

Capacity Planning and Design by Derivation

Load-carrying capability, sizing and redundancy

Draft · 343 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Capacity Planning and Design by Derivation

Volume I

From the planner's questions to probabilistic inference and decisions

Ricardo Marquez

October 31, 2026 (draft for review)

How to read this book

This book is written for the capacity planner, the mechanical design engineer and the facilities planner of a data center. Its questions are the ones a design review asks: whether the plant will carry the load, how much to build, where the bottleneck is, what the uncertainty in the design inputs costs, how the design holds when a unit fails, and what the design's risk is worth. A design sheet answers them by summing nameplate capacities, setting a design-day load against them and applying a margin, and returns one number with no band. It carries handbook values as if they were exact, reads one constraint at a time, and keeps availability and capacity in separate tools. Each question in this book is instead a probability statement over the unknowns of one model of the plant, and the book derives and computes each one.

The plant grows with the book. Part I works on one equipment hall cooled by a single unit, a state of two and then three entries, so that every object of the method can be seen whole. Part II takes one data hall with three cooling units and five racks, calibrated on a night of its own telemetry, and asks what load it can carry. From Part III the plant is the whole site: two halls on one chilled-water plant with its chillers, pump, shared supply pipe and power feed, an economizer that is drawn but not yet built, a year of monthly operating records, and the design alternatives the planner must choose between. On each plant one fitted model serves every question: a what-if is that fit read at a load, an ambient or a unit state not yet seen, and a build decision is an enumeration of such what-ifs with a cost attached. The chapters say which fit they read.

Every chapter that answers a question follows the same progression, and the progression is the method.

Step 0: the plant graph. The plant as it is, before any question is asked: its nodes, paths and sources; which quantities are pinned and which are free; the physical relations that must hold among them. The free quantities form the state vector $\mathbf{z}$.

Step 1: the question. The planner's question in the planner's words, then as a probability or an expectation of a function of the state, with the unknown on the left of the conditioning bar and everything treated as known on the right.

Step 2: the graph extended for the question. Scenario nodes for a load, an ambient, a unit out of service or a design alternative not yet built, and the record attached as rows.

Step 3: the posterior of the state. Derived and computed once, before any question is answered: the MAP state, the information, the covariance, and the exact posterior where one exists.

Step 4: the answers. Each question read from that posterior down to its number: the load at which a limit binds, a joint tail over the constrained entries, an availability from an enumeration of states, or an expected cost.

Step 5: evaluation and checks. The gaps between the reported values and the exact values or draws collected, the approximations tested and, because the record is constructed from a stated truth, the answers tested against that truth.

Step 6: the report. The headroom, the probability, the date or the design the planner signs, with its band and the exact value beside it where one exists.

Step 7: the next measurement. If the band is too wide to decide the build, the tests and surveys on the existing plant ranked by the width each would take off the quantity that decides it.

Five kinds of question cover what the book asks, distinguished by the mathematical object that answers them rather than by the word the design review uses: *estimation*, *prediction under a threshold*, *hypothesis comparison*, *diagnosis* and *design*. Beneath them sit five computations: the *posterior* of the state, the *trend* of an entry across windows, the *evidence* of a drawing of the plant, the *preposterior* after a measurement not yet taken, and the *expected consequence* of a decision. Capacity planning is mostly the posterior read at a scenario and the expected consequence of a design, with the preposterior deciding what to measure before committing. The notation chapter that follows fixes the symbols, and the appendix collects the planner's questions in an atlas that maps each to its mathematical form, its computation and the section that answers it.

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 method and its notation are those of the companion books *Reliability Engineering by Derivation* and *Operations and Control by Derivation*, which ask the reliability engineer's and the operator's questions of the same plant.

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 **vii**

I Thinking in Probabilities **1**

1 Planned plants as graphs **2**
1.1 One system and several questions . . . 2
1.2 The posterior of the state (step 3) . . . 15
1.3 The answers (step 4) . . . 20
1.4 The checks (step 5) . . . 25
1.5 The report (step 6) . . . 28
1.6 The next measurement (step 7) . . . 29
1.7 What to retain . . . 33
1.8 Exercises . . . 34

2 The language of capacity questions **37**
2.1 Planning questions and their mathematical form . . . 37
2.2 The grammar of probability statements . . . 40
2.3 The general forms and the five computations . . . 42
2.4 Estimation, prediction and joint events . . . 54
2.5 Model comparison, diagnosis and decisions . . . 66
2.6 Measurements, classification and workflow . . . 70
2.7 What to retain . . . 76
2.8 Exercises . . . 79

3 Uncertainty and conditional probability **82**
3.1 Uncertainty and evidence in design questions . . . 82
3.2 Conditional probability and Bayes' theorem . . . 87
3.3 Bayesian diagnosis of a capacity shortfall . . . 90
3.4 Marginalization and dependence . . . 93
3.5 Conditioning in a physical model . . . 97
3.6 Predictive variance and the posterior covariance . . . 106
3.7 What to retain . . . 112
3.8 Exercises . . . 115

4 The anatomy of probabilistic inference **117**
4.1 From a plant graph to a design probability . . . 117
4.2 The posterior of the state (step 3) . . . 125
4.3 The answers (step 4) . . . 138

4.4 The checks (step 5) . . . 145
4.5 The report (step 6) . . . 150
4.6 The next measurement (step 7) . . . 152
4.7 What to retain . . . 155
4.8 Exercises . . . 159

II Load-Carrying Capability 161

5 Thermal headroom 162
5.1 How much more load hall B can take . . . 162
5.2 The posterior of the state (step 3) . . . 169
5.3 The answers (step 4) . . . 172
5.4 The checks (step 5) . . . 179
5.5 The report (step 6) . . . 180
5.6 The next measurement (step 7) . . . 182
5.7 What to retain . . . 185
5.8 Exercises . . . 186

6 Load growth scenarios 188
6.1 When hall B leaves its envelope . . . 188
6.2 The posterior of the state (step 3) . . . 191
6.3 The answers (step 4) . . . 191
6.4 The checks (step 5) . . . 198
6.5 The report (step 6) . . . 199
6.6 The next measurement (step 7) . . . 200
6.7 What to retain . . . 202
6.8 Exercises . . . 203

7 Design-day assessment 205
7.1 Hall B on the design day . . . 205
7.2 The posterior of the state (step 3) . . . 209
7.3 The answers (step 4) . . . 210
7.4 The checks (step 5) . . . 216
7.5 The report (step 6) . . . 217
7.6 The next measurement (step 7) . . . 219
7.7 What to retain . . . 220
7.8 Exercises . . . 221

8 Load diversity and coincident peaks 223
8.1 The coincident peak of hall B at its design population . . . 223
8.2 The posterior of the two states (step 3) . . . 228
8.3 The answers (step 4) . . . 231
8.4 The checks (step 5) . . . 237
8.5 The report (step 6) . . . 238
8.6 The next measurement (step 7) . . . 239
8.7 What to retain . . . 242
8.8 Exercises . . . 243

III Sizing and Redundancy **244**

9 Redundancy level **245**
9.1 N+1 or N+2 for the site's chillers . . . 245
9.2 The posterior of the state (step 3) . . . 251
9.3 The answers (step 4) . . . 254
9.4 The checks (step 5) . . . 260
9.5 The report (step 6) . . . 262
9.6 The next measurement (step 7) . . . 263
9.7 What to retain . . . 265
9.8 Exercises . . . 266

10 Unit sizing **268**
10.1 Two quotes for one chiller plant . . . 268
10.2 The posterior of the state (step 3) . . . 274
10.3 The answers (step 4) . . . 275
10.4 The checks (step 5) . . . 282
10.5 The report (step 6) . . . 283
10.6 The next measurement (step 7) . . . 284
10.7 What to retain . . . 286
10.8 Exercises . . . 287

11 Expansion timing **289**
11.1 When the fourth chiller has to be ordered . . . 289
11.2 The posterior of the state (step 3) . . . 294
11.3 The answers (step 4) . . . 294
11.4 The checks (step 5) . . . 301
11.5 The report (step 6) . . . 302
11.6 The next measurement (step 7) . . . 304
11.7 What to retain . . . 305
11.8 Exercises . . . 306

12 Oversizing and stranded capacity **308**
12.1 The build-out drawing at the ten-year load . . . 308
12.2 The posterior of the state (step 3) . . . 313
12.3 The answers (step 4) . . . 315
12.4 The checks (step 5) . . . 321
12.5 The report (step 6) . . . 322
12.6 The next measurement (step 7) . . . 324
12.7 What to retain . . . 326
12.8 Exercises . . . 327

A Atlas of planning questions **329**

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. Hall subscripts are $h \in \{A, B\}$, cooling-unit subscripts $u \in \{A1, A2, A3, B1, B2, B3\}$, rack subscripts $k \in \{A1, \ldots, A5, B1, \ldots, B5\}$ and chiller subscripts $i \in \{C1, C2, C3\}$; 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 at once; its i-th entry.
$\mathbf{z}_t$, t	varies, –	The state of window t and the window index: the five-minute windows of the night record in Part II, the monthly windows of the site's operating year from Part III.
$\mathbf{y}$, y	varies	The evidence: all readings and prior information treated as known; a single reading, subscripted by what it reads (y_{T_s}, $y_{P_{\mathrm{ch}}}$).
$r_k(\mathbf{z})$	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 (30 kW now, 60 kW the design load); effective conductance of its air path to the supply air.
g, G_{ref}	–, kW/K	Log-conductance $\ln(G/G_{\mathrm{ref}})$, 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 halls of Part II and the site		
T_s	°C	Chilled-water supply temperature, 14°C at the set point.
Q_h, $Q_{o,h}$	kW	IT heat load of hall h; its other heat (fans, lighting), pinned.
G_u, $\bar{G}_u$	kW/K	Effective conductance of the coil of cooling unit u; its rated value, pinned.

Symbol	Unit	Meaning
g_u	–	Log-derate $\ln(G_u/\bar{G}_u)$ of unit u.
C_u, Q_u	kW/K, kW	Air capacity rate of unit u, pinned; the heat it rejects to the chilled-water loop.
$T_{r,h}$, $T_{c,h}$, $T_{d,u}$	°C	Return and cold-aisle air temperature of hall h; discharge air temperature of unit u.
T_k, β_k	°C, –	Inlet temperature of rack k; its recirculation fraction.
The chilled-water plant, from Part III		
Q_{IT}, Q_L	kW	Site IT load; heat to the chilled-water loop, Q_{IT} plus the halls' other heat.
T_{wb}	°C	Outdoor wet-bulb temperature, a known input of each window; 22°C at the design day.
N_i, κ_i	kW, –	Nameplate capacity of chiller i at the design day; its as-built log-ratio, the entry that carries the capacity $N_i e^{\kappa_i}$.
d	1/K	Capacity derate per kelvin of wet bulb above the rating condition.
G_e, g_e	kW/K, –	Conductance of the economizer path, drawn and not yet built; its log-ratio to the handbook value.
COP, c	–, –	Efficiency of the chiller plant; the log of its value at 18°C wet bulb.
V, ΔP	L/s, kPa	Loop flow; pressure drop across the shared supply pipe.
k_p, η	varies, –	Flow coefficient of the shared pipe in L/s per $\mathrm{kPa}^{1/2}$, $\Delta P = (V/k_p)^2$; pump efficiency.
P_{ch}, P_{pump}, P_{feed}	kW	Chiller power, pump power, power through the feed.
Scenarios, capacity and design		
x^*	varies	A starred quantity is read at a scenario node: the plant re-solved at a load, ambient, unit state or design alternative not yet seen, with $\mathbf{\Sigma}$ carried through. A starred limit ($T_{\max}$, f^*) is a requirement level; the context tells the two apart, as in the companion books.
$T_{\max}$	°C	A limit on a temperature: the hall's 27.5°C at the design load in Part I; a rack inlet's 27°C, the recommended envelope, from Part II.
$Q_e^{\dagger}$	kW	Binding load of element e: the site or hall IT load at which its limit binds.
H	kW	Headroom: the lowest binding load less the load today.
m, m_0	–	Capacity margin, capacity over the design load less one; the drawing's margin, 0.20.
A	–	Service availability, the probability of the states in which the load is carried; never equated with a capacity margin.
EUE	kWh/yr	Expected unserved energy: the IT load shed per year, in expectation.
λ, μ	1/yr, 1/h	Failure rate and repair rate of a unit.
γ_L	1/yr	Growth rate of the IT load, a stated input.
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, φ its density.
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.
$\mathbf{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\mathbf{\Lambda}^{-1}$; 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 $\hat{\mathbf{z}}$.
f_L, f_U	varies	The ends of the credible interval of a derived quantity.
$\mathbf{J}_a$, $\mathbf{\Lambda}_a$, $\mathbf{\Sigma}_a$	varies	The Jacobian rows of a candidate measurement a, its information $\mathbf{J}_a^\mathsf{T}\mathbf{J}_a$, and the preposterior covariance $(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$.
M, B_{AB}	–	A drawing (model) of the plant; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$.
$\mathcal{R}(a)$, $L(a, \cdot)$	varies	Expected cost (risk) of a candidate design or action a; its loss.
U, U_a	–	Unavailability, $1 - A$; of drawing a (Part III).

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.

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 the record, prior rows carry handbook or commissioning knowledge.

Calibration The fit of the existing plant's parameters to its operating record: the MAP state and its covariance. Every what-if reads it.

Scenario (what-if) The plant re-solved at a scenario node, with a load, an ambient, a unit state or a design alternative pinned to a value not yet seen, and the covariance carried through.

Design day The ambient condition the design is drawn for, taken from the site's climate data; here a wet bulb of 22°C.

Headroom, binding load The load the plant can still take before a limit binds; the load at which an element's limit binds.

Capacity margin Capacity at the design day over the design load, less one. A statement about capacity, never about availability.

Coincidence (diversity) factor The ratio of the coincident peak of a group of loads to the sum of their individual peaks.

N+1, N+2 redundancy Installed units beyond the N needed to carry the design load.

Concurrent maintainability Any one unit can be taken out for service without interrupting the load.

Binding constraint, debottlenecking The element whose limit binds first as the load grows; raising that element's rating.

Importance measure The sensitivity of a system measure to one element: here the gradient of headroom with respect to an element's rating.

Minimal cut set, common-mode failure A smallest set of elements whose joint loss fails the service; a failure that defeats redundant units through an element they share.

Cascading failure A failure whose shifted load drives a second element past its envelope.

Expected unserved energy, loss-of-load probability The IT energy shed per year in expectation; the probability that the load is not carried.

Stranded (oversized) capacity Capacity built and paid for that the load does not use.

Laplace approximation, delta method Replacing the posterior by the Gaussian at the MAP state with the inverse curvature as covariance; linearizing a derived quantity so that its variance is $\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}$.

Variance budget The share of a derived quantity's variance carried by each entry of the state.

Preposterior, value of information The posterior that would result from a measurement not yet taken; the expected improvement in the decision that measurement buys.

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

Part I

Thinking in Probabilities

Chapter 1

Planned Plants as Graphs

1.1 One system and several questions

A design review has one data center equipment hall on its agenda. The hall carries 30 kW of servers today, and its design load is 60 kW, twice that; at the design load the hall air must stay below 27.5°C. The record is thin: one hall sensor, with an error of about half a kelvin, read 23.2°C at commissioning, under today's load, and the design handbook gives the conductance of the hall's air path to the supply air, held at 18°C by the cooling unit's discharge controller, as 7.5 kW/K, right to within a factor of about 1.35. The design sheet carries two numbers for the hall at its design load. The first is the nameplate forward solve, which works from the handbook value. Heat leaves the hall in proportion to its excess over the supply air, so at 7.5 kW/K the hall sits at 22°C under 30 kW, comfortably below its present limit of 24°C, and at 26.0°C under 60 kW, well below 27.5°C: the sheet marks the design load "pass" with room to spare.

The second number works from the commissioning reading and tells another story. Taken as exact, 23.2°C at 30 kW means a conductance of 5.77 kW/K, and at that conductance the design load puts the hall at 28.4°C, over the limit: the sheet marks it "fail". The review must sign off one of the two, and the sheet cannot say how sure either is; each treats its own input as exact, and neither is weighed against the other. In the planner's words the question is

Will the hall carry its 60 kW *design load inside the* 27.5°C *limit, and how sure is that?*

Its answer is a probability, that the hall stays below 27.5°C at a load it has never carried, and that probability can only come from holding the handbook value, the reading and the energy balance that ties them together in one description of the hall, each with its own uncertainty. The question of this chapter is therefore

how a planned plant is represented so that the same representation serves every probabilistic question asked of it, and what must be fixed before any inference is performed.

Jargon: plant graph, node, path, source, pinned and free quantities, state vector, residual, row, observation, prior, posterior, state estimation, model discrepancy, scenario node, Jacobian, normal equations, information matrix, covariance, Laplace approximation, factor graph

A *plant graph* represents a physical system as *nodes*, places where a conserved quantity such as energy or mass must balance; *paths* (edges), along which the quantity flows; and *sources*, which supply it. A quantity is *pinned* when it is treated as known, as a boundary condition, and *free* when it is unknown and must be inferred. The free quantities, collected in one vector, form the *state vector* $\mathbf{z}$. A *residual* is a relation written so that it equals zero when it is satisfied; each residual is a *row* of the model. A physical row states a relation that must hold among entries of the state; an *observation* (a reading) is a measured value related to an entry of the state with an error, and its row compares the two. The *prior* expresses what is known about the state before the readings, for example a handbook value of a conductance; the *posterior* $p(\mathbf{z} \mid \mathbf{y})$ expresses what is known after the evidence $\mathbf{y}$; fitting the state to its rows, which yields the posterior's most probable state and its spread, is what the planner calls the *calibration* of the existing plant (state estimation). A *model discrepancy* is the difference between the physical system and the model of it that would remain if every entry of the state were known exactly. A *scenario node* is a node added to the graph for an operating condition that has not occurred, so that a question can be asked about it. The *Jacobian* $\mathbf{J}$ is the matrix of derivatives of every row with respect to every entry of the state, one row per row of the model and one column per entry. The *normal equations* $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$ give the step that best reduces all the rows at once when each is replaced by its tangent. The *information matrix* $\mathbf{J}^\mathsf{T}\mathbf{J}$ states how sharply the rows together fix the state, direction by direction, and its inverse is the posterior *covariance* $\boldsymbol{\Sigma}$, whose diagonal holds the variances of the entries and whose off-diagonal entries hold how they vary together. The *Laplace approximation* replaces the posterior by the normal distribution centered on its most probable state with that covariance. A *factor graph* draws the entries of the state and the rows that act on them, one node each.

Key idea

A graph describes the plant; a probabilistic question states what is to be learned about it. The plant graph, its pinned and free quantities and its physical rows are built once, and each question is a probability over the state of that one graph, with the unknown on the left of the conditioning bar and the evidence on the right. The posterior of the state is carried by the most probable state $\hat{\mathbf{z}}$, the covariance $\boldsymbol{\Sigma}$ and the Jacobian $\mathbf{J}$ of the rows, and every answer is an entry of the state read off them or a function of the state carried through them.

1.1.1 Several questions about one equipment hall

A planner reviewing the hall's design might ask any of the following:

1. What is the hall air temperature at today's load?
2. What is the effective conductance of the air path?
3. Will the hall carry its 60 kW design load without going over its limit?
4. Is the hall sensor reading true, or is it biased?
5. Would a measurement taken before the sign-off narrow the answer?
6. What happens if the cooling unit loses half its airflow?

The questions are different. The first two ask for the value of an unknown quantity; the third asks for the probability that a future temperature stays within a limit; the fourth asks which of two descriptions of the sensor is true; the fifth asks what a measurement not yet taken would be worth; the sixth asks about an operating condition that has not occurred. All six, however,

concern the same physical system: the same hall, the same air path, the same sensor.

Constructing a separate mathematical model for every question would repeat the same physics six times, and the six models could disagree about it. The method of this book begins instead with one common representation of the equipment and its physical relations, the *plant graph*. It is step 0 of the method stated in *How to read this book*: the system as it is, before any question is asked, built once and kept for every question that follows. Section 1.1.6 writes the questions that this chapter answers over the graph.

1.1.2 The system, the question and the evidence

Three things must be fixed before any probabilistic inference is performed, and the plant graph is where they are fixed. The first is the division of the quantities: which are known and pinned (the supply-air temperature, the heat load), which are unknown and free (the hall temperature, the conductance), and which are observed (the sensor reading). The field often separates the unknowns into *states*, such as temperatures and flows, and *parameters*, such as conductances and rates; in this book both kinds are entries of one state vector, and the inference treats them alike. The second is the set of physical relations that must hold among the quantities, such as conservation of energy at the hall. The third is the evidence: the readings, the records and the prior knowledge, attached to the graph as rows of their own.

The question is a separate object. It is stated over the graph after the graph is built, and it decides what is computed, not what the system is. The separation makes the questions of a design review systematic: when a new question arrives, the graph is already in hand, and the work is to write the question over its state.

> **Principle 1.1.** *The physical model is reusable; the probabilistic question determines what is calculated from it. A plant graph is built once from the equipment and its physics, and every question is a probability or an expectation over the state of that graph.*

1.1.3 Larger plant graphs and factor graphs

The equipment hall of this chapter has one thermal path and two unknown quantities, the smallest graph with every kind of element. The same construction serves a plant of many components, and the graph on which an inference operates is a second object, drawn from the plant graph and the evidence attached to it.

A cooling network. Real plants may include hundreds or thousands of components, each with its own equations and uncertainties. Consider a cooling network with two chillers, two pumps, a common electrical supply and several temperature sensors. Its plant graph may contain temperatures, pressures, flows and electrical states at its nodes; physical components represented by paths or by component relations; heat, mass and power conservation equations; telemetry and historical records; the run and trip states of equipment; and shared parameters and shared dependencies.

Figure 1.1 draws a reduced version. The green boxes are free nodes: the two chillers and the two pumps, each carrying its status and capacity or flow, the supply header with its water temperature T_s and flow, and the hall with its temperature T. The gray box on the left is the pinned outdoor air, to which the chillers reject heat; the white box in the middle is the common electrical supply, which feeds all four machines; the white box at the upper right is the heat load of the hall. Solid green arrows are physical paths of heat and water; thin gray arrows are electrical feeds; dashed arrows are observations, two temperature readings and the maintenance record of a chiller. The common supply is a single node on which four components depend: if it trips, all four stop together, and a graph that omitted it would treat their trips as independent.

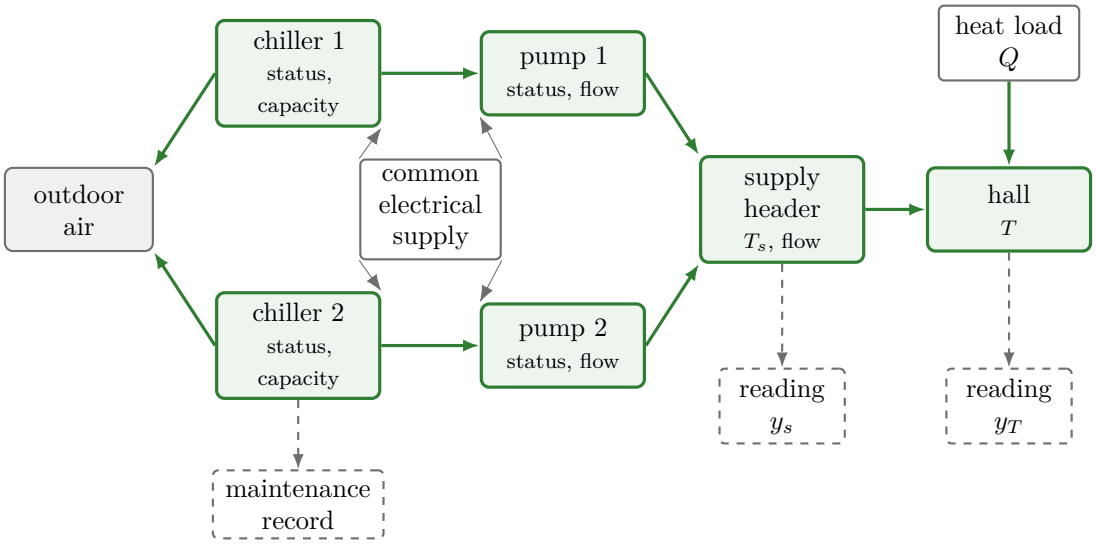


Figure 1.1: Plant graph of a reduced cooling network (illustrative; return piping omitted). Free: the status and capacity of each chiller, the status and flow of each pump, the header water temperature T_s and flow, and the hall temperature T. Pinned: the outdoor air, to which the chillers reject heat, and the heat load Q; the common electrical supply feeds all four machines and is the shared dependency. Observed: T_s and T through two readings, and chiller 2 through its maintenance record.

The mathematical principle is the one the single hall shows below. The network is a collection of physical relations,

$$r_1(\mathbf{z}) = 0, \quad r_2(\mathbf{z}) = 0, \quad \ldots, \quad r_m(\mathbf{z}) = 0,$$

one per balance and per component relation, over a state vector $\mathbf{z}$ that now has many entries. Observations and probability models for the uncertain quantities are attached as rows of their own. What is ultimately sought is the joint posterior of the state,

$$p(\mathbf{z} \mid \mathbf{y}, M),$$

where M represents the assumed system model, including its structure: which components exist, how they are connected and which dependencies are represented. Writing M explicitly matters because the structure itself may be uncertain, and a question about the structure is answered by comparing two such models, not by the posterior inside either.

Physical graphs and factor graphs. A physical graph and a probabilistic factor graph are related, but they are not the same object. The physical graph represents equipment and connections: halls, paths, chillers, pumps, sensors. A *factor graph* represents the variables of the inference and the mathematical relations among them, including the observations, the priors and the dependencies: each entry of the state is one node, each row of the model is one factor node, and each factor is linked to the entries its row involves. It is constructed from the physical graph and the evidence attached to it, and it shows which row acts on which unknown. The distinction becomes important when competing system structures and diagnosis are studied in later chapters. Two physical drawings of the same plant, one with a single air path and one with two, give two factor graphs with different variables and rows, and comparing them is comparing two models M. A sensor bias adds a variable to the factor graph without changing the equipment at all.

1.1.4 The plant graph (step 0)

The simplest plant graph with every kind of element has one free node, one pinned node, one path, one source and one observation, and the equipment hall is that graph. It is built before any question is asked and before any number is fitted.

Nodes, a path, a source and an observation. Consider a simplified equipment hall with one heat-removal path, the hall's air path to the supply air that the cooling unit delivers. Figure 1.2 draws it as a plant graph. The green box on the left is the hall, a free node whose air temperature T is unknown; the gray box on the right is the supply air, a pinned node whose temperature $T_c = 18°\text{C}$ is held by the cooling unit's discharge controller. The solid green arrow between them is the heat-removal path, characterized by its effective thermal conductance G. The box above the hall is the heat load Q supplied by the servers, a source. The dashed box below is the hall sensor, and the dashed arrow from the hall to it is an observation: the sensor reports a value related to T, not T itself. Solid connections are physical relations; the dashed connection is evidence.

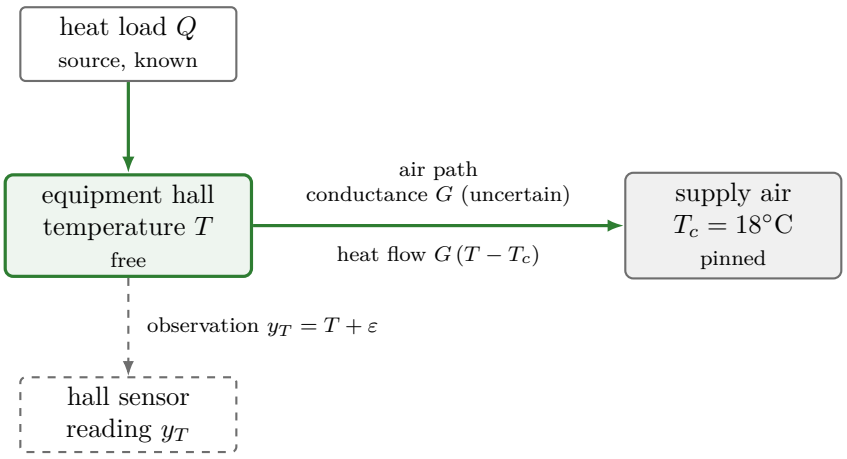


Figure 1.2: Plant graph of the equipment hall in steady state. Free: the hall air temperature T and the conductance G of the air path. Pinned: the supply-air temperature $T_c = 18°\text{C}$, held by the cooling unit's discharge controller, and the heat load Q. Observed: T, through the reading y_T with error ε. Solid arrows are physical relations; the dashed arrow is an observation.

Table 1.1 lists the elements of the graph. The first column names each element, the second its role in the graph, the third its symbol and the fourth its unit; the last column states how the element enters this first example.

Table 1.1: The elements of the plant graph of Figure 1.2.

Element	Role	Symbol	Unit	In the example
Equipment hall	free node	T	°C	air temperature unknown
Supply air	pinned node	T_c	°C	held at 18°C by the discharge controller
Heat load	source	Q	kW	thermal power supplied by the servers, known
Air path	path (edge)	G	kW/K	effective conductance, uncertain
Hall sensor	observation	y_T	°C	a reading related to the hall temperature

For now the supply-air temperature and the heat load are assumed known, and the effective conductance is uncertain. The discharge controller holds T_c at its set point, which is what makes the supply air a pinned boundary; the choice fails if the unit saturates and can no longer hold it (Section 1.6.5). The conductance lumps every mechanism by which the hall's heat reaches the supply air, the airflow through the cooling unit and the mixing of the hall air alike, into one number; how well that lumping represents the hall is a modeling assumption, examined in Section 1.6.5. This is the starting physical model.

Conservation of energy at the hall node. A graph becomes useful for computation when the mathematical relations between its elements are specified. Assume steady state: the hall stores no more and no less heat from one moment to the next, so that its temperature is constant. Conservation of energy at the hall node then requires the heat entering to equal the heat leaving,

$$Q_{\text{in}} = Q_{\text{out}},$$

where Q_{in} and Q_{out} are thermal powers in kilowatts. The heat entering is the load, $Q_{\text{in}} = Q$. The heat leaving is modeled as proportional to the temperature difference across the path,

$$Q_{\text{out}} = G\,(T - T_c),$$

the defining relation of the effective conductance: G is the heat flow per kelvin of difference between the hall and the supply air. Combining the two,

$$Q = G\,(T - T_c), \tag{1.1}$$

where Q is the heat generated by the servers in kW, G the effective thermal conductance in kW/K, T the hall air temperature in °C and T_c the supply-air temperature in °C. A temperature difference has the same value in kelvin and in degrees Celsius, so the units balance: kW/K times K is kW.

Equation (1.1) is a *physical relation*, not a probability statement. It states what must hold if the simplified thermal model is correct, whatever is or is not known about its quantities. Since $G > 0$, dividing by G and adding T_c gives

$$T = T_c + \frac{Q}{G}, \tag{1.2}$$

which gives the temperature when the other three quantities are specified. The hall is warmer than the supply air by Q/G: more load or less conductance raises it.

The forward calculation. The hall carries the load $Q = 30\,\text{kW}$; take as the reference for the forward calculation a conductance of $G = 6\,\text{kW/K}$ (the reference G_{ref}, not the handbook's $7.5\,\text{kW/K}$). With $T_c = 18°\text{C}$, equation (1.2) gives

$$T = 18 + \frac{30\,\text{kW}}{6\,\text{kW/K}} = 18 + 5 = 23°\text{C},$$

the steady-state temperature of the hall at its design conductance. The calculation used no probability, estimated no uncertainty and used no observational evidence. It is a deterministic *forward calculation*: from pinned inputs and a specified conductance to the temperature they imply. Every later chapter needs it, because every inference compares temperatures computed this way with what the sensor reports. Doubling the load doubles the excess $T - T_c$ at every conductance: at $G = 6\,\text{kW/K}$ the hall at 60 kW would be at 28°C.

The state vector. The forward calculation required the conductance to be known. In the inference problems of this book it is not, and neither is the hall temperature. The *state vector* is the collection of the unknown quantities that the inference estimates at once. In the equipment hall the unknowns are the hall temperature and the thermal conductance, and one possible representation is

$$\mathbf{z} = \begin{bmatrix} T \\ G \end{bmatrix}.$$

The state holds everything the inference treats as free; the pinned quantities T_c and Q are not in it, and neither is the reading y_T, which is evidence. The concept recurs throughout the book: the graph fixes the state, the evidence constrains it, and each question is a statement about it. Quantities that are entries of the state, such as T here, are answered directly by the inference; quantities that are functions of the state, such as the heat flow through the path, require one further transformation.

The log-conductance. A conductance is positive, and an estimate of it must be positive too. In practice the conductance is therefore often represented by a logarithmic entry,

$$g = \ln\!\left(\frac{G}{G_{\mathrm{ref}}}\right),$$

where $G_{\mathrm{ref}} > 0$ is a reference conductance in kW/K and g is dimensionless. Inverting,

$$G = G_{\mathrm{ref}}\, \mathrm{e}^{g},$$

which is positive for every real g: an inference free to move g anywhere on the real line can never produce a negative or zero conductance. A change in g is a relative change in G, since $\mathrm{d}G = G_{\mathrm{ref}}\, \mathrm{e}^{g}\, \mathrm{d}g = G\, \mathrm{d}g$, so that $\mathrm{d}g = \mathrm{d}G/G$; uncertainty in g reads directly as a fractional uncertainty in G. The reference value only sets the origin of g. In this chapter it is taken as $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, the design conductance of the forward calculation, so that $g = 0$ there, $G = 7.5\,\mathrm{kW/K}$ corresponds to $g = \ln 1.25 = 0.2231$. The state vector becomes

$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} T \\ g \end{bmatrix}, \tag{1.3}$$

with entries $z_1 = T$ in °C and $z_2 = g$, dimensionless. It is the one state of this chapter: every question below is a statement about these two entries.

The physics row. The thermal relation can now be written in terms of the state. Substituting $G = G_{\mathrm{ref}}\, \mathrm{e}^{g}$ into equation (1.1),

$$Q = G_{\mathrm{ref}}\, \mathrm{e}^{g}\, (T - T_c),$$

or, moving everything to one side,

$$r_{\mathrm{phys}}(\mathbf{z}) = G_{\mathrm{ref}}\, \mathrm{e}^{g}\, (T - T_c) - Q = 0. \tag{1.4}$$

The function $r_{\mathrm{phys}}(\mathbf{z})$ is a *residual*, in kW. A residual measures how far a candidate state is from satisfying a relation: it is zero when the relation holds exactly and grows as the candidate departs from it. With $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, $T_c = 18°\mathrm{C}$ and $Q = 30\,\mathrm{kW}$, the candidate $\mathbf{z} = (23°\mathrm{C},\ 0)$ gives $r_{\mathrm{phys}} = 6 \times 1 \times 5 - 30 = 0$: it is the forward calculation. The candidate $\mathbf{z} = (23.2°\mathrm{C},\ 0)$ gives $6 \times 5.2 - 30 = 1.2\,\mathrm{kW}$: at that temperature and conductance the path would remove 1.2 kW more than the servers supply, so the hall could not be in steady state. The candidate $(23.2°\mathrm{C},\ \ln(30/31.2)) = (23.2°\mathrm{C},\ -0.0392)$ again gives zero: the temperature 23.2°C is consistent with the physics if the conductance is $30/5.2 = 5.769\,\mathrm{kW/K}$. The zero set

of r_{phys} is the curve $T = T_c + Q/(G_{\text{ref}}\,\mathrm{e}^{g})$ in the plane of (g, T); the physics alone allows every point of that curve and no point off it.

The same row has a second form, in the units of the quantity it fixes. Dividing equation (1.4) by $G_{\text{ref}}\,\mathrm{e}^{g} = G$, which is positive, leaves the zero set unchanged and gives the row in kelvin, the departure of a candidate temperature from the temperature the physics implies. Divided further by a *model tolerance* σ_p, the temperature error in kelvin that the physics is allowed, it becomes a dimensionless row,

$$r_1(\mathbf{z}) = \frac{T - T_c - (Q/G_{\text{ref}})\,\mathrm{e}^{-g}}{\sigma_p}, \tag{1.5}$$

which counts the violation of the physics in multiples of its tolerance. This is the form in which the physics row enters every computation of the book. The formulation extends to systems of any size. A larger physical system is represented by a collection of residuals, one per relation: an electrical network has rows for Kirchhoff's current and voltage laws, a fluid network rows for mass conservation at its junctions and for the pressure losses along its pipes, and a cooling plant rows for the energy balance at each of its nodes. Each row has the form of equation (1.4), and the inference treats them alike.

Exact constraints and model discrepancy. Physical rows and reading rows should not automatically be treated as equivalent kinds of noisy measurement. Conservation of energy is a law: if the hall is in steady state and the model has captured every path, equation (1.4) holds exactly, and it can be imposed as an *exact constraint*, $r_{\text{phys}}(\mathbf{z}) = 0$, which restricts the state to the zero set of the row. In a computation the exact constraint is carried as the physics row (1.5) with a model tolerance σ_p far smaller than every other spread of the problem, and exactness is the limit $\sigma_p \to 0$; the limit is checked by showing that the results stop moving as σ_p shrinks, never by using the row to eliminate an entry of the state. A reading row, by contrast, is meant to be violated by an amount of the order of the sensor's error: the reading is not the temperature.

Approximate physics is a third case. The single effective conductance of equation (1.1) is an idealization; the real hall has several paths, a conductance that varies with airflow, and stored heat. When the relation is only approximately true, the honest model is

$$G_{\text{ref}}\,\mathrm{e}^{g}\,(T - T_c) - Q = \delta,$$

with an explicit *model discrepancy* δ in kW, an entry of its own with a stated prior, for example $\delta \sim \mathcal{N}(0, \sigma_\delta^2)$. Algebraically this is a physics row with the weight $1/\sigma_\delta^2$, and in the objective it looks like a reading row. It behaves differently. A measurement error is new at each reading and averages out over many readings; a discrepancy is a property of the model and repeats wherever the same operating condition repeats, so that many readings at one condition do not reduce it. Treating a discrepancy as measurement noise therefore makes the inference overconfident, and treating a law as noisy lets the inference trade physics against data. The choice between an exact constraint and an explicit discrepancy is made when the graph is built, and it is stated with the model. In this chapter conservation of energy is exact.

1.1.5 The record and the handbook value

The record of the hall is one reading. The hall sensor, the one temperature sensor installed in the hall, reported at commissioning, under today's load of 30 kW,

$$y_T = 23.2^{\circ}\mathrm{C}.$$

The sensor is not perfectly accurate. Its measurement model is

$$y_T = T + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_T^2),$$

where ε is the measurement error in kelvin, assumed normal with zero mean, and σ_T its standard deviation, which the sensor's specification gives as

$$\sigma_T = 0.5\,\mathrm{K}.$$

The reading is the only observation: the supply-air temperature and the load are pinned, and nothing measures the conductance.

Knowledge of the conductance that does not come from the reading is also in hand. A design handbook, the commissioning documents or earlier tests of similar halls give an estimate of the effective conductance with an uncertainty; for this hall (illustrative) the handbook gives 7.5 kW/K, with the true conductance within a factor of about 1.35 of it with probability of about 0.68. The handbook value is not a measurement of this hall. It is information held before the reading, and it enters the model as a prior.

1.1.6 The questions (step 1)

The planner's questions can now be written over the graph, each as a probability or an expectation over the state, with the unknown on the left of the bar and the evidence on the right. Table 1.2 lists the five that this chapter answers: the first column is the question in the planner's words; the second the general form that every question of its category takes, written in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it and the posterior density $p(\mathbf{z} \mid \mathbf{y})$; the third its mathematical form for the hall; the fourth its category among the five named at the front of the book, which Chapter 2 develops; and the last the subsection where it is answered. The first two are estimation, of an entry of the state and of a quantity derived from one, the conductance $G = G_{\mathrm{ref}}\,\mathrm{e}^{g}$. The third and fourth are predictions under a threshold, tail probabilities of the posterior: the third at the present load, the fourth at a load of 60 kW, the design load, which the hall has never carried, whose temperature T_2 is a quantity the graph acquires in step 2. The fifth is a design question about a measurement not yet taken; it asks nothing new of the present evidence, and it is answered as the next measurement for the first four.

The questions differ in what they compute and agree in what they are computed over: the same two entries of the state, under the same posterior. Two of the planner's questions are not in the table. Whether the hall sensor is biased compares two complete models of the same reading, the plant with an unbiased sensor and the plant with a bias entry on the reading row, and its answer is a ratio of their evidences rather than a property of either posterior; when every sensor of a larger plant is a candidate, the same comparison run over the candidates is a diagnosis. *Operations and Control by Derivation* answers it in its Chapter 7, and the residual check of Section 1.4.3 takes a first, coarse look at it. What happens if the cooling unit loses half its airflow is a question about a transient that the steady-state graph cannot represent, and Section 1.6.5 states what the graph would need for it.

1.1.7 The evidence as rows (step 2)

Step 2 attaches the reading and the handbook value to the graph as rows on the state $\mathbf{z} = (T, g)^{\mathsf{T}}$, and extends the graph with the operating condition the design-load question names. The three rows, the physics row of Section 1.1.4 among them, are Gaussian, and each is written normalized, with a residual and a Jacobian row.

The modeling choices. One choice makes the physics row, and it buys something and can fail. The row carries the balance of Section 1.1.4 in the limit $\sigma_p \to 0$: this is what confines the posterior to the physics curve, so that a reading of the temperature informs the conductance, and it fails if heat leaves the hall by a mechanism the single conductance does not lump or if the hall is not in steady state (Section 1.6.5).

Table 1.2: The questions asked of the equipment hall, the general form of their category, their mathematical form for the hall, their category and where each is answered. T_2 is the hall temperature at the load $Q_2 = 60\,\mathrm{kW}$ (Section 1.1.7); $\boldsymbol{\Sigma}_a$ the covariance expected after a candidate row with Jacobian row $\mathbf{J}_a$.

Planner's question	General form	Mathematical form	Category	Answered in
What is the hall air temperature at today's load?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(T \mid \mathbf{y})$, $T = z_1$	estimation	1.3.1
What is the effective conductance of the air path?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(G \mid \mathbf{y})$, $G = G_{\mathrm{ref}}\,\mathrm{e}^{z_2}$	estimation of a derived quantity	1.3.1
Is the hall below 24°C now?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T < 24^\circ\mathrm{C} \mid \mathbf{y})$	prediction under a threshold	1.3.2
Will the hall carry its 60 kW design load inside 27.5°C?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2 = T_c + (Q_2/G_{\mathrm{ref}})\,\mathrm{e}^{-z_2}$	prediction under a threshold	1.3.2
Would a measurement before the sign-off narrow the answers?	preposterior covariance $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	$\boldsymbol{\Sigma}_a$ for each candidate row	design	1.6.1

The reading row. The reading of Section 1.1.5 becomes the normalized observation residual

$$r_{\mathrm{obs}}(\mathbf{z}) = \frac{T - y_T}{\sigma_T}, \tag{1.6}$$

which is dimensionless: the difference between the candidate temperature and the reading, in units of the sensor's standard deviation. At $T = 23.2^\circ\mathrm{C}$ it is zero; at the design temperature $T = 23^\circ\mathrm{C}$ it is $(23 - 23.2)/0.5 = -0.4$. The row carries the probability of the reading. Given the state, the reading differs from T only by the error, so its density is the normal density of ε evaluated at $y_T - T$:

$$p(y_T \mid \mathbf{z}) = \frac{1}{\sqrt{2\pi}\,\sigma_T} \exp\left[-\frac{(y_T - T)^2}{2\sigma_T^2}\right] = \frac{1}{\sqrt{2\pi}\,\sigma_T} \exp\left[-\tfrac{1}{2}\, r_{\mathrm{obs}}(\mathbf{z})^2\right].$$

Read as a function of the state for the fixed reading, this is the *likelihood* of the state. A candidate state with a small $|r_{\mathrm{obs}}|$ makes the reading probable; the ratio of the likelihoods of $T = 23^\circ\mathrm{C}$ and $T = 23.2^\circ\mathrm{C}$ is $\mathrm{e}^{-0.4^2/2} = \mathrm{e}^{-0.08} = 0.923$, so the reading alone hardly distinguishes the two. The likelihood does not depend on g: the sensor sees the temperature, and it informs the conductance only through the physics row that ties the two together.

The prior row. The handbook value becomes a row of the same form. Written as a normal distribution on the log-conductance, $g \sim \mathcal{N}(g_0, \sigma_g^2)$, it is the row $(g - g_0)/\sigma_g$: a handbook value is treated as a reading of g taken before the hall's own. The handbook conductance of 7.5 kW/K gives $g_0 = \ln(7.5/6) = 0.2231$, and $\sigma_g = 0.3$ states that the true conductance is within a factor $\mathrm{e}^{0.3} = 1.35$ of it with probability of about 0.68. The prior is written $p(g)$ in general; its form is a modeling choice that must be stated and justified, and Section 1.4.4 tests how much this one weighs.

Three kinds of row. The model now holds three very different kinds of relation. Table 1.3 sets them side by side. Each row of the table is one kind of relation; the columns give its equation for the hall, its residual and its role.

Table 1.3: The three kinds of row in the equipment hall model, with $G = G_{\text{ref}}\,\mathrm{e}^{g}$.

Row	Equation	Residual	Role
Physics	$G\,(T - T_c) - Q = 0$	$G_{\text{ref}}\,\mathrm{e}^{g}(T - T_c) - Q$	constrains what is physically possible
Reading	$y_T = T + \varepsilon$	$(T - y_T)/\sigma_T$	provides imperfect information about the state
Prior	$g \sim \mathcal{N}(g_0, \sigma_g^2)$	$(g - g_0)/\sigma_g$	carries knowledge held before the readings

The distinction between the first two rows is foundational. The physical equation constrains the system: it says which combinations of temperature and conductance can occur at all. The reading provides evidence about the unknown quantities: it says which of the possible combinations are plausible in light of what was measured. The prior row has the form of a reading row and the role of knowledge held in advance. A probabilistic inference method combines the three kinds of relation to obtain a distribution over the plausible states.

Rows and weights. Chapter 3 shows that, for normal errors, the negative logarithm of the posterior is, up to a constant, the objective $\mathcal{L}(\mathbf{z}) = \frac{1}{2}\sum_k w_k\, r_k(\mathbf{z})^2$, a weighted sum of the squared rows with weights $w_k = 1/\sigma_k^2$. A reading row may be written either normalized, as in equation (1.6), with unit weight, or in kelvin, as $T - y_T$ with weight $1/\sigma_T^2$; both put the same term $(T - y_T)^2/\sigma_T^2$ into the objective. Every row of this chapter is written normalized: the physics row in the form (1.5) with the tolerance σ_p, the reading row (1.6) and the prior row, so that every row counts in standard deviations and none dominates because of its unit.

The scenario node. The design-load question doubles the electrical load from 30 kW to 60 kW, an operating condition the hall has never carried. A second operating condition is added to the graph, sharing the same unknown conductance. Figure 1.3 draws it. The upper row is the present condition, at which the hall has been operated and observed: the free node at $Q_1 = Q = 30\,\text{kW}$ with temperature T, its path to the pinned supply air, and its sensor. The lower row is the design-load condition, a scenario node at $Q_2 = 60\,\text{kW}$ with temperature T_2 and its own path to the same supply air. The vertical dashed line between the two paths marks that they carry one and the same conductance. The scenario node has no sensor: it is in the graph only because the question asks about it.

The two operating conditions satisfy the energy balance of equation (1.1) each at its own load,

$$Q_1 = G\,(T - T_c), \tag{1.7}$$

$$Q_2 = G\,(T_2 - T_c), \tag{1.8}$$

with $Q_1 = Q$. The two temperatures are different, but the conductance is shared, and that connection is what makes the second condition predictable from the first. The scenario node adds no evidence and no free quantity: by equation (1.8) its temperature is fixed by the conductance, $T_2 = T_c + (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}$, a function of the state $\mathbf{z} = (T, g)^{\mathsf{T}}$, which keeps its two entries and its three rows. Chapter 4 carries T_2 instead as a third entry with a second physics row, the same formulation; with the same rows it gives the same answers.

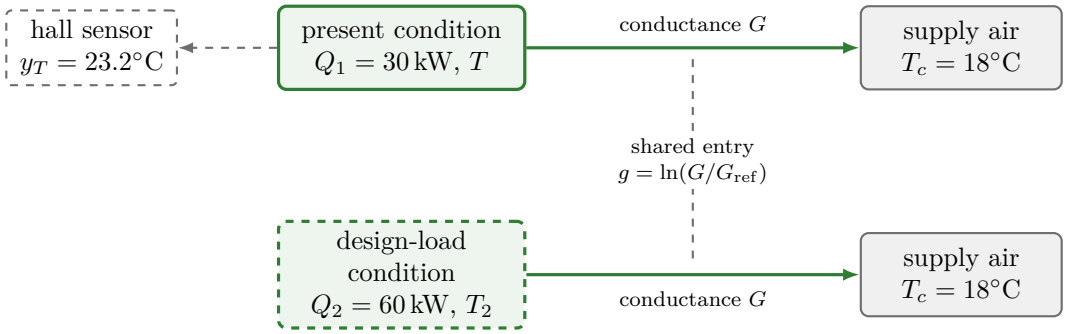


Figure 1.3: Two operating conditions of the equipment hall linked through the same physical conductance. Free: the present-condition temperature T and the shared log-conductance g, the two entries of the state; the design-load temperature T_2 (dashed node) is a function of g. Pinned: the supply-air temperature and the two loads. Observed: T only, through the reading y_T; the scenario node is unobserved and enters the graph because the question asks about it.

Which row acts on which entry. The three rows act on the two entries of the state in a pattern fixed by the graph, before any number is computed. The physics row ties the temperature to the conductance and involves both; the sensor sees the temperature and not the conductance; the handbook speaks about the conductance and not the temperature; the scenario node adds no row. The Jacobian of the rows, the matrix of their derivatives with respect to the entries, can therefore be nonzero only where marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * \\ * & 0 \\ 0 & * \end{bmatrix} \begin{matrix} \text{physics row} \\ \text{reading row,} \\ \text{prior row} \end{matrix} \qquad \text{columns } z_1 = T,\ z_2 = g. \tag{1.9}$$

The zeros are structural: they hold at every state. Only the physics row has entries in both columns, and it is the one row through which a reading of the temperature reaches the conductance.

The factor graph of the hall. The pattern (1.9) is what a factor graph (Section 1.1.3) draws. Figure 1.4 draws it for the hall with its scenario node. The two solid circles are the entries of the state, T and g. The black squares are the rows: the physics row of equation (1.5), linked to both entries; the reading row, linked to T alone; and the prior row, linked to g alone. The dashed circle is the design-load temperature T_2, joined to g by a dashed line: it is a function of g through equation (1.8), not an entry, and no row acts on it. The figure makes visible what the physical graph only implies: T_2 has no reading of its own and is connected to the evidence only through g, and g is connected to the reading only through the physics row. The physical drawing has two supply-air boxes and two paths; the factor graph has no node for the supply air at all, since T_c is pinned and appears only as a constant inside the physics row.

The residual vector. The state is the column vector of equation (1.3),

$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} T \\ g \end{bmatrix},$$

with T in °C and g dimensionless. Three rows act on it, the physics row (1.5), the reading row (1.6) and the prior row, each divided by its standard deviation. Stacked into one residual

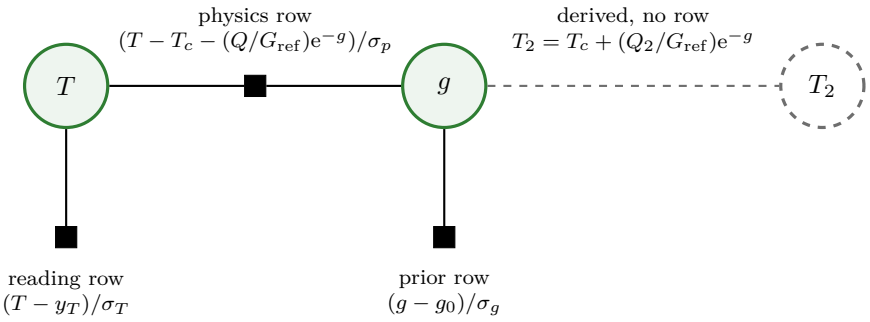


Figure 1.4: Factor graph of the equipment hall with its scenario node. Solid circles: the entries of the state, T and g. Squares: the rows of the model, each linked to the entries it involves, the pattern (1.9) of the Jacobian. Dashed: the design-load temperature T_2, a function of g with no row of its own, which reaches the evidence only through the shared g.

vector, the three rows are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} r_1(\mathbf{z}) \\ r_2(\mathbf{z}) \\ r_3(\mathbf{z}) \end{bmatrix} = \begin{bmatrix} (T - T_c - (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g})/\sigma_p \\ (T - y_T)/\sigma_T \\ (g - g_0)/\sigma_g \end{bmatrix}, \tag{1.10}$$

with $T_c = 18°\mathrm{C}$, $Q = 30\,\mathrm{kW}$, $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, the reading $y_T = 23.2°\mathrm{C}$ with $\sigma_T = 0.5\,\mathrm{K}$, and the handbook prior $g_0 = \ln(7.5/6) = 0.2231$ with $\sigma_g = 0.3$. Each entry of $\mathbf{r}$ is one row: a dimensionless number stating by how many standard deviations a candidate state violates one relation. The objective is half the sum of their squares,

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\,\mathbf{r}(\mathbf{z}) = \tfrac{1}{2}\sum_{k=1}^{3} r_k(\mathbf{z})^2, \tag{1.11}$$

which Chapter 3 shows to be minus the logarithm of the posterior up to a constant, so that the most probable state $\hat{\mathbf{z}}$, the maximum a posteriori (MAP) state, is the minimizer of $\mathcal{L}$. The computation takes $\sigma_p = 0.01\,\mathrm{K}$, fifty times tighter than the sensor, and Section 1.4.1 checks that the limit has been reached.

The Jacobian. The Jacobian $\mathbf{J}$ collects the derivative of every row with respect to every entry of the state: its entry in row k and column i is $\partial r_k/\partial z_i$, so it has one row per row of the model and one column per entry of the state, here three by two. Differentiating equation (1.10) entry by entry,

$$\mathbf{J}(\mathbf{z}) = \frac{\partial \mathbf{r}}{\partial \mathbf{z}} = \begin{bmatrix} 1/\sigma_p & (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}/\sigma_p \\ 1/\sigma_T & 0 \\ 0 & 1/\sigma_g \end{bmatrix}, \tag{1.12}$$

the pattern (1.9) with its entries filled in. Each entry is a sensitivity in standard deviations: how far one row moves when one entry of the state moves by one unit. Only the entry in the first row and second column depends on the state; the others are constants. The first row is the physics row's own sensitivities in its physical unit, rescaled. Differentiating the kilowatt form of equation (1.4),

$$\frac{\partial r_{\mathrm{phys}}}{\partial T} = G_{\mathrm{ref}}\,\mathrm{e}^{g} = G, \qquad \frac{\partial r_{\mathrm{phys}}}{\partial g} = G_{\mathrm{ref}}\,\mathrm{e}^{g}\,(T - T_c) = G\,(T - T_c),$$

in kW/K and kW; on the zero set the second equals the load Q by equation (1.1). A one-kelvin error in the temperature unbalances the hall by G kilowatts, and a unit change in g, a factor e

in the conductance, by about Q kilowatts to first order. Divided by G, to reach the temperature form, and by σ_p, the pair becomes $(1/\sigma_p,\ (Q/G)/\sigma_p)$, the first row of equation (1.12) on the zero set.

1.2 The posterior of the state (step 3)

All the questions of Table 1.2 are read from one posterior: the posterior of the state that the physics row, the reading row and the prior row of step 2 give, approximately the Laplace normal of (T, g) and exactly a density along the physics curve. This section derives it and computes it for the hall; Section 1.3 reads every answer from it.

The objects of the fit recur in the specialized forms, and they are named here by their roles. The *MAP state* $\hat{\mathbf{z}} = (\hat{T}, \hat{g})^{\mathsf{T}}$, the maximum a posteriori state, is the minimum of the objective that the three rows put on the state, the most probable state; the *information matrix* $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the curvature that the rows give that objective at the minimum, with $\mathbf{J}$ the 3×2 Jacobian of the normalized rows; the *Laplace covariance* $\mathbf{\Sigma}$, a 2×2 matrix, is the inverse of that curvature after a term from the physics row's own curvature is added to it, and it makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$ (bold, to tell it from the entry g), with the variance $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. Section 1.2.1 derives them from the rows. Beside them stands the exact posterior that this problem admits in the limit $\sigma_p \to 0$: on the physics curve the state has one free direction, and the posterior is the density of one entry, evaluated by quadrature. The Laplace route, a normal distribution centered on the MAP state, is what lets every answer be read from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ by operations that serve a plant of any size, and it misleads where the posterior is skewed along the curved ridge, which Section 1.4.2 measures against the exact posterior. The chapter reports the Laplace values, sets the exact value beside each answer with the gap, and collects the comparison in Section 1.4.

1.2.1 The posterior from the rows

The answers use the objects of one posterior: the MAP state $\hat{\mathbf{z}}$, the covariance $\mathbf{\Sigma}$ and the Jacobian $\mathbf{J}$ of the rows. The Jacobian is written out in Section 1.1.7; the other two are derived here for the hall. Chapter 3 derives these objects from the probability rules for Gaussian rows over any state, and Chapter 4 dissects the same objects in full on the hall with its scenario node carried as an entry, including the Laplace approximation and its curvature term. Here each is written out once for a reader meeting it for the first time.

The posterior. The ingredients are in hand: the unknown state $\mathbf{z} = (T, g)^{\mathsf{T}}$; the observation $y_T = 23.2°\mathrm{C}$; the physical relation $Q = G\,(T - T_c)$; and the prior information $p(g)$. What is ultimately wanted is

$$p(\mathbf{z} \mid \mathbf{y}), \tag{1.13}$$

the *posterior distribution* of the state, conditional on the evidence $\mathbf{y}$, here the reading y_T and the handbook value, and on the assumed physical model. It states which values of the state remain plausible after the physics, the reading and the prior information have been combined. Chapter 3 derives Bayes' theorem, by which the posterior is proportional to the likelihood times the prior, $p(\mathbf{z} \mid \mathbf{y}) \propto p(y_T \mid \mathbf{z})\, p(\mathbf{z})$; here the reading row plays the part of the likelihood, and the physics row and the prior row on g together play the part of the prior $p(\mathbf{z})$.

The normal equations and the Gauss–Newton step. Near a state $\mathbf{z}$, each row is replaced by its tangent, $\mathbf{r}(\mathbf{z} + \Delta\mathbf{z}) \approx \mathbf{r}(\mathbf{z}) + \mathbf{J}\,\Delta\mathbf{z}$, and the objective by $\frac{1}{2}\,\|\mathbf{r} + \mathbf{J}\,\Delta\mathbf{z}\|^2$, a quadratic in the

step $\Delta\mathbf{z}$. Its gradient with respect to the step is $\mathbf{J}^\mathsf{T}(\mathbf{r} + \mathbf{J}\,\Delta\mathbf{z})$, and setting it to zero gives the *normal equations*

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}, \tag{1.14}$$

a linear system of two equations in the two entries of the step. The Gauss–Newton method solves it, moves to $\mathbf{z} + \Delta\mathbf{z}$, evaluates $\mathbf{r}$ and $\mathbf{J}$ there, and repeats until the step no longer changes the state at the precision wanted. The vector $\mathbf{J}^\mathsf{T}\mathbf{r}$ is the gradient of the objective (1.11), so the iteration stops where that gradient vanishes, at the MAP state. The iteration starts here from the reference state $\mathbf{z}^0 = (23°\mathrm{C},\ 0)^\mathsf{T}$ of the forward calculation, at which the physics row is zero.

The information matrix. The matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$ of the normal equations is the *information matrix* of the rows. It is a sum over the rows: writing $\mathbf{J}_k$ for the k-th row of $\mathbf{J}$,

$$\mathbf{J}^\mathsf{T}\mathbf{J} = \sum_{k=1}^{3} \mathbf{J}_k^\mathsf{T}\mathbf{J}_k,$$

so that each row contributes the outer product of its own sensitivities, and the precisions of independent rows add. For the hall the three outer products are of two kinds. The reading row and the prior row each have one nonzero entry, so each contributes one diagonal entry, $1/\sigma_T^2$ in the T–T position and $1/\sigma_g^2$ in the g–g position. The physics row has two, and its outer product $\mathbf{J}_1^\mathsf{T}\mathbf{J}_1$ fills the whole matrix, off-diagonal entries included: it is the only source of the coupling between the two entries, and it has rank one, with all its information along its own row and none along the tangent of the curve. The eigenvectors of $\mathbf{J}^\mathsf{T}\mathbf{J}$ give the directions in which the rows fix the state, and its eigenvalues how sharply: a large one across the curve, set by $1/\sigma_p^2$, and a small one along it, set by the reading and the prior, as the ridge below makes exact.

The Laplace covariance and the curvature term. The inverse of the information matrix is the covariance of the posterior: exactly so for rows linear in the state (Chapter 3), and as the *Laplace approximation*, the normal distribution fitted at the MAP state, for curved rows (Chapter 4). The Gauss–Newton matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$ drops the curvature of the rows themselves. The full Hessian of the objective (1.11), its matrix of second derivatives, is

$$\frac{\partial^2 \mathcal{L}}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \mathbf{J}^\mathsf{T}\mathbf{J} + \sum_{k=1}^{3} r_k \frac{\partial^2 r_k}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}}, \tag{1.15}$$

and of the three rows only the physics row is curved, through e^{-g}, with the single nonzero second derivative

$$\frac{\partial^2 r_1}{\partial g^2} = -\frac{(Q/G_\mathrm{ref})\,\mathrm{e}^{-g}}{\sigma_p}.$$

The *curvature term* $r_1\,\partial^2 r_1/\partial g^2$ therefore adds to the g–g entry alone. It does not vanish in the exact-physics limit, because r_1 shrinks in proportion to σ_p while $\partial^2 r_1/\partial g^2$ grows as $1/\sigma_p$: at the MAP state the gradient $\mathbf{J}^\mathsf{T}\mathbf{r}$ is zero, its T entry reads $r_1/\sigma_p + r_2/\sigma_T = 0$, so $r_1 = -(\sigma_p/\sigma_T)\,r_2$, and

$$r_1 \frac{\partial^2 r_1}{\partial g^2} = \frac{r_2}{\sigma_T}\frac{Q}{G_\mathrm{ref}}\,\mathrm{e}^{-\hat{g}}$$

at every σ_p. It is the curvature of the physics curve itself, and it acts in the one direction, along the curve, in which the information is small. The *Laplace covariance* is the inverse of the full Hessian at the MAP state,

$$\boldsymbol{\Sigma} = \Big(\mathbf{J}^\mathsf{T}\mathbf{J} + \sum_k r_k \frac{\partial^2 r_k}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}}\Big)^{-1}_{\hat{\mathbf{z}}}, \qquad \mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma}); \tag{1.16}$$

a Newton iteration, which uses the full Hessian in place of $\mathbf{J}^\mathsf{T}\mathbf{J}$ in equation (1.14), reaches the same $\hat{\mathbf{z}}$.

The ridge. The spreads of the two entries are not independent of each other, and the reason is read off the physics row. Its Jacobian row, the derivative of equation (1.5) with respect to (T, g), and the vector it annihilates are

$$\mathbf{J}_1 = \frac{1}{\sigma_p}\begin{bmatrix} 1 & (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g} \end{bmatrix}, \qquad \mathbf{n} = \begin{bmatrix} (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g} \\ -1 \end{bmatrix}, \qquad \mathbf{J}_1\,\mathbf{n} = 0. \tag{1.17}$$

A step along $\mathbf{n}$ leaves the physics row unchanged to first order, so $\mathbf{n}$ is the tangent to the zero set of the row, a change of $(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}$ kelvin of temperature per unit of g with the opposite sign. The information the physics row contributes, $\mathbf{J}_1^{\mathsf{T}}\mathbf{J}_1$, is a matrix of rank one: it grows as $1/\sigma_p^2$ across the curve and is zero along $\mathbf{n}$. Along the curve only the reading row and the prior row inform the state, and by the pattern (1.9) each acts on one entry, with the contributions $\mathbf{J}_2^{\mathsf{T}}\mathbf{J}_2 = \mathrm{diag}(1/\sigma_T^2,\ 0)$ and $\mathbf{J}_3^{\mathsf{T}}\mathbf{J}_3 = \mathrm{diag}(0,\ 1/\sigma_g^2)$. With $n_1 = (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}}$ the first entry of $\mathbf{n}$ at the MAP state and r_2 the reading row there, the physics row's own curvature adds the term $\mathbf{C} = \mathrm{diag}(0,\ n_1 r_2/\sigma_T)$, derived above. With $\mathbf{u} = \mathbf{n}/\|\mathbf{n}\|$ the unit tangent, the information along the curve is

$$\lambda_{\|} = \mathbf{u}^{\mathsf{T}}(\mathbf{J}_2^{\mathsf{T}}\mathbf{J}_2 + \mathbf{J}_3^{\mathsf{T}}\mathbf{J}_3 + \mathbf{C})\mathbf{u} = \frac{n_1^2/\sigma_T^2 + 1/\sigma_g^2 + n_1 r_2/\sigma_T}{n_1^2 + 1}, \tag{1.18}$$

which does not depend on σ_p. The tangent $\mathbf{u}$ is orthogonal to $\mathbf{J}_1$, so in the orthonormal basis $(\mathbf{J}_1^{\mathsf{T}}/\|\mathbf{J}_1\|,\ \mathbf{u})$ the curvature of the objective has the entry $\lambda_{\|}$ along $\mathbf{u}$, an entry across the curve that grows as $1/\sigma_p^2$, and a coupling that stays bounded; its inverse, by the Schur complement of Chapter 3, leaves $1/\lambda_{\|}$ along $\mathbf{u}$ and nothing else as $\sigma_p \to 0$. The covariance therefore tends to the rank-one matrix $\boldsymbol{\Sigma} \to \mathbf{u}\,\mathbf{u}^{\mathsf{T}}/\lambda_{\|}$: the posterior collapses onto a segment of the tangent, a *ridge* of the objective. Then $\sqrt{\Sigma_{11}} \to |u_1|/\sqrt{\lambda_{\|}}$ and $\sqrt{\Sigma_{22}} \to |u_2|/\sqrt{\lambda_{\|}}$ are one spread along the curve seen through the two components of the tangent, and the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$ tends to the sign of $u_1 u_2$, which is -1. The two entries are not known separately: a hall slightly warmer than $\hat{T}$ is a hall with a slightly lower conductance, and the posterior allows the one only together with the other. The reading row has a zero in the column of g, and it narrows the conductance all the same, through the first component of $\mathbf{u}$.

> **Principle 1.2.** *Physical rows constrain the system; reading rows provide evidence about its unknown quantities. The posterior is where the two agree, weighted by the prior, and neither kind of row alone determines it. An exact physical row leaves the posterior only along its own zero set, and its signature in* $\boldsymbol{\Sigma}$ *is a correlation near* -1 *or* $+1$.

The exact posterior along the curve. In the limit $\sigma_p \to 0$ the posterior lies on the physics curve, and along the curve it can be computed without approximation. On the curve the temperature is a function of the conductance entry, $T(g) = T_c + (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}$. The rows define the posterior relative to a base flat in (T, g). Integrated across the curve, the physics row contributes a factor that does not depend on g, because the row is written in kelvin and its slope in T, $1/\sigma_p$, does not depend on g; the base is therefore flat in g along the curve, and what remains is the density of g along the curve, the reading row and the prior row evaluated on it:

$$p(g \mid \mathbf{y}) \propto \exp\left[-\frac{1}{2}\left(\frac{T(g) - y_T}{\sigma_T}\right)^2 - \frac{1}{2}\left(\frac{g - g_0}{\sigma_g}\right)^2\right]. \tag{1.19}$$

One-dimensional quadrature of this density over g gives every probability and moment of a quantity derived from the state exactly, for T through $T(g)$ and for G through $G_{\mathrm{ref}}\,\mathrm{e}^{g}$.

The answer read off. An answer that is an entry of the state is read from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ directly: the entry $\hat{z}_i$ of $\hat{\mathbf{z}}$, with the spread $\sqrt{\Sigma_{ii}}$. The relation between two entries is the correlation $\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$, and the directions of the posterior are the eigenvectors of $\boldsymbol{\Sigma}$. A derived quantity f is carried by its gradient, with the spread $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$, and a threshold on an entry is a normal tail of its marginal. Section 1.2.2 computes $\hat{\mathbf{z}}$, $\mathbf{J}$ and $\boldsymbol{\Sigma}$ for the hall.

1.2.2 The posterior of the hall

The hall's state estimation, the fit of the state to the three rows, has as inputs the pinned values $T_c = 18°\mathrm{C}$, $Q = Q_1 = 30\,\mathrm{kW}$ and $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$; the reading $y_T = 23.2°\mathrm{C}$ with $\sigma_T = 0.5\,\mathrm{K}$; the handbook prior $g_0 = 0.2231$ with $\sigma_g = 0.3$; and the tolerance $\sigma_p = 0.01\,\mathrm{K}$ of the physics row.

The first step. At the reference state $\mathbf{z}^0 = (23°\mathrm{C},\ 0)^{\mathsf{T}}$ the Jacobian (1.12) is

$$\mathbf{J}(\mathbf{z}^0) = \begin{bmatrix} 100 & 500 \\ 2 & 0 \\ 0 & 3.333 \end{bmatrix} :$$

a one-kelvin change in T moves the physics row by 100 and the reading row by 2, and a unit change in g moves the physics row by 500 and the prior row by 3.33. The first row is the pair of kilowatt sensitivities $(G,\ Q) = (6,\ 30)$ divided by $G = 6\,\mathrm{kW/K}$ and by σ_p. The residual vector there is $\mathbf{r} = (0,\ -0.400,\ -0.744)$: the physics holds, the temperature lies 0.4 standard deviations below the reading, and g lies 0.74 standard deviations below the handbook value. The two sides of the normal equations (1.14) are

$$\mathbf{J}(\mathbf{z}^0)^{\mathsf{T}}\mathbf{J}(\mathbf{z}^0) = \begin{bmatrix} 10004 & 50000 \\ 50000 & 250011 \end{bmatrix}, \qquad -\mathbf{J}(\mathbf{z}^0)^{\mathsf{T}}\mathbf{r}(\mathbf{z}^0) = \begin{bmatrix} 0.800 \\ 2.479 \end{bmatrix},$$

and the first step is their solution,

$$\Delta\mathbf{z} = \begin{bmatrix} 0.0684 \\ -0.0137 \end{bmatrix}.$$

The iterations. Table 1.4 lists the iterations. Each row is one iterate $\mathbf{z}^k$; the columns give its two entries, its three rows and the objective of equation (1.11).

Table 1.4: Gauss–Newton iterations for the posterior of the equipment hall (equations (1.10) to (1.14)), with $\sigma_p = 0.01\,\mathrm{K}$, from the reference state $(23°\mathrm{C},\ 0)$.

k	T [°C]	g [–]	r_1 [–]	r_2 [–]	r_3 [–]	$\mathcal{L}$ [–]
0	23.0000	0.0000	0.000	−0.400	−0.744	0.3566
1	23.0684	−0.0137	−0.042	−0.263	−0.789	0.3471
2	23.0701	−0.0139	0.005	−0.260	−0.790	0.3460
3	23.0701	−0.0139	0.005	−0.260	−0.790	0.3460

The first step does almost all the work, because over a step this short the rows are nearly linear. It leaves the physics row slightly violated, $r_1 = -0.042$, through the curvature of e^{-g} that the tangent ignores; the second step restores it, and the third changes nothing at the printed precision. The MAP state is

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{T} \\ \hat{g} \end{bmatrix} = \begin{bmatrix} 23.070°\mathrm{C} \\ -0.0139 \end{bmatrix}, \qquad \hat{G} = 6\,\mathrm{e}^{-0.0139} = 5.917\,\mathrm{kW/K}. \tag{1.20}$$

There the physics holds to $r_1\,\sigma_p = 5\times10^{-5}\,\mathrm{K}$, and the reading and the prior share the disagreement: the temperature lies 0.26 standard deviations below the reading and g lies 0.79 standard deviations below the handbook value.

The information matrix. At the MAP state the Jacobian (1.12) is

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 100 & 507.0 \\ 2 & 0 \\ 0 & 3.333 \end{bmatrix}, \qquad \frac{(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}}}{\sigma_p} = \frac{5\,\mathrm{e}^{0.0139}}{0.01} = 507.0,$$

and the information matrix is the sum of the three outer products,

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \underbrace{\begin{bmatrix} 10000 & 50701 \\ 50701 & 257056 \end{bmatrix}}_{\text{physics}} + \underbrace{\begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix}}_{\text{reading}} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 11.11 \end{bmatrix}}_{\text{prior}} = \begin{bmatrix} 10004 & 50701 \\ 50701 & 257067 \end{bmatrix}.$$

The physics term has rank one: all its information lies along its own row (100, 507.0), across the physics curve, and none along the curve; only the reading's 4 and the prior's 11.11 inform along it. The eigenvalues of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ are 267067 across the curve and 4.27 along it, with the eigenvector $(0.9811,\ -0.1935)^{\mathsf{T}}$ of the small one pointing along the curve; the first two terms of equation (1.18) give it directly, $(5.070^2\times4+11.11)/(5.070^2+1) = 4.27$. The Gauss–Newton covariance is the inverse,

$$\mathbf{\Sigma}_{\mathrm{GN}} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \begin{bmatrix} 0.2256 & -0.0445 \\ -0.0445 & 0.0088 \end{bmatrix},$$

with the standard deviations $\sqrt{0.2256} = 0.475\,\mathrm{K}$ and $\sqrt{0.0088} = 0.094$.

The curvature term and the Laplace covariance. The second derivative of the physics row at the MAP state is $\partial^2 r_1/\partial g^2 = -507.0$; multiplied by $r_1 = 0.0052$, it adds -2.63 to the g–g entry, $257067 \to 257064$, the same value as $(r_2/\sigma_T)(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}} = (-0.2598/0.5)\times5.0701$. Projected on the along-curve eigenvector $(0.9811,\ -0.1935)^{\mathsf{T}}$, the term is $-2.63\times0.1935^2 = -0.099$ against the along-curve information 4.27, a loss of 2.3 % in the one direction in which the information is small; it raises the variances by 2.4 % and the standard deviations by 1.2 %. The Laplace covariance (1.16) is

$$\mathbf{\Sigma} = \begin{bmatrix} 0.23096 & -0.045552 \\ -0.045552 & 0.0089881 \end{bmatrix}, \qquad \sqrt{\Sigma_{11}} = 0.4806\,\mathrm{K}, \quad \sqrt{\Sigma_{22}} = 0.09481. \tag{1.21}$$

These are the Laplace values used in the rest of the chapter, quoted as 0.481 K and 0.0948; Gauss–Newton alone gives 0.475 K and 0.094. The correlation of the two entries, computed from the unrounded $\mathbf{\Sigma}$, is

$$\frac{\Sigma_{12}}{\sqrt{\Sigma_{11}\,\Sigma_{22}}} = \frac{-0.045552}{0.48058\times0.094806} = -0.9998.$$

The ridge. The physics row of $\mathbf{J}(\hat{\mathbf{z}})$, (100, 507.0), is annihilated by $\mathbf{n} = (5.070,\ -1)^{\mathsf{T}}$, equation (1.17) at $\hat{g}$, and the long axis of $\mathbf{\Sigma}$ is $\mathbf{u} = (0.9811,\ -0.1935)^{\mathsf{T}}$, the same direction normalized. By equation (1.18),

$$\lambda_{\parallel} = \frac{5.070^2\times4 + 11.11 + 5.070\times(-0.260)/0.5}{5.070^2+1} = 4.168,$$

END OF THE PREVIEW

Capacity Planning and Design by Derivation

These are the opening pages. The complete book runs 343 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