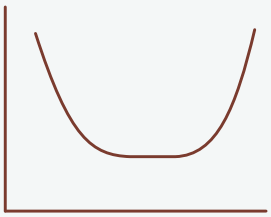
PROBABILISTIC GRAPH MODELING
FOR ENGINEERING DISCIPLINES · VOLUME II



Reliability Engineering by Derivation

Physical models, diagnosis, risk and acceptance

431 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Reliability Engineering by Derivation

Volume II

Physical models, diagnosis, risk and acceptance

A lesson sequence, written up as a book

Ricardo Marquez

October 31, 2026 (work in progress: one chapter per lesson, added as the lessons proceed)

How to read this book

This book teaches one skill: to take a reliability question stated in ordinary engineering language, 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 IV, Chapters 1 to 16: the method itself (a plant graph, a question as a probability over its state, and the five computations beneath every answer: the posterior, the trend, the evidence, the preposterior and the expected consequence), and its application to records of failures, lifetimes and repairs, to predictions of mission reliability, remaining life and future counts, and to the structure of systems (enumeration, cut sets, common cause, Markov chains). This second volume holds Parts V to VII, Chapters 17 to 28. It turns to the physical model itself: the state of a plant fitted to conservation relations and readings, the curvature that tells which entries the readings determine, the propagation of that uncertainty to engineering outputs; then to diagnosis (virtual sensors, bad data, fault isolation, model selection); and to risk, importance measures, acceptance tests and operating margins. Volume III holds Parts VIII and IX, Chapters 29 to 36: the value of information and the engineering decisions it supports, and four integrated case studies that run the whole method on one plant each. Chapter numbers run continuously across the three volumes; a reference to a chapter numbered 16 or below points to Volume I, whose results this volume builds on and does not repeat, and a chapter numbered 29 or above is named in words.

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 evidence attached as rows, with the assumptions they rest on; step 3, the posterior of the state, derived and computed once (the MAP state, the information and the covariance, with the exact posterior where one exists); step 4, the answers, each question read from that posterior down to its number (an entry or a gradient, a ratio of evidences, or an expected consequence); step 5, evaluation and checks against the full calculation; step 6, the report with its band: the probability, the Bayes factor, the expected consequence, or a crossing time with its band by inversion; step 7, the next measurement, ranked by the width it removes. The five kinds of question (estimation, prediction under a threshold, hypothesis comparison, diagnosis, design) and the five computations underneath them (the posterior, the trend, the evidence, the preposterior, the expected consequence) are those of Volume I. The notation chapter is reproduced here so that this volume can be read with Volume I closed, and the atlas of questions in the appendix now covers Volumes I and II.

The register is unchanged: every new quantity is introduced with its name, meaning, units and equation; every derivation proceeds one relation per line; every numerical example is worked with units; a *Jargon* box expands each shorthand term once. The plants of Volume I return: the equipment hall of Chapter 1, now two halls on one chilled-water loop with their sensors and meters; the pumps, valves and generators whose records were the evidence; and the redundant cooling plant, which grows into a data center in the case studies.

The framework the book teaches is physics-constrained probabilistic modeling: a plant written as a graph, its unknowns as a state vector, every piece of evidence 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.

Each chapter closes with a statement of what to retain, a table of its terms with their meaning, units and the section that defines each, and exercises at three levels, formulation, derivation and application, each followed by its answer.

Contents

How to read this book **ii**

Notation and vocabulary **vii**

V Physical Models and Uncertainty Propagation **1**

17 Physics-constrained state estimation **2**
17.1 State estimation on a physical model 2
17.2 The posterior of the state (step 3) 13
17.3 The answers (step 4) 19
17.4 The checks (step 5) 26
17.5 The report (step 6) 29
17.6 The next measurement (step 7) 30
17.7 What to retain 33
17.8 Exercises 35

18 Maximum a posteriori estimation **37**
18.1 Fitting the state of a plant to its readings 37
18.2 The posterior of the state (step 3) 42
18.3 The answers (step 4) 47
18.4 The checks (step 5) 55
18.5 The report (step 6) 60
18.6 The next measurement (step 7) 61
18.7 What to retain 64
18.8 Exercises 66

19 The Hessian, covariance and identifiability **68**
19.1 The spread of a fitted state and what the readings determine 68
19.2 The posterior of the state (step 3) 73
19.3 The answers (step 4) 77
19.4 The checks (step 5) 90
19.5 The report (step 6) 96
19.6 The next measurement (step 7) 98
19.7 What to retain 101
19.8 Exercises 102

20 Propagating uncertainty to outputs **105**
20.1 Hall 1 at twice its present load 105
20.2 The posterior of the state (step 3) 110
20.3 The answers (step 4) 113

20.4 The checks (step 5) . . . 121
20.5 The report (step 6) . . . 126
20.6 The next measurement (step 7) . . . 128
20.7 What to retain . . . 132
20.8 Exercises . . . 134

VI Diagnosis and Competing Hypotheses 137

21 Virtual sensors and unobserved states 138
21.1 Quantities where no sensor is installed . . . 138
21.2 The posterior of the state (step 3) . . . 143
21.3 The answers (step 4) . . . 146
21.4 The checks (step 5) . . . 153
21.5 The report (step 6) . . . 159
21.6 The next measurement (step 7) . . . 160
21.7 What to retain . . . 164
21.8 Exercises . . . 165

22 Sensor bias and bad data 168
22.1 Readings that may be lying . . . 168
22.2 The posterior of the state under each drawing (step 3) . . . 174
22.3 The answers (step 4) . . . 178
22.4 The checks (step 5) . . . 194
22.5 The report (step 6) . . . 198
22.6 The next measurement (step 7) . . . 198
22.7 What to retain . . . 203
22.8 Exercises . . . 204

23 Fault detection and isolation 207
23.1 A plant that has changed . . . 207
23.2 The posterior of the state under each drawing (step 3) . . . 215
23.3 The answers (step 4) . . . 217
23.4 The checks (step 5) . . . 230
23.5 The report (step 6) . . . 234
23.6 The next measurement (step 7) . . . 235
23.7 What to retain . . . 239
23.8 Exercises . . . 241

24 Model selection and Bayes factors 244
24.1 Questions about the drawing . . . 244
24.2 The posterior of the state (step 3) . . . 252
24.3 The answers (step 4) . . . 259
24.4 The checks (step 5) . . . 269
24.5 The report (step 6) . . . 275
24.6 The next measurement (step 7) . . . 277
24.7 What to retain . . . 280
24.8 Exercises . . . 281

VII Risk, Importance and Acceptance 284

25 Risk and expected consequences 285
25.1 The expected loss of the hall's cooling . . . 285
25.2 The posterior of the state (step 3) . . . 291
25.3 The answers (step 4) . . . 292
25.4 The checks (step 5) . . . 309
25.5 The report (step 6) . . . 313
25.6 The next measurement (step 7) . . . 315
25.7 What to retain . . . 319
25.8 Exercises . . . 320

26 Reliability importance measures 322
26.1 The importance of each unit to the loss of cooling . . . 322
26.2 The posterior of the state (step 3) . . . 326
26.3 The answers (step 4) . . . 327
26.4 The checks (step 5) . . . 337
26.5 The report (step 6) . . . 340
26.6 The next measurement (step 7) . . . 342
26.7 What to retain . . . 346
26.8 Exercises . . . 347

27 Reliability demonstration and acceptance tests 349
27.1 Accepting a capacity read through a model . . . 349
27.2 The posterior of the state (step 3) . . . 354
27.3 The answers (step 4) . . . 356
27.4 The checks (step 5) . . . 367
27.5 The report (step 6) . . . 370
27.6 The next measurement (step 7) . . . 370
27.7 What to retain . . . 374
27.8 Exercises . . . 375

28 Operating margins and capacity risk 377
28.1 Margins on the path and on the chiller . . . 377
28.2 The posterior of the state (step 3) . . . 383
28.3 The answers (step 4) . . . 384
28.4 The checks (step 5) . . . 395
28.5 The report (step 6) . . . 400
28.6 The next measurement (step 7) . . . 400
28.7 What to retain . . . 404
28.8 Exercises . . . 405

A Atlas of probabilistic questions 408

Notation and vocabulary

Every symbol below is used with the same meaning in every chapter. The unknown always stands on the left of a conditioning bar and the evidence on the right. Vectors and matrices are bold.

Symbol	Unit	Meaning
The plant graph and its state		
$\mathbf{z}$, z_i	varies	State vector of the plant graph: every free quantity at once (temperatures, flows, conductances, log-rates); its i-th entry.
$\mathbf{y}$, y	varies	The evidence: all readings, counts and prior information treated as known; a single reading.
$r_k(\mathbf{z})$	varies	Row k of the model: a physical relation that must hold ($r_k = 0$), a reading compared with the state ($r_k = z_i - y$, in the units of the reading), or a prior ($r_k = z_i - z_i^0$).
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; both put the same term $\frac{1}{2}w_k r_k^2$ in the objective.
T, T_a	°C	Temperature of the equipment hall; ambient temperature (pinned).
Q, G	kW, kW/K	Heat load supplied to the hall; effective thermal conductance of the heat-removal path.
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	A temperature reading of the hall and the standard deviation of its error.
$T_{\max}$, f^*	varies	A limit on a temperature; a requirement level on a state entry or on a derived quantity $f(\mathbf{z})$: the event $f \geq f^*$ is *the requirement holds*, and an upper limit is the complementary event (Chapter 2). From Part V a starred plant quantity (Q_1^*, T_1^*) is a scenario value: the quantity at a load not yet applied.
ω, ξ	–	Log-rate $\ln(\lambda/\lambda_{\text{ref}})$ and logit of a probability, $\ln(p/(1-p))$: the unbounded entries that carry a rate or a probability in the state.
u, v	–	Log-scale $\ln\eta$ and log-shape $\ln\beta$ of a Weibull lifetime, the two entries of the state in a lifetime fit.
Probability		
$\Pr(A)$, $\Pr(A \mid B)$	–	Probability of event A; conditional on event B.

Symbol	Unit	Meaning
$p(\cdot)$	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.
Beta(a, b), Gamma(a, b)	–	Beta distribution with shape parameters a, b; gamma distribution with shape a and rate b.
Binomial(n, p), Poisson(μ)	–	Number of failures in n trials with failure probability p; count with mean μ.
Inference objects		
$\mathcal{L}(\mathbf{z})$, $\mathcal{L}_k(\mathbf{z})$	–	The objective, minus the log posterior up to a constant, and the term row k puts in it: $\mathcal{L} = \sum_k \mathcal{L}_k$, with $\mathcal{L}_k = \frac{1}{2} w_k r_k(\mathbf{z})^2$ for a Gaussian row and $\mathcal{L}_k = -\ln p(y_k \mid \mathbf{z})$ for a count, lifetime or demand row (Section 4.2.7).
$\mathcal{L}'$, $\mathcal{L}''$, $\mathcal{L}'_k$, $\mathcal{L}''_k$	varies	For a state of one entry, say ω: the slope and the curvature of the objective, $\mathrm{d}\mathcal{L}/\mathrm{d}\omega$ and $\mathrm{d}^2\mathcal{L}/\mathrm{d}\omega^2$, and of row k's term; they add over the rows, and the Newton step is $\Delta\omega = -\mathcal{L}'/\mathcal{L}''$ (Chapter 5).
Working form	varies	The way a count row joins the Jacobian beside Gaussian rows (equation (4.43)): the line $\sqrt{\mu}\,\partial \ln \mu / \partial \mathbf{z}$ with the working residual $(\mu - k)/\sqrt{\mu}$ for a Poisson count k of mean μ, and $\sqrt{np(1-p)}\,\partial \xi / \partial \mathbf{z}$ with $(np - k)/\sqrt{np(1-p)}$ for a binomial count carried on the logit ξ. The pair reproduces the row's gradient and the $\mathbf{J}^\mathsf{T}\mathbf{J}$ part of its Hessian; its half square is not the row's term.
$\hat{\mathbf{z}}$	varies	The maximum a posteriori (MAP) state, the minimizer of $\mathcal{L}$.
$\mathbf{J}$	varies	Jacobian of the rows with respect to the state, evaluated at $\hat{\mathbf{z}}$: for a Gaussian row the line $\partial r_k / \partial z_i$, for a count row the line of its working form (Section 4.2.7).
$\mathbf{\Lambda}$, $\mathbf{H}$, $\mathbf{C}$	varies	Information matrix $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ with $\mathbf{W} = \mathrm{diag}(w_k)$ ($\mathbf{J}^\mathsf{T}\mathbf{J}$ for normalized rows); the Hessian of $\mathcal{L}$, $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{C}$, with $\mathbf{C}$ the curvature term the rows add beyond the information, zero for linear Gaussian rows and for count rows on their canonical parameter (equation (4.44)).
Λ	varies	For a state of one entry, the information, the curvature of the objective at the MAP, $\Lambda = \mathcal{L}''(\hat{\omega})$ (Chapter 5).
$\mathbf{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\mathbf{H}^{-1}$ ($\mathbf{\Lambda}^{-1}$ when $\mathbf{C} = \mathbf{0}$); its entries. The posterior spread of entry i is written $\sqrt{\Sigma_{ii}}$; σ_k always names a row's standard deviation and σ_p a physics row's tolerance.
σ^2_ω, σ_ω	varies	For a state of one entry, say ω or ξ: its Laplace variance $1/\Lambda$ and standard deviation (σ^2_ξ, σ_ξ for the logit). A derived quantity with derivative $g = \mathrm{d}f/\mathrm{d}\omega$ at the MAP has variance $g^2\sigma^2_\omega$.

Symbol	Unit	Meaning
ρ	K/kW	Thermal resistance of a path, $1/G$ (Chapter 4); not a correlation.
f_L, f_U	varies	The ends of the credible interval of a derived quantity $f(\mathbf{z})$ (Chapter 2).
y_{new}, $\tilde{y}_a$	varies	A future observation (the predictive) and the outcome of a candidate measurement a before it is taken (the preposterior).
M, M_j, B_{AB}	–	A drawing (model) of the plant; the members of a family; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$, written B_{10} for M_1 against M_0.
$f_t(\mathbf{z})$, $\mathbf{g}(t)$, $\bar{t}$	varies	A quantity read at time t from a state indexed by time, $f_t = \mathbf{g}(t)^{\mathsf{T}}\mathbf{z}$ with the gradient $\mathbf{g}(t) = (1, t - \bar{t})^{\mathsf{T}}$ about the center $\bar{t}$ of the record, the state holding a level entry α and a trend entry γ (the trend form, Chapter 2).
$U(a)$, u	varies	The expected value of a candidate action a over the outcomes it may produce, and the utility u of the updated posterior it is built from (the design form, Chapter 2).
$\mathbf{J}_a$, $\mathbf{\Lambda}_a$	varies	The Jacobian rows of a candidate measurement and its information $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ (expected counts at $\hat{\mathbf{z}}$ for a count row).
κ	–	Assurance level of an acceptance rule: the posterior probability a requirement must reach, $\Pr(R \geq R_{\min} \mid \mathbf{y}) \geq \kappa$ (Chapter 5).
π_0	–	The prior row's share of the information, $\mathcal{L}_0''/\mathcal{L}''$ for a state of one entry (Chapter 5); in Chapter 16 the stationary probability that no unit is up.
$f(\mathbf{z})$, $\mathbf{g}$	varies	A derived quantity of the state and its gradient $\partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$.
μ_f, σ_f^2	varies	Mean and variance of a derived quantity under the delta method: $f(\hat{\mathbf{z}})$ and $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. From Part V the spread is written $\text{sd}(f \mid \mathbf{y}) = \sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}}$.
$p(\mathbf{y} \mid M)$	–	Evidence (marginal likelihood) of model M.
B_{10}	–	Bayes factor of model M_1 against M_0, $p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$; usually reported as a log in nats.
$\mathbf{\Upsilon}$, υ	varies	The gain of the correction form, $\mathbf{\Sigma}_0\mathbf{A}^{\mathsf{T}}(\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}^{\mathsf{T}} + \mathbf{R})^{-1}$, which moves the prior mean toward the readings; its scalar value for one reading, the weight $\sigma_0^2/(\sigma_0^2 + \sigma_T^2)$ (Chapter 3).
$\mathbf{\Sigma}_a$	varies	Preposterior covariance after a candidate measurement a with information $\mathbf{\Lambda}_a$: $(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$.
EIG	nat or bit	Expected information gain of a candidate measurement.
Reliability quantities (as in the companion book)		
T, $R(t)$, $F(t)$	h, –, –	Time to failure; reliability $\Pr(T > t)$; cumulative failure probability.
$f(t)$, $\lambda(t)$, $\Lambda(t)$	h^{-1}, h^{-1}, –	Failure density; hazard function; cumulative hazard.
λ, μ	h^{-1}	Constant failure rate; constant repair rate.
p, k, n	–	Failure probability per demand; number of failures observed; number of trials.

Symbol	Unit	Meaning
T_{tot}	h	Exposure (total operating time on test) behind a failure count.
β, η	–, h	Weibull shape and scale (characteristic life).
t_i, δ_i, c_i	h, –, h	Observed time of unit i; its event indicator (1 failed, 0 censored); its censoring time.
t_{50}, t_0, T_{new}	h	Median lifetime; an age already reached; the lifetime of a new unit.
ℓ, ℓ_p	–	Log-likelihood; profile log-likelihood (maximized over the other entries).
r, D	–, h	Number of repairs observed; their total duration.
U, A_{req}	–	Unavailability $1 - A$; a required availability.
α, γ	–, yr^{-1}	The log-rate case of the level and trend entries of the trend form: log-rate at the record's center and its trend per year in $\ln \lambda(t) = \alpha + \gamma\,(t - \bar{t})$; t^* the time the rate crosses a limit.
K_f, τ	–, h or yr	A future failure count; the future period or horizon it is counted over.
q_i, $\mathbf{s}$, $\phi(\mathbf{s})$	–	Unavailability of unit i; a configuration of units (up or down); the structure function, 1 when the configuration delivers the service.
R_{sys}, U_{sys}	–	Service availability and unavailability of several units under a drawing, by the 2^n enumeration.
β_{cc}	–	Common-cause share: the fraction of a unit's failures that fail all units sharing the cause (β is the Weibull shape).
A, MTTF, MTTR	–, h, h	Steady-state availability; mean time to failure; mean time to repair.
$L(a, \mathbf{z})$, $\mathcal{R}(a)$, C, C_e, c_a	R$	The loss of action a when the state is $\mathbf{z}$, its consequence; its expected value over the posterior, the risk $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$; the consequence of an event, of event e; the fixed cost of the action (the expected-consequence form, Chapter 2).
Physical models and the state (Part V); a row naming a chapter is that chapter's own symbol		
$\mathbf{z}_{\text{full}}$, $\mathbf{r}_{\text{full}}$, $\mathbf{J}_{\text{full}}$	varies	The full state of a plant graph, every free quantity including the temperatures and meter quantities the balances fix; its residual vector and Jacobian (Chapter 17).
$h_i(\mathbf{z})$, σ_p, $\sigma_{p,P}$	varies, K, kW	The balance of node i as a function of the state; the tolerance of a hall's physics row; the tolerance of the header's (Chapters 17 and 18).
$\mathbf{J}_{\text{e}}$, $\mathbf{J}_{\text{p}}$, $\mathbf{J}_{\text{d}}$, $\mathbf{J}_\pi$	varies	The Jacobian of the evidence rows (readings and priors) and of the physics rows (Chapter 17); the design columns, the reading rows alone, and the prior rows, with information $\mathbf{\Lambda}_{\text{prior}} = \mathbf{J}_\pi^{\mathsf{T}} \mathbf{J}_\pi$ (Chapters 19 and 22).
$\mathbf{N}$, $\mathbf{N}_{\text{b}}$	varies	Null-space basis of the physics rows, $\mathbf{J}_{\text{p}}\mathbf{N} = \mathbf{0}$, whose columns are the free directions of the reduced state $\mathbf{z}$, with reduced Jacobian $\mathbf{J} = \mathbf{J}_{\text{e}}\mathbf{N}$; its block on the balanced entries (Chapter 17).

Symbol	Unit	Meaning
$\tilde{Q}_i$	kW	A load profiled out: the least-squares value with the other entries held at the prior point (Chapter 17).
ν, n_r, n_z	–	Redundancy, the number of rows less the number of entries, $\nu = n_r - n_z$: the degrees of freedom of $2\mathcal{L}(\hat{\mathbf{z}})$.
χ^2_ν, $\chi^2_\nu(c)$	–	Chi-square distribution with ν degrees of freedom; noncentral, with noncentrality c (Chapter 18).
$\mathbf{H}$, $\mathbf{S}$	varies	The exact Hessian of the objective, $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{S}$, and its curvature term $\mathbf{S} = \sum_k r_k \nabla^2 r_k$, which Gauss–Newton drops; a subscript names the variable differentiated ($\mathbf{H_q}$, Chapter 25). $\mathbf{H}$ is a Hessian throughout the book.
$\Delta\mathbf{z}$, tol, ϖ	varies, –, –	The Gauss–Newton step, $\mathbf{J}^\mathsf{T}\mathbf{J}\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$; the stopping tolerance on $\max_i \lvert\Delta z_i\rvert/(1+\lvert z_i\rvert)$; the Levenberg–Marquardt damping (Chapter 18). Chapter 23 writes the step $\mathbf{s}$, since there $\mathbf{\Delta}$ is a change between periods.
$\mathbf{j}_k$, $\mathbf{j}_a$	varies	The Jacobian row of reading k as a column, $\mathbf{j}_k = \mathbf{J}_k^\mathsf{T}$; of a candidate row a.
$\mathbf{K}$, K_{kk}, $1 - K_{kk}$	–	The hat matrix $\mathbf{J}(\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}\mathbf{J}^\mathsf{T}$; the leverage of row k, $\mathbf{j}_k^\mathsf{T}\mathbf{\Sigma}\mathbf{j}_k$; its redundancy number, the share of row k the other rows check, summing to ν (Chapters 18, 21 and 22).
r_k^N, r^*, p_fa	–	The normalized residual $r_k(\hat{\mathbf{z}})/\sqrt{1-K_{kk}}$; its test threshold, by Bonferroni $\Phi^{-1}(1 - p_\mathrm{fa}/2n_r)$; the false-alarm probability of the family of tests (Chapters 18 and 22).
s_j, $\mathbf{v}_j$, $\mathbf{V}$, $\mathbf{F}$	varies	The singular values of $\mathbf{J}$ in decreasing order and its right singular vectors, the columns of $\mathbf{V}$; the left factor, $\mathbf{J} = \mathbf{F}\,\mathrm{diag}(s_j)\mathbf{V}^\mathsf{T}$ ($\mathbf{W}$ is the weight matrix). The weakest direction $\mathbf{v}$, of spread $1/s_\mathrm{min}$ (Chapters 17 and 19).
$\mathrm{cond}(\cdot)$	–	Condition number, s_1/s_n for $\mathbf{J}$ and its square for $\mathbf{\Lambda}$ (Chapter 19); κ is the assurance level.
Λ_{ij}, ϱ_i^2, $\sigma_{0,i}$	varies	The entries of the information $\mathbf{\Lambda}$, $1/\Lambda_{ii}$ the conditional variance of entry i; its squared multiple correlation $1 - 1/(\Lambda_{ii}\Sigma_{ii})$; its prior spread (Chapter 19).
$\mathbf{\Lambda}_{-a}$, $\mathbf{\Sigma}_{-a}$, χ_p	varies	The information and covariance with the rows a removed; the radius of the credible ellipse of probability p (Chapter 19).
Propagation, virtual readings and diagnosis (Parts V and VI)		
$\mathrm{sd}(f \mid \mathbf{y})$	varies	The posterior standard deviation of a derived quantity, $\sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}}$ by the delta method; it replaces σ_f from Part V on.
$[\mathbf{g}]_i$, V_i	varies	Entry i of the gradient; the diagonal term $[\mathbf{g}]_i^2\Sigma_{ii}$ of the variance budget (Chapter 20).
$\mathbf{d}$, $m(x)$, $v(x)$	varies	The deviation $\mathbf{z} - \hat{\mathbf{z}}$ (Chapter 20); the conditional mean and variance of the other entries given one entry x under the Laplace normal, $\Sigma_{5\mid1}$ the conditional variance of entry 5 given entry 1 (Chapters 20 and 21).

Symbol	Unit	Meaning
ζ_j, $\bar{\zeta}_j$, N_{eff}	–	The importance weight of draw j and its normalized value; the effective sample size $(\sum_j \bar{\zeta}_j^2)^{-1}$ (Chapter 20).
$\mathbf{y}_-$, $\mathcal{L}_-$, $\hat{\mathbf{z}}_-$, $\boldsymbol{\Sigma}_-$	varies	The evidence, objective, MAP state and covariance of the refit without one reading row; $\tilde{\boldsymbol{\Sigma}}_-$, $\tilde{\boldsymbol{\Lambda}}_-$ the rank-one downdate of $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}$ at $\hat{\mathbf{z}}$ (Chapter 21).
V_{-k}, V_a	varies	The variance of reading k predicted from the other rows, $\mathbf{g}_k^{\mathsf{T}}\boldsymbol{\Sigma}_-\mathbf{g}_k$; of a reading at a place a without one (Chapters 21 and 22).
$\text{sd}(y_k \mid \mathbf{y}_-)$, d_k	varies, –	The predictive spread of a held-out reading, $\sqrt{\sigma_k^2 + V_{-k}}$; the held-out departure $(y_k - f_k(\hat{\mathbf{z}}_-))/\text{sd}(y_k \mid \mathbf{y}_-)$ (Chapter 21).
T_r, $\dot{m}c_p$	°C, kW/K	The return temperature of the chilled-water loop; its heat-capacity flow (Chapters 21, 22 and 24).
b_k, σ_b, $\hat{b}_k$, b_k^*	K or kW	The bias entry of reading k, its prior spread and its estimate; the minimum detectable bias (Chapter 22).
B_{k0}, ψ_k	–	The Bayes factor of a bias on reading k against no bias; the ratio $\sigma_b^2(1-K_{kk})/\sigma_k^2$ of its closed form (Chapter 22).
$\mathcal{F}$, $\lvert\mathcal{F}\rvert$	–	A family of candidate drawings and its size; the posterior over the family is $\Pr(M_j \mid \mathbf{y})$ for $M_j \in \mathcal{F}$ (Chapters 22 and 23).
N, $\bar{y}_k$	d, varies	The number of days in a record and reading k averaged over them, with spread $\sigma_k/\sqrt{N}$ (Chapter 23; bold $\mathbf{N}$ is the null-space basis).
σ_{h}, σ_{w}	varies	The healthy model's prior spread of an entry, from the earlier period; the widened spread of a candidate drawing (Chapter 23).
$\boldsymbol{\Delta}$, Z_i, χ^2_Δ	varies, –, –	The change of the state between two periods, $\hat{\mathbf{z}}_2 - \hat{\mathbf{z}}_1$; its standardized entry; the joint statistic $\hat{\boldsymbol{\Delta}}^{\mathsf{T}}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1}\hat{\boldsymbol{\Delta}}$ (Chapter 23).
$\Pr(M_d)$, $\ln B_{10}$	–	The prior probability of drawing M_d; the log Bayes factor in nats (Chapter 24).
$\text{KL}(A \,\|\, B)$, ς_d	nat, varies	The Kullback–Leibler divergence of the predictive under A from that under B; the predictive spread of a candidate reading under drawing d (Chapter 24).
c_{f}, Σ_{nn}, U_{c}	varies	The density of a flat prior; the posterior variance of an added entry; the unavailability of the cooling service (Chapter 24).
Risk, importance, acceptance and margins (Part VII)		
$\mathbf{q}$, $\boldsymbol{\Sigma}_{\mathbf{q}}$	–	The unavailabilities of the units as a vector and their posterior covariance (Chapter 25).
t_{yr}, $\mathcal{R}_{\text{yr}}$, L_{yr}	h, R$/yr, R$	The hours of a year, 8760 h; the annual risk $t_{\text{yr}}\mathcal{R}$, the expected loss of a year given $\mathbf{q}$; the loss of one particular year (Chapter 25).
$\mathcal{R}^*$, w_a	R$/yr, decades	A limit on the annual risk; the width of the central 95 % band of the risk after candidate a (Chapter 25).
ν_{loss}, MDT	yr^{-1}, h	The frequency of loss-of-service events; the mean down time of one (Chapter 25).

Symbol	Unit	Meaning
VaR_p, CVaR_p	R$	The p quantile of the predictive year's loss; the tail mean beyond it (Chapter 25).
U_i^+, U_i^-	–	The top event with unit i pinned down ($q_i = 1$) and pinned up ($q_i = 0$) (Chapter 26).
$I_{\mathrm{B},i}$, $\mathbf{I}_\mathrm{B}$	–	Birnbaum importance $U_i^+ - U_i^- = \partial U/\partial q_i$, the probability that unit i is critical; the vector of them (Chapter 26).
$I_{\mathrm{FV},i}$, $I_{\mathrm{CR},i}$	–	Fussell–Vesely importance $(U - U_i^-)/U$, the share a perfect unit removes; criticality importance $q_i I_{\mathrm{B},i}/U$, equal to it for units independent given $\mathbf{q}$ (Chapter 26).
RAW_i, RRW_i	–	Risk achievement worth U_i^+/U and risk reduction worth U/U_i^- (Chapter 26).
V_j, n_j	–, –	Uncertainty importance of shared entry j, its term $(g_j/U)^2\Sigma_{jj}$ in the variance of $\ln U$; the checks of its record (Chapter 26).
G^*, g^*, z_κ	kW/K, –, –	A required conductance and its log-conductance; the standard normal quantile $\Phi^{-1}(\kappa)$ (Chapters 27 and 28).
Q_test, ΔT^*, σ_c	kW, K, K	The held load of a test; the rise at the requirement, Q_test/G^* (the star marks the requirement level); the combined supply spread of sensor and set point (Chapter 27).
$\mathrm{OC}(\mathbf{z}; a)$, $U(a)$	–	The operating characteristic, the probability that test a accepts at a stated truth $\mathbf{z}$; averaged over the posterior, the probability the test ends accepted (Chapter 27).
ε_p, ε_c, π_g	–	The producer's and consumer's risks of a test; the prior's share Σ_{11}/σ_g^2 of the test's posterior (Chapter 27).
m_1^*, Q_cap, Q_κ	K, kW, kW	The margin $T_\mathrm{max} - T_1^*$ at a scenario load; the path's capacity $(T_\mathrm{max} - T_s)\,\mathrm{e}^{g_1}$; the largest load held with assurance κ (Chapter 28).
c_ch, P_pl, P^*, P_hr	–, kW, kW, kW	The chiller's capacity factor; its plate rating; the demand on it at a scenario; its headroom $P_\mathrm{pl}c_\mathrm{ch} - Q_1 - Q_2$ (Chapter 28).
ΔT_s, $(x)^+$, b_P	K, –, kW	The supply's rise when demand exceeds the delivered capacity; the positive part $\max(0, x)$; the bias of the chiller's heat meter (Chapters 23 and 28).

Vocabulary

Plant graph The representation of 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 inference 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; physical rows constrain the state, reading rows compare it with evidence, prior rows carry handbook knowledge.

Prior, likelihood, posterior Knowledge before the evidence; the probability of the evidence given a state; knowledge after the evidence, by Bayes' theorem.

Marginalization Integrating a joint distribution over the unknowns not asked about; the law of total probability.

Aleatory, epistemic Variability in outcomes the model treats as random; incomplete knowledge of quantities that have one true value.

Maximum a posteriori (MAP) estimate The state at which the posterior density is highest; the minimizer of the objective.

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

Information matrix The matrix $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$, the Hessian of the objective without the curvature term $\mathbf{C}$ of $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{C}$ (equation (4.44)); precisions of independent rows add. For a state of one entry it is the number Λ.

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

Change of variables Carrying a density to a new coordinate $u = h(x)$:

$$p_u(u) = p_x(h^{-1}(u))\,|\mathrm{d}x/\mathrm{d}u|,$$

the Jacobian factor keeping the probability of each interval; for several entries the factor is $|\det \partial\mathbf{z}/\partial\mathbf{u}|$. It gives the shape of a derived quantity, where the delta method gives only its spread (Section 4.3.3).

Variance budget The expansion of $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ into one term per state entry plus the cross terms from their covariance.

Threshold probability The probability that a state entry or derived quantity is below (or above) a limit; a normal tail under the Laplace approximation.

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

Credible interval An interval that contains the unknown with a stated posterior probability.

Evidence, Bayes factor The normalizing constant of one model's posterior; the ratio of two models' evidences, which compares the models rather than their parameters.

Preposterior The posterior that would result from a measurement not yet taken; the object that values a measurement before buying it.

Identifiability Whether the evidence determines the state entries separately or only in combination; read from the shape of the posterior, not from marginal spreads.

The five question categories Estimation, prediction under a threshold, hypothesis comparison, diagnosis and design, answered by a marginal, a tail of the posterior or of the trend, an evidence ratio, a ranking of evidences, and a preposterior or an expected consequence.

The five computations The objects that answer every category (Chapter 2): the *posterior* $p(\mathbf{z} \mid \mathbf{y})$, read for an entry or a derived quantity; the *trend*, the posterior over a state indexed by time or window with a trend entry, read as the probability that the trend is positive, as a crossing time with its band, or as a forecast; the *evidence* $p(\mathbf{y} \mid M)$ and its ratios; the *preposterior*, the expected posterior under a measurement not yet taken; and the *expected consequence* $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ of an action over the posterior, with the choice of the action that minimizes it (risk is probability times consequence).

Factor graph A graph of the state's entries and the factors (the rows) that bear on them; the plant graph with its rows is one, and the posterior is the normalized product of its factors (Chapters 1 and 4).

Variable elimination Integrating the entries a question does not name out of the joint distribution one at a time; for a Gaussian posterior it is the Schur complement, and its cost on a large sparse graph is set by the order of elimination (Chapter 3).

Belief propagation Computing the marginals of a graphical model by passing messages between neighboring factors and entries (the sum-product algorithm); exact on a graph without loops, where on a Gaussian graph it returns the Schur complements. Not used in this book, whose posteriors are factored jointly.

Kalman filter and smoother Variable elimination on a state indexed by time, taken in time order: the filter forward, the smoother back.

Importance sampling (reweighted draws) Drawing from an approximation, here the Laplace normal, and weighting each draw by the ratio of the exact posterior to it, so that the draws check the approximation (Chapter 20, in Volume II).

Markov chain Monte Carlo Drawing from a posterior by a Markov chain whose stationary distribution it is (Gibbs sampling, the Metropolis algorithm); the route when neither the Laplace normal nor reweighting is adequate, as for a hierarchical model (Chapter 33, in Volume III).

Explaining away Two causes independent before an effect is seen become dependent once it is seen: evidence for one lowers the probability of the other, as a biased sensor and a changed conductance compete for one disagreeing reading (Chapters 22 and 23, in Volume II).

Physics row, tolerance A balance of the plant graph written as a row with a small standard deviation; in the limit of zero tolerance the balance is exact and the state reduces to its free directions.

Reduced state The free directions left once the balances are exact, the columns of the null-space basis of the physics rows.

Profiling Holding some entries fixed and fitting the others by least squares, to see what the readings say about the fixed ones.

Predictive check Asking whether the readings are probable under the model fitted to them, by the misfit $2\mathcal{L}(\hat{\mathbf{z}})$ against its chi-square law or by a held-out reading against its predictive.

Singular value, weakest direction The scale at which the rows see a direction of the state; the combination of entries with the smallest singular value, which the evidence determines least.

Observability Whether the readings determine an entry at all: an unobservable direction is one the design columns do not see, and only a prior gives it a finite spread.

Condition number The ratio of the largest singular value to the smallest; it bounds the loss of digits in solving with the information matrix.

Virtual sensor The posterior of a quantity no instrument reads, pushed from the state through its function.

Leave-one-out (held-out) check Refitting without a reading and comparing the reading with the predictive of the remaining rows.

Redundancy number The share of a reading the other rows check, $1 - K_{kk}$; zero for a critical measurement, which no other row checks.

Minimum detectable bias The bias on a reading that the residual test flags with a stated probability.

Footprint The pattern a candidate fault leaves on the readings: its column of the Jacobian on the reading rows.

Fault detection and isolation, model-based diagnosis Deciding that the plant has departed from its healthy model, and which component did, by comparing the evidence of candidate drawings generated from the plant graph.

Occam factor The part of a model's evidence that penalizes the prior range its parameters were given and the readings did not use.

Savage–Dickey ratio The Bayes factor of a point value of one entry against the model that frees it, read as the ratio of the posterior to the prior density at that value.

Expected evidence gain The expected log Bayes factor a measurement not yet taken would add between two drawings, the Kullback–Leibler divergence of their predictives.

Value at risk, tail mean A quantile of the predictive loss of a year, and the mean of the loss beyond it.

Importance measures Birnbaum (the slope of the top event in a unit's unavailability), Fussell–Vesely (the share a perfect unit would remove), risk achievement and reduction worths (the factors of an outage and of a perfect unit); each a derived quantity of the state, with a band.

Pivotal decomposition The top event as a straight line in one unit's unavailability, between the unit pinned down and pinned up.

Uncertainty importance A shared entry's term in the variance of the top event: what its record leaves uncertain, as against what its unit carries.

Operating characteristic The probability that a test accepts, as a function of the true value; its complement at a good value is the producer's risk, its value at a bad one the consumer's risk.

Acceptance sampling by variables Accepting on a measured quantity read through a model, rather than on a count of failures (by attributes).

Operating margin, capacity The room between a predicted quantity and its limit at a load; the load at which the limit is reached. Their bands are posteriors, and the load held with an assurance is a tail inverted.

Derating A unit delivering less than its plate rating, estimated as a capacity factor against the plate as a point null.

Part V

Physical Models and Uncertainty Propagation

Chapter 17

Physics-Constrained State Estimation

17.1 State estimation on a physical model

Two data halls are cooled from one chilled-water loop, and on a steady afternoon its five instruments report: the hall sensors 17.654545 and 17.700°C, the supply header 12.10°C, the chiller's heat meter 52.5 kW and the electrical meter on hall 1's equipment 31.0 kW. No instrument reads a conductance, the quantity that falls first when a hall's coil fouls, and none reads hall 2's load. The first reading of the five numbers is a hand calculation. Hall 2's load is the heat meter less hall 1's, $52.5 - 31.0 = 21.5$ kW, and each conductance is its load over its hall's rise above the supply: $31.0/5.5545 = 5.58$ kW/K for hall 1 and $21.5/5.60 = 3.84$ kW/K for hall 2, against a handbook value of 5 kW/K for both. Every reading is reproduced exactly, and hall 2 looks about a quarter below the handbook.

The engineer responsible for the loop asks what the calculation leaves unsaid. It uses each reading once and has none to spare, so it would reproduce a faulty sensor just as exactly; it states no spread, so whether 3.84 differs from 5 by more than the readings' errors allow is not something it can say; and it sets aside the handbook value and the chiller's set point, which were known before any reading. Which entries the readings determine, how well, and whether the readings can be checked against each other are properties of the rows that tie each reading to the unknowns, and they can be read before the state is fitted. The question of this chapter is

how the state of a plant with several free quantities is written as one vector, how each reading is tied to that vector through a physical row, and what the residual vector, its weights and its Jacobian establish about the readings before the state is fitted.

It is answered with the rows of Chapter 4, physics rows with a tolerance beside reading and prior rows on one state vector. What changes from the chapters on records of events is the evidence: readings of temperatures and of heat on a plant whose balances tie them to the unknowns, and with it a state of eight entries of which the balances leave five directions free.

Jargon: state estimation, physics-constrained, measurement model, residual vector, normalized residual, physics row, model tolerance, Jacobian, null directions, design columns, redundancy, prior point, linearization, normal equations, Gauss–Newton step, predictive check, profiling, singular value

State estimation is the inference of the free quantities of a physical system from readings that are related to them through a model; in power systems and process plants it is usually a weighted least-squares fit of the state to the readings. It is *physics-constrained* when the model is the system's own conservation relations, so that the estimate satisfies them. The *measurement model* of a reading states the reading as a function of the state plus an error. The *residual vector* collects every row of the model, each written as predicted minus observed; a *normalized residual* is a row divided by the standard deviation of its error, a dimensionless number of standard deviations. A *physics row* is a balance of the plant written as a residual and divided by a *model tolerance* σ_p, the standard deviation by which the balance is allowed to fail; the exact balance is the limit $\sigma_p \to 0$. The *Jacobian* is the matrix of derivatives of the rows with respect to the entries of the state. The *null directions* of the physics rows are the moves of the state that leave their linearization at zero, the directions along which the balances keep holding; the *design columns* are the columns of the reading rows taken along those directions, which state what the readings alone can see of the states that satisfy the balances. *Redundancy* is the number of rows beyond the number of entries of the state, the degrees of freedom left for a check. The *prior point* is the state at which every entry with a prior sits at its prior mean. *Linearization* replaces each row by its tangent at a point; the *normal equations* are the linear equations whose solution minimizes the sum of squares of linearized rows, and the step they give is the *Gauss–Newton step*. A *predictive check* compares a reading with the distribution the model gives it before it is used. *Profiling* an entry replaces it by the value that minimizes the objective with the other entries held. The *singular values* of a matrix are the lengths by which it scales its orthogonal input directions, the right singular vectors; the smallest says how weakly the least-seen direction is seen.

Key idea

A reading is evidence about the state only through a physical row. In the full state of the plant a hall sensor reads one entry, the hall temperature; it sees the conductance, the load and the supply temperature only because the hall's physics row ties that temperature to all three. Projected onto the null directions of the physics rows, the readings and priors become rows over the free entries alone, with a Jacobian that is the projection of their own and a covariance that is the full state's in the limit of exact balances. The formulation, the rows with their weights and their Jacobian, decides which entries the readings determine, how well, and whether they can be checked, before any state is fitted; a formulation with no rows to spare fits any readings exactly and checks nothing.

17.1.1 From records of events to the state of a plant

In Parts II to IV of Volume I the state was small and the evidence was a count or a list of times. A failure rate was one entry of the state informed by a Poisson row; a lifetime model was two entries informed by a likelihood of observed and censored times; a common-cause share was one more entry informed by the count of simultaneous failures. The physical system stood behind those rows, but no row contained its physics; only the single hall of Part I, with one entry of interest, had a physical row. Part V reverses the emphasis. The plant is the equipment hall of Chapter 1, now two halls on one chilled-water loop, and the evidence is the set of readings that its sensors and meters report in one steady operating condition. The full state has eight entries, two conductances, two loads, three temperatures and the chiller's heat, of which the balances

leave five directions free, and every reading is connected to them by conservation of energy.

This is the setting the field calls state estimation. In electrical networks the state is the set of bus voltages and the readings are power flows and injections; in water networks the state is the set of heads and demands and the readings are pressures and flows; in a cooling plant the state is the set of temperatures, loads and conductances and the readings are the sensors and meters of the plant. The state-estimation chapter of the companion volume *Reliability Engineering* states the method and its vocabulary. Here its formulation is derived for the two halls from the probability rules of Volume I, so that the fit of Chapter 18, the spread and the observability of Chapter 19 and the propagation of Chapter 20 have one formulation to act on.

Part V asks five questions of the fitted plant, and the formulation of this chapter is what they act on. The first asks for the two halls' conductances and loads and the supply temperature, given the readings. By Bayes' theorem the posterior of the state $\mathbf{z}$ is

$$p(\mathbf{z} \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \mathbf{z})\, p(\mathbf{z})}{p(\mathbf{y})}, \tag{17.1}$$

where $p(\mathbf{y} \mid \mathbf{z})$ is the probability of the readings given the state, built from their measurement models (Section 17.1.7), and $p(\mathbf{z})$ is the prior; every entry wanted is an entry of $\mathbf{z}$, so the answer is a marginal of this posterior, its mode and covariance at the fitted state (Chapters 18 and 19). The second asks whether the readings and the physics can be true together, a tail of the minimized objective against the chi-square of the redundant rows (Chapter 18). The third asks which directions of the state the readings resolve at the fitted state, a property of the posterior's shape rather than of any single marginal (Chapter 19). The fourth asks for hall 1's temperature at a future load of 60 kW and whether it stays below 25°C. It extends the graph with a scenario node, hall 1 at $Q_1^* = 60\,\mathrm{kW}$ with the same conductance and the same supply temperature, whose temperature is $T_1^* = T_s + Q_1^*\,\mathrm{e}^{-g_1}/G_{\mathrm{ref}}$, with T_s the supply temperature and g_1 hall 1's log-conductance (Section 17.1.4); the probability is the posterior mass of the states that satisfy the event,

$$\Pr(T_1^* < 25^{\circ}\mathrm{C} \mid \mathbf{y}) = \int \mathbf{1}\{T_s + 60\,\mathrm{e}^{-g_1}/G_{\mathrm{ref}} < 25\}\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{17.2}$$

In Chapter 1 the boundary temperature was pinned and the same event was a threshold on the conductance alone; here it is $g_1 > \ln\!\big(60/(G_{\mathrm{ref}}(25 - T_s))\big)$, a threshold that moves with the free supply temperature, so the answer depends on two entries and on their correlation (Chapter 20). The fifth asks which measurement to take next for that temperature, the preposterior covariance of a candidate reading (Chapter 20). The chapters on bad data, on faults and on model selection, Chapters 22 to 24, ask of the same plant whether a reading is in error, which fault the readings point to and which drawing of the plant they support, from the residuals and from the evidence.

17.1.2 A reading seen through a physical row

Chapter 1 made the central observation on a single hall: the likelihood of the temperature reading does not depend on the conductance, and the reading informs the conductance only through the physics row that ties the temperature to it. With one entry of interest that observation could be made in one line. With eight entries and five readings it becomes the organizing fact of the formulation. Each reading is a reading of one entry: a hall sensor of its hall temperature, the heat meter of the chiller's heat, the load meter of one load. The balances are rows of their own, and a reading sees the other entries only along them. The hall balance moves the hall temperature with the conductance, the load and the supply temperature at once, so a hall sensor sees those three together and none of them alone. The header balance moves the chiller's heat with the sum of the two loads, so the heat meter sees the sum and neither load separately. Which combinations of the entries the readings determine, and which only the

priors determine, is fixed by these physics rows, and it can be read from their derivatives, a column and a null direction at a time, before the state is fitted.

Principle 17.1. *A reading constrains the state only through the physical rows that tie its entry to the others. The pattern of the rows' derivatives, the Jacobian, states which entries each row touches before any reading is processed, and an entry that no reading reaches, directly or along a physics row, is determined by its prior alone.*

17.1.3 The plant graph (step 0)

Two data halls are cooled from one chilled-water loop. The chiller supplies water at the supply temperature T_s to a header that feeds the cooling coils of both halls. Each hall's equipment dissipates a heat load Q_i, in kW, and the hall's heat-removal path, the coil with its air and water flows, carries that heat to the chilled water with an effective conductance G_i, in kW/K; the subscript $i = 1, 2$ names the hall. The warmed water returns to the chiller, which removes the heat from the loop. Figure 17.1 draws the plant graph. The two green boxes on the left and the right are the halls, free nodes whose temperatures T_1 and T_2 follow from the balances derived below; the white boxes above them are the heat loads, sources of unknown strength; the green box in the middle is the supply header, a free node carrying the supply temperature T_s; the white box at the bottom is the chiller, the source of chilled water and the sink of the heat. The solid green arrows are the physical paths of heat: from each load into its hall, from each hall to the header through the conductance G_i, and from the header to the chiller. The drawing routes the heat through the header node, whose temperature is the one the coils see; physically the heat travels back to the chiller in the return line, which the steady balance below does not need. The five dashed boxes are the readings, connected by dashed arrows to the node or source each one measures.

The graph is the system as it is, before any question is asked. Nothing in it is pinned except the reference value of the conductance; in Chapter 1 the ambient temperature was a pinned boundary, and here its place is taken by the supply temperature, which is free, because the chiller's control holds it near a set point but not exactly at it.

17.1.4 The balances of the halls and the header

Chapter 1 derived the steady energy balance of a hall with one heat-removal path, equation (1.1): the heat supplied equals the heat removed, and the heat removed is the conductance times the temperature difference across the path. For hall i the path ends at the chilled water at T_s instead of at the ambient air, and the same steps give

$$Q_i = G_i\,(T_i - T_s), \qquad i = 1, 2, \tag{17.3}$$

with Q_i in kW, G_i in kW/K and the temperatures in °C. Since $G_i > 0$, dividing by G_i and adding T_s,

$$T_i = T_s + \frac{Q_i}{G_i}, \qquad i = 1, 2 : \tag{17.4}$$

each hall sits above the supply temperature by its temperature rise Q_i/G_i. The supply header is a node too. In steady state the water of the loop stores no heat, and if the pump work and the heat gained by the pipes are neglected, the heat the chiller removes equals the heat the two halls deliver,

$$P = Q_1 + Q_2, \tag{17.5}$$

with P the chiller's heat duty in kW. Equations (17.4) and (17.5) are the physical rows of the plant, three relations that must hold among its eight free quantities.

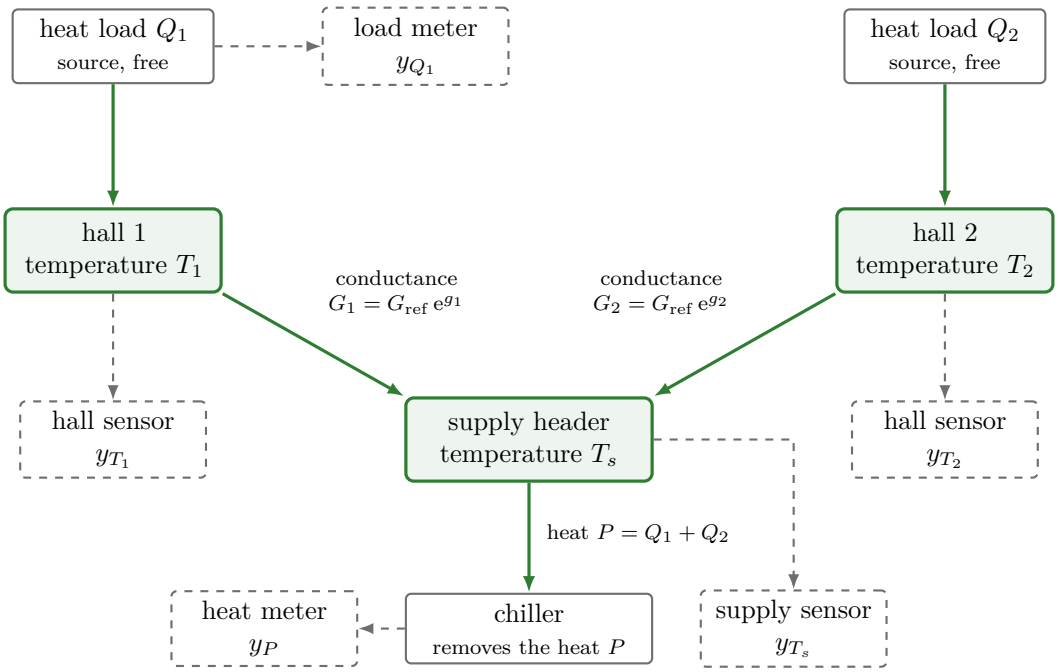


Figure 17.1: Plant graph of two data halls on one chilled-water loop in steady state. Free: the log-conductances g_1, g_2 of the two heat-removal paths, the heat loads Q_1, Q_2, the supply temperature T_s, the hall temperatures T_1, T_2 and the chiller's heat P, the eight entries of the full state; the three balances are physics rows among them. Pinned: the reference conductance $G_{\text{ref}} = 1\,\text{kW/K}$. Observed (dashed): T_1, T_2, T_s, P and Q_1. The three priors, on g_1, g_2 and T_s, are rows of the evidence and are not drawn.

The conductance is carried as the log-conductance of Section 1.1.4, $g_i = \ln(G_i/G_{\text{ref}})$, with $G_{\text{ref}} = 1\,\text{kW/K}$, so that $G_i = G_{\text{ref}}\, e^{g_i}$ is positive for every real g_i and a change of g_i is a relative change of G_i. With the reference at one, e^{g_i} is the conductance in kW/K and e^{-g_i} its reciprocal in K/kW; where a numerical expression below writes $Q_i\, e^{-g_i}$ without the reference, $G_{\text{ref}} = 1\,\text{kW/K}$ has been absorbed and the expression is in kelvin. Every free quantity of the graph is an entry of the full state, as in Chapter 4:

$$\mathbf{z}_{\text{full}} = \begin{bmatrix} g_1 & g_2 & Q_1 & Q_2 & T_s & T_1 & T_2 & P \end{bmatrix}^{\mathsf{T}}, \tag{17.6}$$

eight entries, with g_1, g_2 dimensionless, Q_1, Q_2, P in kW and T_s, T_1, T_2 in °C. The order puts the conductances, the loads and the supply temperature first and the three quantities on the left sides of equations (17.4) and (17.5) last. The balances enter the residual vector as physics rows with a tolerance (Section 17.1.7). When they hold exactly, the full state can move only in five directions, one per entry of

$$\mathbf{z} = \begin{bmatrix} g_1 & g_2 & Q_1 & Q_2 & T_s \end{bmatrix}^{\mathsf{T}}, \tag{17.7}$$

the reduced state. Section 17.2.1 derives it in matrix form, as the null directions of the physics rows, together with its rows and its Jacobian; the inference of the following chapters runs on it.

17.1.5 The readings and the priors

Table 17.1 lists the five readings: each row is one sensor or meter, with its symbol, the value it reports, the standard deviation of its error and the quantity of the plant it measures. The hall sensors and the supply sensor read temperatures; the chiller's heat meter reads the heat it

removes from the loop, which by equation (17.5) is the sum of the two loads; the electrical meter on hall 1's equipment reads its electrical power, which in steady state is the heat it dissipates. The readings are illustrative: they were generated from $G_1 = 5.5\,\mathrm{kW/K}$, $G_2 = 4.0\,\mathrm{kW/K}$, $Q_1 = 30\,\mathrm{kW}$, $Q_2 = 24\,\mathrm{kW}$ and $T_s = 12°\mathrm{C}$ with errors of $+0.2\,\mathrm{K}$, $-0.3\,\mathrm{K}$, $+0.1\,\mathrm{K}$, $-1.5\,\mathrm{kW}$ and $+1.0\,\mathrm{kW}$ added to the five readings in the order of the table.

Table 17.1: The readings of the two-hall plant. The hall-1 reading is printed to six decimals, 17.654545°C, because it is generated as $12 + 30/5.5 + 0.2$; the numbers of this chapter are computed with that value.

Sensor or meter	Symbol	Reading	Error sd	Measures
Hall 1 temperature sensor	y_{T_1}	17.654545°C	$\sigma_T = 0.5\,\mathrm{K}$	T_1, whose balance is $T_s + Q_1/G_1$
Hall 2 temperature sensor	y_{T_2}	17.700°C	$\sigma_T = 0.5\,\mathrm{K}$	T_2, whose balance is $T_s + Q_2/G_2$
Supply temperature sensor	y_{T_s}	12.10°C	$\sigma_s = 0.3\,\mathrm{K}$	T_s
Chiller heat meter	y_P	52.5 kW	$\sigma_P = 2\,\mathrm{kW}$	P, whose balance is $Q_1 + Q_2$
Hall 1 electrical meter	y_{Q_1}	31.0 kW	$\sigma_Q = 1.5\,\mathrm{kW}$	Q_1

Table 17.2 lists the priors: each row is one entry of the state, its prior distribution, and where the prior comes from. The handbook gives a conductance of $5\,\mathrm{kW/K}$ for a hall of this size and coil type, with a spread of 35 %: in the log-conductance this is $g_i \sim \mathcal{N}(g^0, \sigma_g^2)$ with $g^0 = \ln 5 = 1.6094$ and $\sigma_g = 0.3$, since $\mathrm{e}^{0.3} = 1.35$. The supply temperature has the chiller's set point $T_s^0 = 12°\mathrm{C}$ as its prior mean, with a standard deviation $\sigma_0 = 1\,\mathrm{K}$ for the control's offset. The two loads have no prior: nothing is known about them before the readings beyond their being positive, and they enter with a flat prior.

Table 17.2: The priors of the two-hall plant. The loads have none.

Entry	Prior	Source
$g_1 = \ln(G_1/G_{\mathrm{ref}})$	$\mathcal{N}(\ln 5,\ 0.3^2)$	handbook conductance 5 kW/K, spread 35 %
$g_2 = \ln(G_2/G_{\mathrm{ref}})$	$\mathcal{N}(\ln 5,\ 0.3^2)$	the same handbook value; independent of g_1
Q_1, Q_2	none (flat)	determined by the readings alone
T_s	$\mathcal{N}(12°\mathrm{C},\ 1\,\mathrm{K}^2)$	chiller set point and control offset

17.1.6 The questions over the state (step 1)

Table 17.3 lists the four questions of the opening, asked of the formulation before any fit. The first column is the question in engineering language; the second the general form of its category, in the notation of Chapter 2; the third its mathematical form over the reduced state of equation (17.7); the fourth its category; and the last where it is answered. Each is answered from the rows and their Jacobian at the prior point $\mathbf{z}^{(0)}$ (Section 17.2.2).

17.1.7 The evidence as rows (step 2)

Each reading has the measurement model of Chapter 1: the reading equals the quantity it measures plus an error, normal with zero mean and the stated standard deviation, independent of the other readings' errors. For the hall-1 sensor, for example,

$$y_{T_1} = T_1 + \varepsilon_{T_1}, \qquad \varepsilon_{T_1} \sim \mathcal{N}(0, \sigma_T^2),$$

Table 17.3: The questions asked of the two-hall plant before the fit, the general form of their category, their mathematical form over the reduced state $\mathbf{z} = (g_1, g_2, Q_1, Q_2, T_s)$, their category and where each is answered. $\mathbf{y}$ is the five readings with the three priors; $\mathbf{z}^{(0)}$ the prior point; $\tilde{\mathbf{z}}$ the handbook values of g_1, g_2, T_s with the loads that minimize the objective there; $\mathbf{J}$ the Jacobian of the reduced rows and $\mathbf{J}_y\mathbf{N}$ its reading rows; $\mathbf{v}$ the right singular vector of the smallest singular value $s_{\min}$; ν the number of rows beyond entries.

Engineering question	General form	Mathematical form	Category	Answered in
Do the readings agree with the handbook values, the set point and each other, and can that be checked?	a tail, eq. (2.12), of the predictive of a reading, eq. (2.13)	$r_k(\tilde{\mathbf{z}})$, its share $\frac{1}{2}r_k^2/\mathcal{L}(\tilde{\mathbf{z}})$ and $2\Phi(-\lvert r_k(\tilde{\mathbf{z}})\rvert)$; $\nu =$ rows $-$ entries	predictive check of the evidence	17.3.1
Which entries does each reading see, and do the readings determine all five?	$\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, $f = \mathbf{v}^\mathsf{T}\mathbf{z}$	the pattern of $\mathbf{J}_{\text{full}}$; $\mathrm{rank}(\mathbf{J}_y\mathbf{N}) = 5$; $\mathrm{sd}(\mathbf{v}^\mathsf{T}\mathbf{z} \mid \mathbf{y}) = 1/s_{\min}$	observability (estimation along the weakest direction)	17.3.2
How well would the readings fix each entry?	$\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$\sqrt{\Sigma_{ii}}$, $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ at $\mathbf{z}^{(0)}$, against the prior spreads	estimation (the spread)	17.3.3
Where does the first step of the fit take the state?	the location, eq. (2.10), of the posterior with its rows linearized at a point	$\mathbf{z}^{(1)} = \mathbf{z}^{(0)} + \Delta\mathbf{z}$, $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$; $\mathcal{L}(\mathbf{z}^{(1)})$	estimation (the point)	17.3.3

with T_1 the sixth entry of the full state (17.6). The two meters measure loads in kW and the three sensors temperatures in °C, so the five errors have different units, and each row is divided by its own standard deviation. With the five readings and the three priors the evidence has eight rows; with the three physics rows of the balances the residual vector has eleven rows over the eight entries of the full state. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J}_{\text{full}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & * \\ 0 & 0 & * & 0 & 0 & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 & 0 & 0 \\ * & 0 & * & 0 & * & * & 0 & 0 \\ 0 & * & 0 & * & * & 0 & * & 0 \\ 0 & 0 & * & * & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{reading, hall 1} \\ \text{reading, hall 2} \\ \text{reading, supply} \\ \text{reading, heat meter} \\ \text{reading, load meter} \\ \text{prior, } g_1 \\ \text{prior, } g_2 \\ \text{prior, } T_s \\ \text{physics, hall 1} \\ \text{physics, hall 2} \\ \text{physics, header} \end{matrix}, \tag{17.8}$$

with the columns $g_1, g_2, Q_1, Q_2, T_s, T_1, T_2, P$ of equation (17.6). Every reading row and every prior row touches exactly one entry. The columns of g_1, g_2 and Q_2 contain no reading row at all: no sensor reads a conductance, and no meter reads the second load. Those three entries are reached only through the physics rows, the conductance g_i through the balance of its hall and Q_2 through the balance of hall 2 and the header balance. The supply temperature's column has four marks, its own sensor, its prior and both hall balances, and it is the entry that couples the two halls. The header row marks the two loads alike. All eleven rows are Gaussian: the reading and prior rows with the standard deviations of Tables 17.1 and 17.2, the physics rows

with a model tolerance σ_p, as written below.

The modeling choices. Seven modeling choices are made for the whole model, and each buys something and can fail. The reading errors are taken as normal, independent and unbiased, each with the standard deviation of Table 17.1: this is what makes each normalized row a number of standard deviations and lets the squares of the rows add, and it fails for a miscalibrated meter, whose bias a single snapshot cannot separate from the state (Section 17.6.5). The handbook conductance and the chiller's set point are carried as prior rows, and the loads are given none: the priors are what supplies rows beyond entries, and a hall unlike the handbook's would turn its prior row from a check into a source of disagreement, which Section 17.4.4 measures. Every free quantity of the graph is carried as an entry, the hall temperatures and the chiller's heat included, and the balances as physics rows: this is what makes each reading row a comparison of one entry with its reading, so that the pattern of the Jacobian is the graph's own, and it carries three entries the readings could do without, which the projection of Section 17.2.1 removes, valid as far as the limit of Section 17.4.1 holds. The halls and the loop are taken as steady, with one lumped heat-removal path per hall: this is what gives each balance its three terms and fixes the null directions, and a storage term or a second path would add entries and change the columns (Section 17.6.5). The rows are linearized at the prior point, the handbook conductances, an even split of a design load and the set point: this gives a covariance and a step before any fit, and it misleads where the rows curve over the length of the step, which Section 17.4.3 measures. The exact balances are taken as the limit $\sigma_p \to 0$ and the computation runs on the reduced state: this removes the three carried entries from the normal equations without a substitution made by hand, and it holds only if the free-entry block of the full covariance converges to the reduced one, which Table 17.7 checks. The conductance is carried as its logarithm: this keeps every conductance positive and puts the handbook's relative spread in a normal prior row, and on the conductance itself the Jacobian's entries would change as $1/G_i^2$ across the range of conductances (Section 17.1.7).

The rows of the readings, the priors and the balances. The measurement model of the hall-1 sensor states that $y_{T_1} - T_1$ is a normal error of standard deviation σ_T. The hall temperature is the sixth entry of the full state, so the density of the reading given the state,

$$p(y_{T_1} \mid \mathbf{z}) = \frac{1}{\sqrt{2\pi}\,\sigma_T} \exp\left[-\frac{(T_1 - y_{T_1})^2}{2\sigma_T^2}\right],$$

is a function of one entry. The balance of hall 1, equation (17.4), is taken to hold up to a model tolerance σ_p, the device of Chapter 4 (equation (4.16)): the hall temperature is normal about the temperature the balance predicts,

$$p(T_1 \mid g_1, Q_1, T_s) = \mathcal{N}\big(T_1;\; T_s + Q_1\,\mathrm{e}^{-g_1}/G_{\mathrm{ref}},\; \sigma_p^2\big),$$

with σ_p in kelvin, the amount by which the steady balance may fail. Each density is the exponential of minus half the square of a normalized row, and the two rows are

$$r_1(\mathbf{z}) = \frac{T_1 - y_{T_1}}{\sigma_T}, \qquad r_9(\mathbf{z}) = \frac{T_1 - T_s - Q_1\,\mathrm{e}^{-g_1}/G_{\mathrm{ref}}}{\sigma_p}, \tag{17.9}$$

the reading row and the physics row of hall 1, numbered as in the residual vector below. The reading row compares the reading with an entry; the physics row compares that entry with the temperature the balance predicts from (g_1, Q_1, T_s). A state with a large conductance and a small load predicts a hall barely above the supply: if its T_1 stays at the reading, the physics row grows positive, and if T_1 follows the balance down, the reading row grows negative, so the

evidence argues against that state through one row or the other. The rows of hall 2 are the same with the subscript 2, rows 2 and 10.

The heat meter and the header balance pair the same way. The meter's measurement model is $y_P = P + \varepsilon_P$, and the header balance (17.5) holds up to a tolerance $\sigma_{p,P}$ in kW:

$$r_4(\mathbf{z}) = \frac{P - y_P}{\sigma_P}, \qquad r_{11}(\mathbf{z}) = \frac{P - Q_1 - Q_2}{\sigma_{p,P}}.$$

The supply sensor and the load meter read T_s and Q_1, and their rows are $(T_s - y_{T_s})/\sigma_s$ and $(Q_1 - y_{Q_1})/\sigma_Q$. The tolerances are set at $\sigma_p = 0.01\,\mathrm{K}$ and $\sigma_{p,P} = 0.01\,\mathrm{kW}$: tight enough that the answers no longer move as they shrink further (Table 17.7), and within about two decades of the readings' spreads, which keeps the normal equations conditioned well enough for double precision.

Common mistake

Writing rows on incomparable scales. A balance row kept in watts with a standard deviation of a few milliwatts beside meter rows of a few kilowatts makes that row of the Jacobian many decades larger than the others, since a row's derivatives are divided by its standard deviation; the information matrix is then ill-conditioned and an iteration that should converge in a few steps crawls or stalls. Normalizing every row by its standard deviation, and giving a physics row a tolerance that is small against the physics yet within about two decades of the meters' spreads, as $\sigma_p = 0.01\,\mathrm{K}$ and $\sigma_{p,P} = 0.01\,\mathrm{kW}$ are here, keeps the rows on one scale.

Every quantity as an entry. Carrying the hall temperatures and the chiller's heat as entries makes every reading row a comparison of one entry with its reading and puts all of the physics into three rows of its own. The reading rows are then linear in the state, and the whole nonlinearity of the problem sits in the two hall physics rows, which hold the products $Q_i\,\mathrm{e}^{-g_i}$. The same choice turns three extensions into a change of one row. A storage term in the balance, the transient of Section 1.6.5, makes the hall temperature at one time depend on its value at the previous time, and the physics row becomes the discretized balance between two consecutive entries. A model discrepancy (Section 1.1.4) is the tolerance σ_p itself, set to a physical value instead of taken to its limit. And several rows that refer to the same temperature, a second sensor in the same hall or a question about the temperature itself, each touch the one entry, whose mean and variance are then read off the fit as the full-load temperature was in Section 4.3.6.

Common mistake

Substituting a reading for the state. Setting $T_s = y_{T_s} = 12.10\,^{\circ}\mathrm{C}$ and fitting the other four entries treats the supply sensor as exact. The supply temperature then cannot move toward what the two hall readings and the set point say, the hall rows absorb the sensor's error in the conductances, and the fit reports spreads that ignore the supply sensor's 0.3 K. A reading is a row with an error; the quantity it measures stays a free entry of the state.

Three kinds of row and their normalization. Chapter 1 distinguished three kinds of row (Section 1.1.7): physical rows, which state relations that must hold; reading rows, which compare the state with a measured value; and prior rows, which carry knowledge held before the readings. The two-hall plant has all three. Its physical rows are the two hall balances and the header balance, each divided by its tolerance. Its reading rows are the five rows of Table 17.1, each the comparison of one entry with its reading, as in equation (17.9). Its prior rows are the three rows of Table 17.2. A prior row has the form of a reading row, the entry minus its prior mean divided by the prior's standard deviation, because a normal prior $z_i \sim \mathcal{N}(z_i^0, \sigma^2)$ has the

density $\exp[-(z_i - z_i^0)^2/(2\sigma^2)]$ up to a constant: a handbook value is a reading of the entry taken before the test.

Every row is written normalized, divided by the standard deviation of its error: a reading or prior row by the spread of the sensor or the handbook, a physics row by its tolerance. The normalized row is dimensionless, and its weight is one: the term it puts into the objective, $\frac{1}{2}r_k^2$, equals the term $\frac{1}{2}w_k d_k^2$ of the unnormalized difference d_k, in the units of the reading, with weight $w_k = 1/\sigma_k^2$. Normalization therefore changes nothing in the objective. It changes what a row's value means: a normalized residual of 2 says that the state and the reading disagree by twice what the sensor's error allows, whatever the sensor measures, so that a temperature row in kelvin and a meter row in kilowatts can be compared on one scale. With the normalized rows stacked in the residual vector $\mathbf{r}$ and their Jacobian $\mathbf{J}$, the information matrix is $\mathbf{J}^\mathsf{T}\mathbf{J}$; written with the unnormalized differences, whose Jacobian is $\tilde{\mathbf{J}}$, and the weight matrix $\mathbf{W} = \mathrm{diag}(w_k)$, it is the same matrix $\tilde{\mathbf{J}}^\mathsf{T}\mathbf{W}\tilde{\mathbf{J}}$, equation (4.26).

The residual vector. Collecting the eleven rows,

$$\mathbf{r}_{\text{full}}(\mathbf{z}_{\text{full}}) = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \\ r_5 \\ r_6 \\ r_7 \\ r_8 \\ r_9 \\ r_{10} \\ r_{11} \end{bmatrix} = \begin{bmatrix} (T_1 - y_{T_1})/\sigma_T \\ (T_2 - y_{T_2})/\sigma_T \\ (T_s - y_{T_s})/\sigma_s \\ (P - y_P)/\sigma_P \\ (Q_1 - y_{Q_1})/\sigma_Q \\ (g_1 - g^0)/\sigma_g \\ (g_2 - g^0)/\sigma_g \\ (T_s - T_s^0)/\sigma_0 \\ (T_1 - T_s - Q_1\,\mathrm{e}^{-g_1}/G_{\text{ref}})/\sigma_p \\ (T_2 - T_s - Q_2\,\mathrm{e}^{-g_2}/G_{\text{ref}})/\sigma_p \\ (P - Q_1 - Q_2)/\sigma_{p,P} \end{bmatrix} \begin{matrix} \text{reading, hall 1} \\ \text{reading, hall 2} \\ \text{reading, supply} \\ \text{reading, heat meter} \\ \text{reading, load meter} \\ \text{prior, } g_1 \\ \text{prior, } g_2 \\ \text{prior, } T_s \\ \text{physics, hall 1} \\ \text{physics, hall 2} \\ \text{physics, header} \end{matrix} \tag{17.10}$$

a vector of eleven dimensionless numbers for each candidate full state, in three blocks: the five reading rows $\mathbf{r}_\mathrm{y}$, the three prior rows and the three physics rows $\mathbf{r}_\mathrm{p}$. The reading and prior rows together are written $\mathbf{r}_\mathrm{e}$, the rows of the evidence, and their Jacobian $\mathbf{J}_\mathrm{e}$, whose first five rows, the reading rows, are $\mathbf{J}_\mathrm{y}$; the physics rows have the Jacobian $\mathbf{J}_\mathrm{p}$. Rows 1 to 8 and 11 are linear in the full state; rows 9 and 10 are not, through the products $Q_i\,\mathrm{e}^{-g_i}$, and those two physics rows are the whole nonlinearity of the problem.

The objective. For rows with independent normal errors the minus logarithm of the posterior is, up to a constant, half the sum of the squared normalized rows: Chapter 4 derived this for the three-entry state of one hall with its physics rows, equation (4.24), and Chapter 18 derives the multi-row form for this plant. The objective of the full formulation is therefore

$$\mathcal{L}_{\text{full}}(\mathbf{z}_{\text{full}}) = \tfrac{1}{2}\sum_{k=1}^{11} r_k^2 = \tfrac{1}{2}\,\mathbf{r}_{\text{full}}^\mathsf{T}\mathbf{r}_{\text{full}}, \qquad p(\mathbf{z}_{\text{full}} \mid \mathbf{y}) \propto \exp\bigl[-\mathcal{L}_{\text{full}}\bigr], \tag{17.11}$$

with unit weights because the rows are normalized; the flat priors of the loads contribute no row. It is a single number for each candidate state, the sum of the squared disagreements of that state with every reading, every prior and every balance, each in units of its own standard deviation.

The Jacobian and the null directions of the balances. The Jacobian of the full residual vector is the 11×8 matrix

$$\mathbf{J}_{\text{full}} = \frac{\partial \mathbf{r}_{\text{full}}}{\partial \mathbf{z}_{\text{full}}^\mathsf{T}}, \qquad J_{ki} = \frac{\partial r_k}{\partial z_i}, \tag{17.12}$$

with the rows k of equation (17.10) down and the entries i of equation (17.6) across; its nonzero cells are the marks of the pattern (17.8). The hall physics rows carry the physics and their derivatives carry it too. Row $8+i$, for $i = 1, 2$, is $r_{8+i} = (T_i - T_s - Q_i\,\mathrm{e}^{-g_i}/G_{\mathrm{ref}})/\sigma_p$, and the derivative of the exponential is $\partial\,\mathrm{e}^{-g_i}/\partial g_i = -\mathrm{e}^{-g_i}$, so

$$\frac{\partial r_{8+i}}{\partial g_i} = \frac{Q_i\,\mathrm{e}^{-g_i}}{G_{\mathrm{ref}}\,\sigma_p}, \qquad \frac{\partial r_{8+i}}{\partial Q_i} = -\frac{\mathrm{e}^{-g_i}}{G_{\mathrm{ref}}\,\sigma_p} = -\frac{1}{G_i\,\sigma_p}, \qquad \frac{\partial r_{8+i}}{\partial T_s} = -\frac{1}{\sigma_p}, \qquad \frac{\partial r_{8+i}}{\partial T_i} = \frac{1}{\sigma_p}, \tag{17.13}$$

and the derivatives of the row with respect to the other hall's entries and to P are zero. On the balance, $Q_i\,\mathrm{e}^{-g_i}/G_{\mathrm{ref}} = T_i - T_s$, the first entry is the temperature rise of the hall in units of the tolerance: a relative increase of 1 % in the conductance, $\delta g_i = 0.01$, lowers the temperature the balance predicts by 1 % of the rise. The second is the reciprocal conductance: one more kilowatt raises the predicted hall by $1/G_i$ kelvin. The third is the same for both halls: the whole hall moves with the water that cools it. The fourth says that the row compares the entry T_i with that prediction.

The derivatives of the other rows. The other nine rows are linear in the state and their derivatives are constants. Each reading row has a single entry, the reciprocal of its sensor's standard deviation in the column of the entry it reads: $\partial r_1/\partial T_1 = \partial r_2/\partial T_2 = 1/\sigma_T$, $\partial r_3/\partial T_s = 1/\sigma_s$, $\partial r_4/\partial P = 1/\sigma_P$ and $\partial r_5/\partial Q_1 = 1/\sigma_Q$. Each prior row has one entry too, the reciprocal of its prior's standard deviation: $1/\sigma_g$ for g_1 and g_2, and $1/\sigma_0$ for T_s. The header physics row has three, $\partial r_{11}/\partial P = 1/\sigma_{p,P}$ and $\partial r_{11}/\partial Q_1 = \partial r_{11}/\partial Q_2 = -1/\sigma_{p,P}$: the two loads enter the header balance with the same weight, so the balance, and the heat meter that sees the loads through it, cannot tell them apart. A Jacobian entry of a normalized row has the units of the reciprocal of its entry of the state: none for a log-conductance, kW^{-1} for a load or the chiller's heat, K^{-1} for a temperature.

The null directions of the physics rows. Split the full state into its first five entries, the reduced state $\mathbf{z} = (g_1, g_2, Q_1, Q_2, T_s)$ of equation (17.7), and its last three, $\mathbf{z}_{\mathrm{b}} = (T_1, T_2, P)$, and the physics block of the Jacobian into the matching columns, $\mathbf{J}_{\mathrm{p}} = [\,\mathbf{J}_{\mathrm{pa}} \;\; \mathbf{J}_{\mathrm{pb}}\,]$, a 3×5 and a 3×3 block. A move $\Delta\mathbf{z}_{\mathrm{full}}$ keeps the linearized balances when $\mathbf{J}_{\mathrm{p}}\Delta\mathbf{z}_{\mathrm{full}} = \mathbf{0}$, that is $\mathbf{J}_{\mathrm{pa}}\Delta\mathbf{z} + \mathbf{J}_{\mathrm{pb}}\Delta\mathbf{z}_{\mathrm{b}} = \mathbf{0}$. The block $\mathbf{J}_{\mathrm{pb}} = \mathrm{diag}(1/\sigma_p,\ 1/\sigma_p,\ 1/\sigma_{p,P})$ is invertible, so

$$\Delta\mathbf{z}_{\mathrm{b}} = \mathbf{N}_{\mathrm{b}}\,\Delta\mathbf{z}, \qquad \mathbf{N}_{\mathrm{b}} = -\mathbf{J}_{\mathrm{pb}}^{-1}\mathbf{J}_{\mathrm{pa}}, \qquad \mathbf{N} = \begin{bmatrix} \mathbf{I}_5 \\ \mathbf{N}_{\mathrm{b}} \end{bmatrix}, \qquad \mathbf{J}_{\mathrm{p}}\mathbf{N} = \mathbf{0}, \tag{17.14}$$

and the five columns of the 8×5 matrix $\mathbf{N}$ are the null directions of the physics rows, the tangent of the set on which the balances hold, one column per free direction: each moves one entry of the reduced state by one unit and the three last entries as the linearized balances require, by its lower block. The tolerances cancel in that block $\mathbf{N}_{\mathrm{b}}$, which by equation (17.13) and the header row is

$$\mathbf{N}_{\mathrm{b}}(\mathbf{z}) = \begin{bmatrix} -Q_1\mathrm{e}^{-g_1}/G_{\mathrm{ref}} & 0 & 1/G_1 & 0 & 1 \\ 0 & -Q_2\mathrm{e}^{-g_2}/G_{\mathrm{ref}} & 0 & 1/G_2 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{bmatrix},$$

with columns $(g_1, g_2, Q_1, Q_2, T_s)$ and rows (T_1, T_2, P) in °C, °C and kW. Its columns read as the physics: on the balance the null direction of g_i lowers T_i by the hall's rise $T_i - T_s$ per unit of g_i; the direction of Q_i raises T_i by $1/G_i$ kelvin and P by one kilowatt per kilowatt; the direction of T_s raises both halls with the supply. The physics block has rank three, so these five directions are all of its null space.

The log-conductance as the entry. The conductance is carried as $g_i = \ln(G_i/G_{\mathrm{ref}})$ rather than as G_i, and the Jacobian shows three reasons. The first is positivity, as in Chapter 1: every real value of g_i is a positive conductance, so the fit can move the entry freely. The second is the prior. A handbook spread is a relative spread, 35 % of the value, and a normal prior on g_i is a prior of that kind, a factor $\mathrm{e}^{\pm 0.3}$ at one standard deviation; written on G_i the same knowledge is a log-normal density, which no row linear in G_i can carry. The third is the form of the Jacobian entries. In the log-conductance the hall physics row has the entry and the null direction has the component

$$\frac{\partial r_{8+i}}{\partial g_i} = \frac{T_i - T_s}{\sigma_p}, \qquad (\mathbf{N}_{\mathrm{b}})_{i,g_i} = -(T_i - T_s),$$

the temperature rise of the hall, which the hall reading sees as $-(T_i - T_s)/\sigma_T$: it has a physical meaning, it is the same for every conductance that produces the same rise, and it scales with the quantity the sensor sees. In the conductance itself they would be

$$\frac{\partial r_{8+i}}{\partial G_i} = \frac{Q_i}{G_i^2\,\sigma_p}, \qquad (\mathbf{N}_{\mathrm{b}})_{i,G_i} = -\frac{Q_i}{G_i^2},$$

an entry that depends on the units of G_i and that changes as $1/G_i^2$ across the range of conductances, so that the linearization would hold over a narrower range.

The exact balances as a limit. As $\sigma_p \to 0$ the physics rows become the exact balances. A state that misses a hall balance by δ kelvin pays $\frac{1}{2}(\delta/\sigma_p)^2$ in the objective, a term that grows without bound as the tolerance shrinks, so the posterior concentrates on the set of states where the three balances hold. On that set the full state can still move in five directions, the null directions of the physics rows derived in Section 17.1.7, and the readings and priors decide where on it the state lies; the reduced state of Section 17.2.1 is the set described by those five directions. The limit changes no count: the three physics rows bring three rows, and the three directions they remove are the three entries T_1, T_2, P they name, so the redundancy is eleven rows less eight entries, the same as eight readings and priors less five free directions, equation (17.27). The limit is taken at the end of a computation and checked numerically, never used to remove an entry by hand; Section 17.4.1 checks it on the covariance and on the first step of the fit. A tolerance of physical size, a balance that fails by a few tenths of a kelvin, is a different model, one that adds the discrepancy's variance to every quantity the balance determines (equation (4.34)).

17.2 The posterior of the state (step 3)

In the book's notation $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state and $f(\mathbf{z})$ a quantity derived from it. Four objects recur in the specialized forms. The *reduced state* $\mathbf{z}$ of equation (17.7) is the vector of the five free entries, on which the rows act through the full state, whose physics rows hold in the limit $\sigma_p \to 0$. The *MAP state* $\hat{\mathbf{z}}$ is the minimum of the objective that the eight rows of the evidence put on the reduced state; it is not reached here, and every object is evaluated at the prior point $\mathbf{z}^{(0)}$, the linearization point of the first step. The *information matrix* $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$, with $\mathbf{J}$ the 8×5 Jacobian of the normalized rows on the reduced state, is the Gauss–Newton curvature of the objective, and $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$ is the covariance the fit would report if the point were its minimum. No closed form exists, since two of the rows multiply a load by an exponential of a log-conductance.

All four questions are read from one posterior: the posterior of the state that the eight rows of the evidence put on the reduced state, its balances exact. This section derives it, sets out the linearized step that computes it, and evaluates both for the plant at the prior point; Section 17.3 reads every answer from it.

END OF THE PREVIEW

Reliability Engineering by Derivation

These are the opening pages. The complete book runs 431 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