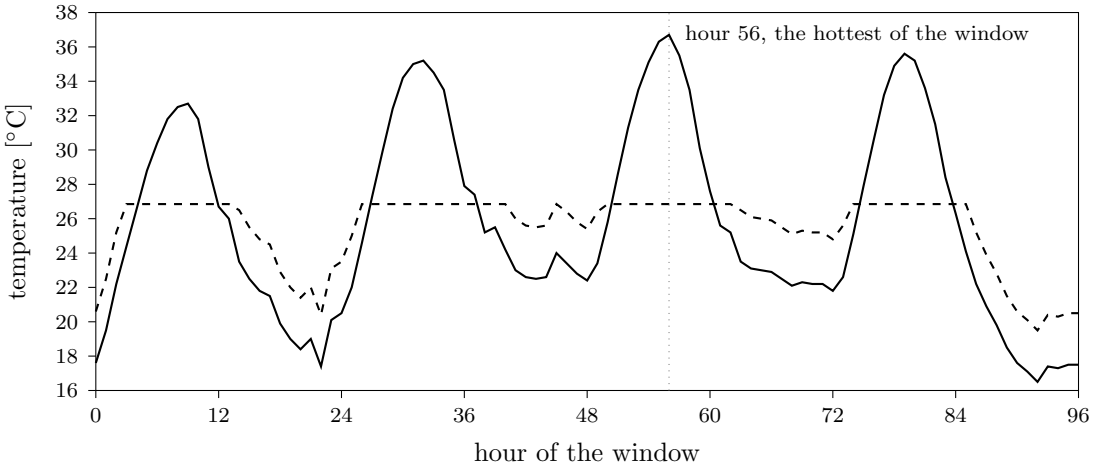
4 Results

4.1 The event

Figure 3 plots the two outdoor drivers over the window. The solid curve is the dry-bulb temperature, the dashed curve the supply-air setpoint the plant delivers under (15), both in degrees Celsius against the hour of the window; the dotted vertical line marks the hottest hour. The window holds three successive hot afternoons, rising to 36.7°C at hour 56 and falling to 16.5°C at night. The setpoint tracks the dry bulb while the weather is cool and sits at its upper bound of 26.85°C through each afternoon, so the racks get no air-side relief exactly when the outdoor temperature is highest, and the plant runs its chillers for those hours. Of the 1,153 time points at which the plant is priced under the cheapest action, the chillers carry some share of the load at 468.

4.2 Rack temperatures with no action

Table 1 gives the peak temperature every rack reaches over the window with no action taken, grouped by the racks that share a peak, with the true peak, the reported peak and each against the limit. Racks in hall A peak at 32.30°C and racks in the two well-contained rows of hall B at 32.71°C, the difference being that hall's reduced room airflow. The other four racks of row B3 peak at 35.97°C, 3.3 K above their neighbors in the same hall, because the recirculation path of (7) returns hot-aisle air to their inlets and the reduced rack-face conductance of (3) makes each rack sit further above the air it draws. Rack B3_r5, which carries 1.7 times their load in the same degraded



solid: outdoor dry bulb; dashed: supply-air setpoint the plant delivers; dotted vertical: the hour of the outdoor peak

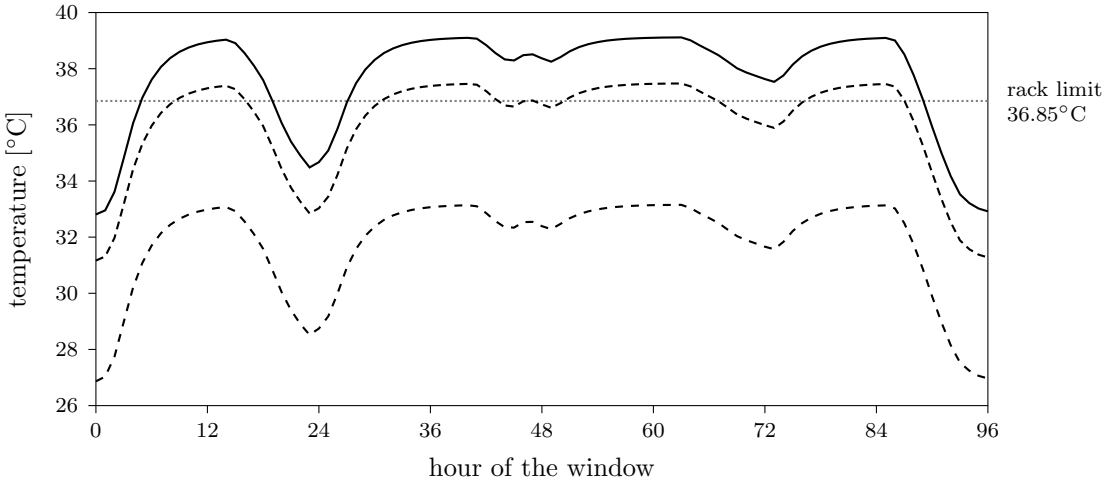
Figure 3: The 96-hour window that drives the study, from 2023-10-04 15:00:00+00:00 to 2023-10-08 14:00:00+00:00. The solid curve is the outdoor dry-bulb temperature at Santa Clara, CA, which reaches 36.7°C at hour 56 and falls to 16.5°C at its coolest; the dashed curve is the supply-air setpoint the plant delivers, the dry bulb plus an approach of 3 K held between 17.85 and 26.85°C. The setpoint sits at its upper bound through the three afternoons of the event, so the racks get no relief from the air side when the outdoor temperature is highest.

Table 1: Peak temperature of every rack over the window with no action taken, grouped by the racks that share a peak. The limit is 36.85°C. The true peak is the maximum of the simulated rack temperature; the reported peak is the same number except for rack B3_r4, whose reading carries a fixed offset of 1.5 K. Of the 30 racks, one exceeds the limit and two readings do.

racks	count	true peak [°C]	reported peak [°C]	true vs limit	reading vs limit
hall A, rows A1–A3	15	32.30	32.30	under	under
hall B, rows B1–B2	10	32.71	32.71	under	under
row B3, its other three racks	3	35.97	35.97	under	under
row B3, rack B3_r4, the biased reading	1	35.97	37.47	under	over
row B3, rack B3_r5, the dense tenant	1	39.12	39.12	over	over

row, reaches 39.12°C and is the one rack above the limit. The mean over all 30 racks peaks at 33.15°C, 3.70 K below the limit, because 29 of the 30 racks are comfortable and the average of a set is insensitive to one member of it.

Figure 4 plots the same result through time: the dashed curve is the mean over all racks, the solid curve the true temperature of rack B3_r5, the dash-dot curve the reading attributed to rack B3_r4, and the dotted horizontal line the limit. The mean stays below the limit for the whole window while the hot rack is above it through each of the three afternoons. Because the hot rack is the hottest in its row and the hottest in its hall, the hottest-per-hall and hottest-per-row traces coincide with the solid curve exactly and are not drawn separately; the resolutions between the mean and the rack differ here not in the peak they read but in where they can place it.



solid: the true temperature of rack B3_r5, which is also the hottest rack of its hall and of its row;
dash-dot: the reading attributed to rack B3_r4; dashed: the mean over all 30 racks

Figure 4: What each resolution reads through the event, with no action taken. The dotted horizontal line is the rack temperature limit, 36.85°C. The dashed curve is the mean over all 30 rack temperatures, which peaks at 33.15°C and never approaches the limit. The solid curve is the true temperature of rack B3_r5, which reaches 39.12°C; because that rack is the hottest in its row and in its hall, the hottest-per-hall and hottest-per-row traces lie exactly on top of it and are not drawn separately. The dash-dot curve is the reading attributed to rack B3_r4, the simulated temperature of that rack plus 1.5 K, which crosses the limit although the rack itself peaks at 35.97°C.

4.3 What each resolution reads

Table 2 gives, for each resolution, the quantity it reads, the peak that quantity reaches, whether that peak crosses the limit and how finely the resolution can place a problem. The coarsest reads 33.15°C and does not cross the limit, against a true peak of 39.12°C somewhere in the facility: it misses the violation entirely. The other three all read 39.12°C and cross it. What separates them is placement and the treatment of the readings. Of the 30 racks, one is truly above the limit and two readings are above it, the extra one being rack B3_r4, whose reading peaks at 37.47°C while the rack itself peaks at 35.97°C. The residual (29) of that reading against the physics is 1.500 K, which is the offset written into it; every other reading has a residual of zero. Only the finest resolution has the rack-level state needed to form that residual at all, so only it can separate the two readings above the limit. The separation is exact here because the comparison is between a constructed reading and the simulation that produced the temperature underneath it.

4.4 The actions and what they cost

Table 3 gives the action each resolution can justify, the worst true rack temperature under it, its cooling electricity over the window, and the false-alarm and missed-violation counts. Figure 5 shows the energies as bars, dark for the cost as priced and light for the same duties priced with the setpoint penalty removed. Doing nothing leaves the worst rack at 39.11°C, above the limit, and misses one violation; it is the only unsafe action and its cost, 1.593 MWh, is the reference the others are read against. Each of the three finer actions holds every rack under the limit, at 35.58, 33.65, 34.65°C respectively. Their costs are 2.259, 1.783, 1.592 MWh: the hall-wide depression costs 41.9 % more than the rack-sized action and the row action with its precautionary trim 12.0 %

Table 2: The four resolutions at which the one simulation is read, the quantity each reads, the peak that quantity reaches over the window, whether that peak crosses the 36.85°C limit, and how far the resolution can place the problem. Only the coarsest, which reads the mean over all 30 racks, fails to cross the limit; the three finer resolutions read the same peak, 39.12°C, and differ only in where they can place it and in whether they can tell a hot rack from a biased reading.

resolution	quantity read	peak read [°C]	crosses the limit	localizes to
L1 building	facility rack-mean temperature	33.15	no	none — one number for the whole site
L2 hall	worst rack per hall (reported)	39.12	yes	hall (which of 2 halls)
L3 row	worst rack per row (reported)	39.12	yes	row (which of 6 rows)
L4 rack	every rack, reconciled against physics	39.12	yes	rack + sensor validity

Table 3: The action each resolution can justify, simulated over the whole facility and priced. The scope column says what the action touches; the worst rack is the maximum simulated rack temperature anywhere in the facility under that action, against a limit of 36.85°C; the energy is the cooling electricity over the 96-hour window. A false alarm is a reading above the limit that the resolution acts on although the rack is below it; a missed violation is a rack above the limit that the resolution does not act on.

resolution	scope of the action	worst rack [°C]	verdict	energy [MWh]	false alarms	missed violations
L1 building	none	39.11	unsafe	1.593	0	1
L2 hall	whole hall (2 rows that did not need it)	35.58	safe	2.259	1	0
L3 row	one row (+ precautionary hall trim)	33.65	safe	1.783	1	0
L4 rack	one row, sized to one rack	34.65	safe	1.592	0	0

more, while the rack-sized action costs 0.0006 MWh less than doing nothing, because restoring the row's airflow lowers every rack in it and so lowers the heat the plant must remove. That the rack-sized action never demands more of the plant than doing nothing holds at every step of the window and in every one of the 192 pricing combinations.

The false-alarm counts separate the two middle resolutions from the finest. The hall resolution acts on the whole of hall B, including the two rows that never needed it, and cannot tell that one of the two readings it responds to belongs to a rack that is below the limit. The row resolution narrows the action to the row that is genuinely degraded but still carries the same unresolved reading, which is what its precautionary 1 K trim pays for; that trim is the whole of the 12.0 % it spends over the rack-sized action. The finest resolution rejects the reading and omits the trim.

4.5 Where the energy ordering comes from

The light bars of Figure 5 price the same solved duties with the setpoint penalty of (26) set to zero. They are 1.5928 MWh for the L1 action, 1.5909 MWh for the L2 action, 1.5916 MWh for the L3 action, 1.5922 MWh for the L4 action. Without the penalty the three safe actions lie within 0.0013 MWh of one another and the ordering reverses: the cheapest becomes L2, the hall-wide depression, and the dearest L4. The reason is that colder supply air lowers every rack temperature and, through

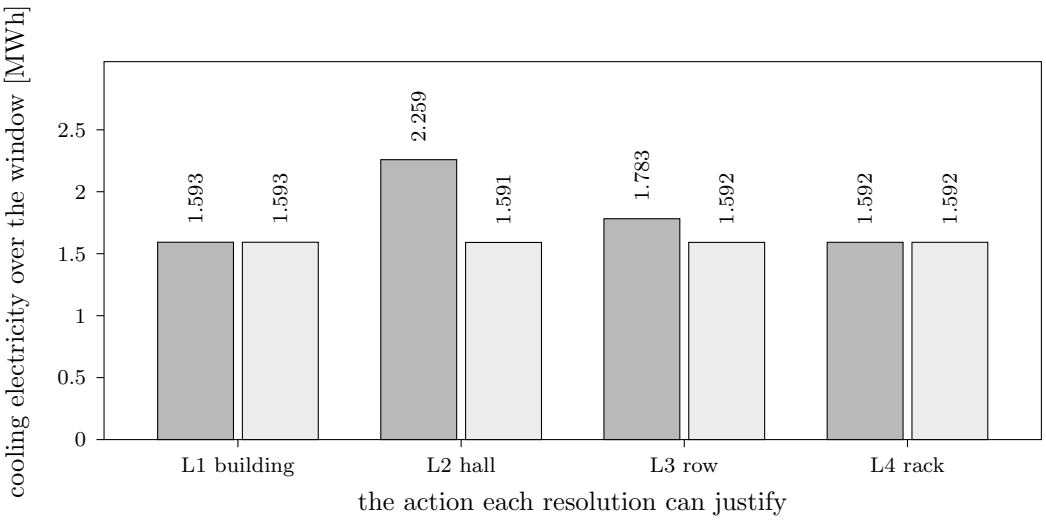


Figure 5: Cooling electricity over the 96-hour window under each of the four actions. Dark bars are the cost as priced: the plant's own electrical draw multiplied by a penalty of 0.12 per kelvin of setpoint depression delivered. Light bars are the same solved duties priced with that penalty set to zero. The ordering the study reports, L2 above L3 above L4, is present only in the dark bars; with the penalty removed the three actions cost within a few thousandths of a MWh of each other and the ordering reverses.

the rack power law (2), lowers the heat the plant has to remove; priced on duty alone, over-cooling is very slightly cheaper than not over-cooling. The entire energy case against the coarser actions therefore lives in the penalty coefficient $\beta = 0.12$ per kelvin, a stated assumption rather than a consequence of the plant model. For comparison, the registered plant's own chiller law gives a rise in chiller power of only 1.57 to 2.33 percent per kelvin of colder chilled water over the running steps of the window, so the assumed penalty is roughly 5 to 8 times the plant's own sensitivity to that one term. Section 5 returns to what this does and does not license.

4.6 Sensitivity of the energy ordering to the plant model

Table 4 lists the two clauses of (30) and (31) and the margin (32) for the registered plant, for each of the nine combinations that differ from it in exactly one choice, and for the combination with the smallest margin among all 192. The ordering holds in every one of them: no row of the table reverses it, and the margin never falls below 0.0611. At the registered plant the margin is 0.1196, which is the second clause, the smaller of the two gaps: measured in units of the rack action's energy, the row action stands 12.0 % above the rack action and the hall action a further 29.9 % above the row action. The five choices that touch only the plant — the loss scale, the staging rule, the unit count, the head-pressure floor and the chilled-water model — move the margin by less than 0.00015 in either direction, because they change what the plant charges for a given duty but not the depressions the actions ask for. The one choice that matters is how the penalty is charged: restricting it to the precooled hall's share of the duty halves the margin, to 0.0615, and that choice is also the one present in the worst combination. The sensitivity table therefore says the same thing as the previous subsection: the ordering is robust to the plant and fragile in the penalty.

Table 4: Sensitivity of the energy ordering of the three safe actions to the choices made in modeling the plant and in charging the setpoint penalty. Each row holds the solved thermal duties fixed and re-prices them with a plant differing from the registered one in the listed choices: the registered plant, each single change, and the combination of all 192 whose margin is smallest. The clauses c_1 and c_2 are the two gaps of the ordering, each expressed as a fraction of the cheapest action's energy, and the margin is the smaller of them; a negative margin would mean the ordering no longer holds. This is a sensitivity analysis over modeling choices, not a statement about how uncertain the result is.

plant choices changed	c_1	c_2	margin g	ordering
none (the registered plant)	0.2992	0.1196	0.1196	holds
loss scale 0.8	0.2992	0.1196	0.1196	holds
load share w $q_B/(q_A+q_B)$	0.1537	0.0615	0.0615	holds
chilled-water model constant_flow	0.2992	0.1196	0.1196	holds
heat-sink floor no floor	0.2991	0.1196	0.1196	holds
heat-sink floor pressure ratio 2.8	0.2992	0.1196	0.1196	holds
heat-sink floor load-independent 296.45 K	0.2991	0.1196	0.1196	holds
unit count 2n	0.2989	0.1195	0.1195	holds
staging rule least_power	0.2992	0.1196	0.1196	holds
loss scale 1.2	0.2992	0.1196	0.1196	holds
joint worst of 192: loss scale 0.8, unit count 2n, heat-sink floor load-independent 296.45 K, load share w $q_B/(q_A+q_B)$	0.1530	0.0611	0.0611	holds

4.7 The range the chiller law is used over

The shape of the chiller law is a fit to the published test data of one water-cooled centrifugal machine, and that fit covers a stated range of conditions: chilled water entering between 279.0 and 289.6 K, condenser water entering between 288.6 and 303.2 K, and unit load fractions between 0.24 and 1.18. Of the 1,153 time points priced under the cheapest action, 685 have no mechanical load at all and 468 run the chillers. Every one of those 468 running points lies outside the fitted range, all of them on the condenser side, and 433 of them also on the lift between the two water temperatures; 143 sit below the lowest fitted load fraction. The head-pressure floor (20) sets the condenser temperature at 386 of them, which is 80.0 % of the chiller energy of that action. The chiller power of this study is therefore an extrapolation of the fitted law throughout, not an interpolation within it.

5 Discussion

The safety result and the cost result rest on different ground, and they should be read differently. That the mean over 30 racks can sit 3.70 K below a limit that one rack exceeds by 2.27 K is arithmetic, not a property of this particular facility: an average is insensitive to one member of the set it averages, and the coarser the reading the more members there are to dilute an excursion. The energy ordering is different. It follows from the assumed penalty for depressing the supply setpoint, and nothing in the plant model produces it: priced on duty alone, the hall-wide depression is the cheapest of the three safe actions, not the dearest. What the study establishes about cost is conditional — if depressing the setpoint costs β per kelvin, then a narrower action is cheaper in proportion to the depression it avoids — and the sensitivity table measures the conditional, not the premise. A reader who does not accept the premise should read only the safety and false-alarm

columns of Table 3.

The facility is constructed, and constructed for the purpose. Its conductances, thermal capacities and leakage coefficient are illustrative rather than derived from measured airflow, and they put well-contained racks near 32°C and the worst rack near 39°C, which read as exhaust-air temperatures rather than as inlet temperatures and are high for either. The three departures from symmetry were placed so that the coarsest reading would be misleading; a facility with its load spread evenly would not produce this result, and the study says nothing about how often real facilities resemble one case or the other. The load is held at the value of the hottest hour for the whole window, so the weather is the only thing that moves and no diurnal load cycle can shift a rack across the limit on its own. One site, one year and one 96-hour window are used.

The treatment of the readings is the part most easily over-read. The simulation produces the temperature of every rack; the reading attributed to rack B3_r4 is that temperature plus a fixed 1.5 K, and the residual of (29) returns that offset exactly because it differences the reading against the very quantity the offset was added to. No noise is present, no sensor model is fitted, and no threshold has to be chosen. The result therefore shows only that the finest resolution is the one that has the rack-level state a residual needs, and that the coarser ones have nothing to check a reading against. It does not measure how large a bias could be detected against real sensor noise, nor how the check would behave when the model itself is wrong — the case that matters in practice, and the one that makes sensor faults hard, as Katipamula and Brambley (2005) report. A study that claimed detection would need readings with a stated noise level, a model error budget, and a decision rule with its false-positive rate.

The plant model has its own limits. Its units are anchored to code minimums rather than to a measured machine, the law's shape comes from a water-cooled centrifugal chiller and is carried by similarity to air-cooled units of 85.3 tons, and every running step of the window lies outside the range the fit covers, most of them with the head-pressure floor setting the condenser temperature rather than the weather. Staged unloading of the compressors and the losses of short cycling are not represented. The supply air is a boundary that the halls cannot move, so a hall that overwhelmed its plant would not see its setpoint rise. The four actions are stated rather than searched: a smaller depression than 3.5 K might also have held the hall, and the report does not claim the hall action is the best action available at hall resolution, only that it is a reasonable one and what it costs.

6 Conclusion

This report asked whether the resolution at which a modeled data-center facility is read changes which cooling action is safe during a heatwave, and what each coarser resolution costs. It does change the action. Read as one mean over 30 racks, the facility peaks at 33.15°C against a limit of 36.85°C and the reading justifies no action, while rack B3_r5 reaches 39.12°C: the coarsest resolution misses one violation. Read at hall, row or rack resolution, the peak is visible and each of the three actions that follow holds every rack under the limit. This report also asked what a reading the physics does not support costs a resolution that cannot check it. Two readings cross the limit and only one rack does; the hall and row resolutions act on both, and the extra reading is what buys the row resolution's precautionary trim and 12.0 % of additional energy, while the rack resolution forms the residual of (29), recovers the 1.500 K offset written into that reading, and omits the trim. The three safe actions cost 2.259, 1.783, 1.592 MWh over the 96-hour window, the hall-wide action 41.9 % more than the rack-sized one, which itself costs 0.0006 MWh less than doing nothing. This report finally asked how far that ordering depends on the modeling choices behind it. It survives all 192 combinations of the registered plant choices with a smallest margin of 0.0611, but it does

not survive the removal of the assumed setpoint penalty, without which the three actions differ by thousandths of a MWh and the ordering reverses. The safety and false-alarm findings stand on the simulation alone; the cost finding stands on the penalty. The next step is to derive the cost of a colder setpoint from the plant and the air side rather than assuming it, and to repeat the reading check against measurements with a stated noise level rather than against the simulation that produced them.

A Reproduction

```
cd docs/terra_graph_engine/case_studies/level3_row_pod/four_resolutions
PYTHONPATH=../../../../../src python run.py
PYTHONPATH=../../../../../src python make_report.py --pdf
cd ../../../../..
python -m pytest -q tests/plugins/datacenter/test_four_resolutions_case_study.py
```

`run.py` simulates the five hall scenarios, reads them at the four resolutions, prices the four actions, runs the sensitivity check over the 192 pricing combinations and writes `results.json`; `make_report.py` reads only `results.json` and writes `report.tex`. No random numbers are drawn, and the only external input is the weather file named below. The test in the last line regenerates the results, compares them with the committed file, pins the verdicts and rebuilds the tex.

B Data and problems

The weather and regional-demand record is `level1_building/real_weather_free_cooling/data/santa_clara_2023_hourly.csv`, 8,751 hourly rows for 2023 at 37.37°N, 121.92°W. The dry-bulb temperature and relative humidity come from the ERA5 reanalysis (Hersbach et al., 2020) through the Open-Meteo historical archive; the regional electricity demand comes from the California Independent System Operator's published system-load history, resampled to hourly means. Both are public and require no credentials. The study uses rows 6,631 to 6,726 of that file. The facility itself, its asymmetry and the offset added to one reading are the study's own construction and are stated in full in Section 2. The chiller law's coefficients are those of a machine anchored to ASHRAE 90.1-2019 minimums at the AHRI 550/590 rating conditions, with the shape of the ASHRAE research project RP-1043 fit to a 90-ton test chiller. The study leaves no corpus problems.

C Notation

Table 5 lists the symbols the report uses.

References

[Masanet et al. 2020] Masanet, E., Shehabi, A., Lei, N., Smith, S., Koomey, J. (2020). Recalibrating global data center energy-use estimates. Science 367(6481), 984–986. doi:10.1126/science.aba3758.

[Barroso et al. 2019] Barroso, L. A., Hölzle, U., Ranganathan, P. (2019). The Datacenter as a Computer: Designing Warehouse-Scale Machines. Synthesis Lectures on Computer Architecture, Springer. doi:10.1007/978-3-031-01761-2.

Table 5: Symbols used in the report, with the values the study gives the constants among them.

symbol	meaning
T_i, P_i, Q_i	rack temperature [K], electrical power [W] and heat flow to the aisle [W]
T_c, T_h	cold-aisle and hot-aisle temperature of a row [K]
T_R, T_B	room and building temperature of a hall [K]
T_{chw}	supply-air temperature the plant delivers [K]
T_{db}	outdoor dry-bulb temperature [K]
α	rise of rack power with rack temperature, 5 W/K
T_{ref}	reference temperature of the rack power law, 300 K
G_{rk}	rack-face conductance, 1,600 W/K at nominal airflow
G_{s}, G_{ret}	row supply and return conductance, 3,000 W/K each
G_{room}, G_{bldg}	room and building conductance, 40,000 and 100,000 W/K at nominal airflow
C_{rk}, C_{a}, C_R, C_B	thermal capacities of rack, aisle, room and building, 5.0×10^5, 1.0×10^6, 1.0×10^6 and 1.0×10^9 J/K
m_r	recirculation ratio of a row, 0.08 when contained
a	approach of the supply setpoint to the dry bulb, 3 K
δ, Δ	setpoint depression an action asks for and the depression actually delivered [K]
β	electrical penalty per kelvin of depression, 0.12 per K
w	share of the plant's duty the penalty is charged on
φ	share of the heat the chillers carry, the rest being carried by outdoor air
b_1, b_2, b_3	entropy generation [W/K], heat leak [W] and exchanger resistance [K/W] of a chiller unit
C, n	capacity of one chiller unit, 300 kW, and the number of units, 2
E	cooling electricity over the window [MWh]
g	margin of the energy ordering in a sensitivity corner

[Schmidt et al. 2005] Schmidt, R. R., Cruz, E. E., Iyengar, M. (2005). Challenges of data center thermal management. IBM Journal of Research and Development 49(4.5), 709–723. doi:10.1147/rd.494.0709.

[Sharma et al. 2005] Sharma, R. K., Bash, C. E., Patel, C. D., Friedrich, R. J., Chase, J. S. (2005). Balance of power: dynamic thermal management for Internet data centers. IEEE Internet Computing 9(1), 42–49. doi:10.1109/MIC.2005.10.

[Katipamula and Brambley 2005] Katipamula, S., Brambley, M. R. (2005). Methods for fault detection, diagnostics, and prognostics for building systems — a review, part I. HVAC&R Research 11(1), 3–25. doi:10.1080/10789669.2005.10391123.

[Crowe 1996] Crowe, C. M. (1996). Data reconciliation — progress and challenges. Journal of Process Control 6(2–3), 89–98. doi:10.1016/0959-1524(96)00012-1.

[Chua et al. 1996] Chua, H. T., Ng, K. C., Gordon, J. M. (1996). Experimental study of the fundamental properties of reciprocating chillers and their relation to thermodynamic modeling and chiller design. International Journal of Heat and Mass Transfer 39(11), 2195–2204. doi:10.1016/0017-9310(95)00335-5.

[Hersbach et al. 2020] Hersbach, H., Bell, B., Berrisford, P., et al. (2020). The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society 146(730), 1999–2049. doi:10.1002/qj.3803.

[ASHRAE 2019] ASHRAE (2019). ANSI/ASHRAE/IES Standard 90.1-2019: Energy Standard for Buildings Except Low-Rise Residential Buildings. ASHRAE, Atlanta, GA.

[AHRI 2023] AHRI (2023). AHRI Standard 550/590 (I-P): Performance Rating of Water-Chilling and Heat Pump Water-Heating Packages Using the Vapor Compression Cycle. Air-Conditioning, Heating and Refrigeration Institute, Arlington, VA.