

Operations and Control by Derivation

Scenarios, interventions and the decisions of one night shift

Draft · 220 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Operations and Control by Derivation

Volume II

Operator actions, model validation and alarms

Ricardo Marquez

November 1, 2026

How to read this book

This book teaches one skill: to take an operator's question stated in the language of the control room, write it as a probability or an expectation over the unknowns of a physical model, derive what that expression requires, compute it, check it, and act on it. Volume I holds Parts I to III, Chapters 1 to 11: the method itself, the state of the plant read from its telemetry (virtual sensors, unit heats, meters, the operating envelope), and the detection and diagnosis of abnormal operation. This second volume holds Parts IV to VI, Chapters 12 to 20. It asks what happens if the operator acts (a set point, a unit out, a load moved, nothing done), then how far the model can be trusted tonight (goodness of fit, sensor loss, the next field measurement), and then where the alarms should sit and how far an alarm should be believed. Chapter numbers run continuously across the two volumes; a reference to a chapter numbered 11 or below points to Volume I, whose results this volume builds on and does not repeat.

Every chapter follows the progression fixed in Volume I's preface and restated here in one line each: step 0, the plant graph; step 1, the question as a probability or expectation over the state; step 2, the graph extended for the question and the telemetry attached as rows; step 3, the posterior of the state, derived and computed once; step 4, the answers, each question read from that posterior down to its number; step 5, evaluation and checks against the full calculation and the constructed truth; step 6, the report with its band; step 7, the next reading, ranked by the width it removes. The five kinds of question and the five computations beneath them are those of Volume I.

The plant of this volume is two data halls on one chilled-water loop, followed through a twelve-hour night shift in five-minute windows: six cooling units, ten racks, sixteen readings. On that plant many of the questions are the same posterior read at a set point not yet applied, and the chapters say so where it happens.

The framework the book teaches is physics-constrained probabilistic modeling: a plant written as a graph, its unknowns as a state vector, every reading as a row, and every question as a probability over that state. TERRA is the author's own implementation of this framework; the book names the framework and the computations, not the software, and every number in it is reproducible from the derivations as written. The notation chapter is reproduced here so that this volume's symbols can be read with Volume I closed, and the atlas of questions in the appendix covers Volumes I and II.

Each chapter closes with a statement of what to retain, a table of its terms, and exercises at three levels, formulation, derivation and application, each followed by its answer.

Contents

How to read this book **ii**

Notation and vocabulary **vi**

IV Predicting the Effect of Operator Actions **1**

12 Set-point changes as scenarios **2**
12.1 A raise of the chilled water in the night . . . 2
12.2 The posterior of the state (step 3) . . . 8
12.3 The answers (step 4) . . . 13
12.4 The checks (step 5) . . . 18
12.5 The report (step 6) . . . 21
12.6 The next reading (step 7) . . . 22
12.7 What to retain . . . 24
12.8 Exercises . . . 25

13 Equipment outages and maintenance windows **26**
13.1 A hall with a unit to spare . . . 26
13.2 The posterior of the state (step 3) . . . 32
13.3 The answers (step 4) . . . 34
13.4 The checks (step 5) . . . 41
13.5 The report (step 6) . . . 43
13.6 The next reading (step 7) . . . 44
13.7 What to retain . . . 46
13.8 Exercises . . . 47

14 Load transfer between halls **48**
14.1 Forty kilowatts between two halls . . . 48
14.2 The posterior of the state (step 3) . . . 51
14.3 The answers (step 4) . . . 52
14.4 The checks (step 5) . . . 59
14.5 The report (step 6) . . . 61
14.6 The next reading (step 7) . . . 62
14.7 What to retain . . . 64
14.8 Exercises . . . 65

15 Look-ahead risk without intervention **66**
15.1 The hours before the morning shift . . . 66
15.2 The posterior of the state (step 3) . . . 71
15.3 The answers (step 4) . . . 74

15.4 The checks (step 5) . . . 83
15.5 The report (step 6) . . . 86
15.6 The next reading (step 7) . . . 87
15.7 What to retain . . . 89
15.8 Exercises . . . 90

V Validating the Model in Operation 92

16 Goodness of fit in operation 93
16.1 A fit-quality light that stays green . . . 93
16.2 The posterior of the state (step 3) . . . 97
16.3 The answers (step 4) . . . 99
16.4 The checks (step 5) . . . 108
16.5 The report (step 6) . . . 109
16.6 The next reading (step 7) . . . 111
16.7 What to retain . . . 112
16.8 Exercises . . . 113

17 Sensor loss and measurement redundancy 114
17.1 A work order for the return sensor . . . 114
17.2 The posterior of the state (step 3) . . . 118
17.3 The answers (step 4) . . . 119
17.4 The checks (step 5) . . . 131
17.5 The report (step 6) . . . 132
17.6 The next reading (step 7) . . . 134
17.7 What to retain . . . 135
17.8 Exercises . . . 136

18 Selecting the next field measurement 137
18.1 One technician and one trip . . . 137
18.2 The posterior of the state (step 3) . . . 140
18.3 The answers (step 4) . . . 141
18.4 The checks (step 5) . . . 148
18.5 The report (step 6) . . . 150
18.6 The next reading (step 7) . . . 151
18.7 What to retain . . . 152
18.8 Exercises . . . 153

VI Alarm Thresholds and Alarm Credibility 154

19 Calibrating alarm thresholds 155
19.1 An alarm that has never fired . . . 155
19.2 The posterior of the state (step 3) . . . 159
19.3 The answers (step 4) . . . 162
19.4 The checks (step 5) . . . 172
19.5 The report (step 6) . . . 173
19.6 The next reading (step 7) . . . 176
19.7 What to retain . . . 177
19.8 Exercises . . . 178

20 Alarm credibility and base rates **179**
20.1 An alarm that has just fired . . . 179
20.2 The posterior of the state (step 3) . . . 183
20.3 The answers (step 4) . . . 184
20.4 The checks (step 5) . . . 193
20.5 The report (step 6) . . . 195
20.6 The next reading (step 7) . . . 195
20.7 What to retain . . . 197
20.8 Exercises . . . 198

A Atlas of operator questions **199**

Notation and vocabulary

Every symbol below has the same meaning in every chapter, except where a chapter's table says otherwise. The unknown always stands on the left of a conditioning bar and the evidence on the right. Vectors and matrices are bold. On the two-unit hall of Parts II and III the units are $u \in \{1, 2\}$ and the racks $k \in \{1, 2, 3\}$. On the shift plant of Parts IV to VI, hall subscripts are $h \in \{A, B\}$, cooling-unit subscripts $u \in \{A1, A2, A3, B1, B2, B3\}$ and rack subscripts $k \in \{A1, \ldots, A5, B1, \ldots, B5\}$; the hall of a unit or a rack is the letter of its name.

Symbol	Unit	Meaning
The plant graph and its state		
$\mathbf{z}$, z_i	varies	State vector of the plant graph: every free quantity of one telemetry window at once; its i-th entry.
$\mathbf{z}_t$, t	varies, –	The state of window t; the window index, each window five minutes long: on the shift plant $t = 1, \ldots, 144$, window 1 starting at 19:00; on the two-unit hall's run $t = 1, \ldots, 36$, window 1 starting at 23:00.
$\mathbf{y}$, y	varies	The evidence: all readings and prior information treated as known; a single reading, subscripted by what it reads (y_{T_s}, y_{Q_A}).
$r_k(\mathbf{z})$	varies	Row k of the model: a physical relation that must hold $(r_k = 0)$, a reading compared with the state, or a prior.
w_k, σ_k	varies	Weight (precision) of row k and the standard deviation it corresponds to, $w_k = 1/\sigma_k^2$. A row may equivalently be written normalized, r_k/σ_k, with unit weight.
σ_p	varies	Tolerance of a physics row; an exact constraint is the limit $\sigma_p \to 0$.
The equipment hall of Part I		
T, T_c	°C	Temperature of the hall air; the supply air held by the cooling unit's discharge controller, pinned.
Q, G	kW, kW/K	Heat load of the hall; effective conductance of its air path to the supply air.
g, G_{ref}	–, kW/K	Log-conductance $\ln(G/G_{\text{ref}})$, the entry of the state that carries the conductance; its reference value.
y_T, σ_T	°C, K	The hall sensor's reading and the standard deviation of its error.
The two-unit hall of Parts II–III and the shift plant		
T_s	°C	Chilled-water supply temperature; the actuator, set by the chiller set point.
Q_h, $Q_{o,h}$	kW	IT heat load of hall h; its other heat (fans, lighting), pinned.

Symbol	Unit	Meaning
G_u, $\bar{G}_u$	kW/K	Effective conductance of the coil of cooling unit u; its rated value, pinned.
g_u, d_u	–	Log-derate $\ln(G_u/\bar{G}_u)$, the entry of the state that carries the unit's conductance; the derate factor $d_u = \mathrm{e}^{g_u}$, one at the rating.
C_u	kW/K	Air capacity rate of unit u (air mass flow times specific heat), pinned at its fixed fan speed.
Q_u	kW	Heat unit u rejects to the chilled-water loop, $G_u(T_{r,h} - T_s)$.
$T_{r,h}$, $T_{c,h}$	°C	Return (hot-aisle) and cold-aisle air temperature of hall h.
$T_{d,u}$	°C	Discharge air temperature of unit u.
T_k, β_k	°C, –	Inlet temperature of rack k; its recirculation fraction, $T_k = T_{c,h} + \beta_k(T_{r,h} - T_{c,h})$.
$T_{\max}$	°C	A limit on a temperature: the hall's 24°C now and 27.5°C at full load in Part I (28°C for the illustrative halls of Chapters 3 and 4); a rack inlet's 27°C, the recommended envelope, from Part II.
f^*, G^*, T^*	varies	A requirement level on a state entry or a derived quantity, as in the companion book: the event $f \geq f^*$ is *the requirement holds*. From Part IV a starred plant quantity (Q_A^*, T_k^*) is also a scenario value: the quantity at a set point, unit state or load not yet applied, as in the companion book's Part V.
b_j	varies	Bias of reading j, an entry freed when a meter is suspected.
Probability		
$\Pr(A)$, $\Pr(A \mid B)$	–	Probability of event A; conditional on event B.
$p(\cdot)$	varies	A probability density; $p(\mathbf{z} \mid \mathbf{y})$ is the posterior of the state, $p(\mathbf{y} \mid \mathbf{z})$ the likelihood, $p(\mathbf{z})$ the prior.
$\mathbb{E}$, Var, Cov	varies	Expectation, variance, covariance.
$\mathcal{N}(\mu, \sigma^2)$	–	Normal distribution with mean μ and variance σ^2; Φ its standard cumulative distribution function, φ its density; $\Phi_2(u_1, u_2; \varrho)$ that of two standard normals of correlation ϱ.
χ^2_ν	–	Chi-square distribution with ν degrees of freedom.
Inference objects		
$\mathcal{L}(\mathbf{z})$	–	The objective, minus the log posterior up to a constant: $\mathcal{L} = \frac{1}{2} r(\mathbf{z})^\mathsf{T} r(\mathbf{z})$ for normalized Gaussian rows.
$\hat{\mathbf{z}}$	varies	The maximum a posteriori (MAP) state, the minimizer of $\mathcal{L}$.
$\mathbf{J}$	varies	Jacobian of the normalized rows with respect to the state, evaluated at $\hat{\mathbf{z}}$.
$\mathbf{\Lambda}$	varies	Information matrix (Gauss–Newton Hessian of $\mathcal{L}$), $\mathbf{J}^\mathsf{T}\mathbf{J}$ for normalized rows, $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ with $\mathbf{W} = \mathrm{diag}(w_k)$ for unnormalized ones.

Symbol	Unit	Meaning
$\boldsymbol{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\boldsymbol{\Lambda}^{-1}$; its entries. The posterior spread of entry i is written $\sqrt{\Sigma_{ii}}$; σ_k always names a row's standard deviation.
$f(\mathbf{z})$, $\mathbf{g}$	varies	A derived quantity of the state and its gradient $\partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$.
$\mathbf{A}$, $\mathbf{A}_{\text{free}}$, $\mathbf{A}_{\text{dep}}$	varies	Jacobian of the physics rows on the full state; its blocks on the free and the dependent entries (Chapter 4).
$\mathbf{N}$, $\mathbf{J}_{\text{rows}}$, $\mathbf{J}_{\text{red}}$	varies	Null-space basis of the physics rows, $[\mathbf{I};\ -\mathbf{A}_{\text{dep}}^{-1}\mathbf{A}_{\text{free}}]$; the Jacobian of the evidence rows (readings and priors); the reduced Jacobian $\mathbf{J}_{\text{rows}}\mathbf{N}$ of the reduced state.
f_L, f_U	varies	The ends of the credible interval of a derived quantity.
y_{new}, $\tilde{y}_a$	varies	A future reading (the predictive) and the outcome of a candidate reading a before it is taken (the preposterior).
e_k	–	Normalized residual of reading k: its residual divided by the standard deviation of that residual.
M, M_j, B_{AB}	–	A drawing (model) of the plant; the members of a family of candidates; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$.
$p(\mathbf{y} \mid M)$	–	Evidence (marginal likelihood) of model M.
$\mathbf{J}_a$, $\boldsymbol{\Lambda}_a$, $\boldsymbol{\Sigma}_a$	varies	The Jacobian rows of a candidate reading, its information $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$, and the preposterior covariance $(\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$.
γ, t^*	varies/h, h	Slope of an entry's trend across windows; the time at which the trend crosses a limit (as in the companion book).
Alarms and decisions		
c	varies	Decision level (alarm threshold) on a monitored quantity.
α	–	False-alarm probability per window.
$U(a)$, $L(a, \cdot)$	varies	Preposterior value of a candidate reading a; the loss of an action a.
Scenario nodes (Part IV); a row naming a chapter is that chapter's own symbol		
T_s^*, T_k^*, $\mathbf{T}^*$	°C	The supply at a scenario node; the inlet of rack k there, a derived quantity of the state; the inlets stacked. A star marks the value at a set point, unit state or load not yet applied.
$T_{r,h}^*$, $T_{c,h}^*$, $T_{d,u}^*$, Q_u^*, Q_h^*	°C, kW	The return, cold aisle, discharge, unit heat and hall load at the node (Chapter 13 writes T_{rA}^* and T_{cA}^* for hall A).
$\mathbf{g}_{T_k}$, $\mathbf{g}_{T_k^*}$	varies	The gradient of the inlet of rack k at the present state and at the node.
$\mathbf{m}^*$, $\mathbf{D}^*$	°C, varies	The inlets at the node evaluated at $\hat{\mathbf{z}}$; their gradients stacked as rows, so that the inlets are jointly $\mathcal{N}(\mathbf{m}^*, \mathbf{D}^*\boldsymbol{\Sigma}\mathbf{D}^{*\mathsf{T}})$.
κ, T_s^κ	–, °C	The assurance a requirement is held to; the largest set point whose joint under the limit is at least κ (Chapter 12).
β_k^{s}	–	The airflow survey's value of the recirculation fraction of rack k, the center of its prior.

Symbol	Unit	Meaning
Q_{loop}	kW	The heat the loop meter reads: both hall loads and both halls' other heat.
$\mathcal{R}$, $\sum G^*$, $\sum C^*$, β_k^*	–, kW/K, kW/K, –	The units of the hall still running with one out; the sums of their conductances and capacity rates; the fraction at the node with the rack air flow held (Chapter 13).
α, γ, $\bar{t}$, d_t	kW, kW per window, –, kW	The line and wander of a load, $Q_{A,t} = \alpha + \gamma(t - \bar{t}) + d_t$: its level, its slope, the center window, and the wander (Chapter 13).
ϕ, s, ρ	–, kW, –	The wander's persistence, its stationary spread, and the part of it left after some windows, $\rho = \phi^{t-73}$ (Chapters 13 and 15).
$Q_f(t)$, $\mathbf{J}_{\text{load}}$, $\mathbf{\Sigma}_{\text{load}}$	kW, varies	The forecast load at window t; the Jacobian and covariance of the second fit, of the load's line and wander (Chapter 13).
Δ, $\Delta_{\max}(\kappa)$	kW	The load removed from a hall (Chapter 13) or moved between halls (Chapter 14), pinned; the largest transfer held at assurance κ.
$\sum G_h$, $\sum C_h$	kW/K	The sums of the conductances and of the capacity rates of the units of hall h.
$T_{k,t}$, $T_{s,t}$, $Q_{h,t}$, $T_{r,h,t}$, $T_{c,h,t}$, $T_{d,u,t}$	°C, kW	The quantities of the plant at window t, when more than one window is in the state.
a_h, b_h, $w_{h,t}$, $\varepsilon_{h,t}$	kW, kW per window, kW, kW	The trend of the load of hall h, $Q_{h,t} = a_h + b_h(t - \bar{t}) + w_{h,t}$: level, slope, the wander $w_{h,t} = \phi w_{h,t-1} + \varepsilon_{h,t}$ and its innovation (Chapter 15).
$g_{\text{B2},t}$, γ, τ, $(x)_+$	–, per window, –, –	The coil's log-derate at window t, $g_{\text{B2}} - \gamma(t - \tau)_+$; the rate of its ramp; the window of its onset; the positive part $\max(x, 0)$ (Chapter 15).
s_j, b_{Q_A}	K, kW	The jitter of each window's set point about one level; the freed bias of the hall-A meter (Chapter 15).
E_s, E, M	–, –, °C	The first-passage event, some inlet at or over the limit at some window through s; the event at the horizon, E_{144}; the largest inlet over the racks and windows (Chapter 15; M_j remains a drawing).
$h(s_1, s_2)$	h^{-1}	The conditional hazard of the first passage between two horizons, given none by the first (Chapter 15).
$\mathbf{G}$, $\widetilde{\mathbf{T}}$	varies	The stacked gradients of the inlets over the horizon; the inlets linearized at $\hat{\mathbf{z}}$, of which the first passage is the complement of an orthant (Chapter 15).
$\mathbf{z}^{(j)}$, w_j	varies, –	A draw of the state and its importance weight (Chapter 15); w_k remains the weight of a row.
m_k, δ_{k,Q_A}, $\mathbf{y}_{1:T}$	varies, –, varies	The prediction of reading k; one on the hall-A meter's row and zero elsewhere; the record through decision window T (Chapter 15).
The model in operation (Part V)		
$\mathbf{H}_t$, h_{kk}	–	The hat matrix $\mathbf{J}_t\mathbf{\Sigma}_t\mathbf{J}_t^{\mathsf{T}}$ of window t; the leverage of row k, its diagonal entry, with $1 - h_{kk}$ the redundancy number.

Symbol	Unit	Meaning
X_t^2, ν_t, r_k^{N}	–	The sum of the squared reading rows at $\hat{\mathbf{z}}_t$; the effective degrees of freedom it is referred to, $\sum_k (1 - h_{kk})$ over the reading rows; the normalized residual $r_k(\hat{\mathbf{z}}_t)/\sqrt{1 - h_{kk}}$.
$\mathbf{d}$, $\boldsymbol{\ell}$, $\boldsymbol{\ell}_{\mathrm{R}}$, λ	–	A stated change on the reading rows; what the fit leaves of it, $(\mathbf{I} - \mathbf{H}_t)$ applied to it; the part on the reading rows; the test's noncentrality, $\|\boldsymbol{\ell}_{\mathrm{R}}\|^2$ (Chapter 16).
$\mathbf{z}_{\mathrm{ref}}$, $\hat{\mathbf{z}}_{\mathrm{s}}$, $\boldsymbol{\Sigma}_{\mathrm{s}}$	varies	The state of the reference fit, shared coils and fractions with each window's supply and loads; the shared entries' MAP values and their covariance (Chapter 16).
$\mathbf{y}_t$, $\mathbf{q}_t$, $\mathbf{q}_0$	varies	The readings of window t; its supply and two loads, free per window against the reference; the center of their prior (Chapter 16).
$\mathbf{A}_{\mathrm{s}}$, $\mathbf{A}_{\mathrm{q}}$, $\mathbf{P}$, $\mathbf{L}$	varies	The gradients of the readings on the shared entries and on $\mathbf{q}_t$; the set point's prior information; $\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}\mathbf{A}_{\mathrm{q}} + \mathbf{P}$ (Chapter 16).
$\mathbf{V}_t$, $\mathbf{e}_t$, $e_{k,t}$	varies	The predictive covariance of the readings, sensors plus the reference's spread; observed minus predicted at the found $\hat{\mathbf{q}}_t$, the sign of the screen; its entry for reading k (Chapter 16).
Q_t, $z_{k,t}$	–	The reference test, $\mathbf{e}_t^{\mathsf{T}}\mathbf{V}_t^{-1}\mathbf{e}_t$ (a statistic, not a load; a load carries a hall, $Q_{h,t}$); the standardized $e_{k,t}$ (Chapter 16).
$\mathbf{a}_k$, $\mathbf{A}_S$	varies	The row of $\mathbf{J}$ for reading k; the rows of a set S of readings (Chapter 17).
$\mathbf{y}_{(k)}$, $\boldsymbol{\Lambda}_{(k)}$, $\boldsymbol{\Sigma}_{(k)}$; $\mathbf{y}_{(S)}$, $\mathbf{J}_{(S)}$, $\boldsymbol{\Sigma}_{(S)}$, $\hat{\mathbf{z}}_{(S)}$	varies	The evidence, information and covariance with reading k removed; the same with the set S removed, the MAP state of its refit (Chapter 17).
c_k, b	varies, –	The coupling $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a}_k^{\mathsf{T}}$ of row k to a stated quantity, whose loss adds variance $c_k^2/(1 - h_{kk})$; the signed departure of a stuck reading in its own spreads (Chapter 17).
m_a, σ_a, r_a, $\mathbf{a}$	varies	What candidate reading a reads, as a function of the state; its spread; its scenario row $(m_a(\mathbf{z}) - \tilde{y}_a)/\sigma_a$; its gradient over its spread at $\hat{\mathbf{z}}$ (Chapter 18).
$\hat{\mathbf{z}}_a$, ρ_a, P_a	varies, –, –	The MAP state after outcome $\tilde{y}_a$; the correlation of the reading, error included, with a worry; the joint at the node after that outcome (Chapter 18).
$T_{\mathrm{in}k}$, T_{rB}, Q_{B2}	°C, °C, kW	The worries by name: the inlet of rack k, written T_k elsewhere; hall B's return, written $T_{r,\mathrm{B}}$ elsewhere; the heat of unit B2 (Chapter 18).
Alarm levels and credibility (Part VI)		
$\mathbf{K}$, $\mathbf{K}_t$	–	The residual map, $\mathbf{I} - \mathbf{A}_{\mathrm{q}}\mathbf{L}^{-1}\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}$, which takes the readings' departures to the residuals with this window's supply and loads refound (Chapter 19; Chapter 20 writes $\mathbf{K}_t$).
e	K	The residual of a monitored reading, its entry of $\mathbf{e}_t$, in kelvin and not in its spreads; in Chapter 20 also the value at which the alarm fired. It is not the normalized e_k of Chapter 17.

Symbol	Unit	Meaning
sd, $\mathrm{sd_n}$, $\mathrm{sd_s}$, x	K	The predictive spread of a monitored quantity; its noise part, new in each window; its shared part, the same in every window; the reference's shared error at that quantity, normal on $\mathrm{sd_s}$ (Chapter 19).
$\mathbf{z}^*$, g_u^*, β_k^*	varies, –, –	The state at the scenario node of an excursion route; the coil's log-derate and the rack's fraction pinned there (Chapters 19 and 20).
s_r, s	K	The shift route r leaves on the residual (Chapter 19); the shift a stated excursion leaves on it (Chapter 20). In Part IV and in the plain reading's spread of Chapter 19, s is the wander's spread.
M	windows	The windows of a month, 8640 (Chapter 19); M_j remains a drawing.
α, π, z_α, z_π	–	In Chapter 19: the requested false-alarm tail per window; the requested catch; their standard normal points.
L_α, L_π	K	The level on the monitored quantity that gives the rate, z_α sd, and the level that gives the catch, $s_r - z_\pi$ sd; both requests hold on the interval $[L_\alpha, L_\pi]$ when it is not empty (Chapter 19).
$\alpha(L)$, $\pi_r(L)$, $\pi_t(L)$	–	The false-alarm tail at level L; the catch of route r there; the catch of the plain reading at window t (Chapter 19).
$\bar{y}_t$, sd_t, $\mathcal{W}$	°C, K, –	The forecast of the plain reading at window t, a superscript ref when it stands on the reference; its spread; the span of windows a fixed level covers (Chapter 19).
a, a_r, σ_T, σ_j	K/kW, K/kW, K, K	The inlet's slope on hall B's load at the one fit and at the reference; the sensor's spread; the set point's jitter (Chapter 19).
π, c, α, L	–, –, –, K	In Chapter 20: the base rate of excursions per window; the catch, the tail past the level under an excursion (elsewhere c is a decision level, and in Chapter 19 π is the catch); the false-alarm tail per window; the level.
$\boldsymbol{\delta}$, $\mathbf{d}$	varies	What a hypothesis moves in the prediction; the same through the map, $\mathbf{K}_t\boldsymbol{\delta}$ (Chapter 20; in Chapter 19 $\mathbf{d}$ is the readings' change before the map).
$B(e)$, S, $w(S)$	–, K, 1/K	The density ratio at the value e; an excursion's size; its stated density (Chapter 20).
d, π_d	windows, 1/window	The windows an excursion lasts; the prior of a window drawn from all the windows that fire (Chapter 20).
γ, σ_γ, $\mathbf{h}_t$, $\mathbf{C}$	1/window, 1/window, varies, varies	The ramp's slope, signed, g_{B2} moved by $\gamma(t-\tau)_+$, so a derate has $\gamma < 0$ (Chapter 15 writes the ramp with a minus sign); its prior spread; its column at window t; the stacked covariance (Chapter 20).
S_{eh}, S_{hh}, Λ	varies	The ramp's two weighted products, $\mathbf{h}^\mathsf{T}\mathbf{C}^{-1}\mathbf{e}$ and $\mathbf{h}^\mathsf{T}\mathbf{C}^{-1}\mathbf{h}$; its information (Chapter 20).
r, α_r, c_r	K, –, –	The ramp's part in the window's residual; the two tails with it in both means (Chapter 20; r_k remains a row).

Vocabulary

Plant graph The physical system as nodes (places where a conserved quantity balances), edges (paths along which it flows) and sources, with the relations that must hold among them. Built before any question is asked.

State vector The collection of free quantities the fit estimates at once; outputs that are entries of the state need no further transformation.

Pinned, free A quantity treated as known (boundary condition) against one carried in the state.

Residual, row A relation written so that it equals zero when satisfied; physics rows constrain the state, reading rows compare it with telemetry, prior rows carry handbook or commissioning knowledge.

Telemetry window The span of readings fitted at once, five minutes here, within which the plant is treated as steady.

State estimation The fit of the state of one window to its readings and priors: the MAP state and its covariance.

Virtual (soft) sensor An unmeasured quantity read from the fitted state with its band, in place of a missing instrument.

Observability Whether the readings determine an entry of the state, or only a combination of entries; read from the Jacobian and from the posterior spread against the prior spread.

Bad-data detection Testing whether a reading is consistent with the others: the chi-square test of the whole residual vector and the largest normalized residual.

Fault detection and isolation Deciding that something changed, and where: a family of candidate faults generated from the graph, ranked by evidence.

What-if (contingency) analysis The plant re-solved at a scenario node, with an actuator or a unit state pinned to a value not yet applied, and the covariance carried through.

Look-ahead A forecast of the plant over the coming windows, with its band.

Operating envelope The region of the state inside which every limit holds.

Alarm rationalization Setting each alarm's decision level from its purpose, its false-alarm rate and its probability of detection.

False-alarm rate, probability of detection The rate at which an alarm fires without a real excursion; the probability that it fires when one of stated size occurs.

Standing alarm, alarm flood An alarm that stays active for long periods; more alarms than an operator can handle in a short time.

Laplace approximation Replacing the posterior by the Gaussian whose covariance is the inverse curvature of the objective at the MAP state.

Delta method Linearizing a derived quantity around the MAP state so that its variance is $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$.

Posterior predictive The distribution of a future reading obtained by averaging its model over the posterior.

Evidence, Bayes factor The normalizing constant of one model's posterior; the ratio of two models' evidences.

Preposterior The posterior that would result from a reading not yet taken; the object that values a reading before it is requested.

The five computations Posterior, trend, evidence, preposterior and expected consequence: the objects every operator question is answered from.

Scenario node A copy of the plant's temperature nodes at values of the pinned or actuated quantities not yet applied. It adds no reading row and no entry of the state; its temperatures are derived quantities with their bands.

Pin, shift A scenario's set point fixed at a value; the same set point moved from its fitted entry by an amount. Under a pin the band at the node is that of the inlet less the supply; under a shift it is the present inlet's.

Assurance The probability a requirement is held to; a largest set point, unit load or transfer is the threshold question inverted at it.

N−1 contingency The plant with one unit out, asked of each unit in turn.

Stop rule, rack air flow held What a stopped unit does to the air: the plant's rule drops its flow; the second form holds the rack air flow.

Forecast peak The load the trend line and the wander give at the horizon's last window.

Load transfer (load migration) Load moved between halls at a scenario node: a donor hall and a receiving hall.

Margin The limit less an inlet's temperature: a capacity margin, not an availability.

First passage Some inlet at or over the limit at some window of the horizon, a union over racks and windows; its conditional hazard is the probability of a first passage in an hour, given none earlier.

Linearized joint (orthant) The state route for a first passage: an orthant of the normal joint of the forecast inlets, integrated by draws.

Load wander A persistent first-order autoregressive departure of a load about its trend line, with its persistence and stationary spread.

Epistemic and aleatory spread The spread from entries the record could fix; the spread of the future's own wander.

Decision window A window at which the answer is read again, using nothing later.

Goodness of fit, effective degrees of freedom The agreement between the readings and the fitted model, tested by the sum of the squared reading rows against the effective degrees of freedom $\sum(1 - h_{kk})$, not rows minus entries.

Fit-quality light A screen's summary of the window's own test of fit.

Noncentrality, power The squared size of a stated change left on the rows a statistic adds up; the noncentral chi-square tail past the level, which is the test's probability of detection.

Reference fit, reference test One fit of a stretch of windows with shared coils and fractions; the test of a later window's readings against its predictive.

First verdict The first window past a level fixed in advance, reported with its spread over fresh sensor errors.

Repaired model The model with an added entry, fitted and tested again.

Sensor loss (removed row) A reading's row taken out of the fit. A sensor that fails and stays connected is a different row, a stuck-at failure, whose band does not widen.

Down-date (deletion formula) The covariance with a row removed, read from $\boldsymbol{\Sigma}$ at $\hat{\mathbf{z}}$; it needs no outcome. Block deletion removes a named group of readings together.

Coupling A row's link to a stated quantity; its square over the redundancy number is the variance its loss adds.

Critical measurement A reading whose loss lowers the rank of the reading rows.

Candidate reading, scenario row An instrument, a place and a spread, not yet used; the reading row it would add, written before its value exists.

Worry A named quantity to narrow, a function of the state; a ranking of candidates is against a worry.

Width taken off The share by which a band falls after a row.

Settled share The probability over a reading's outcomes that a joint ends at or above an upper reference level or below a lower one.

Greedy order The best reading, then the best given it, tested against the best pair.

Decision level The value of a monitored quantity past which an alarm fires.

Monitored quantity, residual alarm What an alarm compares with its level; an alarm on a reading's residual against a reference, with the window's supply and loads refound, as against a fixed level on the plain reading.

Excursion route A physical way a stated rise comes about, injected as a scenario node. A level's catch is stated for a route.

Catch The probability of detection of a stated excursion at a level.

Shared part of a residual's spread The reference's own error, the same in every window, as against the per-window reading noise; it does not average over windows and makes a false-alarm rate itself uncertain.

Base rate How often the thing an alarm is meant for occurs, regardless of any alarm.

Positive predictive value (alarm credibility) The posterior probability of an excursion given the alarm: the two tails are the evidences of excursion and of no excursion, and the base rate gives the prior odds.

Tail ratio, density ratio The likelihood ratio of the observation that the residual passed the level; that of the observation that it took this value.

First alarm The first fired window of a run, as against any window that fires.

Part IV

Predicting the Effect of Operator Actions

Chapter 12

Set-Point Changes as Scenarios

12.1 A raise of the chilled water in the night

It is 01:00, window 73 of the night, the loop is steady, and the operator is considering a raise of the chilled-water set point from 14 to 16°C. Warmer chilled water is commonly asked for because the chiller then does less work, so a raise could be a saving. The screen shows the supply at 13.88°C, the returns at 29.97 and 29.15°C, and the four inlet sensors at 19.51 (rack A1), 22.03 (A3), 22.13 (B2) and 22.67°C (B4). The hottest read inlet, rack B4, is 4.33 K inside the 27°C limit. Adding 2 K to the screen gives 24.67°C, 2.33 K inside, so the dashboard's reading is that the raise is safe.

Chapter 11 closed on a forecast that assumed nobody acts, and here somebody may. That reading adds a number to four sensors, and the plant has ten racks. The plant also carries no chiller-power row, so this chapter prices the risk side of the raise only. The question it answers is

what the probability is that every inlet stays under the limit at the raised set point, including the six racks no sensor reads, and what stands behind that probability.

It is answered on a bigger plant than the two-unit hall, with the method of Chapter 4 and one new object. The quantities asked for lie at a value of the actuator that has not been applied, and no reading can touch them.

Jargon: set point, actuator, what-if (contingency) analysis, scenario node, pinned, carried, assurance, joint tail, pin and shift

The *set point* is the value a controller is told to hold; the *actuator* is what the operator moves to change the plant, here the chilled-water set point. *What-if (contingency) analysis* computes the plant at a value of an actuator or a unit's run state not yet applied (in the framework, a *scenario node*). A quantity at the node is *pinned* if the scenario fixes it, and *carried* if it keeps the posterior of the present state. The *assurance* κ is the probability the operator requires of a requirement, as in Chapter 2. A *joint tail* is the probability of an event on several quantities at once, as in Chapter 8. The *pin* sets T_s^* to a value; the *shift* adds an amount to the fit's entry for T_s.

Key idea

A scenario node adds no reading and no entry to the reduced state: its temperatures are derived quantities of the state, the posterior is unchanged, and $\boldsymbol{\Sigma}$ is carried by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. A raise of the supply moves every temperature by the same amount, so the physics of the raise is one line. The answer is not, because that amount lands on inlets whose fractions no reading has seen.

12.1.1 A scenario node and what it carries

Chapter 5 read a temperature of the present plant at a place with no sensor. A what-if reads a temperature of a plant that does not yet exist: the same racks and coils, at a value of the actuator the operator has not entered. A *scenario node* is a copy of the plant's temperature nodes at values of the pinned or actuated quantities not yet applied. Write $\mathbf{u}^*$ for those values. Each scenario temperature is the forward chain of the physical rows solved at $\mathbf{u}^*$ and at the free entries,

$$T_k^*(\mathbf{z}) = T_k(\mathbf{z};\ \mathbf{u}^*), \qquad T_k^* \mid \mathbf{y} \approx \mathcal{N}(T_k^*(\hat{\mathbf{z}}),\ \mathbf{g}_{T_k^*}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{T_k^*}), \tag{12.1}$$

a function of the state like the virtual sensor of Chapter 5. The node adds no reading row, since nobody can read a plant that has not happened, and no entry to the reduced state, so the posterior is the one the readings already gave. On the full state it adds one copy of the temperatures and one copy of the physical rows, which the limit of exact rows removes again.

Principle 12.1. *What-if analysis needs no new reading and no new fit: pin what the operator would change, re-solve the temperatures, and carry $\boldsymbol{\Sigma}$ through their gradients. What a scenario can know is what the present posterior knows.*

12.1.2 The two-hall shift plant (step 0)

The two-unit hall of Chapter 5 had two units, three racks and a state of six entries. The shift plant has two halls A and B on one chilled-water loop at supply temperature T_s, three cooling units per hall (A1 to A3, B1 to B3) and five racks per hall (A1 to A5, B1 to B5); a unit and a rack may share a name, and the subscript of $T_{d,\mathrm{A1}}$ is a unit's while that of T_{A1} is a rack's.

The physical relations. They are the rows of Chapter 5, repeated per hall h with $G_u = \bar{G}_u \mathrm{e}^{g_u}$:

$$\begin{aligned}
r_{\mathrm{h},h} &= \Big[\sum_{u\in h} G_u\,(T_{r,h} - T_s) - Q_h - Q_{o,h}\Big]/\sigma_p && \text{hall balance,}\\
r_{\mathrm{a},u} &= \big[C_u\,(T_{r,h} - T_{d,u}) - G_u\,(T_{r,h} - T_s)\big]/\sigma_p && \text{air side of unit } u,\\
r_{\mathrm{m},h} &= \Big[\sum_{u\in h} C_u\,T_{c,h} - \sum_{u\in h} C_u T_{d,u}\Big]/\sigma_p && \text{cold-aisle mixing,}\\
r_{\mathrm{i},k} &= \big[T_k - T_{c,h} - \beta_k\,(T_{r,h} - T_{c,h})\big]/\sigma_p && \text{rack inlet } k,
\end{aligned} \tag{12.2}$$

ten rows per hall, twenty in all, with σ_p the tolerances of those rows taken to zero together. The pinned quantities are the ratings $\bar{G}_u = 8\,\mathrm{kW/K}$ in hall A and $7\,\mathrm{kW/K}$ in hall B, the air capacity rates $C_u = 10\,\mathrm{kW/K}$ and the other heat $Q_{o,\mathrm{A}} = 12\,\mathrm{kW}$, $Q_{o,\mathrm{B}} = 10\,\mathrm{kW}$. The free quantities and the twenty temperatures form a full state of 39 entries, and with one scenario copy of the temperatures, 59.

The reduced state. On the reduced state of Chapter 4 the fit runs on the 19 free entries,

$$\mathbf{z}_t = \big(T_s,\ Q_{\mathrm{A}},\ Q_{\mathrm{B}},\ g_{\mathrm{A1}},\dots,g_{\mathrm{B3}},\ \beta_{\mathrm{A1}},\dots,\beta_{\mathrm{B5}}\big)^{\mathsf{T}}, \qquad t = 73, \tag{12.3}$$

in °C, kW and sixteen dimensionless entries. The dependent block $\mathbf{A}_{\mathrm{dep}}$ of the rows (12.2) is triangular with a nonzero diagonal, so it is invertible, and the temperatures of each hall follow in order,

$$\begin{aligned}
T_{r,h} &= T_s + \frac{Q_h + Q_{o,h}}{\sum_{u\in h} G_u}, \qquad T_{d,u} = T_{r,h} - \frac{G_u\,(T_{r,h} - T_s)}{C_u},\\
T_{c,h} &= \frac{\sum_{u\in h} C_u T_{d,u}}{\sum_{u\in h} C_u}, \qquad T_k = T_{c,h} + \beta_k\,(T_{r,h} - T_{c,h}).
\end{aligned} \tag{12.4}$$

As in equation (5.5), the air balance of each hall,

$$\sum_{u \in h} C_u \left(T_{r,h} - T_{c,h}\right) = Q_h + Q_{o,h}, \tag{12.5}$$

holds at every point of the set: the rise from the cold aisle to the return is set by the heat and the airflow alone.

Modeling choices at the graph. The plant is taken in steady state within a window, which makes the rows (12.2) algebraic and lets the node say where the plant settles at T_s^*; it says nothing of how long the loop and the air take to get there or of any controller's overshoot on the way, and the record holds no set-point step from which either could be read, so a statement of the form "safe for the next hour" is not available from this node. The units' airflow is pinned and no controller acts on the coils within a window; this lets a discharge reading inform its coil, and a valve that held a discharge temperature would break it. Each coil carries one effective conductance from the return air to the chilled water; this keeps one entry per coil, and it fails if a coil's heat transfer depends on its discharge temperature in a way one conductance cannot carry.

The graph. Figure 12.1 draws the rows. The two halls stand left and right of the chilled water in the middle (green). In each hall the load (green) enters the return air (blue), the return feeds the three coils (green, each with its log-derate), the coils' discharge mixes in the cold aisle (blue), and the cold aisle feeds five racks, each carrying its own fraction. Dashed arrows leave for the sixteen readings. Rack B3 is drawn in red. The red dashed box at the bottom is the scenario node: a copy of every temperature, with T_s^* pinned, and no dashed arrow to a reading.

12.1.3 The operator's questions (step 1)

The operator asks five questions of window 73, and step 1 writes each over the reduced state (12.3). Table 12.1 lists them. The first column is the question in the operator's words, the first of them the drill question; the second the general form of its category from Chapter 2; the third its mathematical form for the plant; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered.

12.1.4 The evidence and the node as rows (step 2)

Step 2 attaches window 73 to the graph as rows on the reduced state. The sixteen readings are those of Figure 12.1: the supply, of standard deviation 0.1 K; the two returns, six discharges and four inlets, each 0.3 K, the discharges of units A1 to B3 reading 17.87, 16.48, 17.60, 18.07, 19.99 and 18.97°C; the two hall meters, each 4 kW; and the loop meter on $Q_{\mathrm{A}} + Q_{\mathrm{B}} + Q_{o,\mathrm{A}} + Q_{o,\mathrm{B}}$, 8 kW. The priors are the plant's: $g_u \sim \mathcal{N}(0,\ 0.15^2)$ from the commissioning record, $\beta_k \sim \mathcal{N}(\beta_k^{\mathrm{s}},\ 0.08^2)$ from the airflow survey, $T_s \sim \mathcal{N}(14,\ 0.5^2)$. The loads have no prior row. That makes $16+17 = 33$ rows on 19 entries. Each reading is taken as Gaussian, unbiased and independent of the others given the state; this makes the objective a sum of squares, and it fails if a meter carries a bias the rows do not model, which Section 12.2.3 examines. The airflow survey is taken as the only knowledge of an unread rack's fraction. This gives the six empty columns of the Jacobian a proper posterior, and it fails if the survey's center or width is wrong or has gone stale, since no reading can correct it; the answer inherits whatever the survey got wrong, and Section 12.4.3 varies both.

What the node pins, leaves free and carries. Table 12.2 lists each quantity at the node of this chapter, with what it does there and why.

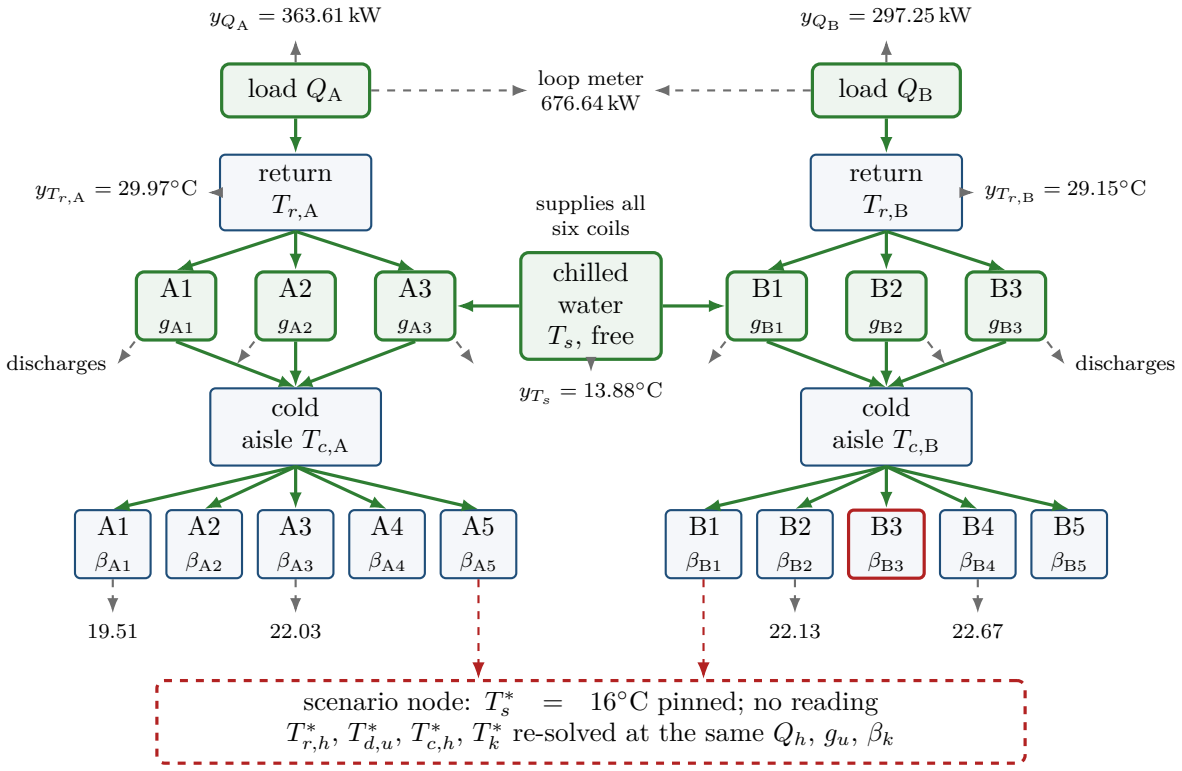


Figure 12.1: The rows of the shift plant at window 73, 01:00. Free (green): the supply T_s, the loads Q_h and the log-derates g_u; the recirculation fractions β_k are free entries carried by the rack nodes, and a rack's mix of cold aisle and return is its inlet row, not drawn from the return. Blue: temperatures fixed by the rows (12.2). Observed (dashed gray): the supply, two returns, six discharges, two hall meters, the loop meter and the inlets of racks A1, A3, B2 and B4, in °C and kW. Not drawn: the pinned ratings, capacity rates and other heat, and the priors. Red: rack B3, and the scenario node with the arrows that give it the state's entries.

Table 12.1: The questions asked of window 73. $\mathbf{y}$ is the sixteen readings and the seventeen priors; T_k^* the inlet of rack k at the node (12.1); the limit is 27°C.

Operator question	General form	Mathematical form	Category	Computation	Answered in
If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit?	(2.12), a joint tail at a scenario node	$\Pr(\max_k T_k^* < 27^\circ\mathrm{C} \mid \mathbf{y},\ T_s^* = 16^\circ\mathrm{C})$	prediction under a threshold	posterior	12.3.1
Which rack carries it, and what stands behind its band?	(2.9), (2.11)	T_k^* and its band; the budget of $T_{\mathrm{B}3}^*$	estimation	posterior	12.3.2
Is the raise seen by the sensors on the screen?	(2.12) on two subsets	the measured four against the unread six; the sensors see the raise and not the tail	prediction under a threshold	posterior	12.3.3
Does the probability fall with the raise, and how large a raise holds at a stated assurance?	(2.12) read along T_s^*	$\Pr(\cdot \mid T_s^*)$; the largest T_s^* with $\Pr \geq \kappa$	prediction under a threshold	posterior	12.3.4
Is the raise a pinned value or a shift?	(2.12), two nodes	$T_s^* = 16^\circ\mathrm{C}$ against $T_s + 2\,\mathrm{K}$	prediction under a threshold	posterior	12.3.5

Table 12.2: The scenario node of the raise.

Quantity	At this node	Why
T_s^*	pinned, 16.0°C	the actuator; the set point the controller is told
Q_{A}, Q_{B}	free, window 73's entries	loads do not respond to the water
g_u	free, window 73's entries, assumed independent of T_s	a choice, below
β_k	free, window 73's entries	the air pattern does not respond to the water
C_u, $Q_{o,h}$	pinned as in the plant	fixed fans, other heat
every temperature	re-solved	$T_{r,h}^*$, $T_{d,u}^*$, $T_{c,h}^*$, T_k^*
loop and unit heats	unchanged at the node	Q_u depends on $T_{r,h} - T_s$, which does not move

The physics before any number. By the chain (12.4), every temperature is T_s plus a quantity that does not depend on T_s, so at the node

$$T_k^* = T_s^* + \frac{Q_h + Q_{o,h}}{\sum_{u \in h} G_u} - (1 - \beta_k)\, \frac{Q_h + Q_{o,h}}{\sum_{u \in h} C_u}, \tag{12.6}$$

the second term the rise from the supply to the return and the third the rise (12.5) scaled. Every row is linear in T_s with coefficient one: a raise moves every temperature of the chain by the same amount and no unit heat at all. A raised supply must therefore make the probability fall. The chapter's physics is that simple, and its content is that the amount is added to a rack that no sensor reads.

Modeling choices at the node. Four choices make this node: the coil and the actuator are varied in Section 12.4.3, the pin in Section 12.3.5; the loads are not varied (Section 12.6.4).

The conductance independent of the supply temperature. This buys the coefficient one above. The record cannot test it: its supply readings span 13.75 to 14.23°C with a standard deviation of 0.102 K against the reading's 0.10 K, so the spread cannot be told from reading noise, and the node lies 1.8 K above the largest of them. It fails if warmer water lowers a coil's effective conductance, through less latent duty or less temperature difference at the valve.

The set point pinned at 16.0°C*, not the entry plus* 2 K. A controller holds a set point, and the operator enters that number. This takes the entry T_s out of the gradient of every scenario temperature; what that does to the band is computed in Section 12.3.5, where the shift stands beside it.

The loads held at window 73's values. The node is the plant at 01:00 with a different set point. This keeps the question to the actuator, and it fails as a statement about any later window, since it holds no load change.

The actuator exact. The node's T_s^* has no spread. This keeps the answer to what the fit knows; it fails to the extent the loop misses the value it is told.

The limit and the assurance. The limit, 27°C, is the plant's. The assurance κ in equation (2.12) belongs to the operator. The chapter computes at the stated values 0.90, 0.95 and 0.99 and names none of them as the rule.

Which row acts on which entry. Figure 12.2 draws the sixteen reading rows of the Jacobian on the 19 columns: each row of the figure is a reading, each column an entry, and a filled square a nonzero; the seventeen prior rows, each a single entry on its own column, are not drawn. Each hall's temperature rows see T_s, that hall's load and its three coils; a measured inlet adds its own fraction; the meters see the loads only. The six shaded columns are empty.

The block structure reads as observability. The six fractions of the unread racks are not observable from the readings, and only their prior rows hold them; the two halls share the column of T_s and, through the loop meter, those of the loads. The question's gradient is not a row. It has the pattern

$$\mathbf{g}_{T^*_{\mathrm{B3}}}^{\mathsf{T}} = (0,\ 0,\ *,\ 0,\ 0,\ 0,\ *,\ *,\ *,\ \underbrace{0,\ \ldots,\ 0}_{\beta_{\mathrm{A1}},\ldots,\beta_{\mathrm{B2}}},\ *,\ 0,\ 0), \tag{12.7}$$

with entries on Q_{B}, the three coils of hall B and β_{B3}, one of the empty columns, and none on T_s: the pin removes it.

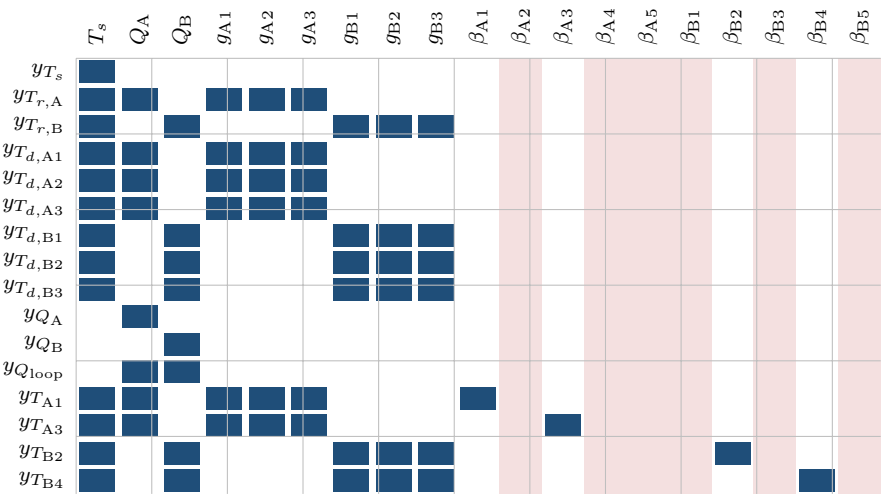


Figure 12.2: The sparsity of the reading rows of $\mathbf{J}$. Shaded: the columns of β_{A2}, β_{A4}, β_{A5}, β_{B1}, β_{B3} and β_{B5}, in which no reading has a nonzero.

12.2 The posterior of the state (step 3)

This is the fit of window 73, 01:00, computed by the plant's own rows, and it is the one fit Part IV reads. No concrete reason against this window was found: its test passes. Everything an answer below uses from the record's events is computed here too.

12.2.1 The rows and the normal equations

Stacked in the order of Figure 12.2, each divided by its standard deviation, the rows on the reduced state are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_s - y_{T_s})/0.1 \\ (T_{r,h}(\mathbf{z}) - y_{T_{r,h}})/0.3 \quad (h = \mathrm{A}, \mathrm{B}) \\ (T_{d,u}(\mathbf{z}) - y_{T_{d,u}})/0.3 \quad (\text{six units}) \\ (Q_h - y_{Q_h})/4 \quad (h = \mathrm{A}, \mathrm{B}) \\ (Q_{\mathrm{A}} + Q_{\mathrm{B}} + Q_{o,\mathrm{A}} + Q_{o,\mathrm{B}} - y_{Q_{\mathrm{loop}}})/8 \\ (T_k(\mathbf{z}) - y_{T_k})/0.3 \quad (k = \mathrm{A1}, \mathrm{A3}, \mathrm{B2}, \mathrm{B4}) \\ (T_s - 14)/0.5 \\ g_u/0.15 \quad (\text{six units}) \\ (\beta_k - \beta_k^{\mathrm{s}})/0.08 \quad (\text{ten racks}) \end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \tag{12.8}$$

with the temperatures from the chain (12.4) and the reduced Jacobian $\mathbf{J}$ of equation (4.47). The MAP state minimizes $\mathcal{L}$ by the normal equations of equation (4.29),

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \qquad \boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}\big|_{\hat{\mathbf{z}}}, \tag{12.9}$$

from the set point, the hall meters, the ratings and the survey, the Gauss–Newton form of equation (4.38); the posterior is approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. The empty columns of Figure 12.2 separate the objective: the terms in the six unread fractions are their prior terms alone, so their posterior is the survey's exactly, not only under the linearization, and they are uncorrelated with every other entry, as in equation (5.10).

The exact posterior. The density $p(\mathbf{z} \mid \mathbf{y}) \propto \exp(-\mathcal{L}(\mathbf{z}))$ has no closed form on 19 entries. It is set beside the Laplace normal by importance sampling: 400 000 draws from that normal, each weighted by the ratio of the two densities, with an effective size of 396 168 and a largest weight of 7.58×10^{-6}. Each draw is pushed through the chain with T_s replaced by the pin.

12.2.2 The fit of window 73

The iterations. Table 12.3 runs equation (12.9). Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 31.8994 to 1.3134 in the first step and settles at 1.2748 by the second, and the norm falls to 8.46×10^{-9} by the fifth.

Table 12.3: Gauss–Newton iterations of window 73, equation (12.9).

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	31.8994	4.37×10^{2}
1	1.3134	1.98×10^{1}
2	1.2748	1.43×10^{-2}
3	1.2748	6.52×10^{-5}
4	1.2748	3.22×10^{-7}
5	1.2748	8.46×10^{-9}

The Jacobian at the MAP state. Table 12.4 gives the reading rows of $\mathbf{J}(\hat{\mathbf{z}})$ on the thirteen observed columns: each row a reading over its standard deviation, each cell its derivative in that column. The six unread fractions' columns hold zeros in these rows. The prior rows are single entries on their own columns: 2 for T_s, 1/0.15 for each coil and 1/0.08 for each fraction. An inlet's row repeats its hall's return row in the coil columns, because by the air balance (12.5) a derate leaves the rise unchanged.

Table 12.4: The reading rows of $\mathbf{J}(\hat{\mathbf{z}})$ on the observed columns, window 73; the Q columns in kW^{-1}, the others per unit of the entry, T_s per K.

[–]	T_s	Q_{A}	Q_{B}	g_{A1}	g_{A2}	g_{A3}	g_{B1}	g_{B2}	g_{B3}	β_{A1}	β_{A3}	β_{B2}	β_{B4}
y_{T_s}	10.000	0	0	0	0	0	0	0	0	0	0	0	0
$y_{T_{r,\mathrm{A}}}$	3.333	0.142	0	−16.977	−18.907	−17.358	0	0	0	0	0	0	0
$y_{T_{r,\mathrm{B}}}$	3.333	0	0.166	0	0	0	−18.487	−15.409	−17.038	0	0	0	0
$y_{T_{d,\mathrm{A1}}}$	3.333	0.036	0	−44.140	−4.745	−4.357	0	0	0	0	0	0	0
$y_{T_{d,\mathrm{A2}}}$	3.333	0.024	0	−2.815	−47.549	−2.878	0	0	0	0	0	0	0
$y_{T_{d,\mathrm{A3}}}$	3.333	0.033	0	−3.976	−4.428	−44.839	0	0	0	0	0	0	0
$y_{T_{d,\mathrm{B1}}}$	3.333	0	0.045	0	0	0	−42.072	−4.208	−4.653	0	0	0	0
$y_{T_{d,\mathrm{B2}}}$	3.333	0	0.066	0	0	0	−7.287	−36.932	−6.716	0	0	0	0
$y_{T_{d,\mathrm{B3}}}$	3.333	0	0.055	0	0	0	−6.102	−5.086	−39.746	0	0	0	0
$y_{Q_{\mathrm{A}}}$	0	0.250	0	0	0	0	0	0	0	0	0	0	0
$y_{Q_{\mathrm{B}}}$	0	0	0.250	0	0	0	0	0	0	0	0	0	0
$y_{Q_{\mathrm{loop}}}$	0	0.125	0.125	0	0	0	0	0	0	0	0	0	0
$y_{T_{\mathrm{A1}}}$	3.333	0.050	0	−16.977	−18.907	−17.358	0	0	0	41.689	0	0	0
$y_{T_{\mathrm{A3}}}$	3.333	0.072	0	−16.977	−18.907	−17.358	0	0	0	0	41.689	0	0
$y_{T_{\mathrm{B2}}}$	3.333	0	0.090	0	0	0	−18.487	−15.409	−17.038	0	0	34.001	0
$y_{T_{\mathrm{B4}}}$	3.333	0	0.096	0	0	0	−18.487	−15.409	−17.038	0	0	0	34.001

The MAP state. Table 12.5 gives the fit: each row an entry, the columns its value, its posterior spread, its prior spread and the rows that inform it. Four fractions move and narrow under their inlet sensors; the other six stay at the survey's values with the survey's spread. At the MAP state $T_{r,\mathrm{A}} = 29.8674$, $T_{r,\mathrm{B}} = 29.1749$, $T_{c,\mathrm{A}} = 17.3607$ and $T_{c,\mathrm{B}} = 18.9745$°C; the conductances sum to 23.4901 and 20.0268 kW/K; and the rises from the supply to the return are 15.9727 and 15.2802 K.

The covariance. On the observed entries the variances, the diagonal of $\boldsymbol{\Sigma}$, are 9.5167×10^{-3} for T_s; 11.794 and 11.766 for Q_{A} and Q_{B}; 5.4524, 4.6074, 5.2627, 6.2183, 8.4368 and 7.1135×10^{-4} for the six coils in order; and 6.4333, 6.4406, 9.1888 and 9.1992×10^{-4} for the four read fractions. Table 12.6 gives their correlations, each cell the correlation of its row's entry and its column's.

Table 12.5: The MAP state of window 73, 01:00, with the posterior and prior spreads.

Entry [unit]	$\hat{z}_i$ [entry's]	$\sqrt{\Sigma_{ii}}$ [entry's]	Prior sd [entry's]	Informed by
T_s [°C]	13.8947	0.0976	0.5	supply reading
Q_{A} [kW]	363.2006	3.4343	none	hall and loop meters, hall A's temperatures
Q_{B} [kW]	296.0123	3.4302	none	hall and loop meters, hall B's temperatures
g_{A1} [–]	−0.0659	0.0234	0.15	hall A's temperatures
g_{A2} [–]	0.0418	0.0215	0.15	hall A's temperatures
g_{A3} [–]	−0.0437	0.0229	0.15	hall A's temperatures
g_{B1} [–]	0.0377	0.0249	0.15	hall B's temperatures
g_{B2} [–]	−0.1444	0.0290	0.15	hall B's temperatures
g_{B3} [–]	−0.0439	0.0267	0.15	hall B's temperatures
β_{A1} [–]	0.1698	0.0254	0.08	inlet sensor and survey
β_{A3} [–]	0.3715	0.0254	0.08	inlet sensor and survey
β_{B2} [–]	0.3082	0.0303	0.08	inlet sensor and survey
β_{B4} [–]	0.3665	0.0303	0.08	inlet sensor and survey
β_{A2}, β_{A4}, β_{A5} [–]	0.25, 0.30, 0.45	0.0800	0.08	survey only
β_{B1}, β_{B3}, β_{B5} [–]	0.20, 0.50, 0.35	0.0800	0.08	survey only

Table 12.6: Correlations of $\boldsymbol{\Sigma}$ on the observed entries, window 73. The six unread fractions have a variance of 0.08^2 and no correlation with any entry.

[–]	T_s	Q_{A}	Q_{B}	g_{A1}	g_{A2}	g_{A3}	g_{B1}	g_{B2}	g_{B3}	β_{A1}	β_{A3}	β_{B2}	β_{B4}
T_s	1.00	−0.00	−0.00	0.26	0.28	0.26	0.25	0.22	0.24	−0.00	−0.00	−0.00	−0.00
Q_{A}	−0.00	1.00	−0.18	0.18	0.15	0.17	−0.03	−0.04	−0.04	0.02	−0.04	0.00	0.01
Q_{B}	−0.00	−0.18	1.00	−0.03	−0.03	−0.03	0.19	0.23	0.21	−0.00	0.01	−0.02	−0.04
g_{A1}	0.26	0.18	−0.03	1.00	−0.10	−0.12	0.06	0.05	0.05	0.17	0.16	−0.00	0.00
g_{A2}	0.28	0.15	−0.03	−0.10	1.00	−0.10	0.07	0.05	0.06	0.18	0.17	−0.00	0.00
g_{A3}	0.26	0.17	−0.03	−0.12	−0.10	1.00	0.06	0.05	0.06	0.17	0.16	−0.00	0.00
g_{B1}	0.25	−0.03	0.19	0.06	0.07	0.06	1.00	−0.17	−0.15	−0.00	0.00	0.15	0.15
g_{B2}	0.22	−0.04	0.23	0.05	0.05	0.05	−0.17	1.00	−0.18	−0.00	0.00	0.13	0.12
g_{B3}	0.24	−0.04	0.21	0.05	0.06	0.06	−0.15	−0.18	1.00	−0.00	0.00	0.14	0.14
β_{A1}	−0.00	0.02	−0.00	0.17	0.18	0.17	−0.00	−0.00	−0.00	1.00	0.18	0.00	0.00
β_{A3}	−0.00	−0.04	0.01	0.16	0.17	0.16	0.00	0.00	0.00	0.18	1.00	−0.00	−0.00
β_{B2}	−0.00	0.00	−0.02	−0.00	−0.00	−0.00	0.15	0.13	0.14	0.00	−0.00	1.00	0.17
β_{B4}	−0.00	0.01	−0.04	0.00	0.00	0.00	0.15	0.12	0.14	0.00	−0.00	0.17	1.00

The rows fix what to read in it. The supply correlates with every coil at +0.22 to +0.28: the return is read, so a lower supply leaves a larger difference to the return, and the same heat then needs a smaller conductance. The two loads correlate at −0.18 through the loop meter, which reads their sum. Across the two halls the other correlations are at most 0.07 in size, so the halls are otherwise nearly separate. The fractions read by an inlet sensor correlate with their own hall's coils; the six unread fractions correlate with nothing.

12.2.3 What the record shows of its events

These are the methods of Parts II and III computed at 01:00, stated once here for Part IV. A one-window fit has its own loads, so a load that drifts between windows is no misspecification of it; it would matter only to a forecast. Fitted one window at a time, the entries do drift: per window, 0.28135 (0.02802) kW on Q_{A} and 0.19141 (0.02404) kW on Q_{B} over windows 1 to 73, and −0.00333 (0.00044) on g_{B2} over windows 37 to 73, by least squares whose standard errors ignore autocorrelation.

The window's own test. The chi-square test of Chapter 7, on the reading rows with the effective degrees of freedom ν of equation (7.14), gives $X^2 = 0.755$ on 3.691, an upper tail of 0.925, and a largest normalized residual, equation (7.15), of −0.718 on the discharge of unit B2. The hall-A meter's own is −0.20. The window's test sees nothing.

The coil of unit B2. The refit stacks windows 1 to 36, with the shared coils and fractions as one entry each and T_s, the two loads and g_{B2} free per window. The mean of g_{B2} over those first three hours is −0.0022. Window 73's entry, −0.1444, lies −4.90 of its spreads from that mean, a difference of −0.1423. A second refit of window 73 with g_{B2} pinned at that mean is misspecified: its reading rows give $X^2 = 28.421$ on effective degrees of freedom of 4.644, with an upper tail below 0.001 and a largest normalized residual of −5.078 on the discharge of B2, and the supply reading's is +4.12. The free entry absorbs the derate, so the plain fit needs no further term. The derate is therefore already in every scenario read from this fit, and what is in it is the coil's level at 01:00, not a slope.

The hall-A meter. The meter reconciliation of Section 7.2.5 does not act on one window; it acts on the record stacked to window T, the same stacked rows with one constant bias entry on one meter, prior $\mathcal{N}(0,\ 10^2)$ in kW, and the Laplace Bayes factor of equation (7.17) against no bias. Each window of that stack has eighteen rows: its sixteen readings, its T_s prior and its g_{B2} prior, with one prior row for each shared entry. Table 12.7 gives the results. Each row is a stack; each of the three meters has its logarithm of the Bayes factor, its bias estimate and that estimate's spread. Figure 12.3 draws the logarithms against T, with the level $\ln 20 = 2.9957$ dashed.

Only the hall-A meter passes the level. It is past it at $T = 24$, under it at 36, past it again at 48, and above 10 at 72 and 73, where it stands 8.742 above the loop meter and 12.209 above the hall-B meter. This is one night's draw. Over 100 nights of fresh sensor errors on the same truth, $\ln B$ of this meter has medians 1.512, 3.085, 4.379, 5.986 and 7.725 at $T = 24$, 36, 48, 61 and 73, and the fractions of nights past the level are 0.33, 0.52, 0.73, 0.86 and 0.86. At $T = 73$ the range from the 10th to the 90th percentile is 2.681 to 13.968: this night's 10.139 is above the median, and 0.14 of nights are not yet past the level at 01:00. So no first window is reported as a date; by 01:00 this night's record shows a high hall-A meter, and the one-window fit absorbs a part of it.

The fit carrying the bias. A refit of window 73 subtracts one bias entry from the meter's reading, its prior the stack of windows 1 to 72, 4.636 (0.918) kW. It gives $Q_{\mathrm{A}} = 359.815$ (3.499) kW

Table 12.7: The meter reconciliation on the record stacked to window T: $\ln B$ against no bias [–], bias estimate and its spread [kW].

		y_{Q_A}		y_{Q_B}		$y_{Q_{loop}}$	
T [–]	Clock [h:min]	$\ln B$ [–]	Bias (sd) [kW]	$\ln B$ [–]	Bias (sd) [kW]	$\ln B$ [–]	Bias (sd) [kW]
12	19:55	−1.198	1.74 (2.20)	−0.670	2.73 (2.20)	−1.204	−1.34 (2.67)
24	20:55	4.048	5.43 (1.58)	−1.273	1.58 (1.57)	−0.144	−3.37 (1.92)
36	21:55	2.881	4.08 (1.29)	−2.043	0.05 (1.29)	−1.184	−1.85 (1.58)
48	22:55	4.065	3.99 (1.12)	−2.142	−0.46 (1.12)	−1.271	−1.67 (1.37)
61	00:00	5.227	3.90 (1.00)	−2.294	0.09 (1.00)	−0.720	−2.05 (1.22)
72	00:55	10.197	4.64 (0.92)	−2.114	0.59 (0.92)	1.239	−2.96 (1.12)
73	01:00	10.139	4.60 (0.91)	−2.070	0.64 (0.91)	1.396	−3.01 (1.11)

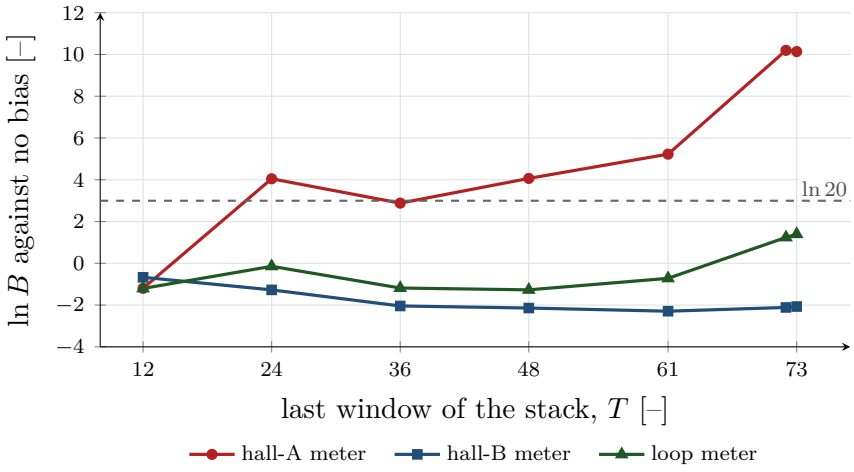


Figure 12.3: The logarithm of the Bayes factor of a bias on each meter against no bias, on the record stacked to window T (Table 12.7).

against the plain fit's 363.201 (3.434) kW, a difference of −3.385 kW, −0.99 plain spreads, and through the loop meter $Q_{\mathrm{B}} = 296.614$ against 296.012 kW.

The fit the answers read. The reported fit is the plain fit of the plant's rows, for three reasons: it passes its window's test, the one fit of a Part should be the one its plant states, and the bias is the size of the meter's noise. The fit carrying the bias is a second named route, set beside each answer.

12.3 The answers (step 4)

Every answer is read from the pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 12.2.2, and that value is the one reported. A scenario temperature is carried by its gradient at $\hat{\mathbf{z}}$, equation (12.1). The exact posterior's draws stand beside each reported value where they were run, which is at the 16.0°C node.

12.3.1 The inlets at the raised set point

Operator question. If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit? The question is about all ten inlets at once, at the set point the operator would enter, at this window's loads and coils.

General form. A prediction under a threshold, equation (2.12), here on a maximum as in equation (8.6), and read on the requirement's side.

Specialized form. The question is equation (2.12) with $\mathbf{z}$ the reduced state (12.3), $f(\mathbf{z}) = -\max_k T_k^*(\mathbf{z})$ at $T_s^* = 16°\mathrm{C}$, $f^* = -27°\mathrm{C}$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal (12.9). The assurance κ is the operator's. By equation (12.6) the inlet's gradient has the entries

$$\frac{\partial T_k^*}{\partial T_s} = 0, \quad \frac{\partial T_k^*}{\partial Q_h} = \frac{1}{\sum_{u\in h} G_u} - \frac{1-\beta_k}{\sum_{u\in h} C_u}, \quad \frac{\partial T_k^*}{\partial g_u} = -\frac{G_u\,(T_{r,h} - T_s)}{\sum_{v\in h} G_v}, \quad \frac{\partial T_k^*}{\partial \beta_k} = T_{r,h} - T_{c,h}, \tag{12.10}$$

with every difference $T_{r,h} - T_s$ and $T_{r,h} - T_{c,h}$ at $\hat{\mathbf{z}}$, and they do not depend on the pin. Stacking the ten gradients as the rows of $\mathbf{D}^*$, as in equation (8.9), the inlets are jointly normal, and the requirement is the corner of that normal,

$$\begin{gathered} \mathbf{T}^* \mid \mathbf{y} \approx \mathcal{N}(\mathbf{m}^*,\ \mathbf{D}^*\boldsymbol{\Sigma}\mathbf{D}^{*\mathsf{T}}), \\ \Pr(\max_k T_k^* < 27°\mathrm{C} \mid \mathbf{y},\ T_s^* = 16°\mathrm{C}) = \Pr(T_k^* < 27°\mathrm{C} \text{ for all ten} \mid \mathbf{y}), \end{gathered} \tag{12.11}$$

with $\mathbf{m}^*$ the ten inlets at $\hat{\mathbf{z}}$. One choice enters here: the joint is the delta method's normal, which gives the corner from $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and one matrix, and whose linearity the draws test. A corner of ten dimensions has no closed form; it is evaluated by seeded numerical integration of the normal, with an absolute error requested of 10^{-6}. The union bracket (8.13) applies to its complement with ten terms.

Building blocks. The MAP state and $\boldsymbol{\Sigma}$ (Section 12.2.2); the gradients (12.10); the integration of the corner.

The joint at the node. By equation (12.11),

$$\Pr(\text{all ten} < 27°\mathrm{C} \mid \mathbf{y},\ T_s^* = 16°\mathrm{C}) = 0.80983,$$

and the exact posterior's draws give 0.81026 (se 0.00062), a gap of +0.00043. The exceedance is 0.19017, between the largest marginal exceedance, 0.16300 at rack B3, and the union bound,

0.19587. No stated assurance from 0.90 up is met, on either value. The fit carrying the bias gives 0.81013.

12.3.2 The rack that carries the tail

Operator question. Which rack carries it, and what stands behind its band? The question is which inlet comes nearest the limit at the node, and which entries of the state make its spread.

General form. An estimation of a derived quantity, equation (2.9), with its central interval, equation (2.11), and its variance split by entry, as in equation (5.14).

Specialized form. The question is equations (2.9) and (2.11) with $f(\mathbf{z}) = T_k^*(\mathbf{z})$ and the Laplace normal, so the band is $T_k^*(\hat{\mathbf{z}}) \pm 1.960\,\mathrm{sd}$ by equation (12.1). The budget of T_{B3}^* has the gradient's five nonzero entries,

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \sum_i [\mathbf{g}]_i^2\,\Sigma_{ii} + \sum_{i \neq j} [\mathbf{g}]_i [\mathbf{g}]_j \Sigma_{ij}, \tag{12.12}$$

with no cross term in β_{B3}, whose column is empty. The gradient's signs fix those of the others. On hall B's coils it is negative, and the coils' mutual correlations are negative, so each such term is a negative covariance on two same-signed gradients. On Q_{B} the gradient is positive and its correlations with the coils are positive, so those terms are a positive covariance on opposite-signed gradients. Every cross term is negative.

Building blocks. The gradients (12.10); $\boldsymbol{\Sigma}$ (Section 12.2.2); the budget (12.12).

The ten inlets. Table 12.8 gives all ten: each row a rack, the columns its inlet at the present state with its spread, its inlet at the node with its spread and band, its probability of staying under the limit, and the exact posterior's median and probability. Figure 12.4 draws them, with the horizontal axis the racks, the vertical axis the inlet temperature, a gray bar for the present band, a blue bar for the band at the node, a red line at the limit, and a dotted tick for the stated truth at the node, a check marker for Section 12.4.4. The unread racks' bars are longer than the read racks', and the bars move up together.

Table 12.8: The ten inlets at the present state and at the node $T_s^* = 16\,^{\circ}\mathrm{C}$, window 73. Racks A1, A3, B2 and B4 have sensors. Medians and probabilities beside them from the exact posterior.

Rack [–]	T_k [°C]	sd [K]	T_k^* [°C]	sd [K]	2.5 % [°C]	97.5 % [°C]	Pr(< 27) [–]	Exact median [°C]	Exact Pr [–]
A1	19.4840	0.2876	21.5893	0.3037	20.9941	22.1845	1.00000	21.5881	1.00000
A2	20.4874	1.0112	22.5927	1.0157	20.6018	24.5835	0.99999	22.5888	0.99999
A3	22.0065	0.2876	24.1118	0.3037	23.5166	24.7070	1.00000	24.1108	1.00000
A4	21.1127	1.0112	23.2180	1.0158	21.2271	25.2089	0.99990	23.2175	0.99991
A5	22.9887	1.0114	25.0940	1.0160	23.1026	27.0854	0.96967	25.0899	0.97002
B1	21.0145	0.8288	23.1198	0.8344	21.4844	24.7553	1.00000	23.1165	0.99999
B2	22.1185	0.2821	24.2238	0.2985	23.6389	24.8088	1.00000	24.2223	1.00000
B3	24.0747	0.8293	26.1800	0.8349	24.5436	27.8164	0.83700	26.1767	0.83720
B4	22.7126	0.2821	24.8179	0.2985	24.2330	25.4029	1.00000	24.8175	1.00000
B5	22.5446	0.8289	24.6499	0.8345	23.0144	26.2855	0.99757	24.6458	0.99758

Rack B3. Rack B3 is nearest the limit, and the upper end of its band in Table 12.8 lies past it; rack A5 comes next and rack B5 after it. The fit carrying the bias gives rack B3

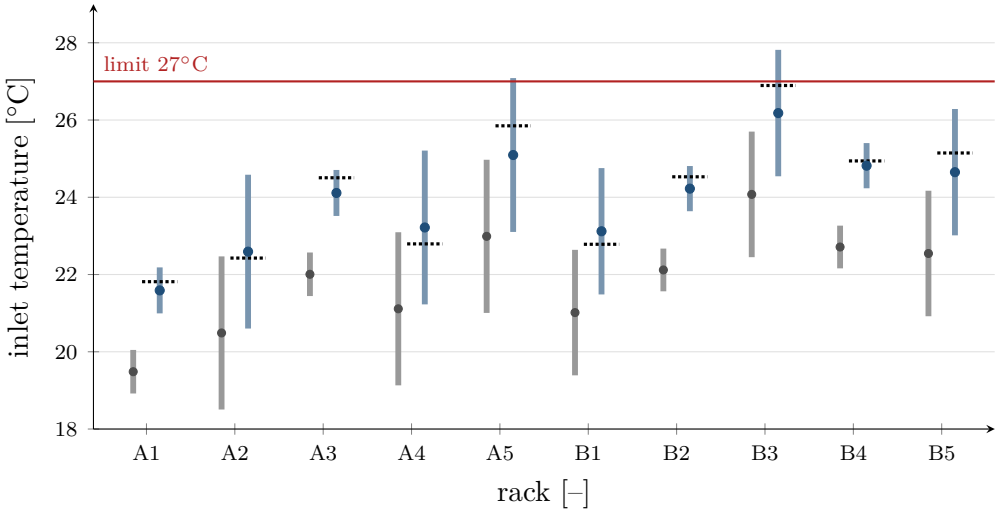


Figure 12.4: The ten inlets with their central 95 % bands. Gray: the present state. Blue: the node $T_s^* = 16°\mathrm{C}$. Dotted: the stated truth at the node, a check marker (Section 12.4.4).

26.1849 (0.8365)°C and 0.83509, and rack A5 25.0707 (1.0071)°C and 0.97229. By equation (12.10) the gradient of T^*_{B3} is 0.0333 K/kW on Q_{B}, −5.5461, −4.6226, −5.1115 K on g_{B1}, g_{B2}, g_{B3}, and 10.2004 K on β_{B3}.

The budget. Figure 12.5 draws the budget (12.12), one bar per entry and one for the cross terms, the horizontal axis in percent of the variance $0.69705\,\mathrm{K}^2$. The terms are $0.66591\,\mathrm{K}^2$ for β_{B3}, 0.01913, 0.01859 and 0.01803 for g_{B1}, g_{B3} and g_{B2}, 0.01302 for Q_{B}, and −0.03762 for the cross terms. With β_{B3} known the spread is 0.1765 K. The answer is one survey's fraction, 0.50 ± 0.08.

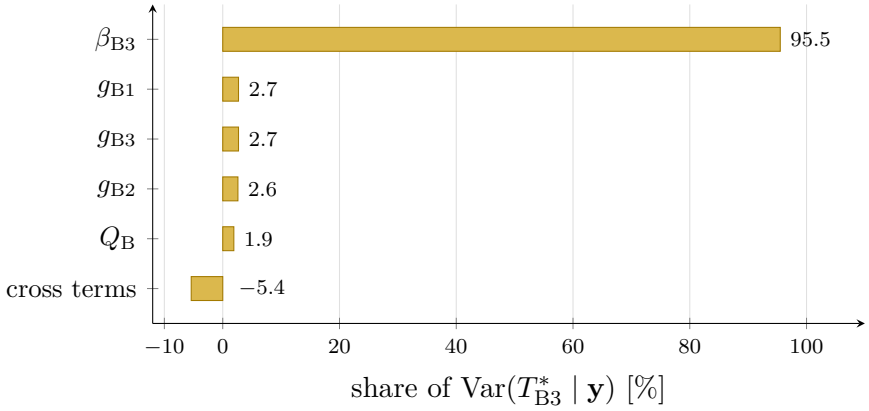


Figure 12.5: The variance budget of the inlet of rack B3 at the node, equation (12.12), in percent of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$.

The correlations at the node. The correlation of T^*_{A5} and T^*_{B3} is 0.0108, and that of T^*_{B3} and T^*_{B4} is 0.0479. Hall A alone has the joint 0.969566 and hall B alone 0.835050, and their product, independence of the halls, is 0.809636 against the joint's 0.80983; the fit carrying the bias gives the halls 0.97221 and 0.83310. Here the joint adds little to the marginals, because the racks that carry the tail share little beyond entries the readings have fixed, while Section 8.3.4 found otherwise on a poorly read plant.

END OF THE PREVIEW

Operations and Control by Derivation

These are the opening pages. The complete book runs 220 pages.

The full PDF opens with a token. Write to **rmarquez123@gmail.com**, name the book, and I will send you one.

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com