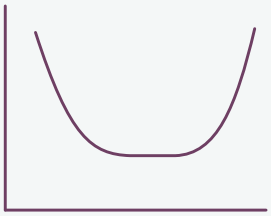
LESSON BOOK — VOLUME II



Reliability Engineering Volume II

Degradation and prognostics, power-system adequacy, observability and decisions

Sixteen lessons · 479 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Reliability Engineering

Volume II

Degradation and prognostics, power-system adequacy, observability and engineering decisions

A lesson sequence, written up as a book

Ricardo Marquez

October 14, 2026

How to read this book

This book is about one question: how an engineered system fails, how likely that is, what it costs, and what should be done about it. Each chapter expands one lesson of a sequence that builds the answer in the order the discipline itself uses,

function and failure → lifetime models → system logic → dependencies and risk → repair and availability → degradation and prognosis → power-system adequacy → measurements and credibility → decisions.

The sequence has nine modules and thirty-six lessons, published in two volumes. Volume I holds Parts I to V, Chapters 1 to 20: the classical reliability theory as the textbooks state it, from failure definitions and lifetime models through system logic, probabilistic risk assessment and repairable systems. This second volume holds Parts VI to IX, Chapters 21 to 36, which extend that theory to degradation and prognostics, power-system adequacy, the sufficiency of measurements, and engineering decisions. Chapter numbers run continuously across the two volumes; a reference to a chapter numbered 20 or below points to Volume I, whose results this volume builds on and does not repeat. The notation chapter is reproduced here so that this volume can be read with Volume I closed.

The book is written for a reader who builds or uses physics-based models of engineered systems and wants the concepts and the vocabulary of reliability engineering stated precisely. Every new quantity is therefore introduced in the same order: its name, its meaning, its units, its equation, and how it is estimated. Where the field uses a shorthand term, a *Jargon* box expands it once and the text then uses the expanded meaning. The vocabulary is the standard one of the discipline, as fixed by Rausand, Barros and Høyland's *System Reliability Theory*, Billinton and Allan's *Reliability Evaluation of Power Systems*, Abur and Expósito's *Power System State Estimation*, the prognostics and health management literature and the other references each chapter names; no term is coined where a standard one exists.

Three systems recur throughout, with growing complexity as the methods allow: an industrial cooling-water pump, a data center cooling plant with three chillers of which two are required, and an electric transmission network whose line ratings depend on the weather. The numbers in the examples are illustrative, chosen to make each relation concrete, and are labeled as such.

One distinction recurs in every chapter: the physical system, which fails or performs; the model, which represents it under stated assumptions; and the evidence, which is never complete. Three kinds of uncertainty are kept apart throughout: aleatory variability in which unit fails and when, epistemic uncertainty about the parameters of the model, and uncertainty about whether the model's structure is adequate at all.

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 four exercises that can be solved from the chapter alone.

Contents

How to read this book **ii**

Notation and vocabulary **vii**

VI Prognostics and Health Management **1**

21 Degradation and damage accumulation **2**
21.1 Failure as a threshold crossing of a degradation state 2
21.2 Degradation mechanisms and deterministic models 5
21.3 Fatigue, crack growth and thermal aging 9
21.4 Stochastic degradation processes 15
21.5 Condition-based reliability 21
21.6 Operating history, decisions and credibility 25
21.7 What to retain 29
21.8 Exercises 31

22 Condition monitoring and fault diagnosis **33**
22.1 Detection, diagnosis and prognosis 33
22.2 Physical models and condition indicators 36
22.3 Signal processing for condition monitoring 39
22.4 Bayesian diagnosis with competing hypotheses 44
22.5 Sensor faults, observability and the next measurement 50
22.6 Validation and application of diagnosis 56
22.7 What to retain 59
22.8 Exercises 61

23 Remaining useful life estimation **62**
23.1 Remaining useful life as a conditional random variable 62
23.2 Wiener-process remaining useful life 67
23.3 State and parameter uncertainty 69
23.4 Gamma processes and particle filters 75
23.5 Calibration and validation of RUL predictions 80
23.6 Decisions and credibility 82
23.7 What to retain 86
23.8 Exercises 87

24 Predictive maintenance and PHM decisions **89**
24.1 Predictive maintenance as a decision problem 89
24.2 Probability-based maintenance triggers 92
24.3 The cooling-pump replacement decision 94

24.4 Inspection value and scheduling . . . 97
24.5 Maintenance logistics and system availability . . . 103
24.6 Sequential decisions under uncertainty . . . 106
24.7 What to retain . . . 111
24.8 Exercises . . . 113

VII Power-System Reliability and Adequacy 115

25 Generation adequacy and reliability indices 116
25.1 Generation adequacy and the adequacy event . . . 116
25.2 Forced outages and the COPT . . . 119
25.3 Annual adequacy and uncertain demand . . . 125
25.4 Dependent outages and the construction of the COPT . . . 130
25.5 Adequacy, security and sensitivity . . . 134
25.6 Limits and interpretation of adequacy indices . . . 138
25.7 What to retain . . . 140
25.8 Exercises . . . 141

26 State enumeration and capacity outage models 143
26.1 Multistate generating units . . . 143
26.2 Exact enumeration of a multistate fleet . . . 147
26.3 COPT recursion and generating functions . . . 151
26.4 Correlated generator outages . . . 156
26.5 Monte Carlo adequacy simulation . . . 160
26.6 Choice of method and validation . . . 166
26.7 What to retain . . . 169
26.8 Exercises . . . 171

27 Transmission reliability and contingency analysis 173
27.1 Contingencies, security and transmission reliability . . . 173
27.2 The network model and single-outage analysis . . . 176
27.3 Distribution factors and AC security . . . 181
27.4 Probabilistic contingency analysis . . . 184
27.5 Uncertain ratings and corrective actions . . . 189
27.6 Risk attribution and the limits of the analysis . . . 195
27.7 What to retain . . . 200
27.8 Exercises . . . 202

28 Time-varying reliability and cascading failures 204
28.1 Cascades and static contingency analysis . . . 204
28.2 Time-dependent and weather-correlated failures . . . 207
28.3 Conductor thermal dynamics . . . 210
28.4 Protective controls and the cascade event tree . . . 213
28.5 Branching models and chronological simulation . . . 220
28.6 Risk-based decisions and model credibility . . . 225
28.7 What to retain . . . 228
28.8 Exercises . . . 229

VIII State Estimation, Observability, and Model Credibility 231

29 Observability and measurement sufficiency 232
29.1 States, measurements and observability . . . 232
29.2 Observability of the three-bus network . . . 235
29.3 Uncertainty, redundancy and bad data . . . 238
29.4 Sensor placement and measurement sufficiency . . . 245
29.5 Nonlinear, dynamic and parameter observability . . . 247
29.6 Observability in reliability decisions . . . 253
29.7 What to retain . . . 256
29.8 Exercises . . . 257

30 State estimation and bad-data detection 259
30.1 State estimation and its two outputs . . . 259
30.2 Weighted least-squares estimation . . . 262
30.3 Residual analysis and bad-data detection . . . 266
30.4 Critical measurements and topology errors . . . 272
30.5 Robust state estimation . . . 276
30.6 Bayesian state estimation and reliability decisions . . . 279
30.7 What to retain . . . 284
30.8 Exercises . . . 286

31 Identifiability, uncertainty and model adequacy 288
31.1 Parameter estimation and identifiability . . . 288
31.2 Sensitivity matrices and excitation . . . 293
31.3 Fisher information and Bayesian estimation . . . 296
31.4 Model adequacy and model discrepancy . . . 302
31.5 Bayesian model comparison . . . 305
31.6 Applications and the modeling workflow . . . 309
31.7 What to retain . . . 312
31.8 Exercises . . . 314

32 Bayesian updating and value of information 315
32.1 Bayesian updating of competing hypotheses . . . 315
32.2 Entropy and expected information gain . . . 319
32.3 Value of information for decisions . . . 322
32.4 The pump replacement decision . . . 324
32.5 Measurements for parameters and models . . . 329
32.6 Sequential decisions and model credibility . . . 336
32.7 What to retain . . . 342
32.8 Exercises . . . 344

IX Engineering Applications and Decisions 346

33 Reliability standards and engineering practice 347
33.1 Engineering assessments and the role of standards . . . 347
33.2 Transmission planning standards . . . 350
33.3 Functional safety and asset management . . . 353
33.4 Requirements and a two-pump assessment . . . 359
33.5 Verification, validation and traceability . . . 366
33.6 Risk registers, uncertainty and acceptance . . . 369

33.7 What to retain . . . 375
33.8 Exercises . . . 376

34 Data center power and cooling reliability 378
34.1 Data center service reliability . . . 378
34.2 Electrical infrastructure reliability . . . 381
34.3 Cooling infrastructure reliability . . . 384
34.4 Coupled failures and the integrated case . . . 387
34.5 Thermal ride-through and the coupled event tree . . . 392
34.6 Importance, decisions and a unified model . . . 397
34.7 What to retain . . . 402
34.8 Exercises . . . 403

35 Risk importance and decision optimization 405
35.1 Component importance and intervention value . . . 405
35.2 Risk worth measures and parameter sensitivity . . . 409
35.3 Counterfactual risk reduction . . . 413
35.4 An investment portfolio . . . 418
35.5 Value of an upgrade and value of information . . . 424
35.6 Uncertainty and the decision framework . . . 427
35.7 What to retain . . . 431
35.8 Exercises . . . 432

36 Integrated reliability assessment 435
36.1 The integrated assessment . . . 435
36.2 Physical model and measurement audit . . . 438
36.3 From diagnosis to risk and action . . . 443
36.4 The next measurement . . . 448
36.5 Credibility and the assessment report . . . 453
36.6 Synthesis of the two volumes . . . 458
36.7 What to retain . . . 461
36.8 Exercises . . . 462

Notation and vocabulary

Every symbol below is used with the same meaning in every chapter. Time is measured in operating hours unless a chapter states otherwise.

Symbol	Unit	Meaning
T	h	Random time to failure of a nonrepairable item, measured from a defined starting condition; $T \geq 0$.
t	h	A specified operating duration, or the running time variable.
H	h	Operating horizon or mission time over which a requirement must hold.
$R(t)$	–	Reliability (survivor) function, $\Pr(T > t)$.
$F(t)$	–	Cumulative failure probability, $\Pr(T \leq t) = 1 - R(t)$.
$f(t)$	h^{-1}	Failure density, $\mathrm{d}F/\mathrm{d}t$.
$\lambda(t)$	h^{-1}	Hazard function (instantaneous failure rate), $f(t)/R(t)$.
$\Lambda(t)$	–	Cumulative hazard, $\int_0^t \lambda(u)\,\mathrm{d}u$, so that $R(t) = \exp[-\Lambda(t)]$.
λ	h^{-1}	Constant failure rate of the exponential model.
μ	h^{-1}	Constant repair rate, 1/MTTR.
$R(t + \Delta t \mid t)$	–	Conditional reliability: survival over a further Δt given survival to t, $R(t + \Delta t)/R(t)$.
MTTF, MTBF, MTTR	h	Mean time to failure (nonrepairable, or operating duration between restorations); mean time between failures (repairable, including restoration); mean time to repair or restore.
MRL(t)	h	Mean residual life at age t, $\mathbb{E}[T - t \mid T > t]$.
A, $A(t)$, A_{mission}	–	Steady-state availability MTTF/(MTTF + MTTR); instantaneous (point) availability at time t; point availability at the end of a mission of length H, $A(H)$, for an item that starts in the operating state.
R_{system}, R_i	–	Reliability of the system; of component i.
k-out-of-n	–	A system of n units that functions when at least k of them function.
$\mathcal{R}$	R$/yr	Risk, $\sum_i \Pr(S_i)\, C(S_i)$ over a partition of outcomes S_i; expected annual loss when C is monetary.
$C(F)$	R$	Consequence of failure event F, conditional on its occurrence.
$X(t)$	varies	Uncertain system state or inputs at time t (load, weather, equipment condition).
$g(x)$	varies	Performance-margin function; $g \geq 0$ means the requirement is met.

Symbol	Unit	Meaning
$L(t)$, $C(t)$	varies	Required load or demand; available capacity, in the units of the function (m^3/h, kW, A).
$C_{\text{available}}(t)$	varies	Available capacity of a set of units, $\sum_i S_i(t)\, C_i(t)$.
$S_i(t)$	–	Status of unit i at time t: 1 available, 0 unavailable.
$T_{\max}$	°C	Upper temperature limit of a conductor, a room or a component.
I, $R(T)$	A, Ω	Current; temperature-dependent resistance of a conductor (never the reliability function, which is R).
mc_p	J/K	Heat capacity of the modeled conductor length.
q_{solar}, q_{conv}, q_{rad}	W	Solar heat gain; convective and radiative heat loss of a conductor.
S, O, D	–	FMEA severity, occurrence and detection ratings on a 1–10 scale.
RPN	–	Risk priority number, $S \cdot O \cdot D$.
β, η	–, h	Shape and scale parameters of the Weibull distribution.
σ	varies	Standard deviation; stress amplitude, by context.
N	–	Number of components, of load cycles, of trials, by context.
Statistical estimation (Chapters 6–8)		
$\hat{R}(t)$, $\hat{\lambda}$	–, h^{-1}	Estimates from data (Kaplan–Meier survival estimate; maximum likelihood estimate of the failure rate).
r, T_{tot}	–, h	Number of observed failures; total exposure (time on test) of all units, censored ones included.
$L(\theta)$, $\ell(\theta)$, $\mathcal{I}(\theta)$	varies	Likelihood, log-likelihood and Fisher information of a parameter θ.
Gamma(a, b)	–	Gamma distribution with shape a and rate b; prior and posterior of a failure rate.
μ, σ (lognormal)	–	Mean and standard deviation of $\ln T$ in the lognormal model; in that use μ is not the repair rate.
$\Gamma(\cdot)$, $\Phi(\cdot)$	–	Gamma function; standard normal cumulative distribution function.
System logic and probabilistic risk assessment (Chapters 9–15)		
$\phi(\mathbf{x})$, x_i	–	Structure function and component state indicators (1 operating, 0 failed; or the failure indicator, as stated).
q_i	–	Probability of basic event i (component failure or unavailability).
Q_{top}, Q_{sys}	–	Probability of the top event; system failure probability or unavailability.
C_j, P_j	–	Minimal cut set j; minimal path set j (chapter text says when C and P name components instead).
I_i^{B}, I_i^{FV}	–	Birnbaum and Fussell–Vesely importance of component i.
RAW_i, RRW_i	–	Risk achievement worth; risk reduction worth.
β (beta factor), α_k	–	Common-cause beta factor (fraction of a component's failure rate that is common cause; not the Weibull shape); alpha factor, the fraction of failure events involving exactly k components.

Symbol	Unit	Meaning
λ_T, λ_I, λ_C	h^{-1}	Total, independent and common-cause failure rates of one component in a common-cause group.
f_{IE}, f_{S_k}	yr^{-1}	Initiating-event frequency; frequency of accident sequence k.
$\Pr(S_k \mid I)$	–	Probability of sequence k given the initiating event.
LOLE, EUE	h/yr, MWh/yr	Loss-of-load expectation; expected unserved energy.
Repairable systems and availability (Chapters 17–20)		
$N(t)$, T_j, X_j	–, h, h	Number of failures of a repairable item by time t; the j-th failure epoch; the j-th interfailure time.
$W(t)$, $w(t)$	–, h^{-1}	Cumulative intensity (mean function) $\mathbb{E}[N(t)]$ of a counting process; its rate of occurrence of failures (ROCOF), $\mathrm{d}W/\mathrm{d}t$, which is not the hazard of one lifetime.
$M(t)$, $m(t)$	–, h^{-1}	Renewal function (mean number of renewals) and renewal density of a perfect-repair process.
b, η (power law)	–, h	Shape and scale of the power-law intensity $w(t) = (b/\eta)(t/\eta)^{b-1}$ of a nonhomogeneous Poisson process.
V_n	h	Virtual age of an item after its n-th repair (Kijima models).
$P_i(t)$, $\boldsymbol{\pi}$	–	Probability of being in state i at time t of a Markov model; its stationary distribution.
$\mathbf{Q}$, q_{ij}	h^{-1}	Generator (transition-rate) matrix of a continuous-time Markov chain and its off-diagonal rates; not the top-event probability Q_{top} nor the basic-event probability q_i.
ρ	–	Ratio of failure rate to repair rate, λ/μ; by context, a discount rate.
f_{F}, MDT	yr^{-1}, h	Frequency of entering the failed state; mean down time per failure, so that unavailability equals $f_{\mathrm{F}} \cdot \mathrm{MDT}$.
T_p, $C(T_p)$	h, R$/h	Preventive replacement age; long-run cost rate of the age-replacement policy.
c_p, c_f, c_D	R$, R$, R$/h	Cost of a preventive replacement, of a failure replacement, of one hour of downtime.
s, $\nu\tau$	–, –	Spare-parts stock level; expected spare demand over a lead time τ at demand rate ν.
Degradation, prognostics and health management (Chapters 21–24)		
$D(t)$, D_0, D_{crit}	varies	Degradation (damage) state, its initial value and the failure threshold; failure is the first crossing. Not the FMEA detection rating D.
μ_D, σ_D, $B(t)$	varies	Drift and diffusion of the Wiener degradation model $D(t) = D_0 + \mu_D t + \sigma_D B(t)$; standard Brownian motion. μ_D is not the repair rate.
α_D, β_D	–, varies	Shape rate and scale of a gamma degradation process.
T_{f}, $\mathrm{RUL}(t)$, $\mathrm{RUL}_\alpha(t)$	h, d	First-passage (failure) time; remaining useful life at time t, $T_{\mathrm{f}} - t$ given survival; its α-quantile.

Symbol	Unit	Meaning
e, z	varies, –	Residual, observed minus model-predicted value; standardized residual.
$H_0, H_1, \ldots$	–	Diagnostic hypotheses (healthy, each fault, each sensor fault) with posterior probabilities $\Pr(H_i \mid \mathbf{y})$.
$\mathcal{H}$	bit	Shannon entropy of the hypothesis probabilities; expected information gain of a test.
Y^+, Y^-	–	Positive and negative outcome of an inspection, characterized by sensitivity and specificity.
τ_{L}, p^*	d, –	Action lead time; decision threshold on the failure probability.
EVSI, EVPI	R$	Expected value of sample (test) information; of perfect information.
Power-system adequacy (Chapters 25 onward)		
L_t, C_t, S_t	MW	Hourly demand; available capacity; shortfall $(L_t - C_t)^+$.
FOR_i, U_i	–	Forced outage rate of unit i: its unavailability, a probability, not a rate.
LOLP_t, LOLH	–, h/yr	Loss-of-load probability in interval t; loss-of-load hours, the LOLE on an hourly basis.
X, $P(X \geq x)$	MW, –	Capacity on outage and its cumulative probability in a capacity outage probability table (COPT).
K_i, $\mathrm{PTDF}_{\ell,k}$, $\mathrm{LODF}_{\ell,m}$	MW, –, –	Rating of unit i; power transfer and line outage distribution factors of line ℓ.
θ_i, V_i, x_{ij}, P_{ij}, $P_{ij}^{\max}$	rad, pu, pu, MW	Bus voltage angle and magnitude; line reactance; DC line flow $(\theta_i - \theta_j)/x_{ij}$; line rating.
$\mathbf{B}$	pu	DC susceptance matrix of the network.
λ_{b}, N_{c}	–, –	Branching ratio (mean offspring number) of a cascade model; cascade size.
τ_{th}	min	Thermal time constant of a conductor.
State estimation and identifiability (Chapters 29–31)		
$\mathbf{x}$, $\mathbf{y}$	varies	State vector; measurement vector (written $\mathbf{z}$ in the power-system literature), $\mathbf{y} = h(\mathbf{x}) + \boldsymbol{\varepsilon}$.
$\mathbf{R}$, $\mathbf{W}$	varies	Measurement error covariance (bold; not the reliability function) and its inverse, the weight matrix.
$\mathbf{H}$, $\mathbf{G}$	varies	Measurement Jacobian $\partial h/\partial \mathbf{x}$ (bold; not the horizon) and the gain (information) matrix $\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}$.
$J(\mathbf{x})$, $\hat{\mathbf{x}}$	–, varies	Weighted least-squares objective; its minimizer, the state estimate.
$\mathbf{e}$, $\mathbf{S}$, $\boldsymbol{\Omega}$, e_i^{N}	varies	Residual $\mathbf{y} - h(\hat{\mathbf{x}})$; residual sensitivity matrix; residual covariance $\mathbf{SR}$; normalized residual $e_i/\sqrt{\Omega_{ii}}$.
$\boldsymbol{\theta}$, $\mathbf{S}_\theta$	varies	Parameter vector (bold; not the bus angles) and its sensitivity matrix $\partial h/\partial \boldsymbol{\theta}$.
$\sigma_1 \geq \sigma_2 \geq \ldots$, κ	–	Singular values of the scaled sensitivity matrix; condition number $\sigma_1/\sigma_{\min}$.
$\delta(\cdot)$, B_{12}	varies, –	Model discrepancy term; Bayes factor between models 1 and 2.
Value of information and engineering decisions (Chapters 32–36)		
$L(a, H)$, $\rho(a)$	R$	Loss of action a when hypothesis H holds; its expected loss under the current probabilities.

Symbol	Unit	Meaning
C_{now}, C_{sample}, C_{perfect}	R$	Minimum expected loss with current information; with the outcome of a test; with the true hypothesis known.
EIG, $I(H;Y)$, D_{KL}	bit or nat	Expected information gain of a test, equal to the mutual information between hypothesis and outcome; Kullback–Leibler divergence.
$\mathcal{R}(a)$, $\Delta\mathcal{R}(a)$	varies	Risk after intervention a; its reduction from the baseline $\mathcal{R}_0$.
x_i, c_i, B	–, R$/yr, R$/yr	Selection variable of project i (1 selected, 0 not); its annualized cost; the budget.
v, H_{rel}	R$/MWh, h/yr	Consequence value of unserved energy; hours per year for which an analyzed operating condition applies.
CVaR_α	varies	Conditional value at risk: the expected loss in the upper $(1-\alpha)$ tail.
A_{IT}, $Y(t)$	–	Availability of the required IT service; its indicator at time t.
C_{eff}, t_{crit}	MJ/K, min	Effective thermal capacitance of a room; time for its temperature to reach the limit after a loss of cooling.
λ_{DU}, T_{p}	h^{-1}, h	Dangerous undetected failure rate of a safety function; proof-test interval (not the replacement age of Chapter 18).
PFD_{avg}, PFH	–, h^{-1}	Average probability of dangerous failure on demand; average frequency of dangerous failure per hour.

Vocabulary

The book uses the standard vocabulary of reliability engineering and probabilistic risk assessment, as fixed by the textbooks named in the preface. Where a term has a common shorthand, the shorthand is given in parentheses once and the full term is used thereafter.

Item A component, subsystem or system treated as a unit of analysis.

Required function What an item must do for the system to count as operating; stated as a measurable performance criterion.

Failure The termination of an item's ability to perform its required function; a functional event, which need not coincide with physical breakage.

Fault The state of an item after failure, or a defect that may lead to failure; a fault can exist without a failure having occurred.

Failure mode The manner in which the failure is observed (no flow, low flow, leak, spurious trip).

Failure mechanism The physical, chemical or electrical process that produces the failure (fatigue, corrosion, wear, electromigration).

Failure cause The circumstance that activates the mechanism (overload, contamination, misuse).

Failure effect The consequence of the mode at the next level of the system hierarchy.

System boundary The set of items and interfaces inside the reliability model; what is outside is an input or an assumption.

Reliability Probability that an item performs its required function under specified conditions for a specified period.

Availability Probability that an item is in an operating state at a given time, or the long-run fraction of time it is operating, restoration included.

Maintainability Probability that a failed item is restored within a specified time under specified procedures.

Risk The combination of the probability of outcomes and their consequences.

RAMS Reliability, availability, maintainability and safety, considered together.

Hazard function (failure rate) The conditional failure intensity at age t given survival to t; constant for the exponential model, increasing with wear-out, decreasing with early failures.

Bathtub curve The hazard shape with an early decreasing part, a constant middle and an increasing wear-out part.

Mission reliability Reliability over a stated mission time without credit for repair.

Repairable system An item that is restored to operation after failure; described by counting and renewal processes rather than a single lifetime.

FMEA, FMECA Failure mode and effects analysis; the same with a criticality analysis.

Fault tree analysis (FTA) A deductive logic model of the combinations of events that produce a defined top event.

Minimal cut set A smallest set of component failures sufficient to cause the top event.

Common-cause failure Failure of several items from one shared cause, which defeats the independence that redundancy assumes.

Importance measure A quantity that ranks components by their contribution to system unreliability (Birnbaum, Fussell–Vesely, risk achievement worth, risk reduction worth).

Probabilistic risk assessment (PRA) The scenario-based quantification of risk: initiating events, event trees, fault trees and consequences.

Aleatory and epistemic uncertainty Variability inherent in the process, and lack of knowledge about its parameters or structure.

Censoring A lifetime observation that ends before failure, so that only a lower bound of T is known (right-censoring); interval and left censoring bound T from both sides or from above.

Kaplan–Meier estimator The nonparametric product-limit estimate of the reliability function from censored data.

Maximum likelihood estimate (MLE) The parameter value that maximizes the likelihood of the observed failures and censored records.

Credible interval A Bayesian interval that contains the parameter with stated posterior probability; distinct from a confidence interval.

Posterior predictive reliability The reliability of a new unit obtained by averaging the model over the posterior of its parameters.

Structure function The Boolean function that gives the system state from the component states.

Reliability block diagram A success-logic diagram in which the system operates when a path of operating blocks connects input to output.

Cut set, path set A set of component failures that causes system failure; a set of operating components that keeps the system operating. Minimal when no proper subset qualifies.

Rare-event approximation The sum of minimal-cut-set probabilities as an estimate of the top-event probability; an upper bound that is accurate when the probabilities are small.

Inclusion–exclusion The exact expansion of the probability of a union of events in alternating sums over intersections.

Coherent system A system in which no component failure can restore operation and every component matters.

Common-cause component group A set of redundant, similar components exposed to the same common causes; the unit to which beta and alpha factors apply.

Coupling factor The shared property (design, environment, procedure) through which one root cause affects several components.

Initiating event The event that starts an accident sequence; characterized by a frequency per year.

Event tree An inductive model that branches on the success or failure of each protection function after an initiating event.

Accident sequence, end state One path through an event tree; the consequence category in which it ends.

Frequency–consequence curve The annual frequency of exceeding each consequence level; the risk profile beyond its expectation.

Integrated PRA A risk assessment that links initiating-event frequencies, fault trees, event trees, consequences and uncertainty into one traceable model; organized by the risk triplet (what can go wrong, how likely, with what consequences).

Counting process, ROCOF The number of failures of a repairable item by time t; the rate of occurrence of failures, the derivative of its mean, as distinct from the hazard of one lifetime.

Homogeneous and nonhomogeneous Poisson process (HPP, NHPP) Counting processes with a constant, respectively time-varying, intensity; the NHPP with a power-law intensity is the Crow–AMSAA model.

Perfect, minimal and imperfect repair Repair that returns an item to as-good-as-new, to as-bad-as-old, or to a state between the two (virtual age).

Renewal process The counting process of a perfectly repaired item: independent identically distributed interfailure times.

Age replacement The maintenance policy that replaces an item at failure or at age T_p, whichever comes first; optimized through the renewal-reward theorem.

Continuous-time Markov chain (CTMC) A state model with exponential holding times and a generator matrix of transition rates; solved by the Kolmogorov forward equations or its stationary distribution.

Frequency and duration The description of unavailability as the product of the failure frequency and the mean down time.

Degradation, failure threshold The gradual loss of a property that fixes an item's ability to function; failure is the first crossing of the threshold by the degradation state.

Prognostics and health management (PHM) Detection, diagnosis and prognosis of equipment condition from measurements, and the maintenance decisions built on them.

Remaining useful life (RUL) The time from now until the degradation state reaches its threshold, a random variable conditional on the present evidence.

Condition indicator, residual A quantity computed from measurements that tracks condition; the difference between a measurement and its model prediction, which detection and diagnosis test.

Fault detection and isolation (FDI) Deciding that a fault is present, and which one; with bad-data detection, the discipline's names for separating physical faults from sensor faults.

Value of information The expected reduction in decision loss from a test or inspection before deciding; sample information (EVSI) and perfect information (EVPI).

Generation (resource) adequacy Whether installed resources suffice to meet demand at every hour under random outages; measured by LOLP, LOLE and EUE.

Forced outage rate (FOR) The unavailability of a generating unit, the fraction of time it is out of service on forced outage; a probability.

Capacity outage probability table (COPT) The distribution of capacity on outage for a fleet, built by recursive convolution of the units' two-state models.

Multistate unit, derated state A generating unit modeled with partial-capacity states between full availability and outage.

Contingency, N−1 criterion The outage of one or more network elements; the requirement that the system withstand the loss of any single element without violating limits.

Cascading failure A sequence of dependent outages in which each failure raises the stress on, and the failure probability of, the remaining elements.

Observability Whether the states of interest can be inferred from the available measurements; the rank condition on the measurement Jacobian.

State estimation The weighted least-squares (or robust, or Bayesian) inference of the system state from redundant measurements.

Critical measurement A measurement whose removal makes the system unobservable; its residual is always zero, so its error cannot be detected.

Bad-data detection and identification The chi-square test on the weighted residual sum, and the largest normalized residual test that names the suspect measurement.

Topology error A mismatch between the network model's switching state and the physical network, which a residual test can confuse with a sensor fault.

Identifiability Whether the parameters of a model can be determined from the measurements, structurally (from the model) and practically (given the noise and the design).

Model discrepancy The systematic difference between a model's prediction and the physical system, distinct from measurement error.

Expected information gain The expected reduction in the entropy of the hypothesis probabilities from a test; a measure of knowledge, not of decision value.

Counterfactual risk reduction The change in risk from a hypothetical intervention on the system, as distinct from conditioning on an observed state.

Portfolio optimization The selection of projects under a budget, with system risk recomputed for each combination because project benefits interact.

Verification, validation Whether a model is implemented and solved correctly; whether it represents the physical system adequately for its intended use.

Traceability The documented connection from a requirement through the model, its inputs and evidence to the result and the decision.

Safety integrity level (SIL) The band of required failure measures for a safety function, from 1 to 4.

Thermal ride-through The time a room's thermal inertia allows between a loss of cooling and the crossing of the temperature limit.

Part VI

Prognostics and Health Management

Chapter 21

Degradation and Damage Accumulation

21.1 Failure as a threshold crossing of a degradation state

Chapters 17 to 20 of Volume I represented equipment as operating, failed or under repair. They modeled failure and restoration rates, availability and the optimization of maintenance, and in all of them a failure was an event whose rate was estimated from the history of earlier failures. Those methods leave one question unanswered: what happens physically inside a component before it fails. A bearing gradually loses material; the paper insulation of a transformer accumulates thermal damage; a battery loses capacity; a welded joint develops a fatigue crack that grows with every load cycle. Such processes often evolve over months or years before they reach a critical condition, and if their evolution can be modeled, the failure can be predicted, the remaining life of the component estimated, and maintenance planned on the condition of the equipment rather than on its age. This chapter, which opens the part of the book on prognostics and health management, asks

whether equipment failure can be predicted from the evolving physical condition of the equipment, rather than only from historical failure frequencies.

The answer is a change of the object that is modeled. Instead of the failure event, the model describes the physical process that produces it: a degradation state that evolves under the operating conditions and a failure threshold that it eventually crosses. The chapter first states this framing; it then describes the physical mechanisms and the deterministic models of degradation, works the physics-based models of fatigue, crack growth and thermal aging, introduces the stochastic degradation processes whose first-passage times are lifetime distributions, conditions the reliability of one particular component on its measured condition, and ends with operating history, engineering decisions and the credibility of degradation models.

Jargon: degradation, degradation state, failure threshold, first-passage time, damage accumulation, PHM, condition monitoring, diagnosis, prognostics, remaining useful life, condition-based maintenance

Degradation is a progressive change in the physical or functional condition of an item that reduces its ability to perform its required function. The *degradation state* $D(t)$ is a measurable or inferred quantity that describes that condition at time t: a crack length, a wall penetration, a wear depth, a normalized damage index or a fraction of capacity lost. The *failure threshold* D_{crit} is the value of the degradation state at which the item no longer meets its required function. The *first-passage time* T_{f}, also called the threshold-crossing time, is the first time at which $D(t)$ reaches D_{crit}. *Damage accumulation* is the growth of a degradation state that does not recover, such as fatigue damage or metal loss. *Prognostics and health management* (PHM) is the discipline that monitors the condition of equipment, diagnoses its faults, predicts its future degradation and failure, and uses the prediction for maintenance and operating decisions. Its three activities have standard names: *condition monitoring*, the observation of indicators of condition; *diagnosis*, the determination of the mechanism responsible for an observed change; and *prognostics*, the prediction of future degradation and of the *remaining useful life* (RUL), the time from the present until the item can no longer perform its required function. *Condition-based maintenance* is maintenance whose timing is decided from the observed condition of the item rather than from its age or a fixed calendar.

Key idea

A failure mechanism that accumulates damage turns the time to failure into the first time a degradation state crosses a threshold. Modeling the state, its dependence on the operating history and its random variation gives a distribution of the time to failure for each individual component, conditioned on what has been measured about it.

21.1.1 From failure events to degradation trajectories

In the lifetime models of Chapter 5 of Volume I, the time to failure of an item is a random variable and the reliability function is the probability that it exceeds a stated duration,

$$R(t) = \Pr(T_{\text{f}} > t), \tag{21.1}$$

estimated from the failure times of a population of similar items. Nothing in that description refers to the physical condition of the item between its installation and its failure. A degradation model introduces that condition explicitly as a degradation state $D(t)$, measured directly or inferred from measurements. Examples are the length of a fatigue crack, the wear of a bearing, the depth of corrosion penetration in a pipe wall and the cumulative damage of an insulation. (D here is the degradation state of Chapter 4 of Volume I; it is not the FMEA detection rating D of that chapter.) The item fails when the degradation reaches a critical threshold, so that the time to failure is the first time at which the state reaches it:

$$T_{\text{f}} = \inf\{\, t \geq 0 : D(t) \geq D_{\text{crit}} \,\}. \tag{21.2}$$

The infimum is the greatest lower bound of the set of times at which the state is at or above the threshold, which for a continuous path is the first such time. T_{f} is called a *first-passage time* or *threshold-crossing time*. When failure is defined by the threshold, T_{f} is the random time to failure T of Volume I, and equation (21.1) gives its reliability function.

Chapter 4 stated the deterministic version of this model in Section 4.3.3: with a damage rate $\mathrm{d}D/\mathrm{d}t = k\, s(t)^n$, in which k is a degradation coefficient, $s(t)$ the intensity of the stress and n the stress sensitivity exponent, a constant stress gives the time to failure $(D_{\text{crit}} - D_0)/(ks^n)$,

and the cumulative failure probability over a horizon H is $\Pr(D(H) \geq D_{\text{crit}})$ when the damage never decreases (equation (4.11)). This chapter extends that model in three directions: to physical laws for particular mechanisms, to degradation that is random along its path, and to conditioning on measurements of one particular item.

21.1.2 Failure as a threshold crossing

Figure 21.1 draws the framing. The horizontal axis is operating time and the vertical axis the degradation state. The blue curve is one degradation path that starts at an initial value D_0 and increases, slowly at first and faster later; the dashed red line is the critical threshold D_{crit}; the dot marks the first crossing and the dashed vertical line below it the failure time T_{f}. The arrow at an intermediate time t is the remaining distance to the threshold, $D_{\text{crit}} - D(t)$, which a later section uses to predict the remaining life. Failure occurs where the path meets the threshold, not at a predefined age of the component: a component whose path rises faster fails earlier at the same age.

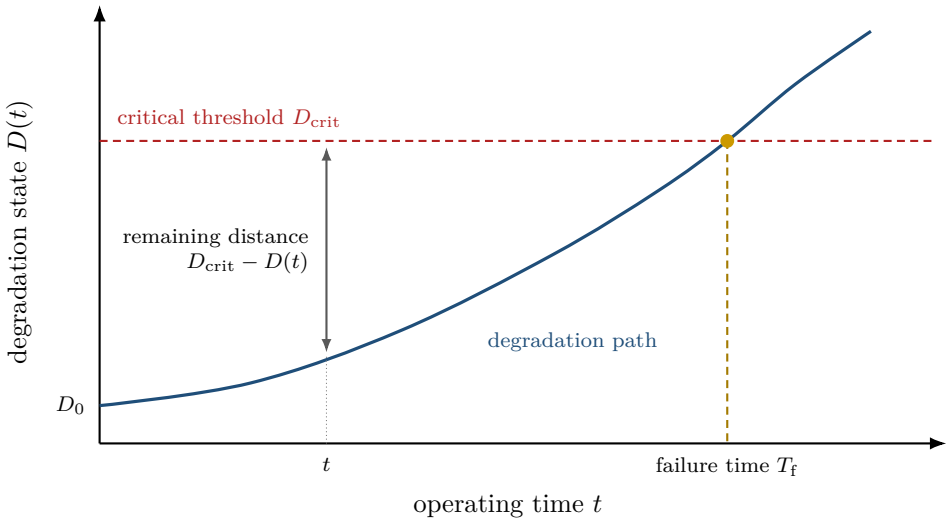


Figure 21.1: An illustrative increasing degradation path. Failure occurs at the first crossing of the critical threshold D_{crit}, the first-passage time T_{f} of equation (21.2), not at a predefined component age. The double arrow is the remaining distance to the threshold at an intermediate time t.

This changes the reliability problem. Instead of modeling the failure events directly, the model describes the physical process that produces them, in the chain

$$\text{operating conditions} \;\rightarrow\; \text{damage evolution} \;\rightarrow\; \text{threshold crossing} \;\rightarrow\; \text{failure}.$$

The first arrow is a physical law that turns loads, temperatures and cycles into a rate of damage; the second is the accumulation of that rate up to the threshold; the third is the definition of failure against the required function. The chain is the stress-strength model of Section 4.3.1 with a strength that falls over time. There, an item with strength C under a load L fails when the margin $g = C - L$ is negative, with probability $p_f = \Pr(L > C)$; here, degradation lowers the capacity of the item until the margin is consumed, and the threshold D_{crit} is the degradation state at which the margin reaches zero. Each link of the chain can be modeled deterministically, which gives a predicted failure time, or stochastically, which gives a distribution of failure times. The sections that follow develop both.

Principle 21.1. *Under a degradation model, the time to failure is the first-passage time of the degradation state through the failure threshold. The reliability function is the probability that the whole degradation path stays below the threshold up to the stated time.*

21.2 Degradation mechanisms and deterministic degradation models

This section states what degradation is relative to a required function, lists the physical mechanisms that produce it with their degradation states, and develops the deterministic models in which the state follows an ordinary differential equation: linear, nonlinear and dependent on the operating conditions.

21.2.1 Degradation and the required function

Degradation is a progressive change in a component's physical or functional condition. It becomes relevant to reliability when it reduces the component's ability to satisfy its required function, in the sense of Chapter 2 of Volume I. Much degradation happens while the component is still operating. A battery remains in service while it loses capacity, and it fails only when its available energy or power falls below what its application requires. A transmission-line conductor accumulates fatigue damage at its suspension clamps while it continues to carry its load. A cooling pump continues to deliver flow while the wear of its bearings and its vibration increase, and it fails, for the cooling-water system of Chapter 4 of Volume I, when it can no longer deliver the required $100\,\mathrm{m}^3/\mathrm{h}$ or when the bearing seizes.

The degradation variable must therefore connect to a meaningful failure criterion. A crack length matters because a critical length exists at which the remaining section fractures; a capacity fraction matters because the application sets a minimum energy; a vibration velocity matters only through the bearing condition it indicates and the seizure or flow loss that condition leads to. A degradation state without a threshold tied to the required function is a monitored quantity, not yet a reliability model.

21.2.2 Common physical mechanisms

The mechanisms that produce degradation were surveyed in Chapter 4; five of them recur in prognostics. *Fatigue*: repeated stress cycles initiate and propagate cracks, in rotating shafts, turbine components, transmission structures and mechanical connections. *Corrosion*: chemical or electrochemical reactions gradually remove material or degrade protective layers, and eventually threaten pressure containment or structural strength. *Wear*: contact, lubrication conditions, particle contamination and repeated loading damage surfaces and change friction, vibration and clearances. *Thermal and electrical aging*: temperature and electrical stress accelerate the deterioration of insulation, polymers and other materials, in transformers, motors and power electronics. *Electrochemical degradation*: batteries lose usable capacity and their internal resistance grows, through calendar aging, cycling, temperature and the operating conditions. Table 21.1 lists, for each mechanism, a degradation state that describes it, the stresses that drive it, a failure criterion and the model of this chapter that represents it.

The mechanisms differ, but many of them can be described by one mathematical structure, a first-order evolution law for the degradation state,

$$\frac{\mathrm{d}D}{\mathrm{d}t} = g_D(D, \mathbf{u}(t), \mathbf{z}(t), \boldsymbol{\theta}), \tag{21.3}$$

where $D(t)$ is the degradation state, $\mathbf{u}(t)$ the operating inputs and loading (speed, current, flow, cycles), $\mathbf{z}(t)$ the environmental conditions (temperature, humidity, water chemistry), $\boldsymbol{\theta}$ the

Table 21.1: Degradation mechanisms, their degradation states, driving stresses, failure criteria and the models of this chapter.

Mechanism	Degradation state $D(t)$	Driving stresses	Failure criterion	Model in this chapter
Fatigue	Miner damage sum; crack length a	stress range, number of cycles	$D = 1$; $a = a_c$	S-N curve and Miner rule; Paris law
Corrosion	wall penetration (mm)	water chemistry, temperature, flow, coating	maximum acceptable penetration	linear degradation; gamma process
Wear	bearing damage index; clearance	contact load, speed, lubrication, particles	seizure; flow below requirement	Wiener process; speed-dependent rate
Thermal and electrical aging	insulation damage; equivalent aging time	temperature, electric stress	end of insulation life	Arrhenius rate, cumulative thermal damage
Electro-chemical	capacity loss $1 - Q_{\mathrm{cap}}/Q_{\mathrm{cap}}(0)$; resistance growth	calendar time, cycles, temperature, state of charge	80 % capacity retention; power limit	calendar and cycling model

physical parameters of the model, and g_D the degradation evolution law, with the units of D per unit time. (g_D is not the performance margin g of Volume I.) The law should be supported by a physical mechanism or by adequate empirical evidence; equation (4.8) of Chapter 4 is the special case $g_D = k\, s(t)^n$, which depends on the stress alone.

Not every observed change in a sensor signal is evidence of progressive damage. The vibration of a pump may increase because its operating speed changed, not because its bearing deteriorated. The degradation state is a property of the component; a measurement responds to the state and to the operating conditions together. The distinction is taken up in Section 21.5.2 and, in full, in the treatment of fault diagnosis in Chapter 22.

21.2.3 Linear degradation

The simplest degradation model assumes a constant rate of damage accumulation,

$$\frac{\mathrm{d}D}{\mathrm{d}t} = \kappa, \qquad \kappa > 0, \tag{21.4}$$

where κ is the degradation rate in the units of D per unit time. With the initial condition $D(0) = D_0$, integration gives

$$D(t) = D_0 + \kappa t. \tag{21.5}$$

Failure occurs when $D(t) = D_{\mathrm{crit}}$, so $D_0 + \kappa\, T_{\mathrm{f}} = D_{\mathrm{crit}}$, and solving for the failure time,

$$T_{\mathrm{f}} = \frac{D_{\mathrm{crit}} - D_0}{\kappa}. \tag{21.6}$$

Example 21.1 (Corrosion penetration)**.** A pipe wall corrodes at a penetration rate $\kappa = 0.08\,\mathrm{mm/yr}$. The penetration measured at the latest inspection is $D_0 = 0.60\,\mathrm{mm}$, and the maximum acceptable penetration, set by the wall thickness needed to contain the operating pressure, is $D_{\mathrm{crit}} = 2.0\,\mathrm{mm}$ (illustrative values). Then

$$T_{\mathrm{f}} = \frac{2.0\,\mathrm{mm} - 0.60\,\mathrm{mm}}{0.08\,\mathrm{mm/yr}} = \frac{1.40}{0.08}\,\mathrm{yr} = 17.5\,\mathrm{yr},$$

counted from the inspection. The deterministic model predicts another 17.5 years until the threshold is reached.

That prediction is not yet a probabilistic reliability estimate. It assumes that the corrosion rate and the threshold are known exactly and remain constant, whereas the rate of corrosion depends on the water chemistry, the temperature, the flow velocity, the condition of protective coatings and other factors that vary along the pipe and over the years. To compute a failure probability, the uncertain quantities and the random variation of the process must be represented, which Section 21.4 does for this same pipe with a gamma process.

21.2.4 Nonlinear degradation

Degradation often accelerates as damage accumulates, because the damage itself raises the rate: a longer crack concentrates more stress at its tip, a worn bearing vibrates more and wears faster. The simplest model of such acceleration makes the rate proportional to the state,

$$\frac{\mathrm{d}D}{\mathrm{d}t} = \kappa D, \tag{21.7}$$

with κ now in units of inverse time. Separating the variables, $\mathrm{d}D/D = \kappa\,\mathrm{d}t$, and integrating from D_0 at time zero, $\ln(D/D_0) = \kappa t$, so

$$D(t) = D_0\,\mathrm{e}^{\kappa t}. \tag{21.8}$$

Failure occurs when $D_0\,\mathrm{e}^{\kappa T_\mathrm{f}} = D_\mathrm{crit}$, that is, at

$$T_\mathrm{f} = \frac{1}{\kappa}\,\ln\!\left(\frac{D_\mathrm{crit}}{D_0}\right). \tag{21.9}$$

The model requires $D_0 > 0$: a state that starts at zero never grows under equation (21.7), which is a reminder that accelerating models describe the growth of damage that already exists, not its initiation.

A more general law raises the state to a damage exponent n_D,

$$\frac{\mathrm{d}D}{\mathrm{d}t} = \kappa D^{n_D}. \tag{21.10}$$

For $n_D > 1$ the rate increases strongly with damage. Separating the variables, $D^{-n_D}\,\mathrm{d}D = \kappa\,\mathrm{d}t$, and integrating,

$$\frac{D_0^{1-n_D} - D(t)^{1-n_D}}{n_D - 1} = \kappa t, \qquad n_D \neq 1,$$

so that the failure time is

$$T_\mathrm{f} = \frac{D_0^{1-n_D} - D_\mathrm{crit}^{1-n_D}}{(n_D - 1)\,\kappa}. \tag{21.11}$$

For $n_D > 1$ the solution $D(t) = [D_0^{1-n_D} - (n_D - 1)\kappa t]^{-1/(n_D-1)}$ grows without bound as t approaches $D_0^{1-n_D}/[(n_D - 1)\kappa]$, a finite time; in a physical system the threshold, or a different mechanism such as unstable fracture, is reached first. Behavior of this kind appears in some crack-growth and wear processes, although the actual laws are often more complicated: the Paris law of Section 21.3.4 is equation (21.10) with the crack length as the state and $n_D = m/2$.

Figure 21.2 compares the three laws on one example. The horizontal axis is operating time in hours and the vertical axis a normalized damage. All three paths start at $D_0 = 0.10$ and reach the threshold $D_\mathrm{crit} = 1$ at the same time, 900 h (the example of Section 4.3.3), which fixes their rates: $\kappa = 0.001\,\mathrm{h}^{-1}$ for the linear law, $\kappa = \ln(10)/900\,\mathrm{h} = 2.558 \times 10^{-3}\,\mathrm{h}^{-1}$ for the exponential law, and, from equation (21.11) with $n_D = 2$, $\kappa = (10 - 1)/900 = 0.01\,\mathrm{h}^{-1}$ for the quadratic law. The dotted horizontal line marks half the distance to the threshold, $D = 0.55$. The linear path reaches it at 450 h, the exponential path at 666 h and the quadratic path at 818 h. Models that agree on the failure time disagree strongly on the state along the way: under

the quadratic law, a component that shows less than a fifth of the critical damage at 450 h (the path is at 0.18) has only half of its life left, and the last half of the damage accumulates in the last 9 % of the life. An inspection program designed under the linear law would inspect too rarely for an accelerating mechanism.

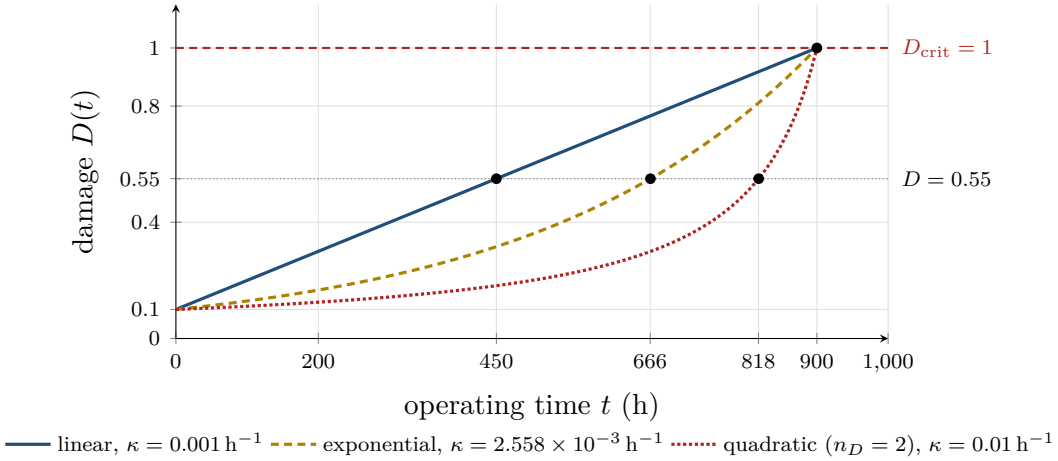


Figure 21.2: Linear, exponential and quadratic degradation laws, $D_0 + \kappa t$, $D_0 \mathrm{e}^{\kappa t}$ and $[D_0^{-1} - \kappa t]^{-1}$, with $D_0 = 0.10$ and rates chosen so that all three reach $D_{\mathrm{crit}} = 1$ at 900 h (illustrative). The marked points are the times at which each path reaches half the distance to the threshold: 450, 666 and 818 h.

Whether an accelerating law is physically appropriate depends on the mechanism. Uniform corrosion of a pipe wall at a steady water chemistry proceeds at a roughly constant rate; a fatigue crack accelerates; wear can do either, depending on whether the worn surface degrades the lubrication. The choice of g_D is a statement about the mechanism and has to be justified as such.

21.2.5 Damage that depends on operating conditions

A constant degradation rate may be reasonable under stable operating conditions, but industrial systems experience changing load and environment. A more useful model makes the rate a function of the conditions,

$$\frac{\mathrm{d}D}{\mathrm{d}t} = g_D\big(\sigma(t), T(t), I(t), \ldots\big), \tag{21.12}$$

where $\sigma(t)$ is the mechanical stress, $T(t)$ the absolute temperature and $I(t)$ the electric current. (In this chapter T with an argument or a subscript other than f is a temperature in kelvin; the failure time is always T_{f}.) Writing $\mathbf{z}(u)$ for the vector of conditions at time u and integrating from the initial state, the accumulated damage is

$$D(t) = D_0 + \int_0^t g_D\big(\mathbf{z}(u)\big)\,\mathrm{d}u. \tag{21.13}$$

The damage at time t depends on the whole history of the conditions, not on the elapsed time alone. Two nominally identical components can therefore age differently. One cooling pump may operate near its maximum load continuously while another runs only intermittently; their chronological ages are equal, but their accumulated physical damage can differ by a large factor. Section 21.6.1 returns to this point with a quantitative example.

Chapter 5 of Volume I reached a similar conclusion for the hazard function. There, Section 5.6.3 wrote a condition-dependent hazard $\lambda(t \mid \mathbf{z}(t)) = \lambda_0(t) \exp[\boldsymbol{\gamma}^{\mathsf{T}} \mathbf{z}(t)]$ and obtained the reliability by integrating it along the history of the conditions, equation (5.22). The two models

answer different questions. A condition-dependent hazard changes the probability of a failure event without describing the state of the component; a condition-dependent degradation law changes the state, and the probability of failure follows from where the state is relative to its threshold. The second can use a measurement of the state; the first cannot.

21.3 Physics-based damage models: fatigue, crack growth and thermal aging

Three mechanisms have physical damage laws developed far enough to predict failure from the loading history: fatigue, through the stress-life curve and the Miner rule; fatigue crack growth, through the Paris law; and thermal aging, through the Arrhenius rate. This section restates each from Chapter 4 of Volume I, applies it to a new example, and shows how each becomes a degradation state that can be updated as the operating history is recorded.

21.3.1 Fatigue as damage driven by load cycles

Fatigue damage arises from repeated loading, and it differs from ordinary wear in two respects. A metal component can fail under cyclic stresses that are each below its static strength: a transmission tower connection may withstand a particular wind-induced load once, but repeated cycles of the same load can initiate a crack and grow it. And the relevant exposure is often not elapsed time but the number of loading cycles N and their amplitudes, so that a fatigue model needs information about the amplitude and the number of stress cycles, not only about the operating hours.

An S-N curve relates the stress amplitude to the fatigue life. One common empirical relation is Basquin's law,

$$\sigma_a = \sigma'_f \, (2N_f)^{b_{\mathrm{B}}}, \tag{21.14}$$

where σ_a is the stress amplitude (half the stress range of a cycle), σ'_f the fatigue strength coefficient in the units of stress, N_f the number of cycles to failure, and $b_{\mathrm{B}} < 0$ the fatigue strength exponent (written b in fatigue texts; the subscript separates it from the S-N exponent b below). The factor $2N_f$ is the number of stress reversals, two per cycle; some formulations use cycles directly and absorb the factor into the fitted coefficient. Solving equation (21.14) for the life, $2N_f = (\sigma_a/\sigma'_f)^{1/b_{\mathrm{B}}}$, gives

$$N_f = \tfrac{1}{2}\, \sigma_f'^{\,b}\, \sigma_a^{-b}, \qquad b = -\frac{1}{b_{\mathrm{B}}} > 0,$$

which is the simpler engineering form used in Section 4.4.1,

$$N(\Delta\sigma) = K\, \Delta\sigma^{-b}, \tag{21.15}$$

with $K > 0$ and $b > 0$, written with the stress range $\Delta\sigma = 2\sigma_a$ of a fully reversed cycle and the factor 2^b absorbed into K. Basquin exponents of smooth metal specimens are typically between -0.05 and -0.12, which gives b between about 8 and 20; welded details, whose life is mostly crack growth, have b near 3, as Section 4.4.2 showed. The important result is that a higher cyclic stress produces fewer cycles to failure within the regime in which the curve applies.

21.3.2 The Miner rule and a cyclically loaded connection

Real components experience many different stress amplitudes. The Miner rule of Section 4.4.1, the standard linear damage accumulation rule, assigns to each cycle at stress level i the fraction

$1/N_i$ of the life at that level, where N_i is the constant-amplitude fatigue life. After n_i cycles at each of the levels, the accumulated damage is

$$D = \sum_i \frac{n_i}{N_i}, \tag{21.16}$$

and nominal fatigue failure is assumed when D reaches $D_{\text{crit}} = 1$. Each fraction n_i/N_i is the fraction of the nominal fatigue life consumed at load level i, and the sum is the degradation state of a fatigue model.

Example 21.2 (A cyclically loaded structural connection)**.** A structural connection has experienced, over its service so far, the stress cycles of Table 21.2, counted in three classes of stress range. For each class, testing or an appropriate design fatigue curve gives the number of cycles to failure at constant amplitude (illustrative values). The table reads row by row: the second column is the number of cycles n_i experienced in the class, the third the constant-amplitude life N_i, and the fourth the damage contribution n_i/N_i. The accumulated damage is

$$D = \frac{100\,000}{1\,000\,000} + \frac{40\,000}{200\,000} + \frac{10\,000}{50\,000} = 0.100 + 0.200 + 0.200 = 0.50.$$

Under the nominal Miner model, the connection has consumed 50 % of its critical damage. If the future load spectrum repeats the historical spectrum, another identical block of cycles adds $\Delta D = 0.50$, and the nominal prediction is failure at the end of one additional block.

Table 21.2: Miner damage of a cyclically loaded structural connection (illustrative values).

Stress range	Cycles experienced n_i	Cycles to failure N_i	Damage n_i/N_i
50 MPa	100 000	1 000 000	0.100
75 MPa	40 000	200 000	0.200
100 MPa	10 000	50 000	0.200
Total			0.500

The three lives of the table do not lie on one straight line of equation (21.15): between 50 and 75 MPa they imply an exponent $b = \ln 5/\ln 1.5 = 3.97$, and between 75 and 100 MPa an exponent $\ln 4/\ln(4/3) = 4.82$. A design fatigue curve may legitimately bend in this way; the Miner sum needs only the life at each level, not a single power law. Unlike the welded joint of Section 4.4.1, where the many small cycles dominated the damage, here the two higher classes contribute equally and together 80 % of it.

The rule has an important limitation: it does not account for the order of the loading cycles. High-amplitude loading early in life can affect crack initiation and growth differently from the same loading late in life, and the rule does not represent residual stresses, crack closure, corrosion fatigue or an uncertain threshold. The value $D = 0.50$ should therefore not be read as an experimentally verified statement that 50 % of the useful life remains. It is a model-dependent damage index, and measured damage sums at failure scatter widely around one.

21.3.3 Fatigue under time-varying operating conditions

In practice the stress cycles are not given in classes; they are extracted from a recorded or simulated stress history by a cycle-counting method, of which rainflow counting is the standard one. Rainflow counting pairs each local stress reversal with the reversal that closes a hysteresis loop, so that an irregular history becomes a list of complete cycles, each with a range and a mean. The damage calculation is then the chain

$$\text{stress history } \sigma(t) \;\rightarrow\; \text{stress cycle identification} \;\rightarrow\; \{n_i, \Delta\sigma_i\} \;\rightarrow\; D.$$

Figure 21.3 draws the chain for a wind-exposed structure such as a transmission tower. The boxes in the main row are the stages of the calculation, from left to right: the stress history at the critical detail, the cycle counting, the histogram of cycles by stress range, the S-N curve with the Miner sum, and the damage with the remaining life it implies. The note above the first box lists the inputs from which the stress history is computed: the mean wind speed, the turbulence intensity, the wind direction, the natural frequencies of the structure and its aerodynamic response. The note above the fourth box is the material fatigue property, the S-N curve of the detail. The dashed arrow below the row is the update: each new record of stress, or of wind from which stress is computed, adds its increment of damage to the state.

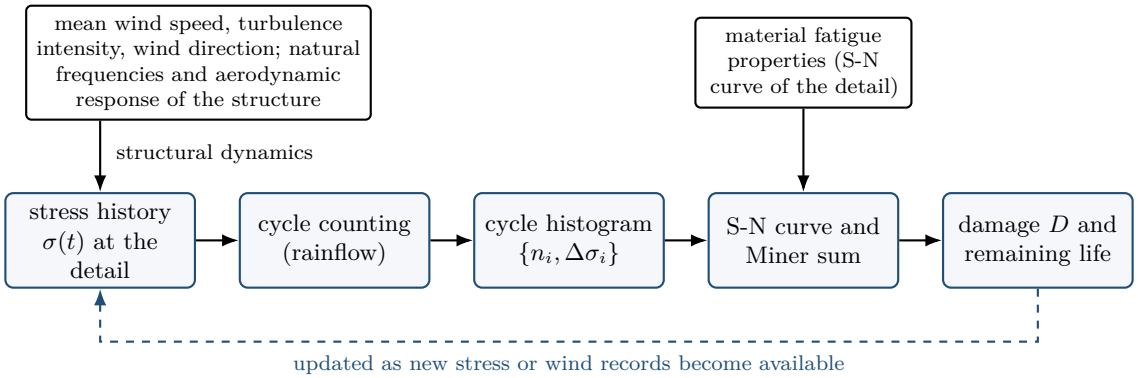


Figure 21.3: Fatigue damage of a wind-exposed structure computed from its operating history: the stress history at the critical detail is reduced to counted cycles, and the cycles to a Miner damage, which is updated with every new record.

Because the damage is a sum, it is updated by adding the increment of each new record. For the connection of Example 21.2, suppose that the next block of service, a windier period than the first, brings 30 % more cycles in every class (illustrative). Its damage increment is $1.3 \times 0.50 = 0.65$ instead of 0.50, and the threshold $D = 1$ is reached when $0.50/0.65 = 77\,\%$ of that block has elapsed, not at its end. This is a direct connection between meteorological modeling, structural dynamics and reliability analysis: the weather that sets the dynamic ratings of a transmission line also sets the fatigue damage of its structures.

21.3.4 Crack growth and the cycles to a critical length

The Miner rule summarizes damage as an accumulated scalar. When a crack is already present, its growth can be modeled directly. Let $a(N)$ be the crack length after N cycles. The Paris law of Section 4.4.2 states that the growth per cycle is a power of the stress-intensity-factor range,

$$\frac{\mathrm{d}a}{\mathrm{d}N} = C_{\mathrm{P}}\,(\Delta K)^m, \qquad \Delta K = Y\,\Delta\sigma\,\sqrt{\pi a}, \tag{21.17}$$

where C_{P} and m are material and environmental parameters determined by experiment, ΔK is the stress-intensity-factor range in $\mathrm{MPa}\sqrt{\mathrm{m}}$, Y a dimensionless geometry correction factor and $\Delta\sigma$ the stress range. Substituting, $\mathrm{d}a/\mathrm{d}N = C_{\mathrm{P}}\,(Y\Delta\sigma\sqrt{\pi a})^m$: the growth rate depends on the current crack length, and as the crack grows the stress-intensity factor increases and the crack accelerates. The law is a physical explanation of nonlinear degradation.

Suppose that Y and $\Delta\sigma$ remain constant over the interval analyzed, and define

$$K_0 = Y\,\Delta\sigma\,\sqrt{\pi}, \tag{21.18}$$

in MPa. Then $\Delta K = K_0\sqrt{a}$ and

$$\frac{\mathrm{d}a}{\mathrm{d}N} = C_{\mathrm{P}}\, K_0^m\, a^{m/2},$$

which is the power law of equation (21.10) with the crack length as the state, cycles in place of time and $n_D = m/2$. Rearranging, $\mathrm{d}N = \mathrm{d}a/(C_{\mathrm{P}} K_0^m a^{m/2})$, and integrating from the present crack length a_0 to the critical length a_c gives the remaining number of cycles,

$$N_{\mathrm{rem}} = \int_{a_0}^{a_c} \frac{\mathrm{d}a}{C_{\mathrm{P}}\,(K_0\sqrt{a})^m}. \tag{21.19}$$

For $m \neq 2$ the antiderivative of $a^{-m/2}$ is $a^{1-m/2}/(1-m/2)$, so

$$N_{\mathrm{rem}} = \frac{a_c^{\,1-m/2} - a_0^{\,1-m/2}}{C_{\mathrm{P}}\, K_0^m\,(1-m/2)}, \qquad m \neq 2, \tag{21.20}$$

which is equation (4.20) of Chapter 4 with numerator and denominator multiplied by -1. For $m = 2$ the integrand is $1/(C_{\mathrm{P}} K_0^2 a)$ and

$$N_{\mathrm{rem}} = \frac{\ln(a_c/a_0)}{C_{\mathrm{P}}\, K_0^2}, \qquad m = 2. \tag{21.21}$$

The expression connects the current crack length, the applied cyclic stress and the remaining fatigue life. With the illustrative steel parameters of Section 4.4.2, $C_{\mathrm{P}} = 10^{-11}$ meters per cycle per $(\mathrm{MPa}\sqrt{\mathrm{m}})^3$, $m = 3$, $Y = 1.12$, $\Delta\sigma = 100\,\mathrm{MPa}$ and $a_c = 10\,\mathrm{mm}$,

$$K_0 = 1.12 \times 100 \times 1.7725 = 198.5\,\mathrm{MPa}, \qquad C_{\mathrm{P}} K_0^3 = 7.823 \times 10^{-5}\,\mathrm{m}^{-1/2} \text{ per cycle},$$

and with $1 - m/2 = -\tfrac{1}{2}$,

$$N_{\mathrm{rem}}(a_0 = 1\,\mathrm{mm}) = \frac{(0.010)^{-1/2} - (0.001)^{-1/2}}{7.823 \times 10^{-5} \times (-0.5)} = \frac{10.00 - 31.62}{-3.912 \times 10^{-5}} = 5.53 \times 10^5 \text{ cycles},$$

the value of Chapter 4.

21.3.5 Crack length as a predictor of remaining life

The crack-growth model can outperform a model based on age alone. Two identical components of equal calendar age may have different crack lengths, for instance $a_0 = 1\,\mathrm{mm}$ and $a_0 = 4\,\mathrm{mm}$. For the second, $(0.004)^{-1/2} = 15.81\,\mathrm{m}^{-1/2}$ and

$$N_{\mathrm{rem}}(a_0 = 4\,\mathrm{mm}) = \frac{15.81 - 10.00}{0.5 \times 7.823 \times 10^{-5}} = 1.49 \times 10^5 \text{ cycles},$$

3.72 times fewer than for the first, although an age-only model gives both the same reliability. A degradation model recognizes the distinction because its state carries the history that the age discards.

Figure 21.4 shows the remaining life as a function of the present crack length. The horizontal axis is the crack length a_0 found at an inspection, in millimeters, and the vertical axis, on a logarithmic scale, the remaining number of cycles to the critical length of 10 mm, in units of 10^5, from equation (21.20). The three curves are three stress ranges, and the two marked points are the components of the example. Each curve falls by more than two orders of magnitude between 0.5 and 9.5 mm, and in absolute terms it falls fastest at small crack lengths: a crack of 2 mm at 100 MPa leaves 3.16×10^5 cycles, 57 % of the life of a 1 mm crack. The curves for different stress ranges are parallel because the remaining life scales as $\Delta\sigma^{-m} = \Delta\sigma^{-3}$ at every

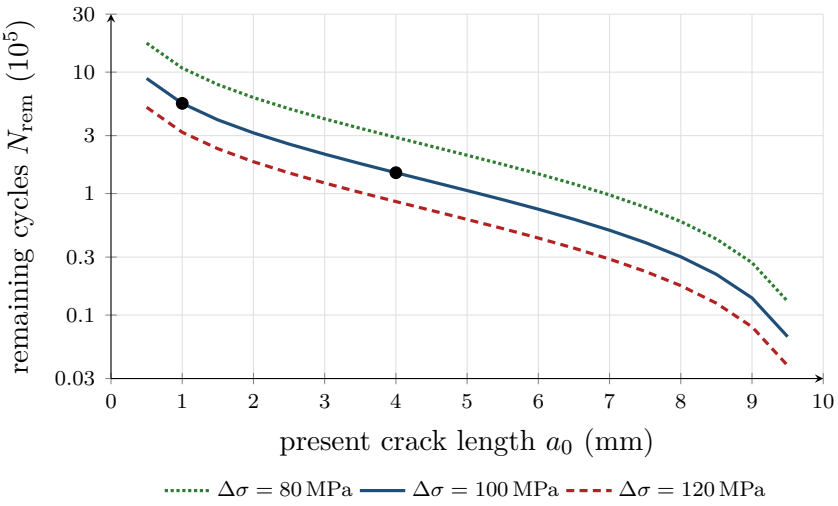


Figure 21.4: Remaining fatigue life against the present crack length, from the Paris law, equation (21.20), with illustrative steel parameters $C_{\mathrm{P}} = 10^{-11}$, $m = 3$, $Y = 1.12$ and a critical length $a_c = 10\,\mathrm{mm}$. Marked: 5.53×10^5 cycles for $a_0 = 1\,\mathrm{mm}$ and 1.49×10^5 cycles for $a_0 = 4\,\mathrm{mm}$ at 100 MPa.

crack length. A measurement of the crack length therefore moves the prediction more than any refinement of the age distribution could.

The Paris law applies only within its crack-growth regime. Near the threshold stress-intensity range below which cracks do not propagate, and near unstable fracture as the crack approaches its critical size, other fracture-mechanics relations are needed, and the critical length itself depends on the fracture toughness of the material and the peak load.

21.3.6 Thermal aging and the Arrhenius acceleration factor

Many insulation materials and polymers age faster at higher temperatures. Section 4.3.4 wrote the rate of a thermally activated degradation process as $k(T) = k_0 \exp[-E_a/(k_B T)]$, with the activation energy E_a in electronvolts per molecule and the Boltzmann constant k_B. The same phenomenological law is often written per mole,

$$k(T) = k_0 \exp\!\left(-\frac{E_a}{R_{\mathrm{g}} T}\right), \tag{21.22}$$

where $k(T)$ is the degradation reaction rate, k_0 the pre-exponential factor, E_a the activation energy in J/mol, $R_{\mathrm{g}} = 8.314\,\mathrm{J/(mol\,K)}$ the molar gas constant (neither the reliability function R nor a resistance), and T the absolute temperature. The temperature must be in kelvin. The two forms are the same law, since E_a/R_{g} with E_a per mole equals E_a/k_B with E_a per molecule.

For two temperatures T_1 and T_2 the pre-exponential factor cancels in the ratio of the rates,

$$\frac{k(T_2)}{k(T_1)} = \frac{k_0\, \mathrm{e}^{-E_a/(R_{\mathrm{g}} T_2)}}{k_0\, \mathrm{e}^{-E_a/(R_{\mathrm{g}} T_1)}},$$

so the acceleration factor of T_2 relative to T_1 is

$$\mathrm{AF} = \exp\!\left[\frac{E_a}{R_{\mathrm{g}}}\left(\frac{1}{T_1} - \frac{1}{T_2}\right)\right]. \tag{21.23}$$

If $T_2 > T_1$, then $1/T_1 > 1/T_2$ and the acceleration factor exceeds one for a positive activation energy.

Example 21.3 (Transformer insulation aging)**.** Take an insulation aging mechanism with $E_a = 80\,000\,\mathrm{J/mol}$ (illustrative; 0.829 eV per molecule) and compare the operating temperatures $T_1 = 333.15\,\mathrm{K}$ (60°C) and $T_2 = 343.15\,\mathrm{K}$ (70°C). Then

$$\begin{aligned} \frac{E_a}{R_\mathrm{g}} &= \frac{80\,000}{8.314}\,\mathrm{K} = 9622\,\mathrm{K}, \\ \frac{1}{T_1} - \frac{1}{T_2} &= \frac{343.15 - 333.15}{333.15 \times 343.15}\,\mathrm{K}^{-1} = 8.747 \times 10^{-5}\,\mathrm{K}^{-1}, \\ \mathrm{AF} &= \exp(9622 \times 8.747 \times 10^{-5}) = \exp(0.8417) = 2.32. \end{aligned}$$

Under this illustrative model, the insulation degrades 2.32 times as fast at 70°C as at 60°C.

The result can be compared with the doubling rules used for insulation. The general rule of thumb that the aging rate of an insulation doubles for every 10 K corresponds, near 60°C, to an activation energy of about 66 kJ/mol (0.68 eV). The loading guides for oil-immersed power transformers use a steeper rule for non-thermally-upgraded paper: the relative aging rate doubles for every 6 K of hot-spot temperature, $V = 2^{(\theta_h - 98)/6}$ with the hot-spot temperature θ_h in degrees Celsius and $V = 1$ at 98°C, which near 98°C corresponds to about 134 kJ/mol (1.39 eV). Over a 10 K rise the three give factors of 2, $2^{10/6} = 3.17$ and, for the example, 2.32; at 80 kJ/mol the rate doubles for every 8.2 K near 60°C, between the two rules. The doubling interval of an Arrhenius law is not a constant: it grows as T^2, to 8.7 K near 70°C. A doubling rule is a local summary of an activation energy, valid near the temperature at which it was established.

An acceleration factor of 2.32 does not mean that the transformer fails 2.32 times as often. It means that the modeled degradation mechanism proceeds 2.32 times as fast. Connecting the accumulated damage to equipment failure still requires a failure criterion, such as the loss of mechanical strength of the paper at which a short-circuit force tears it, and a model of the material condition at which that occurs, as in Section 4.3.4.

21.3.7 Cumulative thermal damage

Suppose that the damage accumulates at a rate proportional to the temperature-dependent reaction rate, $\mathrm{d}D/\mathrm{d}t = k(T(t))$, where the temperature follows a history $T(t)$. Integrating, as in equation (21.13),

$$D(t) = D_0 + \int_0^t k(T(u))\,\mathrm{d}u. \tag{21.24}$$

Because k_0 is rarely known, the damage is normalized by the rate at a reference temperature T_ref,

$$D_\mathrm{equiv}(t) = \int_0^t \frac{k(T(u))}{k(T_\mathrm{ref})}\,\mathrm{d}u, \tag{21.25}$$

in which the integrand is the acceleration factor of equation (21.23) at each instant, with $T_1 = T_\mathrm{ref}$. The result has units of time: it is the equivalent exposure time at the reference temperature, which transformer loading guides call the loss of life. One hour at a temperature with acceleration factor 2.32 contributes the same nominal damage as 2.32 hours at the reference temperature. The normalization is useful for equipment exposed to strongly time-varying temperatures, whose aging cannot be read from the calendar.

As an illustration, let the transformer of Example 21.3 operate every day for 16 h at 60°C and 8 h at 70°C, with $T_\mathrm{ref} = 60$°C. Its equivalent aging per day is

$$16\,\mathrm{h} \times 1 + 8\,\mathrm{h} \times 2.32 = 34.56\,\mathrm{h},$$

1.44 times the calendar time. Evaluating the acceleration factor at the time-averaged temperature, $(16 \times 60 + 8 \times 70)/24 = 63.3$°C, would give 1.33 and understate the aging by 8 %. Because

the rate is a convex function of temperature, the hot hours contribute more than the average temperature suggests: the damage rate must be integrated along the temperature history, not evaluated at its mean, as the hazard was in Section 5.6.4 of Chapter 5 of Volume I.

Principle 21.2. *A physics-based damage law turns the recorded operating history, counted cycles, crack lengths, temperatures, into a degradation state. The state is updated by integrating the damage rate along the actual history; evaluating the rate at average conditions understates the damage whenever the rate is convex in the stress.*

21.4 Stochastic degradation processes

Deterministic models predict a threshold-crossing time but not a failure probability. This section introduces random variation along the degradation path, develops the Wiener process with its first-passage distribution and the gamma process with its failure probability, and compares the families of stochastic degradation models.

21.4.1 Why deterministic models are insufficient

Two identical pumps operating under apparently identical conditions do not follow the same degradation trajectory. The differences come from

- manufacturing variability;
- different microscopic material defects;
- unobserved loading fluctuations;
- lubrication and environmental differences;
- random damage initiation;
- uncertain parameters of the degradation model.

These effects are represented by adding stochastic variation to the evolution law. A general formulation is the stochastic differential equation

$$\mathrm{d}D(t) = g_D(D, t, \mathbf{u}, \boldsymbol{\theta})\,\mathrm{d}t + \sigma_D(D, t)\,\mathrm{d}B(t), \tag{21.26}$$

where g_D is the deterministic degradation drift of equation (21.3), $\sigma_D(D, t)$ the intensity of the stochastic variability, and $B(t)$ a standard Brownian motion, also called a standard Wiener process. (B is used for the Brownian motion because $W(t)$ is the cumulative intensity of a counting process in Chapter 17 of Volume I.) The stochastic term lets different components follow different trajectories under the same drift, and it makes the first-passage time of equation (21.2) a random variable even when the parameters are known.

Two families of degradation processes are the most used, the Wiener process and the gamma process. Which of them is appropriate depends in part on whether the physical damage must be monotonic: a Wiener path can decrease, a gamma path cannot.

21.4.2 The Wiener degradation process

The basic Wiener degradation model is equation (21.26) with a constant drift and a constant diffusion, integrated from D_0:

$$D(t) = D_0 + \mu_D t + \sigma_D B(t), \tag{21.27}$$

where $\mu_D > 0$ is the drift, the average degradation rate in units of D per unit time, and $\sigma_D > 0$ is the diffusion parameter in units of D per square root of time. (μ_D is the drift of the degradation state, not the repair rate μ of Volume I.) The standard Brownian motion $B(t)$ has three defining properties: $B(0) = 0$; its increments over disjoint intervals are independent;

and the increment over an interval of length Δt is normal with mean zero and variance Δt, so that $B(t) \sim \mathcal{N}(0, t)$, in units of the square root of time. Its paths are continuous but nowhere differentiable, which is why equation (21.26) is written in differentials.

The mean and the variance of the state follow from these properties. Since D_0 and $\mu_D t$ are constants and $\mathbb{E}[B(t)] = 0$,

$$\mathbb{E}[D(t)] = D_0 + \mu_D t + \sigma_D \, \mathbb{E}[B(t)] = D_0 + \mu_D t. \tag{21.28}$$

Adding a constant does not change a variance, and multiplying a random variable by σ_D multiplies its variance by σ_D^2, so

$$\operatorname{Var}[D(t)] = \sigma_D^2 \, \operatorname{Var}[B(t)] = \sigma_D^2 \, t. \tag{21.29}$$

A linear function of a normal variable is normal; under known parameters, therefore,

$$D(t) \sim \mathcal{N}\big(D_0 + \mu_D t, \; \sigma_D^2 t\big). \tag{21.30}$$

The expected degradation follows the straight line of the deterministic linear model, equation (21.5), and the individual trajectories fluctuate around it with a standard deviation that grows as the square root of time.

The Wiener process has an important limitation: its increments can be negative, so the modeled degradation state can decrease. That can be appropriate for a health indicator that fluctuates with the operating conditions, or for a damage measure whose small decreases represent measurement and modeling noise. It is inappropriate for a physically irreversible quantity such as accumulated metal loss, which never decreases; for strictly increasing damage, a gamma process (Section 21.4.5) may be preferable.

21.4.3 Stochastic bearing degradation

Example 21.4 (Stochastic bearing degradation)**.** The degradation of a bearing of the cooling-water pump is expressed as a normalized damage index with failure threshold $D_{\text{crit}} = 1$ and present damage $D_0 = 0$. Assume a Wiener model with $\mu_D = 0.01\,\text{day}^{-1}$ and $\sigma_D = 0.03\,\text{day}^{-1/2}$ (illustrative), so that

$$D(t) = 0.01\,t + 0.03\,B(t),$$

with t in days. By equation (21.28), the mean degradation after 100 days is $\mathbb{E}[D(100)] = 0.01 \times 100 = 1$, exactly the threshold. By equation (21.29), its standard deviation is

$$\sqrt{\operatorname{Var}[D(100)]} = 0.03 \times \sqrt{100} = 0.30.$$

At day 100 the damage of an individual bearing is uncertain, normal with mean 1 and standard deviation 0.30, even though its expected value has just reached the threshold.

Figure 21.5 shows what this means for individual bearings. The horizontal axis is time in days and the vertical axis the damage index. The five thin colored lines are sample paths of the process, simulated with daily increments drawn from $\mathcal{N}(0.01, 0.03^2)$; the straight black line is the mean $0.01\,t$, and the two dashed gray curves are the mean plus and minus two standard deviations, $0.01\,t \pm 0.06\sqrt{t}$. The dashed red line is the threshold, and the dots mark where each path first reaches it: on days 83, 91, 115, 143 and 145. Two features stand out. The paths spread out as $\sqrt{t}$, so the failure times spread around the deterministic 100 days. And some paths cross the threshold and return below it: the path that first reaches 1 on day 91 is at 0.975 on day 100, below the threshold, although under the definition of equation (21.2) the bearing failed nine days earlier.

The probability that the degradation is above the threshold at exactly day 100 is therefore not the probability that the threshold has been crossed at some time during the first 100 days. The first, from equation (21.30), is $\Pr(D(100) \geq 1) = 1 - \Phi\big((1-1)/0.30\big) = 0.5$, where Φ is the standard normal cumulative distribution function. The second is a first-passage probability, and it is larger, because it also counts the paths that crossed earlier and came back.

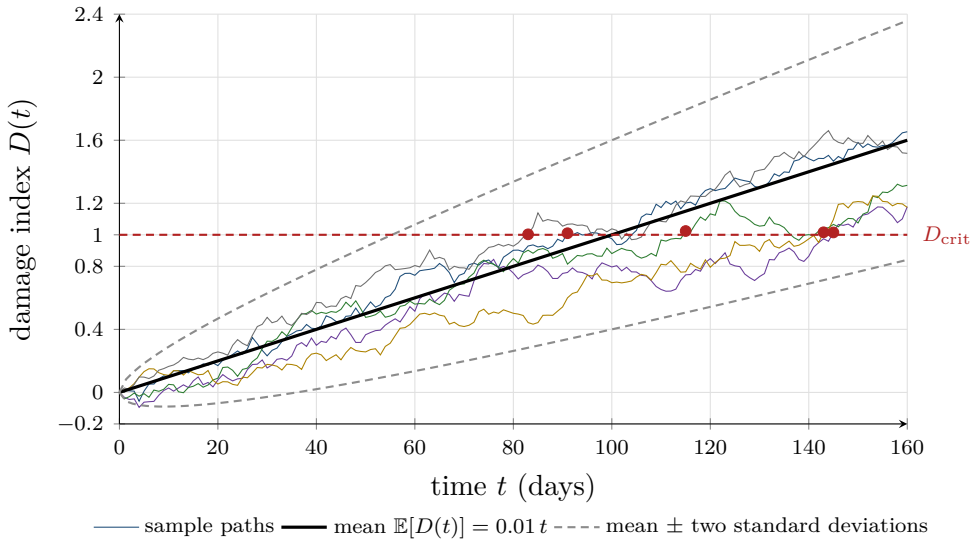


Figure 21.5: Five simulated sample paths of the Wiener bearing model $D(t) = 0.01\,t + 0.03\,B(t)$, with t in days (illustrative), the mean $0.01\,t$ and the band $0.01\,t \pm 0.06\sqrt{t}$. The dots mark the first crossings of $D_{\text{crit}} = 1$, on days 83, 91, 115, 143 and 145. The path that crosses on day 91 is back below the threshold on day 100.

21.4.4 The first-passage time distribution

For the Wiener model of equation (21.27), define the remaining distance to the threshold,

$$\delta = D_{\text{crit}} - D_0 > 0, \tag{21.31}$$

in the units of D. The failure time is the first-passage time of equation (21.2). For a positive constant drift, it is a standard result of the theory of Brownian motion that T_{f} follows the inverse Gaussian distribution, with the cumulative distribution function

$$F_{T_{\text{f}}}(t) = \Pr(T_{\text{f}} \le t) = \Phi\left(\frac{\mu_D t - \delta}{\sigma_D\sqrt{t}}\right) + \exp\left(\frac{2\mu_D\delta}{\sigma_D^2}\right)\Phi\left(-\frac{\mu_D t + \delta}{\sigma_D\sqrt{t}}\right), \qquad t > 0, \tag{21.32}$$

and, differentiating with respect to t, the density

$$f_{T_{\text{f}}}(t) = \frac{\delta}{\sigma_D\sqrt{2\pi t^3}}\exp\left[-\frac{(\delta - \mu_D t)^2}{2\sigma_D^2 t}\right], \qquad t > 0. \tag{21.33}$$

The first term of equation (21.32) is the probability that the free process is above the threshold at time t; the second, obtained by reflecting the paths at the threshold, adds the probability of the paths that crossed before t and are below it at t. In the usual two-parameter form the inverse Gaussian distribution has mean δ/μ_D and shape parameter δ^2/σ_D^2, in units of time. Because $\mu_D > 0$, the process crosses the threshold with probability one: $F_{T_{\text{f}}}(t) \to 1$ as $t \to \infty$.

The mean and the variance can be derived without integrating the density, from two properties of Brownian motion stopped at a first-passage time whose mean is finite: $\mathbb{E}[B(T_{\text{f}})] = 0$ and $\mathbb{E}[B(T_{\text{f}})^2] = \mathbb{E}[T_{\text{f}}]$ (Wald's identities). Because the paths are continuous, the state at the crossing is exactly the threshold, $D(T_{\text{f}}) = D_{\text{crit}}$, so from equation (21.27)

$$\delta = \mu_D T_{\text{f}} + \sigma_D B(T_{\text{f}}). \tag{21.34}$$

Taking expectations, $\delta = \mu_D\,\mathbb{E}[T_{\text{f}}]$. Squaring $\sigma_D B(T_{\text{f}}) = \delta - \mu_D T_{\text{f}} = -\mu_D\,(T_{\text{f}} - \delta/\mu_D)$ and taking expectations, $\sigma_D^2\,\mathbb{E}[T_{\text{f}}] = \mu_D^2\,\mathrm{Var}(T_{\text{f}})$. Hence

$$\mathbb{E}[T_{\text{f}}] = \frac{\delta}{\mu_D}, \qquad \mathrm{Var}(T_{\text{f}}) = \frac{\delta\,\sigma_D^2}{\mu_D^3}. \tag{21.35}$$

The mean is the deterministic failure time of the linear model, equation (21.6) with $\kappa = \mu_D$; the variance grows with the diffusion and falls steeply with the drift.

For the bearing of Example 21.4, $\delta = 1$, $\mu_D = 0.01\,\mathrm{day}^{-1}$ and $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$, so

$$\mathbb{E}[T_\mathrm{f}] = \frac{1}{0.01} = 100 \text{ days}, \qquad \mathrm{SD}(T_\mathrm{f}) = \sqrt{\frac{1 \times 0.03^2}{0.01^3}} = \sqrt{900} = 30 \text{ days}.$$

This prediction is more informative than the deterministic estimate of exactly 100 days: it is a distribution of possible failure times.

The reliability is the probability of surviving without crossing the threshold,

$$R(t) = \Pr(T_\mathrm{f} > t) = 1 - F_{T_\mathrm{f}}(t). \tag{21.36}$$

At $t = 100$ days, the first argument of equation (21.32) is $(1 - 1)/0.30 = 0$, so the first term is $\Phi(0) = 0.5$; the exponent of the second is $2 \times 0.01 \times 1/0.0009 = 22.22$ and its argument $-(1+1)/0.30 = -6.667$, so the second term is $\mathrm{e}^{22.22}\,\Phi(-6.667) = 4.48 \times 10^9 \times 1.31 \times 10^{-11} = 0.059$. Hence $F_{T_\mathrm{f}}(100) = 0.559$ and $R(100) = 0.441$. The probability that the bearing has failed by day 100 is 0.559, not the 0.5 that the distribution of $D(100)$ alone suggests. The calculation answers the physically relevant question, the probability that the component reaches the critical damage before time t, rather than the probability that its damage happens to exceed the threshold at the instant t.

Figure 21.6 shows the distribution. In panel (a) the horizontal axis is time in days and the vertical axis the first-passage density of equation (21.33); the solid curve is the new bearing, $\delta = 1$, and the dashed curve a bearing whose damage is already 0.40, $\delta = 0.60$, the case of Section 21.5.1. In panel (b) the vertical axis is a probability of failure by time t; the solid blue curve is the first-passage probability of equation (21.32) for $\delta = 1$, the dotted gray curve the fixed-time probability $\Pr(D(t) \geq 1)$, and the dashed green curve the gamma process of Section 21.4.5 with the same mean and variance. Both densities are skewed to the right: for $\delta = 1$ the mode is at 87.4 days and the median at 95.7 days, below the mean of 100, and the 5 % quantile is 59.1 days. The first-passage curve lies above the fixed-time curve at every time, by up to 0.059, which is the probability of the paths that crossed and returned. The gamma curve is discussed below.

21.4.5 Gamma degradation processes

Many physical degradation mechanisms are, to a good approximation, monotonic: cumulative corrosion penetration, progressive material loss, accumulated irreversible wear, and some forms of cumulative fatigue damage. For these a gamma process is a natural model. Let $D(t) \geq D_0$ be the cumulative damage. A stationary gamma process has independent increments, and the increment over any interval of length Δt is gamma distributed,

$$D(t + \Delta t) - D(t) \sim \mathrm{Gamma}\big(\alpha_D\,\Delta t,\ \beta_D\big), \tag{21.37}$$

with shape $\alpha_D \Delta t$ and *scale* β_D: the density of an increment $x > 0$ is

$$\frac{x^{\alpha_D \Delta t - 1}\,\mathrm{e}^{-x/\beta_D}}{\Gamma(\alpha_D \Delta t)\,\beta_D^{\alpha_D \Delta t}}.$$

Here α_D is the shape rate, per unit time, and β_D the scale, in the units of D. (The gamma prior $\mathrm{Gamma}(a, b)$ of Chapter 8 of Volume I is written with a rate b; the gamma process is conventionally written with a scale, and the two parameterizations are related by $b = 1/\beta_D$.) The shape grows in proportion to the length of the interval, which is what makes increments over adjacent intervals add up consistently: the sum of independent gamma variables with a

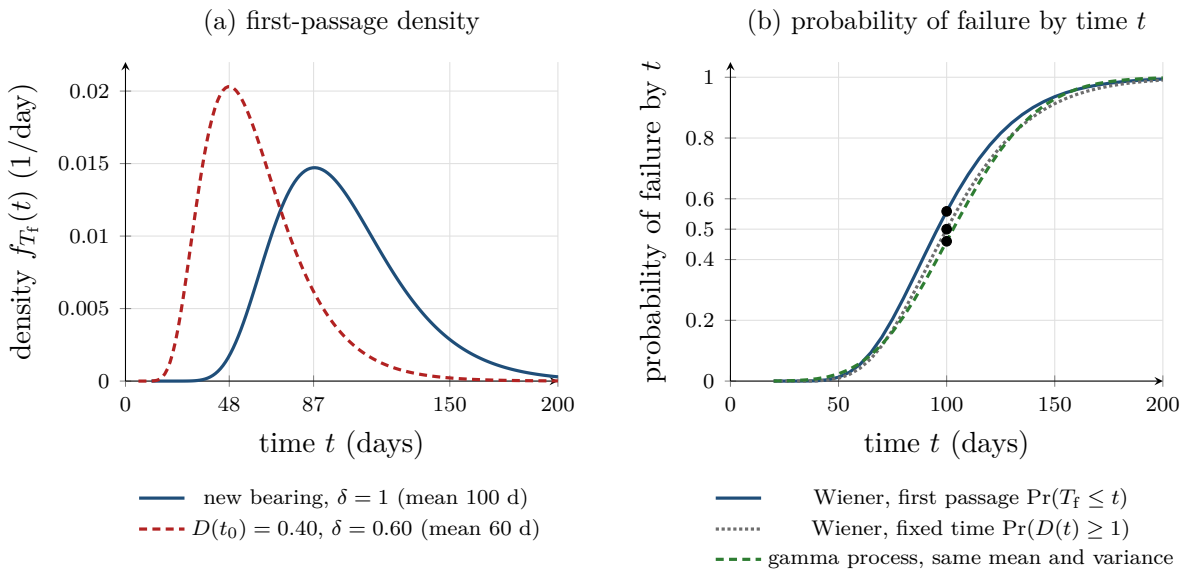


Figure 21.6: First-passage time of the Wiener bearing model with $\mu_D = 0.01\,\text{day}^{-1}$ and $\sigma_D = 0.03\,\text{day}^{-1/2}$ (illustrative). (a) Inverse Gaussian density, equation (21.33), for $\delta = 1$ (mode 87.4 days) and $\delta = 0.60$ (mode 48 days). (b) For $\delta = 1$: the first-passage probability, equation (21.32), the fixed-time probability $\Pr(D(t) \ge 1)$, and the failure probability of a gamma process with shape rate $\alpha_D = 0.111\,\text{day}^{-1}$ and scale $\beta_D = 0.09$; marked at day 100: 0.559, 0.500 and 0.460.

common scale is gamma with the shapes added. From the mean $\alpha\beta$ and variance $\alpha\beta^2$ of a gamma variable with shape α and scale β,

$$\mathbb{E}[\Delta D] = \alpha_D\,\beta_D\,\Delta t, \qquad \operatorname{Var}(\Delta D) = \alpha_D\,\beta_D^2\,\Delta t, \tag{21.38}$$

so the mean degradation rate is $\alpha_D\beta_D$ and the variance grows linearly in time, as for the Wiener process. Most importantly, $\Delta D \ge 0$ with probability one: the degradation never decreases.

Failure probability under a gamma process. For an initial damage D_0 and threshold D_{crit}, let $\delta = D_{\text{crit}} - D_0$ as in equation (21.31), and let the future increments follow the stationary gamma process. Failure occurs when $D(t) \ge D_{\text{crit}}$. Since the gamma process is nondecreasing, crossing the threshold before time t is equivalent to being at or above it at time t, and

$$\Pr(T_f \le t) = \Pr\big[D(t) - D_0 \ge \delta\big], \tag{21.39}$$

where $D(t) - D_0 \sim \text{Gamma}(\alpha_D t, \beta_D)$. The reliability is therefore the gamma cumulative distribution function with shape $\alpha_D t$, evaluated at δ,

$$R(t) = F_{\text{Gamma}(\alpha_D t,\,\beta_D)}(\delta) = \frac{\gamma(\alpha_D t,\ \delta/\beta_D)}{\Gamma(\alpha_D t)}, \tag{21.40}$$

where $\gamma(s, x) = \int_0^x u^{s-1}\mathrm{e}^{-u}\,\mathrm{d}u$ is the lower incomplete gamma function, so that the ratio is the regularized incomplete gamma function, available in every numerical library. This is a simplification compared with the Wiener process: because the damage cannot decrease, there is no possibility of crossing the threshold and returning below it, and no reflection term. When the shape $\alpha_D t = k$ is an integer, the gamma variable is a sum of k exponential variables with mean β_D, and it reaches δ if and only if fewer than k events of a Poisson process of rate $1/\beta_D$ occur in an interval of length δ; then

$$\Pr(T_f \le t) = \sum_{j=0}^{k-1} \mathrm{e}^{-\delta/\beta_D}\,\frac{(\delta/\beta_D)^j}{j!}, \qquad \alpha_D t = k \in \{1, 2, \ldots\}, \tag{21.41}$$

which can be evaluated by hand.

Example 21.5 (Corrosion as a gamma process)**.** Return to the pipe of Example 21.1, with $\delta = 2.0 - 0.60 = 1.40\,\text{mm}$ of penetration left, and model the corrosion as a gamma process with shape rate $\alpha_D = 2\,\text{yr}^{-1}$ and scale $\beta_D = 0.04\,\text{mm}$ (illustrative). The mean rate is $\alpha_D\beta_D = 0.08\,\text{mm/yr}$, the rate of the deterministic example, and the variance rate is $\alpha_D\beta_D^2 = 0.0032\,\text{mm}^2/\text{yr}$. After 17.5 years the penetration added has shape 35, mean 1.40 mm and standard deviation $0.04\sqrt{35} = 0.237\,\text{mm}$, and equation (21.39) gives the probability that the threshold has been reached: 0.002 at 10 years, 0.032 at 12.5, 0.177 at 15, 0.478 at 17.5, 0.780 at 20 and 0.941 at 22.5 years. The median failure time is 17.7 years, slightly after the deterministic 17.5 years because the gamma distribution is skewed to the right, and the probability of reaching the threshold within 15 years, 2.5 years before the deterministic prediction, is about 18 %.

The gamma process of Figure 21.6(b) has the same mean rate and variance rate as the Wiener bearing model: $\alpha_D\beta_D = 0.01$ and $\alpha_D\beta_D^2 = 0.0009$ per day, so $\beta_D = 0.0009/0.01 = 0.09$ and $\alpha_D = 0.01/0.09 = 0.111\,\text{day}^{-1}$. Its failure probability at day 100 is 0.460, below both Wiener curves. At that time its state has mean 1 and standard deviation 0.30, like the Wiener state, but the gamma distribution is skewed to the right, so its median is below the threshold, and it has no crossings that return. In the early tail the order reverses: at day 40 the gamma probability is 0.0077 against 0.0011 for the Wiener first passage, because the right skew puts more probability on rapid early damage. Two models with the same first two moments thus give different failure probabilities, most of all in the early tail on which maintenance decisions depend.

21.4.6 Comparison of stochastic degradation models

Table 21.3 compares the families of degradation models by the sign of their increments, the uses for which they are appropriate and their main limitation. The rows run from the deterministic evolution law of Section 21.2, which has no intrinsic random variation, through the Wiener and gamma processes of this section, to the inverse Gaussian process, a monotone process whose increments follow the inverse Gaussian distribution, and to physics-based stochastic models in which the parameters of a damage law such as the Paris law are random.

Table 21.3: Families of degradation models, their increments, appropriate uses and main limitations.

Model	Degradation increments	Appropriate use	Main limitation
Deterministic evolution law	fixed, given parameters and inputs	well-understood physical laws	no intrinsic random variation of trajectories
Wiener process	positive or negative	fluctuating health indicators, stochastic drift	damage can decrease
Gamma process	nonnegative	monotone cumulative damage	the independent-increment assumption may be unrealistic
Inverse Gaussian process	nonnegative	monotone damage with flexible increment variability	more complicated estimation and interpretation
Physics-based stochastic model	mechanism-dependent	fatigue, thermal damage, battery aging	more states and parameters; harder calibration

The model should be chosen from the physical damage mechanism and the observations, not because one mathematical family is convenient. Corrosion penetration and accumulated wear call for a monotone process; a vibration-based health indicator, which fluctuates with speed

and load, may be represented adequately by a Wiener process; and where the mechanism is known, a stochastic version of its own law carries the most information.

Principle 21.3. *A stochastic degradation model turns a deterministic crossing time into a distribution of first-passage times. For a process that can decrease, the probability of failure by time t is the first-passage probability, which exceeds the probability that the state is above the threshold at t; for a monotone process the two are equal.*

21.5 Condition-based reliability

The models so far describe a population of components or a component whose state is known. This section conditions the reliability on the observed state of one particular component, distinguishes the measurement from the state it informs, and estimates the state from noisy measurements with a Bayesian filter.

21.5.1 From population reliability to equipment-specific reliability

Traditional lifetime reliability calculates $R(t) = \Pr(T_\mathrm{f} > t)$ for a population or a modeled class of components. Suppose instead that the current condition of one particular pump is observed: an inspection at time t_0 estimates the normalized damage of its bearing as $D(t_0) = 0.40$, with critical damage $D_\mathrm{crit} = 1$. The question becomes the probability that this pump, given its current damage and operating history, survives another 30 days. This is a conditional reliability. Define the remaining useful life at t_0 as

$$\mathrm{RUL}(t_0) = T_\mathrm{f} - t_0, \tag{21.42}$$

conditional on the component having survived until t_0. Its reliability over a further duration τ is

$$R(\tau \mid t_0) = \Pr\big[\mathrm{RUL}(t_0) > \tau \mid \mathbf{y}_{1:k},\ T_\mathrm{f} > t_0\big], \tag{21.43}$$

where $\mathbf{y}_{1:k} = (y_1, \dots, y_k)$ are all the condition measurements observed up to $t_k = t_0$. The expression includes the uncertainty about the present degradation state and, where they are estimated, about the model parameters. Its quantile $\mathrm{RUL}_\alpha(t_0)$, defined by $\Pr[\mathrm{RUL}(t_0) \le \mathrm{RUL}_\alpha(t_0)] = \alpha$, is the remaining life that is reached with probability α; a small α, such as 0.05, gives a conservative planning horizon.

Damage known exactly. For a Wiener process with independent future increments and known μ_D and σ_D, the future depends on the past only through the present state, and the remaining life is the first-passage time from $D(t_0)$, with the remaining distance

$$\delta = D_\mathrm{crit} - D(t_0) = 1 - 0.40 = 0.60.$$

With the parameters of Example 21.4, $\mu_D = 0.01\,\mathrm{day}^{-1}$ and $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$, equation (21.35) gives

$$\mathbb{E}[\mathrm{RUL}] = \frac{0.60}{0.01} = 60 \text{ days}, \qquad \mathrm{SD}(\mathrm{RUL}) = \sqrt{\frac{0.60 \times 0.03^2}{0.01^3}} = \sqrt{540} = 23.24 \text{ days},$$

and equation (21.32) with $t = \tau = 30$ days gives $\Pr(\mathrm{RUL} \le 30) = 0.047$, so the probability that this pump survives another 30 days is $R(30 \mid t_0) = 0.953$. The 5 % quantile of the remaining life is $\mathrm{RUL}_{0.05} = 30.3$ days. Compare this with the population prediction of 100 days from a zero-damage state: the current condition changes the prediction, but it does not eliminate the uncertainty.

The difference from the conditional survival of Chapter 5 of Volume I is worth computing. There, Section 5.2.4 conditioned only on survival, $R(t+\Delta t \mid t) = R(t+\Delta t)/R(t)$. Suppose the inspection took place at $t_0 = 40$ days on a bearing installed new, when the mean damage of the population is $0.01 \times 40 = 0.40$, the measured value. The population conditional survival over the next 30 days is $R(70)/R(40) = (1-0.1442)/(1-0.0011) = 0.857$. The equipment-specific value for the bearing measured at 0.40 is 0.953; for a bearing of the same age measured at 0.80 it would be 0.172. The population value averages over all the bearings that survived to day 40, whose damage then is spread with a standard deviation of $0.03\sqrt{40} = 0.19$; the measurement places the particular bearing within that spread.

21.5.2 Measurements are not the degradation state

In most applications the physical damage is not observed directly. Sensors produce measurements that are related to it, through an observation model

$$y_k = h(D(t_k), \mathbf{u}_k) + \varepsilon_k, \tag{21.44}$$

where y_k is the sensor measurement at time t_k, $D(t_k)$ the latent degradation state, $\mathbf{u}_k$ the operating conditions at that time, h the measurement relationship, and ε_k the measurement error. The vibration amplitude of a pump, for instance, depends on the damage of its bearing, but also on the rotational speed, the load and the mounting conditions. A simple linear observation model is

$$y_k = c_y\, D(t_k) + \varepsilon_k, \qquad \varepsilon_k \sim \mathcal{N}(0, \sigma_y^2), \tag{21.45}$$

with a sensitivity c_y in units of the measurement per unit of damage and a measurement standard deviation σ_y. If σ_y is large, a single measurement carries little information about the damage; if c_y is poorly known, the damage is difficult to identify at all, because a change in y can be explained by a change in c_y as well as by a change in D.

Example 21.6 (Pump vibration)**.** The vibration velocity of the cooling-water pump is measured at $y = 4.0\,\mathrm{mm/s}$. Suppose that a calibration suggested $c_y = 5\,\mathrm{mm/s}$ per unit of damage (illustrative). Inverting the observation model would give $D = 4.0/5 = 0.80$, but that conclusion is not justified without a credible mapping between vibration and bearing condition: the same vibration can arise from misalignment, imbalance, looseness, cavitation or other mechanisms, and from a change of speed or load with no damage at all. The number 0.80 is a hypothesis about the state, not a measurement of it.

The example separates the three activities of PHM named at the start of the chapter. *Condition monitoring* observes indicators, here the vibration velocity. *Diagnosis* determines the mechanism responsible for a change in the indicators, here bearing damage rather than misalignment or imbalance. *Prognostics* predicts the future degradation and failure from the diagnosed state. Diagnosis is the subject of Chapter 22 and the estimation of the remaining useful life that of Chapter 23; the next subsection shows the inference that links a measurement to the state once the mechanism and the observation model are accepted.

Common mistake

Treating a sensor indicator as a direct measurement of damage. A vibration velocity, a temperature rise or a resistance responds to the operating conditions and to other fault mechanisms as well as to the damage of interest. The link from the indicator to the state is an observation model with its own uncertainty, and the state is inferred through it.

21.5.3 Bayesian estimation of the degradation state

Let $p(D_k \mid \mathbf{y}_{1:k})$ be the posterior distribution of the degradation state $D_k = D(t_k)$ after the measurements up to time t_k, the probability distribution of the state given the evidence, in the sense of Chapter 8 of Volume I. A Bayesian filter updates it recursively with two operations.

Prediction. The previous distribution of the state is propagated through the process model, the probability $p(D_k \mid D_{k-1})$ of reaching D_k from D_{k-1}, by the law of total probability:

$$p(D_k \mid \mathbf{y}_{1:k-1}) = \int p(D_k \mid D_{k-1})\, p(D_{k-1} \mid \mathbf{y}_{1:k-1})\, \mathrm{d}D_{k-1}. \tag{21.46}$$

Measurement update. The predicted distribution is updated with the new measurement by Bayes' theorem, with the likelihood $p(y_k \mid D_k)$ given by the observation model:

$$p(D_k \mid \mathbf{y}_{1:k}) \propto p(y_k \mid D_k)\, p(D_k \mid \mathbf{y}_{1:k-1}). \tag{21.47}$$

The result is a probability distribution over the current damage, not a single precise estimate. Future degradation trajectories are then propagated from this posterior, and the probability of failure within a further τ,

$$\Pr(T_\mathrm{f} \le t_k + \tau \mid \mathbf{y}_{1:k}) = \int \Pr\big(\mathrm{RUL} \le \tau \mid D_k = d\big)\, p(d \mid \mathbf{y}_{1:k})\, \mathrm{d}d, \tag{21.48}$$

integrates the threshold-crossing probability of each possible present state over the posterior uncertainty in that state (and, where they are estimated, in the parameters). This connects degradation modeling directly to the probabilistic inference methods of Volume I.

The scalar Kalman filter. For the Wiener model and the linear observation model, both operations have closed forms. Sampled at interval Δt, equation (21.27) gives the process model

$$D_k = D_{k-1} + \mu_D \Delta t + w_k, \qquad w_k = \sigma_D\,[B(t_k) - B(t_{k-1})] \sim \mathcal{N}(0, \sigma_D^2 \Delta t),$$

with w_k independent of D_{k-1} by the independence of the increments. Let the posterior at t_{k-1} be normal with mean m_{k-1} and variance v_{k-1}. In the prediction step, D_k is the sum of two independent normal variables, $D_{k-1} + \mu_D \Delta t$ and w_k, so the integral of equation (21.46), a convolution, is normal with the means and the variances added:

$$m_{k|k-1} = m_{k-1} + \mu_D \Delta t, \qquad v_{k|k-1} = v_{k-1} + \sigma_D^2 \Delta t. \tag{21.49}$$

The prediction moves the mean by the drift and widens the distribution by the process variance. In the update step, the likelihood of equation (21.45) and the predicted density are both Gaussian in D_k, and their product is proportional to

$$\exp\left[-\frac{(y_k - c_y D_k)^2}{2\sigma_y^2} - \frac{(D_k - m_{k|k-1})^2}{2 v_{k|k-1}}\right].$$

The exponent is quadratic in D_k. The coefficient of $-D_k^2/2$ is $c_y^2/\sigma_y^2 + 1/v_{k|k-1}$, which is the inverse of the posterior variance, and the coefficient of D_k is $c_y y_k/\sigma_y^2 + m_{k|k-1}/v_{k|k-1}$, which is the posterior mean divided by the posterior variance. Completing the square therefore gives a normal posterior with

$$\frac{1}{v_k} = \frac{1}{v_{k|k-1}} + \frac{c_y^2}{\sigma_y^2}, \qquad m_k = v_k \left(\frac{m_{k|k-1}}{v_{k|k-1}} + \frac{c_y\, y_k}{\sigma_y^2}\right): \tag{21.50}$$

the precisions (inverse variances) add, and the posterior mean is a precision-weighted combination of the prediction and the measurement. Writing $S_k = c_y^2 v_{k|k-1} + \sigma_y^2$ for the variance of the predicted measurement, the same result takes the standard form of the scalar Kalman filter,

$$G_k = \frac{v_{k|k-1}\, c_y}{S_k}, \qquad m_k = m_{k|k-1} + G_k\, e_k, \qquad v_k = (1 - G_k c_y)\, v_{k|k-1}, \tag{21.51}$$

where $e_k = y_k - c_y m_{k|k-1}$ is the innovation, the residual of the measurement against its prediction, and G_k is the Kalman gain (usually written K, here G to keep K for the S-N constant). The two forms agree: $v_k = v_{k|k-1}\sigma_y^2/S_k = (1 - G_k c_y) v_{k|k-1}$, and expanding m_k of equation (21.50) over the common denominator S_k gives $m_{k|k-1} + v_{k|k-1} c_y (y_k - c_y m_{k|k-1})/S_k$. The gain is large when the prediction is uncertain relative to the measurement and small when the measurement is noisy.

Example 21.7 (Filtering the pump vibration)**.** Let the bearing of Example 21.6 have been estimated at the previous inspection as $D \sim \mathcal{N}(0.40, 0.05^2)$, and let the vibration of $y = 4.0\,\mathrm{mm/s}$ be measured $\Delta t = 10$ days later. Take the Wiener parameters of Example 21.4, $c_y = 5\,\mathrm{mm/s}$ per unit of damage, and $\sigma_y = 1.0\,\mathrm{mm/s}$ to represent the effect of speed, load and mounting on the reading (illustrative). The prediction is

$$m_{k|k-1} = 0.40 + 0.01 \times 10 = 0.50, \qquad v_{k|k-1} = 0.05^2 + 0.03^2 \times 10 = 0.0025 + 0.0090 = 0.0115,$$

a standard deviation of 0.107. The predicted measurement is $5 \times 0.50 = 2.5\,\mathrm{mm/s}$, so the innovation is $e_k = 4.0 - 2.5 = 1.5\,\mathrm{mm/s}$, with $S_k = 25 \times 0.0115 + 1 = 1.2875\,(\mathrm{mm/s})^2$. The gain is $G_k = 0.0115 \times 5/1.2875 = 0.0447$ per mm/s, and

$$m_k = 0.50 + 0.0447 \times 1.5 = 0.567, \qquad v_k = (1 - 0.0447 \times 5) \times 0.0115 = 0.00893,$$

a posterior standard deviation of 0.095. The check by equation (21.50) gives $1/v_k = 1/0.0115 + 25/1 = 86.96 + 25 = 111.96$, the same v_k. The posterior places the damage at 0.57 ± 0.09, not at the 0.80 of the naive inversion, which lies 2.8 predicted standard deviations above the prediction and is therefore discounted by the filter.

Figure 21.7 shows the four distributions of the example. The horizontal axis is the damage D and the vertical axis a probability density. The narrow gray curve is the estimate at the previous inspection; the dashed gold curve the prediction ten days later, moved to the right by the drift and widened by the process variance; the dotted red curve the likelihood of the measurement as a function of D, centered at the naive value 0.80 with standard deviation $\sigma_y/c_y = 0.20$ and normalized to unit area; and the solid blue curve the posterior, the normalized product of the prediction and the likelihood. The posterior lies between the prediction and the likelihood, closer to the narrower of the two, and is narrower than either.

The posterior then gives the failure probability of equation (21.48). Integrating the first-passage probability of equation (21.32) over the posterior $\mathcal{N}(0.567, 0.095^2)$ gives a probability of failure within 30 days of 0.303. Using the posterior mean alone, as if the state were known, gives 0.271; using the naive state 0.80 gives 0.828. The posterior uncertainty raises the probability of failure relative to the point estimate, because in this range the first-passage probability rises more than linearly with the damage, so that the states above the mean add more than the states below it remove; the misinterpretation of the measurement inflates it by a factor of almost three.

Principle 21.4. *Condition-based reliability conditions the remaining life on the measured state of one component. Measurements inform the state through an observation model; a Bayesian filter turns them into a distribution of the state, and the probability of failure is the first-passage probability averaged over that distribution.*

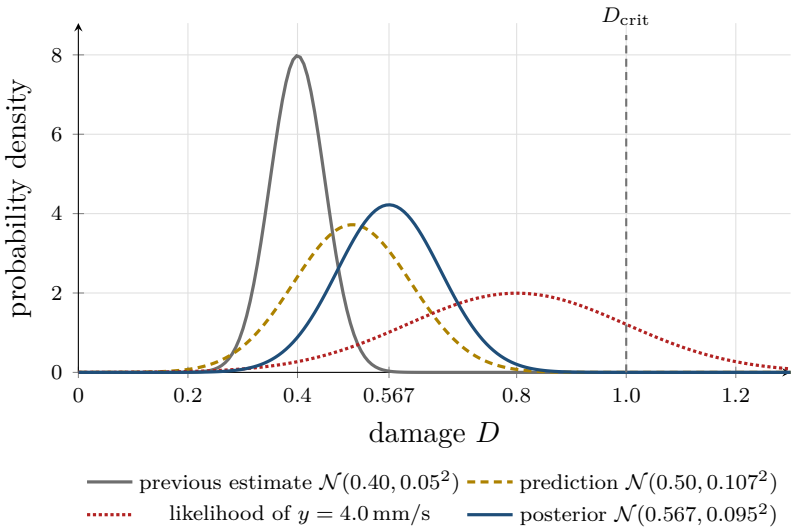


Figure 21.7: One step of the scalar Kalman filter for the bearing of the cooling-water pump (illustrative): the previous estimate, the prediction after 10 days of Wiener drift and diffusion ($\mu_D = 0.01\,\text{day}^{-1}$, $\sigma_D = 0.03\,\text{day}^{-1/2}$), the likelihood of a vibration reading of 4.0 mm/s under $y = 5D + \varepsilon$ with $\sigma_y = 1.0\,\text{mm/s}$, and the posterior.

21.6 Operating history, decisions and model credibility

A degradation model relates the operating history of a component to its condition, and so it can inform how equipment is operated as well as when it is maintained. This section shows the effect of operating history, works examples from the data center cooling plant and from battery storage, lists the common mistakes of degradation modeling and the questions that make a model credible, and ends with a worked table of remaining-life probabilities.

21.6.1 Operating history and equal chronological age

Consider two transmission components of the same chronological age. Component A has experienced mild loading and moderate temperatures; component B has experienced repeated high loads and severe thermal cycling. An age-only model assigns them the same failure probability. A physical degradation model instead accumulates the damage of each along its own history,

$$D_i(t) = \int_0^t g_D(\mathbf{u}_i(s), \mathbf{z}_i(s), \boldsymbol{\theta}_i)\,\mathrm{d}s, \qquad i \in \{\mathrm{A}, \mathrm{B}\}, \tag{21.52}$$

where each component has its own operating inputs $\mathbf{u}_i$, environmental conditions $\mathbf{z}_i$ and, possibly, parameters $\boldsymbol{\theta}_i$. Then $D_{\mathrm{A}}(t) \neq D_{\mathrm{B}}(t)$ even when the chronological ages are equal. This is one reason why condition-based reliability provides information that population lifetime distributions cannot: the distributions describe the average history of the population, and the integral describes the history of the component. It matters most for equipment whose degradation depends strongly on load and environmental exposure, such as overhead conductors, transformers and the pumps and chillers of a cooling plant.

21.6.2 Speed-dependent bearing damage in a data center cooling plant

The chilled-water pumps of a data center cooling plant run at variable speed. Suppose that the degradation rate of a pump's bearing depends on its rotational speed, with the illustrative

model

$$\frac{\mathrm{d}D}{\mathrm{d}t} = \kappa_{\mathrm{ref}} \left(\frac{\omega(t)}{\omega_{\mathrm{ref}}} \right)^{n}, \tag{21.53}$$

where $\omega(t)$ is the rotational speed, ω_{ref} a reference speed, κ_{ref} the damage rate at the reference speed, and n an empirically calibrated load sensitivity. This is the damage-rate law of Section 4.3.3 of Chapter 4 of Volume I with the stress measured by the relative speed. With $n = 3$, operation at 120 % of the reference speed gives

$$\frac{\dot{D}}{\dot{D}_{\mathrm{ref}}} = 1.2^{3} = 1.728,$$

so damage accumulates 72.8 % faster. At 80 % of the reference speed the ratio is $0.8^{3} = 0.512$. The exponent is not a universal law of pump bearings; it must be justified for the actual failure mechanism.

The convexity of the law has operational consequences. A pump that alternates equal periods at 80 % and 120 % of the reference speed accumulates damage at $(0.512 + 1.728)/2 = 1.12$ times the reference rate, 12 % more than a pump running steadily at the same average speed. And if a required flow, which under the pump affinity laws and neglecting static head is proportional to the speed of identical pumps in parallel, can be delivered either by two pumps at $1.2\,\omega_{\mathrm{ref}}$ or by three at $0.8\,\omega_{\mathrm{ref}}$ ($2 \times 1.2 = 3 \times 0.8$), the total bearing damage rate of the running pumps is $2 \times 1.728 = 3.46$ in the first case and $3 \times 0.512 = 1.54$ in the second, 2.25 times less, under this illustrative model. A degradation-aware operating strategy weighs

- energy efficiency;
- pump loading;
- redundancy;
- expected future damage;
- maintenance cost.

Running a third pump consumes the standby margin of a redundant plant and changes its energy use, which depends on the system curve; the damage model supplies one of the terms of that trade-off. This extends reliability engineering from the prediction of failure to the optimization of operation.

21.6.3 Battery capacity fade

For a battery energy storage system, a relevant degradation state is often the usable capacity $Q_{\mathrm{cap}}(t)$, in ampere-hours or kilowatt-hours. The capacity-loss degradation is

$$D(t) = 1 - \frac{Q_{\mathrm{cap}}(t)}{Q_{\mathrm{cap}}(0)}, \tag{21.54}$$

so that a 10 % capacity loss is $D = 0.10$. If end of life is defined as 80 % capacity retention, the threshold is $D_{\mathrm{crit}} = 0.20$. A simple empirical model adds a calendar term and a cycling term,

$$D(t) = a_{\mathrm{cal}} \sqrt{t} + b_{\mathrm{cyc}}\, n_{\mathrm{cyc}}(t), \tag{21.55}$$

where $a_{\mathrm{cal}}\sqrt{t}$ is an illustrative calendar-aging contribution, with a_{cal} in $\mathrm{yr}^{-1/2}$, and $b_{\mathrm{cyc}} n_{\mathrm{cyc}}(t)$ an illustrative cycling contribution, with $n_{\mathrm{cyc}}(t)$ the number of equivalent full cycles to time t and b_{cyc} the loss per cycle. With $a_{\mathrm{cal}} = 0.03\,\mathrm{yr}^{-1/2}$, $b_{\mathrm{cyc}} = 2 \times 10^{-5}$ per cycle and one cycle a day, $n_{\mathrm{cyc}}(t) = 365\,t$ with t in years (illustrative), the threshold is reached when $0.03\sqrt{t} + 0.0073\,t = 0.20$. This is a quadratic equation in $x = \sqrt{t}$, $0.0073\,x^{2} + 0.03\,x - 0.20 = 0$, whose positive root is

$$x = \frac{-0.03 + \sqrt{0.03^{2} + 4 \times 0.0073 \times 0.20}}{2 \times 0.0073} = \frac{-0.03 + 0.0821}{0.0146} = 3.568,$$

so $T_\mathrm{f} = 3.568^2 = 12.7$ years. With two cycles a day the cycling coefficient doubles to $0.0146\,\mathrm{yr}^{-1}$ and the life falls to 7.9 years, not to half of 12.7, because the calendar term does not depend on the cycling.

Real battery degradation depends on the chemistry, the temperature, the state of charge, the depth of discharge, the current and the interactions between the calendar and the cycling mechanisms; the square-root and linear terms are the simplest empirical shapes, not a law. Also, 80 % capacity retention is only one possible end-of-life criterion. A battery may become unsuitable for its application earlier because of resistance growth, power limits, thermal constraints or operational requirements: a frequency-regulation service that needs a stated power for a stated duration fails when the battery can no longer deliver it, whatever its capacity. The failure threshold must correspond to the required function.

21.6.4 Model credibility and common mistakes

Degradation models fail in recognizable ways. Table 21.4 lists the common mistakes in the left column and, in the right, why each matters; most of them have appeared in the examples of this chapter.

Table 21.4: Common mistakes in degradation modeling.

Mistake	Why it matters
Treating a sensor indicator as direct damage	measurements depend on load, environment and other fault mechanisms (Example 21.6)
Assuming all damage is linear	fatigue cracks, corrosion and thermal aging may be nonlinear (Figure 21.2)
Using a Wiener process for necessarily monotone physical damage	its trajectories can decrease
Treating expected degradation as a deterministic failure time	stochastic paths cross the threshold earlier or later; at the deterministic time the failure probability already exceeds one half
Ignoring uncertain failure thresholds	the critical condition varies across components
Assuming model parameters never change	degradation rates depend on the operating regime and on maintenance
Extrapolating beyond observed damage	a model calibrated during early wear may fail near end of life
Equating physical damage with functional failure	the required system function may depend on redundancy and operating conditions

A credible prognostics model addresses three separate questions. *Model adequacy*: does the chosen evolution law capture the relevant degradation mechanism? *Identifiability*: are the available measurements sufficient to estimate the damage states and the parameters? The filter of Section 21.5.3 cannot separate the damage from the sensitivity c_y if both are unknown and only y is measured. *Predictive validation*: does the model produce calibrated and useful predictions of the future condition and of failure, checked against outcomes it was not fitted to? A physically plausible equation alone does not establish that its predictions are trustworthy. Guidance on condition-based maintenance accordingly recommends combining several diagnostic technologies, and using the observed condition of equipment to support maintenance decisions, rather than relying on one indicator for every failure mechanism.

21.6.5 A worked table of remaining-life probabilities

The deterministic damage-to-threshold model predicts the time until the threshold is crossed from the current damage $D(t_0)$, the threshold D_crit, a reference degradation rate and an

operating-stress acceleration factor that multiplies it:

$$\text{time to threshold} = \frac{D_{\text{crit}} - D(t_0)}{\mu_D}, \qquad \mu_D = (\text{reference rate}) \times (\text{acceleration factor}).$$

It assumes a constant future degradation rate. The Wiener model with the same drift adds the random variation of the path and gives the probability that the remaining life is shorter than a horizon τ, from equation (21.32) with $\delta = D_{\text{crit}} - D(t_0)$. Table 21.5 evaluates both for the bearing model, with $D_{\text{crit}} = 1$, $\sigma_D = 0.03\,\text{day}^{-1/2}$, a reference rate of $0.01\,\text{day}^{-1}$, acceleration factors of one and two ($\mu_D = 0.01$ and $0.02\,\text{day}^{-1}$), current damage levels from 0.2 to 0.8 and horizons of 10, 30 and 60 days. The table depends on the current damage and the threshold only through their difference δ, so a lower threshold is read from the row of a higher damage: $D(t_0) = 0.40$ with $D_{\text{crit}} = 0.80$ is the row $D(t_0) = 0.60$.

For the row $D(t_0) = 0.40$, $\mu_D = 0.01\,\text{day}^{-1}$ and $\tau = 30$ days: $\delta = 0.60$, $\sigma_D\sqrt{\tau} = 0.03\sqrt{30} = 0.1643$, the first term of equation (21.32) is $\Phi\big((0.30 - 0.60)/0.1643\big) = \Phi(-1.826) = 0.0339$, and the second is $\exp(2 \times 0.01 \times 0.60/0.0009)\,\Phi(-0.90/0.1643) = \mathrm{e}^{13.33}\,\Phi(-5.477) = 6.17 \times 10^5 \times 2.16 \times 10^{-8} = 0.0133$, which sum to 0.047.

Table 21.5: Remaining-life prediction for the Wiener bearing model with $D_{\text{crit}} = 1$ and $\sigma_D = 0.03\,\text{day}^{-1/2}$ (illustrative): deterministic time to threshold δ/μ_D, standard deviation of the remaining life from equation (21.35), and probability of failure within τ days from equation (21.32).

$D(t_0)$	μ_D (1/day)	δ/μ_D (days)	SD (days)	Pr(RUL $\leq$ 10)	Pr(RUL $\leq$ 30)	Pr(RUL $\leq$ 60)
0.2	0.01	80	26.8	$< 10^{-6}$	0.0017	0.239
0.2	0.02	40	9.49	$< 10^{-6}$	0.134	0.968
0.4	0.01	60	23.2	$< 10^{-6}$	0.047	0.575
0.4	0.02	30	8.22	1.9×10^{-5}	0.554	0.997
0.6	0.01	40	19.0	0.0013	0.346	0.866
0.6	0.02	20	6.71	0.024	0.919	> 0.999
0.8	0.01	20	13.4	0.213	0.828	0.982
0.8	0.02	10	4.74	0.590	0.997	> 0.999

Figure 21.8 draws the deterministic projections of four rows. The horizontal axis is days from the present and the vertical axis the damage; each line starts at its current damage and rises at its drift until it meets the dashed threshold, where a dot marks the predicted crossing. Doubling the acceleration factor halves the predicted life under the linear law, from 60 to 30 days at $D(t_0) = 0.40$, and raising the current damage from 0.40 to 0.80 cuts it from 60 to 20 days, although chronological age appears nowhere in the prediction.

The probabilities of Table 21.5 add what the projections omit. At the deterministic crossing time the probability of failure is always above one half (0.554 for the row 0.40 at factor two, at 30 days), because the median of the inverse Gaussian distribution is below its mean. A faster drift shortens the life and also narrows its distribution, since the standard deviation $\sqrt{\delta\sigma_D^2/\mu_D^3}$ falls as $\mu_D^{-3/2}$: at $D(t_0) = 0.20$ the standard deviation is 26.8 days at factor one and 9.5 days at factor two. And a short horizon is safe even at high damage only when the drift is low: at $D(t_0) = 0.80$ the probability of failure within 10 days is 0.21 at factor one and 0.59 at factor two. The current physical state is more informative than age alone, but the inference of that state from measurements, which this table takes as known, is the subject of the next chapter on condition monitoring and fault diagnosis.

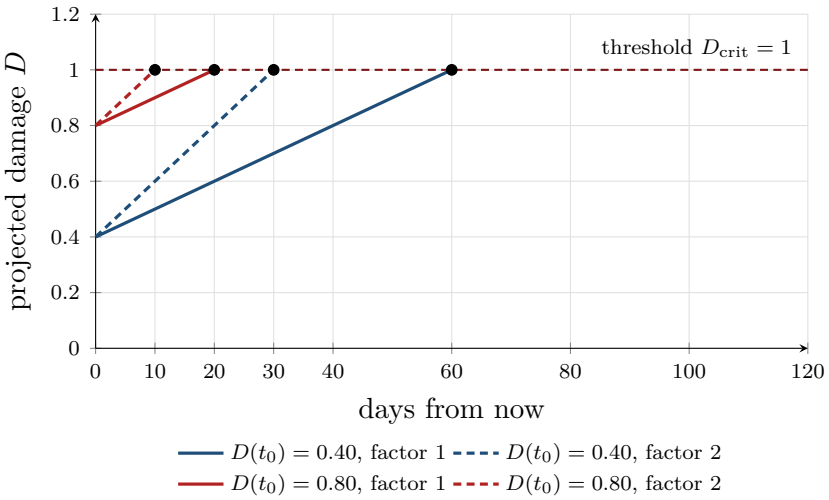


Figure 21.8: Deterministic damage projections, $D(t_0) + \mu_D t$ with $\mu_D = 0.01\,\text{day}^{-1}$ times an acceleration factor of one or two, for current damage 0.40 and 0.80 (illustrative). The predicted crossings of $D_{\text{crit}} = 1$ are at 60, 30, 20 and 10 days.

21.7 What to retain

Definition 21.1 (Degradation and damage accumulation)**.** A degradation model describes the physical state of a component as it evolves under its operating history and defines failure as the first passage of that state through a threshold tied to the required function; with random variation along the path and conditioning on measurements, it gives a distribution of the remaining life of each individual component rather than a single failure time for the population.

Table 21.6 lists the chapter's terms with their meaning, their expression and units, and the subsection that defines each.

Table 21.6: Terms of degradation and damage accumulation.

Term	Meaning	Expression, units	Defined in
Degradation state	measurable or inferred condition	$D(t)$, units of the quantity or normalized	21.1.1
Failure threshold	state at which the required function is lost	D_{crit}	21.1.1
First-passage time	first crossing of the threshold	$T_{\mathrm{f}} = \inf\{t : D(t) \geq D_{\mathrm{crit}}\}$	21.1.1
Degradation evolution law	rate of change of the state	$\mathrm{d}D/\mathrm{d}t = g_D(D, \mathbf{u}, \mathbf{z}, \boldsymbol{\theta})$	21.2.2
Linear degradation	constant rate	$D_0 + \kappa t$; $T_{\mathrm{f}} = (D_{\mathrm{crit}} - D_0)/\kappa$	21.2.3
Nonlinear degradation	rate grows with damage	$\mathrm{d}D/\mathrm{d}t = \kappa D^{n_D}$	21.2.4
Condition-dependent damage	damage integrated along the history	$D_0 + \int_0^t g_D(\mathbf{z}(u))\,\mathrm{d}u$	21.2.5
Basquin's law, S-N curve	fatigue life against stress	$\sigma_a = \sigma_f'(2N_f)^{b_{\mathrm{B}}}$; $N = K\Delta\sigma^{-b}$	21.3.1
Miner damage	sum of consumed life fractions	$D = \sum_i n_i/N_i$, failure at 1	21.3.2
Rainflow counting	cycles from a stress history	$\sigma(t) \to \{n_i, \Delta\sigma_i\}$	21.3.3
Remaining crack-growth life	cycles from a_0 to a_c	N_{rem}, eq. (21.20), cycles	21.3.4
Arrhenius acceleration factor	ratio of aging rates	$\exp[(E_a/R_{\mathrm{g}})(1/T_1 - 1/T_2)]$	21.3.6
Equivalent aging time	exposure at the reference temperature	$\int k(T(u))/k(T_{\mathrm{ref}})\,\mathrm{d}u$, h	21.3.7
Wiener degradation process	drift plus Brownian diffusion	$D_0 + \mu_D t + \sigma_D B(t)$; $\mathcal{N}(D_0 + \mu_D t, \sigma_D^2 t)$	21.4.2
Inverse Gaussian first passage	failure time of the Wiener model	mean δ/μ_D, variance $\delta\sigma_D^2/\mu_D^3$	21.4.4
Gamma degradation process	monotone, independent gamma increments	Gamma$(\alpha_D\Delta t, \beta_D)$; $R(t) = \gamma(\alpha_D t, \delta/\beta_D)/\Gamma(\alpha_D t)$	21.4.5
Remaining useful life	time left to the threshold	$\mathrm{RUL}(t_0) = T_{\mathrm{f}} - t_0$; quantile $\mathrm{RUL}_\alpha(t_0)$, h or days	21.5.1
Conditional reliability	survival given the measured condition	$\Pr[\mathrm{RUL}(t_0) > \tau \mid \mathbf{y}_{1:k}, T_{\mathrm{f}} > t_0]$	21.5.1
Observation model	measurement as a function of the state	$y_k = h(D(t_k), \mathbf{u}_k) + \varepsilon_k$	21.5.2
Condition monitoring, diagnosis, prognostics	observe, explain, predict	–	21.5.2
Bayesian filter, Kalman filter	recursive estimate of the state	prediction and update; gain G_k, innovation e_k	21.5.3

Three ideas carry forward. Degradation provides a physical explanation of failure, connecting the operating conditions to progressive damage and eventual loss of function, and fatigue, wear, corrosion, thermal aging and battery deterioration each require their own physical model. A deterministic degradation model predicts a threshold-crossing time but not a failure probability; stochastic models represent the variability of trajectories and failure times, and first-passage analysis is essential, because the probability of exceeding a threshold at one instant differs from the probability of having crossed it earlier. And the current condition can be more informative than chronological age, provided the measured signals are connected to the underlying damage through credible observation models, with physical plausibility, identifiable parameters, adequate measurements and predictive validation; in short, operating history

determines physical degradation, which determines the conditional failure probability.

Further reading. Suresh, *Fatigue of Materials* (fatigue crack initiation, propagation and life prediction). Dowling, *Mechanical Behavior of Materials* (stress-life, strain-life and fracture-mechanics approaches). *Degradation Data Analysis and Remaining Useful Life Estimation: Wiener-Process Methods* (a survey of Wiener-based degradation modeling and prognostics). *Condition-Based Maintenance Using the Inverse Gaussian Degradation Model* (stochastic degradation and maintenance decisions). NASA, *Predictive Testing, Inspection, and Condition-Based Maintenance* (condition-monitoring technologies and maintenance practice).

21.8 Exercises

Exercise 21.1. Answer each part in one or two sentences. (a) How is the failure time defined in a threshold degradation model? (b) What does the Miner rule calculate? (c) Why can a gamma process be preferable to a Wiener process for corrosion penetration? (d) What does a first-passage reliability calculation determine? (e) Why is a vibration reading not necessarily a direct measurement of bearing damage? (Answers: (a) as the first time the degradation reaches the critical threshold, the first-passage time of the degradation process; (b) the sum of the nominal fatigue-life fractions n_i/N_i consumed by the cycles of the different stress classes; (c) because its increments are nonnegative, so its paths are monotone, consistent with irreversible cumulative metal loss; (d) the probability that the threshold has not been crossed by a given time, which depends on the whole path remaining below the threshold; (e) because vibration also responds to load, speed, imbalance, alignment and other mechanisms, so an observation model is needed to connect the measurement to the latent physical condition.)

Exercise 21.2. (a) A pipe wall has a measured corrosion penetration of 0.90 mm, a penetration rate of 0.12 mm/yr and a maximum acceptable penetration of 2.0 mm. Compute the deterministic time to the threshold. (b) The connection of Example 21.2, at $D = 0.50$, next experiences 120 000 cycles at 50 MPa, 30 000 at 75 MPa and 15 000 at 100 MPa. Compute the Miner damage of this block and the fraction of the block, assuming its cycles are spread uniformly, at which $D = 1$ is reached. (c) With the Paris-law parameters of Section 21.3.4, compute the remaining cycles from a crack of 2 mm to 10 mm at 100 and at 120 MPa. (d) For $E_a = 80\,\mathrm{kJ/mol}$, compute the acceleration factor of 80°C relative to 60°C, and the equivalent aging per day at 60°C of a transformer that spends 20 h at 60°C and 4 h at 80°C. (Answers: (a) 9.17 years; (b) $0.12 + 0.15 + 0.30 = 0.57$, reached at $0.50/0.57 = 88\,\%$ of the block; (c) 3.16×10^5 and 1.83×10^5 cycles; (d) AF = 5.13 and $20 + 4 \times 5.13 = 40.5$ h per day, 1.69 times the calendar time.)

Exercise 21.3. A degradation process follows the Wiener model with $D_0 = 0.20$, $D_{\mathrm{crit}} = 1$, $\mu_D = 0.005\,\mathrm{day}^{-1}$ and $\sigma_D = 0.02\,\mathrm{day}^{-1/2}$. (a) Compute the mean and the standard deviation of the first-passage time. (b) At the mean first-passage time, compute $\Pr(D(t) \geq D_{\mathrm{crit}})$ and the first-passage probability $\Pr(T_{\mathrm{f}} \leq t)$, and explain the difference. (c) Compute the reliability at 100 days. (d) A pipe has 1.40 mm of penetration left before its threshold and corrodes as a gamma process with shape rate $\alpha_D = 0.5\,\mathrm{yr}^{-1}$ and scale $\beta_D = 0.16$ mm. Compute its mean corrosion rate and, with equation (21.41), the probability that it reaches the threshold within 16 years. (Answers: (a) 160 days and $\sqrt{0.8 \times 0.0004/0.005^3} = 50.6$ days; (b) 0.5 and $0.5 + \mathrm{e}^{20}\,\Phi(-6.325) = 0.5 + 4.85 \times 10^8 \times 1.27 \times 10^{-10} = 0.562$, the difference being the paths that crossed and returned below the threshold; (c) $R(100) = 1 - 0.086 = 0.914$; (d) 0.08 mm/yr, and with shape 8 and $\delta/\beta_D = 8.75$, $\sum_{j=0}^{7} \mathrm{e}^{-8.75}\, 8.75^j/j! = 0.354$.)

Exercise 21.4. The bearing of a cooling-water pump was estimated at the last inspection as $D \sim \mathcal{N}(0.30, 0.04^2)$. Five days later a vibration of $y = 2.2$ mm/s is measured. The Wiener parameters are $\mu_D = 0.01\,\mathrm{day}^{-1}$ and $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$; the observation model is $y = c_y D + \varepsilon$

with $c_y = 5\,\mathrm{mm/s}$ and $\sigma_y = 0.5\,\mathrm{mm/s}$ (illustrative). (a) Compute the predicted mean and variance of the damage. (b) Compute the innovation, the Kalman gain and the posterior mean and standard deviation, and compare the posterior mean with the naive inversion y/c_y. (c) Treating the posterior mean as the known damage, compute the expected remaining life and the probability of failure within 30 days. (d) Explain why the probability in (c) understates the probability of failure computed by integrating over the posterior. (Answers: (a) 0.35 and $0.0016 + 0.0045 = 0.0061$; (b) $e = 2.2 - 1.75 = 0.45\,\mathrm{mm/s}$, $S = 0.4025$, $G = 0.0758$, posterior mean 0.384 and standard deviation 0.062, against 0.44 by inversion; (c) $(1 - 0.384)/0.01 = 61.6$ days and 0.038; (d) the first-passage probability is convex in the damage over this range, so averaging it over the posterior spread gives more than its value at the mean, here 0.049.)

Chapter 22

Condition Monitoring and Fault Diagnosis

22.1 Detection, diagnosis and prognosis

Chapter 21 modeled the deterioration of equipment as a physical state, the degradation $D(t)$, that grows through wear, fatigue, thermal aging, corrosion and the other mechanisms of Chapter 4 of Volume I. That state is rarely observed directly. What is measured are temperatures, currents, pressures, vibrations, voltages and flows, and each of them responds not only to damage but also to operating conditions, control actions, sensor errors and the environment. A cooling pump that draws more electrical power may have a mechanical fault; it may also be delivering more flow, pumping against a higher pressure, or reporting its power through an inaccurate meter. The question of this chapter is

how it is determined whether equipment is malfunctioning, which physical cause is the most likely, and how actual degradation is distinguished from sensor errors and from changing operating conditions.

The chapter answers it in five steps. It builds a physical model of the healthy pump and turns the difference between measurement and model into residuals with a known uncertainty and a statistical test. It extracts condition indicators from raw signals in the time and frequency domains and accumulates small persistent changes. It works a complete example in which an excess of electrical power is first detected and then diagnosed by Bayes' theorem over three competing hypotheses, one of them a biased sensor. It examines when the measurements can distinguish the hypotheses at all, and which further measurement is worth taking. Finally, it states how a diagnostic system is validated and how its result enters reliability and maintenance decisions.

Jargon: condition monitoring, condition indicator, residual, fault detection, fault isolation, FDI, fault diagnosis, prognostics, PHM, false alarm, missed detection, sensor bias, analytical redundancy, bad-data detection

Condition monitoring is the observation and summary of the condition of equipment through measurements, operating data and quantities derived from them. A *condition indicator* (also called a feature) is such a derived quantity whose value changes with the condition of interest, for example a vibration amplitude or an apparent efficiency. A *residual* e is the difference between an observed quantity and the value a model of the healthy system predicts for it. *Fault detection* decides whether the observed behavior is inconsistent with normal operation; *fault isolation* decides which component or fault is responsible; *fault identification* estimates the size of the fault. Detection and isolation together are *fault detection and isolation* (FDI); *fault diagnosis* covers isolation and identification, and in this chapter means the comparison of competing explanations of the evidence. *Prognostics* estimates future condition, failure probability or remaining useful life, and *prognostics and health management* (PHM) is the discipline that combines monitoring, diagnosis, prognosis and the maintenance decisions they support. A *false alarm* is a detection declared when no fault is present; a *missed detection* is a fault present and not detected. A *sensor bias* is a persistent offset between a sensor's reading and the true value. *Analytical redundancy* is the use of a model, such as an energy balance, to check measurements against one another instead of duplicating sensors; in power-system state estimation the resulting check is called *bad-data detection*.

Key idea

Detection tests whether the measurements are consistent with a model of the healthy system; diagnosis compares competing explanations, physical faults and measurement faults alike, by how well each predicts the measurements and how plausible it was beforehand. Neither can resolve explanations whose effects on the measurements are indistinguishable.

This section explains why degradation must be inferred from indirect measurements, separates the four functions of monitoring, detection, diagnosis and prognosis, and introduces the variable-speed cooling-water pump on which the chapter's methods are worked.

22.1.1 From degradation to measurements

The degradation state $D(t)$ of Chapter 21 is a property of the physical item: the depth of wear on a bearing race, the loss of material from an impeller, the growth of a crack. (The symbol D denotes degradation in this volume; it is not the FMEA detection rating D of Chapter 4.) A condition monitoring system sees the item through sensors, and every sensor reading is the joint outcome of several influences: the condition of the item, the operating point at which it runs, the control actions applied to it, the environment, and the behavior of the sensor itself. The distinction between the physical system, the model of it and the evidence is therefore explicit from the start. The pump and its damage are the system; the equations that predict what a healthy pump would show at the present operating point are the model; the sensor readings, with their own errors, are the evidence.

Electrical power illustrates the difficulty. A rise in the power drawn by a cooling pump is consistent with a mechanical fault that wastes energy in friction, with an increase in delivered flow, with an increase in the pressure the pump works against, and with a power meter that reads high. The reading alone cannot separate these explanations. Condition monitoring detects that something is unusual; fault diagnosis determines which explanations the evidence supports. The distinction is the foundation of prognostics and health management, since a prognosis built on a wrong diagnosis predicts the evolution of damage that is not there.

22.1.2 Monitoring, detection, diagnosis and prognosis

The functions that lead from data to a prediction are usually separated into four, each answering a different question. Figure 22.1 draws them in order. The four boxes, left to right, are condition monitoring, fault detection, fault diagnosis and prognosis; the text under each box states its task and gives an example for a pump; the arrow at the bottom indicates that each function needs more information and a more complete model than the one before it.

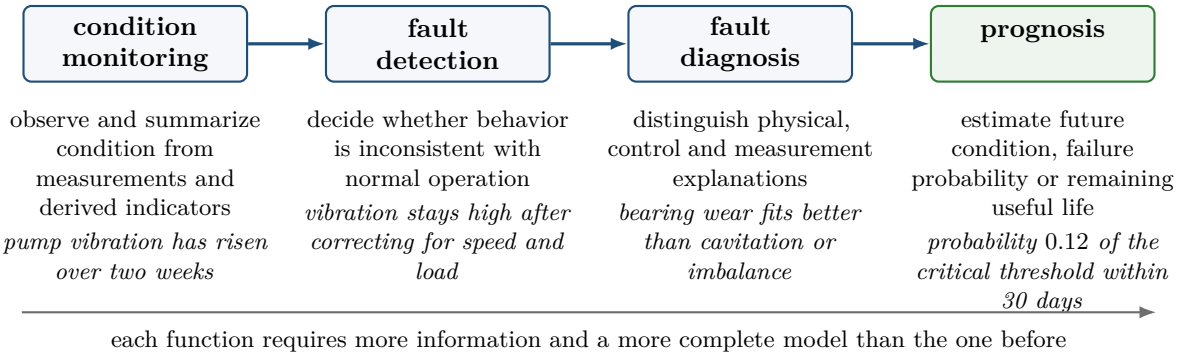


Figure 22.1: The four functions of condition-based maintenance, with an example for a pump under each (illustrative). An unusual signal may suffice for detection; diagnosis needs a model that distinguishes candidate causes; prognosis needs a model of how the diagnosed damage evolves.

Condition monitoring observes and summarizes, for example by reporting that pump vibration has increased over the past two weeks. *Fault detection* decides that the observed behavior is inconsistent with acceptable normal operation, for example that vibration remains unusually high after its dependence on rotational speed and load has been accounted for. *Fault diagnosis* distinguishes among competing physical, control and measurement explanations, for example that bearing wear is more consistent with the evidence than cavitation or imbalance. *Prognosis* estimates future condition, failure probability or remaining useful life, for example that, given the inferred wear state, the probability of reaching a critical threshold within 30 days is 0.12; that function is the subject of Chapter 23.

The four functions require progressively more information. An unusual signal may be enough to detect an anomaly. Identifying the fault behind it requires a model able to distinguish the candidate causes, that is, one in which different causes predict different measurements. Predicting the remaining life requires, in addition, a credible model of how the diagnosed damage evolves.

22.1.3 The cooling-water pump example

The methods of the chapter are worked on the cooling-water pump that recurs in this book, here in the form it takes in a data center cooling-water loop: a variable-speed centrifugal pump driven by an electric motor through a variable-speed drive. Figure 22.2 draws it with its instruments. The drive feeds the motor from the left; the motor turns the pump through a shaft whose bearing housing carries a vibration sensor and a temperature sensor; water enters the pump from the return header on the right and leaves upward to the cooling coils. The circles are the instruments, labeled with the symbols of Table 22.1.

The abnormal conditions of interest are bearing wear, cavitation, increased hydraulic resistance in the loop, reduced pump efficiency and sensor bias. The last belongs on the list on equal terms with the others: Section 4.5.7 of Chapter 4 showed that a flow sensor reading $20\,\mathrm{m}^3/\mathrm{h}$ high hides a real flow deficiency, and the same failure mode of a power meter can imitate

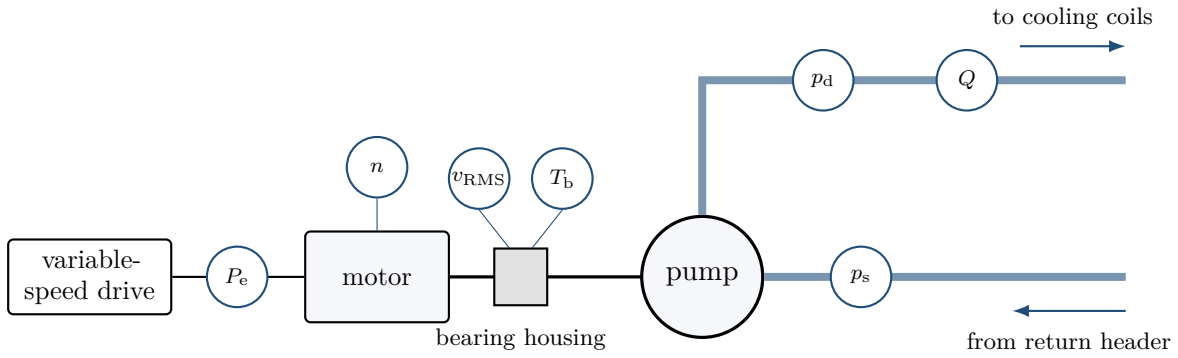


Figure 22.2: The variable-speed cooling-water pump and its instruments: electrical input power P_e, rotational speed n, bearing vibration v_{RMS}, bearing temperature T_b, suction pressure p_s, discharge pressure p_d and volumetric flow Q (schematic).

Table 22.1: Measurements available on the variable-speed cooling-water pump.

Measurement	Symbol	Typical units
Motor rotational speed	n	rpm
Electrical input power	P_e	kW
Volumetric flow	Q	m^3/s
Suction pressure	p_s	kPa
Discharge pressure	p_d	kPa
Bearing vibration velocity	v_{RMS}	mm/s
Bearing temperature	T_b	°C

a real fault. Before any hypothesis is evaluated, the physical meaning of the measurements must be established, which is the subject of the next section.

22.2 Physical models and condition indicators

A measurement becomes evidence about condition only when it is compared with what a healthy item would show under the same operating conditions. This section derives the energy relation of the pump, defines residuals and standardizes them by their uncertainty, and derives the chi-square test that combines several residuals into one detection statistic.

22.2.1 The pump energy relation

For incompressible flow, the hydraulic power transferred to the water is the product of the volumetric flow and the pressure rise across the pump,

$$P_h = Q\,\Delta p, \qquad \Delta p = p_d - p_s, \tag{22.1}$$

with units $m^3/s \cdot Pa = m^3/s \cdot N/m^2 = N\,m/s = W$. The pressure rise is often expressed as the pump's total head H_p, the height of a water column with the same pressure, $\Delta p = \rho_w\, g\, H_p$, where ρ_w is the density of water (about $1000\,kg/m^3$), $g = 9.81\,m/s^2$ is the gravitational acceleration and H_p is in meters (the head is not the operating horizon H). In that form $P_h = \rho_w\, g\, Q\, H_p$, with units $kg/m^3 \cdot m/s^2 \cdot m^3/s \cdot m = kg\,m^2/s^3 = W$.

The overall wire-to-water efficiency η is the ratio of the hydraulic power delivered to the electrical power drawn by the drive, motor and pump together, $\eta = P_h/P_e$ (here η is an efficiency,

not the Weibull scale). Solving for the electrical power,

$$P_{\mathrm{e}} = \frac{Q\,\Delta p}{\eta} = \frac{\rho_{\mathrm{w}}\, g\, Q\, H_{\mathrm{p}}}{\eta}. \tag{22.2}$$

The relation is a simplified energy balance, useful when the kinetic and elevation corrections between the suction and discharge pressure taps are small or handled consistently; a complete pump energy balance uses the total head, including the velocity and elevation terms.

Equation (22.2) shows why electrical power is an ambiguous condition indicator by itself. It rises when the flow Q rises, when the pressure rise Δp rises, or when the efficiency η falls. Only the third is deterioration of the pump; the first two are changes of the operating point imposed by the loop and its controls. A pump drawing more power is not necessarily deteriorating, and power becomes a useful indicator only after it is conditioned on the flow and pressure rise at which it was measured.

22.2.2 Residuals: observed against expected behavior

A model-based diagnostic system predicts each measurement under the assumption that the system is healthy. Let k index the sampling instants, $\mathbf{x}_k$ be the modeled physical state, $\mathbf{u}_k$ the known operating inputs (speed, flow set point, valve positions) and $\boldsymbol{\theta}_0$ the parameters of the healthy system, such as the healthy efficiency. The predicted measurement is

$$\hat{y}_k = h(\mathbf{x}_k, \mathbf{u}_k, \boldsymbol{\theta}_0), \tag{22.3}$$

where h is the measurement model, for the power channel equation (22.2) evaluated at the measured flow and pressure rise. The *residual* is the observation minus the prediction,

$$e_k = y_k - \hat{y}_k, \tag{22.4}$$

in the units of the measurement. It measures the discrepancy between the observation and what the healthy model expects; under the healthy model and with an unbiased prediction its mean is zero.

If the expected motor power is 20 kW and the measured power is 24 kW, the power residual is

$$e_P = 24\,\mathrm{kW} - 20\,\mathrm{kW} = 4\,\mathrm{kW}.$$

Whether 4 kW is significant cannot be decided from the number alone. It depends on how large a residual the healthy system produces by itself, through sensor noise, uncertainty in the healthy parameters, uncertainty in the operating inputs and the imperfection of the model. A discrepancy cannot be interpreted without an uncertainty model.

22.2.3 Standardized residuals

Suppose that under healthy operation the residual is approximately Gaussian with mean zero and standard deviation σ_e, $e \sim \mathcal{N}(0, \sigma_e^2)$. The *standardized residual* is the residual in units of its own standard deviation,

$$z = \frac{e}{\sigma_e}, \tag{22.5}$$

dimensionless, and standard normal, $z \sim \mathcal{N}(0, 1)$, under the healthy model. With $e_P = 4\,\mathrm{kW}$ and $\sigma_e = 1\,\mathrm{kW}$, $z = 4$: the observation lies four residual standard deviations from the healthy prediction, and a healthy pump produces a residual at least that extreme in either direction with probability $2[1 - \Phi(4)] = 6.3 \times 10^{-5}$. With $\sigma_e = 5\,\mathrm{kW}$ the same discrepancy gives $z = 0.8$, which a healthy pump exceeds in magnitude with probability 0.42: the same 4 kW is unremarkable. Residual magnitude alone does not establish the significance of a fault.

The standard deviation σ_e must collect every source of healthy discrepancy: the noise of the sensor, the uncertainty of the healthy parameters $\boldsymbol{\theta}_0$, the uncertainty of the operating inputs from which the prediction is computed, and the discrepancy between the model and the physics it simplifies. A σ_e that includes only sensor noise makes every model imperfection look like a fault.

Principle 22.1. *A residual is evidence of a fault only in relation to its healthy uncertainty. The standardized residual $z = e/\sigma_e$, with σ_e including model and input uncertainty as well as sensor noise, is the quantity a detection rule compares with a threshold.*

22.2.4 Residual covariance and the multivariate chi-square test

Several quantities are usually monitored together. Take three: the electrical power, the pressure rise and the bearing vibration, with observation and prediction vectors

$$\mathbf{y} = \begin{bmatrix} P_{\mathrm{e}} \\ \Delta p \\ v_{\mathrm{RMS}} \end{bmatrix}, \qquad \mathbf{e} = \mathbf{y} - \hat{\mathbf{y}},$$

where the predicted pressure rise comes from the pump's characteristic curve at the measured speed and flow. Under the healthy model the residual vector has mean zero and covariance matrix

$$\boldsymbol{\Sigma}_e = \mathbb{E}\big[\mathbf{e}\,\mathbf{e}^{\mathsf{T}}\big], \tag{22.6}$$

whose diagonal holds the variances $\sigma_{e,j}^2$ of the individual residuals and whose off-diagonal entries are their covariances. Residuals are correlated when they share an error source: an error in the measured flow, for example, enters the predicted power and the predicted pressure rise together.

Assume $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_e)$ with a known positive-definite $\boldsymbol{\Sigma}_e$ of dimension $m \times m$. The common detection statistic is the squared Mahalanobis distance of the residual vector from zero,

$$J = \mathbf{e}^{\mathsf{T}} \boldsymbol{\Sigma}_e^{-1}\, \mathbf{e}, \tag{22.7}$$

dimensionless. Its distribution follows in three steps. First, a positive-definite matrix has a Cholesky factorization $\boldsymbol{\Sigma}_e = \mathbf{A}\mathbf{A}^{\mathsf{T}}$ with $\mathbf{A}$ lower triangular and invertible. Second, the whitened residual $\mathbf{w} = \mathbf{A}^{-1}\mathbf{e}$ is Gaussian with mean zero and covariance $\mathbf{A}^{-1}\boldsymbol{\Sigma}_e\mathbf{A}^{-\mathsf{T}} = \mathbf{A}^{-1}\mathbf{A}\mathbf{A}^{\mathsf{T}}\mathbf{A}^{-\mathsf{T}} = \mathbf{I}$; its m components are uncorrelated standard normal variables, and being jointly Gaussian they are independent. Third, since $\boldsymbol{\Sigma}_e^{-1} = \mathbf{A}^{-\mathsf{T}}\mathbf{A}^{-1}$,

$$J = \mathbf{e}^{\mathsf{T}}\mathbf{A}^{-\mathsf{T}}\mathbf{A}^{-1}\mathbf{e} = \mathbf{w}^{\mathsf{T}}\mathbf{w} = \sum_{j=1}^{m} w_j^2,$$

a sum of m squared independent standard normal variables, which is by definition the chi-square distribution with m degrees of freedom:

$$J \sim \chi^2_m \quad \text{under the healthy model.} \tag{22.8}$$

The result requires a known positive-definite covariance; the residuals themselves may be correlated, because the whitening removes the correlation. When $\boldsymbol{\Sigma}_e$ is diagonal, $J = \sum_j z_j^2$ is the sum of the squared standardized residuals.

For three monitored quantities $J \sim \chi^2_3$, and the rule

$$J > \chi^2_{3,0.99} = 11.34 \tag{22.9}$$

declares an anomaly with a nominal false-alarm probability of 1 % per test, $\chi^2_{\nu,p}$ being the p quantile of the chi-square distribution with ν degrees of freedom. As an illustrative case, let the residuals be $e_P = 4\,\mathrm{kW}$, $e_{\Delta p} = 5\,\mathrm{kPa}$ and $e_v = 0.3\,\mathrm{mm/s}$, with healthy standard deviations 1.5 kW, 10 kPa and 1.0 mm/s. With uncorrelated residuals,

$$J = \left(\frac{4}{1.5}\right)^2 + \left(\frac{5}{10}\right)^2 + \left(\frac{0.3}{1.0}\right)^2 = 7.111 + 0.250 + 0.090 = 7.451,$$

below the threshold, with $\Pr(\chi^2_3 > 7.451) = 0.059$. If the power and pressure-rise residuals have correlation coefficient 0.6, because both depend on the measured flow, the off-diagonal entry is $0.6 \times 1.5 \times 10 = 9\,\mathrm{kW\,kPa}$ and, with the standardized residuals $z_1 = 2.667$ and $z_2 = 0.5$,

$$J = \frac{z_1^2 - 2(0.6) z_1 z_2 + z_2^2}{1 - 0.6^2} + z_3^2 = \frac{7.111 - 1.600 + 0.250}{0.64} + 0.090 = 9.092,$$

still below 11.34 but larger: a high power residual is more unusual when the pressure-rise residual, which normally moves with it, does not follow. The covariance is part of the healthy model, and getting it wrong changes the test.

Figure 22.3 draws the test. The horizontal axis is the statistic J and the vertical axis the χ^2_3 density, $\sqrt{J/(2\pi)}\,\mathrm{e}^{-J/2}$; the shaded tail to the right of the threshold has area 0.01, and the two vertical lines are the uncorrelated and correlated values of the example. Both lie in the body of the healthy distribution.

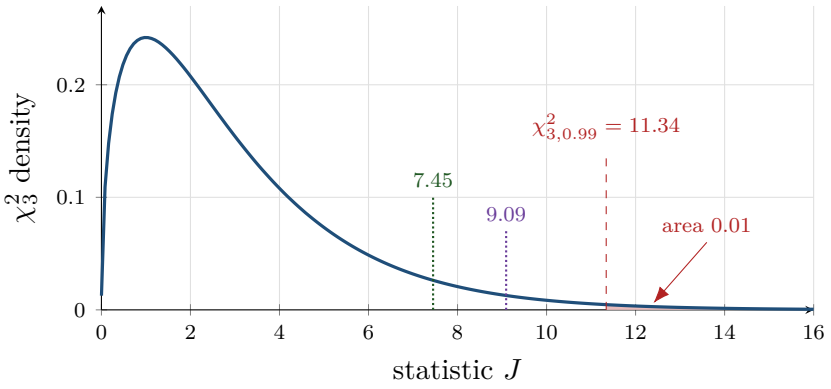


Figure 22.3: The multivariate chi-square test for three residuals: density of χ^2_3, $\sqrt{J/(2\pi)}\,\mathrm{e}^{-J/2}$, the 1 % threshold 11.34 (shaded tail) and the statistic of the illustrative example with uncorrelated residuals (7.45) and with correlation 0.6 between the power and pressure-rise residuals (9.09).

The nominal 1 % holds per test and under idealized conditions. A test repeated once a minute makes 1440 decisions a day and, if the tests were independent, would raise on average 14.4 false alarms a day on a healthy pump. Real residuals are also autocorrelated in time, their covariance is estimated rather than known, and the healthy model is never exact; each of these changes the actual false-alarm rate, and the rate must be verified on data, as Section 22.6 describes. The next section turns from model-based residuals to the indicators that signal processing extracts from the raw measurements.

22.3 Signal processing for condition monitoring

Raw sensor records are rarely usable as condition indicators as they stand. This section explains why, defines the standard time-domain features, treats frequency-domain analysis and the characteristic frequencies of rotating machinery, and derives the two recursive statistics, the exponentially weighted moving average and the cumulative sum, that detect gradual and persistent changes.

22.3.1 Why raw data are insufficient

The pump's sensors produce data of very different kinds. An accelerometer on the bearing housing may record thousands of samples per second; the bearing temperature may be reported once a minute; a supervisory control and data acquisition (SCADA) system may deliver electrical power or pressure every few seconds. These records differ in sampling rate, in noise and in physical meaning, and most of what they contain is unrelated to the condition of the pump: the operating point, the control actions and the noise of acquisition. Signal processing extracts from them the quantities that do respond to the condition of interest. A typical sequence is

$$\text{raw signal} \;\to\; \text{quality checks} \;\to\; \text{operating-state normalization}$$
$$\to\; \text{feature extraction} \;\to\; \text{detection}.$$

Quality checks remove missing, frozen, out-of-range and misaligned samples. Operating-state normalization removes the dependence of a measurement on the operating point, as equation (22.2) does for power, or by comparing vibration only within bands of speed and load. Feature extraction reduces a record to a few condition indicators, and detection applies a statistical rule to them.

22.3.2 Time-domain features

For a sampled vibration signal $x_1, x_2, \ldots, x_N$, for example vibration velocity in mm/s over a record of N samples, the mean is $\bar{x} = \frac{1}{N}\sum_{k=1}^{N} x_k$. Three features are standard. The *root-mean-square (RMS) amplitude*

$$x_{\mathrm{RMS}} = \sqrt{\frac{1}{N}\sum_{k=1}^{N} x_k^2}, \tag{22.10}$$

in the units of the signal, measures the overall energy of the signal and is the indicator most often trended and compared with severity limits. The *crest factor*

$$\mathrm{CF} = \frac{\max_k |x_k|}{x_{\mathrm{RMS}}}, \tag{22.11}$$

dimensionless, compares the largest excursion with the RMS level; it is $\sqrt{2} = 1.414$ for a pure sine wave, and a large value indicates impulsive behavior. The *kurtosis*

$$\mathrm{Ku} = \frac{\frac{1}{N}\sum_{k=1}^{N}(x_k - \bar{x})^4}{\left[\frac{1}{N}\sum_{k=1}^{N}(x_k - \bar{x})^2\right]^2}, \tag{22.12}$$

dimensionless, is the normalized fourth central moment. It equals 3 for Gaussian noise and 1.5 for a sine wave, and it rises when a few large excursions dominate the fourth power, which is how the impacts of a rolling element over a spall on a bearing race appear in a signal.

Table 22.2 computes the three features for three illustrative records of $N = 100$ samples. The first is a periodic signal sampled at its peaks, alternating $+1$ and -1 mm/s; the second replaces one consecutive pair of samples by $+5$ and -5 mm/s, an impact and its rebound; the third has two such pairs. One row of the arithmetic, for the single impact: the squares sum to $98 \times 1 + 2 \times 25 = 148$, so $x_{\mathrm{RMS}} = \sqrt{1.48} = 1.217$ mm/s and $\mathrm{CF} = 5/1.217 = 4.11$; the mean is zero and the fourth powers sum to $98 + 2 \times 625 = 1348$, so $\mathrm{Ku} = 13.48/1.48^2 = 6.15$.

The single impact raises the RMS by 22 % but the crest factor fourfold and the kurtosis sixfold: impulsive features respond to a few impacts that barely move the energy. With a second impact the RMS keeps rising while the crest factor falls, because the maximum is unchanged and the RMS has grown. The features measure different properties and none of them is specific to one mechanism: a large crest factor may come from a bearing defect, but also from cavitation, from looseness or from an electrical spike in the acquisition chain. No single feature identifies every failure mechanism.

Table 22.2: Time-domain features of three illustrative vibration records of 100 samples: a periodic signal of amplitude 1 mm/s, the same with one impact pair of amplitude 5 mm/s, and with two.

Record	x_{RMS} (mm/s)	CF	Ku
Periodic, ±1	1.000	1.00	1.00
One impact pair, ±5	1.217	4.11	6.15
Two impact pairs, ±5	1.400	3.57	6.76

22.3.3 Frequency-domain analysis

Rotating machinery produces vibration at frequencies set by its speed and geometry, and the frequency at which energy appears often says more than the amount. The discrete Fourier transform of N samples taken at the sampling frequency f_{s} is

$$X_j = \sum_{k=0}^{N-1} x_k \, \mathrm{e}^{-\mathrm{i}2\pi jk/N}, \qquad j = 0, 1, \ldots, N-1, \tag{22.13}$$

where i is the imaginary unit; coefficient X_j belongs to the frequency $f_j = j f_{\mathrm{s}}/N$, so that the frequency resolution is $\Delta f = f_{\mathrm{s}}/N$, the reciprocal of the record length. A sinusoid of amplitude a at a frequency that falls on a bin gives $|X_j| = aN/2$. A record of one second at $f_{\mathrm{s}} = 10\,\mathrm{kHz}$, for example, has $N = 10\,000$ and a resolution of 1 Hz.

The reference frequency of a pump is its rotational frequency. At $n = 1800\,\mathrm{rpm}$,

$$f_{\mathrm{r}} = \frac{n}{60} = \frac{1800}{60\,\mathrm{s}} = 30\,\mathrm{Hz}.$$

An imbalance of the rotor produces a strong component near $1 \times f_{\mathrm{r}}$; misalignment and looseness add harmonics at $2\times$ and $3 \times f_{\mathrm{r}}$; impeller vanes produce the vane-pass frequency, the number of vanes times f_{r}; cavitation produces broadband vibration; and rolling-element bearings produce characteristic defect frequencies. The last follow from the bearing kinematics. Let the bearing have N_{b} rolling elements of diameter d_{b} on a pitch diameter D_{p} with contact angle φ, the inner race turning with the shaft at f_{r} and the outer race fixed. The cage turns at the mean of the surface speeds of the two races at the contact points, which gives the fundamental train (cage) frequency, and the defect frequencies follow from how often elements pass a point on each race:

$$\begin{aligned}
f_{\mathrm{FTF}} &= \frac{f_{\mathrm{r}}}{2}\left(1 - \frac{d_{\mathrm{b}}}{D_{\mathrm{p}}}\cos\varphi\right), \\
f_{\mathrm{BPFO}} &= N_{\mathrm{b}}\, f_{\mathrm{FTF}} = \frac{N_{\mathrm{b}} f_{\mathrm{r}}}{2}\left(1 - \frac{d_{\mathrm{b}}}{D_{\mathrm{p}}}\cos\varphi\right), \\
f_{\mathrm{BPFI}} &= N_{\mathrm{b}}\,(f_{\mathrm{r}} - f_{\mathrm{FTF}}) = \frac{N_{\mathrm{b}} f_{\mathrm{r}}}{2}\left(1 + \frac{d_{\mathrm{b}}}{D_{\mathrm{p}}}\cos\varphi\right), \\
f_{\mathrm{BSF}} &= \frac{D_{\mathrm{p}} f_{\mathrm{r}}}{2 d_{\mathrm{b}}}\left[1 - \left(\frac{d_{\mathrm{b}}}{D_{\mathrm{p}}}\cos\varphi\right)^2\right],
\end{aligned} \tag{22.14}$$

the ball-pass frequencies of the outer race (BPFO) and the inner race (BPFI) and the ball spin frequency (BSF). These are the standard kinematic results, exact for rolling without slip. For an illustrative deep-groove ball bearing with nine balls of diameter 7.94 mm on a pitch diameter of 39.04 mm and zero contact angle, $d_{\mathrm{b}}/D_{\mathrm{p}} = 0.2034$, and at 30 Hz Table 22.3 gives the four frequencies. One row of the arithmetic: $f_{\mathrm{BPFO}} = 4.5 \times 30 \times (1 - 0.2034) = 107.5\,\mathrm{Hz}$.

Table 22.3: Characteristic bearing frequencies, equation (22.14), for an illustrative bearing with $N_\mathrm{b} = 9$, $d_\mathrm{b}/D_\mathrm{p} = 0.2034$ and $\varphi = 0$, on a shaft at $f_\mathrm{r} = 30\,\mathrm{Hz}$.

Frequency	Symbol	Value (Hz)	Multiple of f_r
Fundamental train (cage)	f_FTF	11.95	0.398
Ball spin	f_BSF	70.70	2.357
Ball pass, outer race	f_BPFO	107.54	3.585
Ball pass, inner race	f_BPFI	162.46	5.415

The defect frequencies are not integer multiples of the running speed, which is what lets a spectrum separate a bearing defect from imbalance or misalignment. An outer-race spall appears at BPFO and its harmonics; an inner-race defect, which turns with the shaft through the load zone, appears at BPFI modulated at f_r, as sidebands at BPFI $\pm f_\mathrm{r}$. In practice the impacts excite structural resonances at much higher frequencies, and the defect frequencies are found in the spectrum of the envelope of the band-passed signal.

Figure 22.4 draws an illustrative spectrum. The horizontal axis is frequency and the vertical axis spectral amplitude, both without numerical scale; the bars are spectral lines, and the tallest, in red, is the component at the rotational frequency of 30 Hz. A peak at 30 Hz is consistent with imbalance but does not prove it: load, coupling condition, mounting stiffness, speed variation and the location of the sensor all change the amplitude of the 1× component, and a misaligned coupling also produces it. Spectral peaks become diagnostic when their frequencies and the operating conditions are physically interpretable.

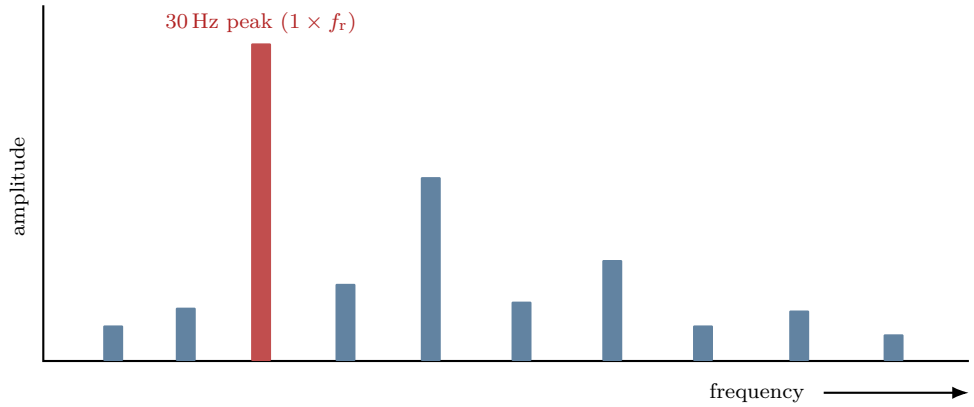


Figure 22.4: Illustrative vibration spectrum (schematic, not measured data). The red line is the component at the rotational frequency $f_\mathrm{r} = 30\,\mathrm{Hz}$ of a pump at 1800 rpm; the blue lines are other components. Spectral peaks become diagnostically useful when their frequencies and the operating conditions are physically interpretable.

22.3.4 Filtering and gradual trend detection

Condition monitoring looks mostly for persistent changes rather than isolated spikes: degradation shifts the level of an indicator slowly, while noise moves it from sample to sample. The *exponentially weighted moving average* (EWMA) of a sequence y_k is the recursion

$$s_k = \alpha\, y_k + (1 - \alpha)\, s_{k-1}, \qquad 0 < \alpha \le 1, \tag{22.15}$$

started at the in-control mean, $s_0 = 0$ for a standardized residual. The smoothing constant α is dimensionless. Unrolling the recursion,

$$s_k = \alpha \sum_{j=0}^{k-1} (1-\alpha)^j \, y_{k-j} + (1-\alpha)^k s_0,$$

a weighted average in which the weights $\alpha(1-\alpha)^j$ decay geometrically with the age j of the sample; for $\alpha = 0.2$ they are 0.200, 0.160, 0.128, 0.102, …. If the y_k are independent with variance σ^2, the variance of s_k is $\sigma^2\alpha^2 \sum_j (1-\alpha)^{2j}$, which tends to

$$\operatorname{Var}(s_k) \to \frac{\alpha^2\sigma^2}{1-(1-\alpha)^2} = \frac{\alpha}{2-\alpha}\,\sigma^2.$$

For $\alpha = 0.2$, $s_k = 0.2\,y_k + 0.8\,s_{k-1}$ and the variance factor is $0.2/1.8 = 0.111$, so the standard deviation of the average is one third of that of a single sample. A smaller α smooths more strongly and responds more slowly. Applied to a noisy bearing temperature, the EWMA reveals a gradual drift that the scatter of the raw readings hides; its alarm limit is set at a multiple of its own standard deviation, $L\sigma\sqrt{\alpha/(2-\alpha)}$ with L near 3.

22.3.5 The cumulative-sum statistic

A sustained small shift may never produce one large residual but accumulates evidence over many samples. The *cumulative-sum* (CUSUM) statistic collects that evidence and follows from a likelihood ratio. Let the standardized residuals z_k be independent and $\mathcal{N}(0,1)$ while the pump is healthy, and $\mathcal{N}(\delta,1)$ after a change of unknown onset, $\delta > 0$ being the shift to be detected in standard deviations. The log-likelihood ratio of one sample for the shifted against the healthy distribution is

$$\ln \frac{\exp[-(z_k-\delta)^2/2]}{\exp[-z_k^2/2]} = \delta z_k - \frac{\delta^2}{2} = \delta\left(z_k - \frac{\delta}{2}\right).$$

If the change began at sample j, the log-likelihood ratio of the samples $j,\dots,k$ is the sum of these terms. The onset is unknown, so the evidence for a change by time k is the largest such sum over all candidate onsets, with the empty sum, zero, for "no change yet". Dividing by δ and writing $\kappa = \delta/2$,

$$S_k^+ = \max\Bigl(0,\, \max_{1\le j\le k} \sum_{i=j}^{k} (z_i - \kappa)\Bigr).$$

Every candidate sum ending at k is the term $z_k - \kappa$ plus a candidate sum ending at $k-1$ (or the empty one), and the largest of the latter is S_{k-1}^+. The maximum therefore obeys the recursion

$$S_k^+ = \max(0,\; S_{k-1}^+ + z_k - \kappa), \qquad S_0^+ = 0, \tag{22.16}$$

the one-sided CUSUM. The reference value κ, in standard deviations, is half the shift the chart is designed to detect; samples below κ pull the statistic down, and the floor at zero discards evidence from before the change. An alarm is raised when

$$S_k^+ > h_{\mathrm{C}}, \tag{22.17}$$

which is equivalent to the likelihood ratio of the most likely onset exceeding $\mathrm{e}^{\delta h_{\mathrm{C}}}$. The threshold h_{C} trades false alarms against detection delay, both measured by the *average run length* (ARL), the mean number of samples to an alarm: ARL_0 while healthy, which should be long, and ARL_1 after a shift of size δ, which should be short. A standard diffusion approximation for the one-sided CUSUM gives

$$\mathrm{ARL} \approx \frac{\mathrm{e}^{-2\Delta b} + 2\Delta b - 1}{2\Delta^2}, \qquad \Delta = \delta_{\mathrm{true}} - \kappa, \quad b = h_{\mathrm{C}} + 1.166,$$

where δ_{true} is the actual shift. With $\kappa = 0.5$ and $h_{\text{C}} = 4$ it gives $\text{ARL}_0 \approx 338$ samples and, for a shift of one standard deviation, $\text{ARL}_1 \approx 8.3$; with $h_{\text{C}} = 5$, $\text{ARL}_0 \approx 938$ and $\text{ARL}_1 \approx 10.3$. A rule that alarms on a single sample above $z = 3$ has $\text{ARL}_0 = 741$ but needs 44 samples on average to see the same shift.

Table 22.4 works an illustrative sequence of fifteen standardized residuals in which the mean shifts from zero to about one standard deviation at sample 6. The columns are the sample index, the residual, the increment $z_k - \kappa$ with $\kappa = 0.5$, the CUSUM of equation (22.16), and the EWMA of equation (22.15) with $\alpha = 0.2$. One row of the arithmetic, sample 8: $S_8^+ = \max(0, 0.8 + 1.4 - 0.5) = 1.7$ and $s_8 = 0.2 \times 1.4 + 0.8 \times 0.3086 = 0.5269$.

Table 22.4: CUSUM ($\kappa = 0.5$, $h_{\text{C}} = 4$) and EWMA ($\alpha = 0.2$) of an illustrative sequence of standardized residuals with a shift of about one standard deviation from sample 6.

Sample k	z_k	$z_k - \kappa$	S_k^+	s_k	Alarm
1	0.3	−0.2	0.0	0.060	
2	−0.8	−1.3	0.0	−0.112	
3	0.5	0.0	0.0	0.010	
4	−0.2	−0.7	0.0	−0.032	
5	0.1	−0.4	0.0	−0.005	
6	1.2	0.7	0.7	0.236	
7	0.6	0.1	0.8	0.309	
8	1.4	0.9	1.7	0.527	
9	0.9	0.4	2.1	0.602	
10	1.6	1.1	3.2	0.801	
11	0.4	−0.1	3.1	0.721	
12	1.3	0.8	3.9	0.837	
13	1.1	0.6	4.5	0.889	$S_{13}^+ > 4$
14	0.8	0.3	4.8	0.872	
15	1.5	1.0	5.8	0.997	

Figure 22.5 draws the sequence. In panel (a) the horizontal axis is the sample index and the vertical axis the standardized residual; the stems are the residuals, the dashed red line is a single-sample limit at $z = 3$, and the dotted line marks the onset of the shift. No residual comes near the limit; the largest is 1.6. In panel (b) the vertical axis is the CUSUM, the dashed line the threshold $h_{\text{C}} = 4$: the statistic stays at zero for five samples, climbs once the shift begins, and crosses the threshold at sample 13, eight samples after the onset, close to the ARL_1 of 8.3. The EWMA of the table rises from about zero to about 0.9, tracking the new level of the residuals. The thresholds of both charts must be chosen for the required false-alarm rate and detection delay, and with temporal correlation in the residuals the independent-sample run lengths above no longer hold; they must then be calibrated by simulation or on healthy data.

22.4 Bayesian diagnosis with competing hypotheses

Detection establishes that the healthy model is contradicted; it does not say why. This section works a complete example: an excess of electrical power is detected on the pump, three explanations are proposed, each is turned into a Gaussian likelihood for two residuals, and Bayes' theorem converts the likelihoods and the prior probabilities into posterior probabilities that change when the evidence changes.

22.4.1 Baseline performance and the power residual

The pump operates at the illustrative point

$$Q = 0.050\,\text{m}^3/\text{s}\ (180\,\text{m}^3/\text{h}), \qquad \Delta p = 300\,000\,\text{Pa} = 300\,\text{kPa},$$

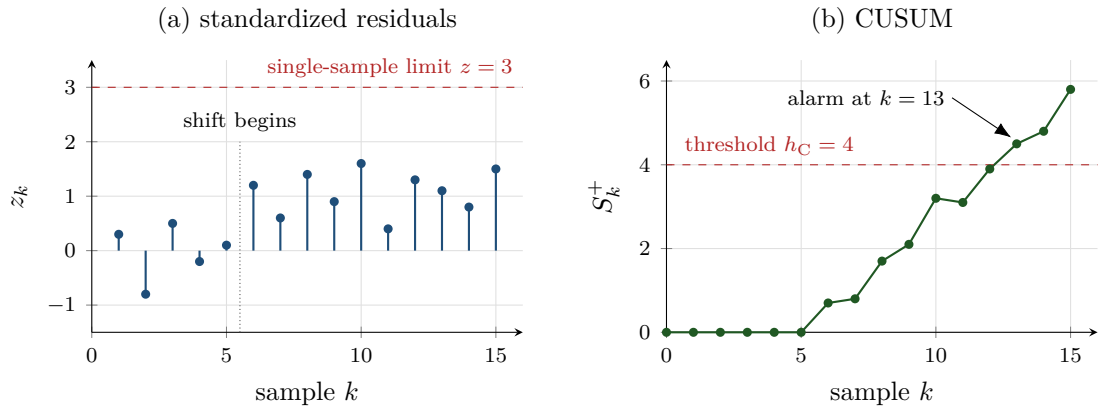


Figure 22.5: Detection of a sustained shift of about one standard deviation beginning at sample 6 (illustrative data of Table 22.4). (a) No single standardized residual approaches a 3σ limit. (b) The one-sided CUSUM, equation (22.16) with $\kappa = 0.5$, accumulates the small positive deviations and crosses $h_C = 4$ at sample 13.

a head of $H_\mathrm{p} = 300\,000/(1000 \times 9.81) = 30.6\,\mathrm{m}$. By equation (22.1) the hydraulic power is

$$P_\mathrm{h} = Q\,\Delta p = 0.050\,\mathrm{m^3/s} \times 300\,000\,\mathrm{Pa} = 15\,000\,\mathrm{W} = 15\,\mathrm{kW}.$$

With a healthy wire-to-water efficiency $\eta_0 = 0.75$, the expected electrical power is, by equation (22.2),

$$P_{\mathrm{e},0} = \frac{15\,\mathrm{kW}}{0.75} = 20\,\mathrm{kW}.$$

The measured power is $P_\mathrm{e} = 24\,\mathrm{kW}$. If the flow, pressure and power measurements are accurate and simultaneous, the apparent efficiency is

$$\eta_\mathrm{obs} = \frac{15\,\mathrm{kW}}{24\,\mathrm{kW}} = 0.625,$$

62.5 % against the healthy reference of 75 %. Five explanations are consistent with that number, and each requires a definite size of error: an actual loss of pump efficiency from 0.75 to 0.625; a flow meter that reads low, the true flow being $0.060\,\mathrm{m^3/s}$, so that $0.060 \times 300\,000/0.75 = 24\,\mathrm{kW}$; a pressure measurement that reads low, the true pressure rise being 360 kPa; a power meter that reads 4 kW, or 20 %, high; or a healthy-efficiency reference that does not apply at the current operating point, since pump efficiency varies along the characteristic curve. An inconsistency with the reference model has been detected. Its cause has not yet been diagnosed.

The power residual is

$$e_P = P_\mathrm{e,obs} - P_\mathrm{e,model} = 24\,\mathrm{kW} - 20\,\mathrm{kW} = 4\,\mathrm{kW}.$$

Let the combined standard deviation of the model prediction error and the measurement error be $\sigma_e = 1.5\,\mathrm{kW}$, an illustrative value. Then

$$z = \frac{4\,\mathrm{kW}}{1.5\,\mathrm{kW}} = 2.667,$$

and under a calibrated Gaussian healthy-residual model the probability of a residual at least this extreme in either direction is $2[1 - \Phi(2.667)] = 2 \times 0.00383 = 0.0077$, about 0.77 %. At a 1 % two-sided level, whose threshold is $|z| > 2.576$, the healthy model is rejected. The residual is evidence against the healthy reference; it cannot, by itself, distinguish a loss of efficiency from a biased measurement, because both produce exactly this residual. That requires fault hypotheses and further evidence.

22.4.2 Competing hypotheses and the observation vector

Three explanations of the excess power are compared:

$$H_0 : \text{healthy pump and sensors},$$
$$H_1 : \text{mechanical or hydraulic deterioration},$$
$$H_2 : \text{electrical power measurement bias}.$$

Within this simplified diagnostic model they are assumed mutually exclusive and collectively exhaustive. In reality several faults may coexist and the true explanation may be none of the three; both limitations return in Section 22.6.3. The prior probabilities, before the present measurements, are

$$\Pr(H_0) = 0.70, \qquad \Pr(H_1) = 0.15, \qquad \Pr(H_2) = 0.15.$$

They are hypothetical, chosen for the example, and not estimated from pump failure data; in practice they would come from failure and calibration records such as those of Chapter 8 of Volume I.

The evidence is two condition indicators, collected in the observation vector

$$\mathbf{y} = \begin{bmatrix} e_P \\ e_v \end{bmatrix},$$

where e_P is the electrical power residual in kW and e_v the vibration residual in mm/s, the measured RMS vibration minus its healthy value at the current speed and load. The observed residuals are

$$\mathbf{y} = \begin{bmatrix} 4.0\,\text{kW} \\ 0.3\,\text{mm/s} \end{bmatrix}: \tag{22.18}$$

the power is 4 kW above expectation, the vibration only slightly elevated. Each hypothesis predicts a different mean residual vector $\boldsymbol{\mu}_i$, listed in Table 22.5: the healthy pump predicts no excess in either; deterioration that wastes energy in the bearings or the rotor raises the power and the vibration together; a power-meter bias raises the power reading and leaves the vibration untouched. The magnitudes are illustrative.

Table 22.5: Mean residual vector $\boldsymbol{\mu}_i$ predicted by each hypothesis (illustrative).

Hypothesis	Expected power residual	Expected vibration residual
H_0: healthy	0 kW	0 mm/s
H_1: deterioration	+4 kW	+2 mm/s
H_2: power-sensor bias	+4 kW	0 mm/s

Around its mean, the residual vector is assumed conditionally Gaussian under every hypothesis, with the common covariance

$$\boldsymbol{\Sigma} = \begin{bmatrix} 1.5^2 & 0 \\ 0 & 1.0^2 \end{bmatrix}, \tag{22.19}$$

in kW^2 and $(\text{mm/s})^2$. The zero off-diagonal entry is an assumption; shared model errors, such as an error in the measured speed that enters both predictions, would correlate the two residuals.

22.4.3 The Gaussian likelihood under each hypothesis

Under hypothesis H_i the observation vector is $\mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma})$, so its likelihood is the multivariate Gaussian density with $m = 2$:

$$p(\mathbf{y} \mid H_i) = \frac{1}{(2\pi)^{m/2}\,|\boldsymbol{\Sigma}|^{1/2}} \exp\!\Big[-\tfrac{1}{2}(\mathbf{y}-\boldsymbol{\mu}_i)^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{y}-\boldsymbol{\mu}_i)\Big]. \tag{22.20}$$

For the diagonal covariance of equation (22.19), $|\boldsymbol{\Sigma}| = 1.5^2 \times 1.0^2$ and the quadratic form splits into one term per residual, so the density is the product of two univariate Gaussian densities,

$$p(\mathbf{y} \mid H_i) = \frac{1}{\sqrt{2\pi}\,1.5}\,\mathrm{e}^{-(e_P-\mu_{i,P})^2/(2\cdot 1.5^2)} \cdot \frac{1}{\sqrt{2\pi}\,1.0}\,\mathrm{e}^{-(e_v-\mu_{i,v})^2/(2\cdot 1.0^2)},$$

as independence of the two residuals requires. Define the squared normalized discrepancy of the observation from the prediction of H_i,

$$d_i^2 = (\mathbf{y}-\boldsymbol{\mu}_i)^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{y}-\boldsymbol{\mu}_i) = \frac{(e_P-\mu_{i,P})^2}{1.5^2} + \frac{(e_v-\mu_{i,v})^2}{1.0^2}. \tag{22.21}$$

The normalizing factor $1/(2\pi \times 1.5 \times 1.0) = 0.1061$ per kW mm/s is the same for the three hypotheses, because they share the covariance; it cancels in Bayes' theorem, and only the relative likelihoods

$$L_i \propto \mathrm{e}^{-d_i^2/2} \tag{22.22}$$

matter. For the observation of equation (22.18):

$$\begin{aligned}
d_0^2 &= \frac{4^2}{1.5^2} + \frac{0.3^2}{1^2} = 7.1111 + 0.0900 = 7.2011, & L_0 &= \mathrm{e}^{-3.6006} = 0.0273,\\
d_1^2 &= \frac{(4-4)^2}{1.5^2} + \frac{(0.3-2)^2}{1^2} = 2.8900, & L_1 &= \mathrm{e}^{-1.4450} = 0.2357,\\
d_2^2 &= \frac{(4-4)^2}{1.5^2} + \frac{(0.3-0)^2}{1^2} = 0.0900, & L_2 &= \mathrm{e}^{-0.0450} = 0.9560.
\end{aligned}$$

The bias hypothesis H_2 is the most consistent with the residual pattern: it explains the excess power and predicts the near-zero vibration. The deterioration hypothesis explains the power but predicts a vibration excess of 2 mm/s, which is 1.7 standard deviations away from what was observed. The healthy hypothesis predicts neither. Multiplied by 0.1061, the absolute likelihoods are 0.00290, 0.0250 and 0.1014 per kW mm/s. The discrepancy d_0^2 is also the two-dimensional version of the detection statistic J: under H_0 it is χ^2_2, whose tail $\Pr(\chi^2_2 > x) = \mathrm{e}^{-x/2}$ makes $L_0 = 0.0273$ its significance level.

22.4.4 Posterior probabilities

Bayes' theorem for a discrete set of hypotheses, stated in Chapter 8 for competing models, gives the posterior probability of each hypothesis as its likelihood times its prior, normalized over the set:

$$\Pr(H_i \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid H_i)\,\Pr(H_i)}{\sum_j p(\mathbf{y} \mid H_j)\,\Pr(H_j)}. \tag{22.23}$$

The relative likelihoods L_i may replace $p(\mathbf{y} \mid H_i)$ because the common factor cancels. The unnormalized weights $w_i = \Pr(H_i)\,L_i$ and the posterior are collected in Table 22.6, whose columns are the hypothesis, the prior, the discrepancy, the relative likelihood, the weight and the posterior, for the observation of equation (22.18) and for the second observation of Section 22.4.5. For the first, $w_0 = 0.70 \times 0.0273 = 0.0191$, $w_1 = 0.15 \times 0.2357 = 0.0354$, $w_2 = 0.15 \times 0.9560 = 0.1434$; their sum is 0.1979, and dividing each weight by it gives the posterior.

Table 22.6: Bayesian diagnosis of the excess power for two observations. Priors 0.70, 0.15, 0.15; means of Table 22.5; covariance of equation (22.19).

Hypothesis	Prior	d_i^2	L_i	w_i	$\Pr(H_i \mid \mathbf{y})$
Observation $\mathbf{y} = (4.0\,\text{kW},\ 0.3\,\text{mm/s})$					
H_0: healthy	0.70	7.2011	0.0273	0.0191	0.0966
H_1: deterioration	0.15	2.8900	0.2357	0.0354	0.1787
H_2: power-sensor bias	0.15	0.0900	0.9560	0.1434	0.7247
Observation $\mathbf{y} = (4.0\,\text{kW},\ 2.0\,\text{mm/s})$					
H_0: healthy	0.70	11.1111	0.0039	0.0027	0.0156
H_1: deterioration	0.15	0.0000	1.0000	0.1500	0.8670
H_2: power-sensor bias	0.15	4.0000	0.1353	0.0203	0.1173

The power-sensor bias hypothesis has posterior probability 0.725. It explains the elevated power reading without predicting a large increase in vibration. The deterioration hypothesis also explains the high power, but its prediction of clearly elevated vibration is less consistent with the observation; and although the healthy hypothesis started with the largest prior, 0.70, its likelihood is so small that its posterior falls to 0.097. This is how several measurements distinguish competing hypotheses. The probabilities are conditional on the three hypotheses, the assumed priors and the Gaussian likelihood model; they are illustrative and are not validated field probabilities of these faults.

The comparison of two hypotheses is clearest in posterior odds. Taking the logarithm of the ratio of equation (22.23) for H_1 and H_2,

$$\ln \frac{\Pr(H_1 \mid \mathbf{y})}{\Pr(H_2 \mid \mathbf{y})} = \ln \frac{\Pr(H_1)}{\Pr(H_2)} - \frac{d_1^2 - d_2^2}{2}. \tag{22.24}$$

Because the two hypotheses share the covariance, the quadratic terms in $\mathbf{y}$ cancel in $d_1^2 - d_2^2$, which is linear in the observation; here $d_1^2 - d_2^2 = (e_v - 2)^2 - e_v^2 = 4 - 4e_v$, and the log odds are $0 + 2e_v - 2$, independent of e_P. At $e_v = 0.3\,\text{mm/s}$ they are -1.4, odds of $\mathrm{e}^{-1.4} = 0.247$, which is $0.1787/0.7247$. Only the vibration residual separates deterioration from the power-sensor bias, because only in vibration do their predictions differ.

22.4.5 A different vibration observation

Suppose the power residual stays at $e_P = 4\,\text{kW}$ but the vibration residual is $e_v = 2.0\,\text{mm/s}$, so that $\mathbf{y} = (4\,\text{kW}, 2\,\text{mm/s})$. The deterioration hypothesis predicts exactly this pattern, $d_1^2 = 0$; the bias hypothesis predicts no vibration increase, $d_2^2 = (2-0)^2 = 4$; the healthy hypothesis predicts neither, $d_0^2 = 7.1111 + 4 = 11.1111$. The lower half of Table 22.6 carries the computation through: the weights are 0.0027, 0.1500 and 0.0203, their sum is 0.1730, and the posterior probabilities are 0.016 for the healthy pump, 0.867 for deterioration and 0.117 for the power-sensor bias. One additional measurement dimension reverses the diagnosis.

This is *diagnostic discrimination*. The vibration residual is useful not because it detects abnormal behavior, since the power residual already did, but because it responds differently to the competing explanations. The second observation is also detected by the two-dimensional test, $d_0^2 = 11.11 > \chi^2_{2,0.99} = 9.21$, whereas the first, $d_0^2 = 7.20$, is not, although its power residual alone rejects the healthy model at the 1 % level: a multivariate statistic spreads its false-alarm budget over every direction, and a quiet channel dilutes a loud one.

22.4.6 Posterior probabilities over the residual plane

The model of the previous subsections maps every pair of residuals (e_P, e_v) to three posterior probabilities. Table 22.7 evaluates the map for the two observations already worked and for

ten others: six with the power residual held at 4 kW and the vibration residual raised from 0.3 to 3.0 mm/s, and five with the vibration held at 0.3 mm/s and the power residual moved from 0 to 6 kW, plus one near the point where all three hypotheses are equally probable. Each row gives the residuals, the three discrepancies of equation (22.21), the posterior probabilities of equation (22.23) and the most probable hypothesis. One row of the arithmetic, $\mathbf{y} = (4, 1.5)$: $d^2 = (9.361, 0.250, 2.250)$, weights $0.70\,\mathrm{e}^{-4.681} = 0.00649$, $0.15\,\mathrm{e}^{-0.125} = 0.13237$ and $0.15\,\mathrm{e}^{-1.125} = 0.04870$, sum 0.18756, posterior $(0.035, 0.706, 0.260)$.

Table 22.7: Posterior probabilities of the three hypotheses for pairs of residuals. Priors 0.70, 0.15, 0.15; means $(0, 0)$, $(4, 2)$, $(4, 0)$; standard deviations 1.5 kW and 1.0 mm/s. The rows other than $(4, 0.3)$ and $(4, 2.0)$ are added to show the trend.

					$\Pr(H_i \mid \mathbf{y})$			
e_P (kW)	e_v (mm/s)	d_0^2	d_1^2	d_2^2	H_0	H_1	H_2	Most probable
4.0	0.3	7.20	2.89	0.09	0.097	0.179	0.725	H_2
4.0	0.5	7.36	2.25	0.25	0.089	0.245	0.666	H_2
4.0	1.0	8.11	1.00	1.00	0.063	0.469	0.469	H_1, H_2 tie
4.0	1.5	9.36	0.25	2.25	0.035	0.706	0.260	H_1
4.0	2.0	11.11	0.00	4.00	0.016	0.867	0.117	H_1
4.0	3.0	16.11	1.00	9.00	0.002	0.980	0.018	H_1
0.0	0.3	0.09	10.00	7.20	0.992	0.002	0.006	H_0
1.0	0.3	0.53	6.89	4.09	0.957	0.009	0.035	H_0
2.0	0.3	1.87	4.67	1.87	0.789	0.042	0.169	H_0
3.0	0.3	4.09	3.33	0.53	0.388	0.121	0.491	H_2
6.0	0.3	16.09	4.67	1.87	0.003	0.197	0.800	H_2
2.87	1.0	4.66	1.57	1.57	0.332	0.334	0.334	near tie

Three trends appear. With the power residual at 4 kW, raising the vibration residual moves the diagnosis from the sensor bias to deterioration, the two being equally probable at exactly $e_v = 1$ mm/s, where the log odds $2e_v - 2$ of equation (22.24) vanish. With the vibration at 0.3 mm/s, reducing the power residual toward zero makes the healthy hypothesis dominant; above about 2.9 kW the bias overtakes it, and at 6 kW the healthy pump is excluded while the split between bias and deterioration is again set by the vibration alone. Where the probabilities are close, near the last row, the present evidence does not discriminate, and a further measurement is especially valuable.

Figure 22.6 draws the whole map. The horizontal axis is the power residual and the vertical axis the vibration residual, over the ranges -2 to 7 kW and -1 to 4 mm/s; each shaded region is the set of observations for which one hypothesis is the most probable; the squares are the hypothesis means of Table 22.5, and the dots the two observations of Table 22.6. The boundaries are straight lines, because with a shared covariance each pairwise log odds is linear in $\mathbf{y}$: the boundary between H_0 and H_2 is $e_P = 2 + (4.5/8)\ln(0.70/0.15) = 2.87$ kW, the boundary between H_1 and H_2 is $e_v = 1$ mm/s, the boundary between H_0 and H_1 is $e_P = 3.99 - 1.125\,e_v$, and the three meet at $(2.87, 1.0)$. The dashed ellipse is the boundary of the healthy 99 % region, $d_0^2 = 9.21$: the first observation lies inside it and the second outside.

The map also shows that a fault classification is not a statement of certainty. Every region contains observations near its boundaries, where the most probable hypothesis has a posterior barely above that of a neighbor, and the boundaries themselves move when the priors, the fault signatures or the covariance change: raising the prior of H_2 to twice that of H_1, for example, shifts the H_1–H_2 boundary to $e_v = 1 + \frac{1}{2}\ln 2 = 1.35$ mm/s. The diagnosis depends on the model as much as on the data, and the next section asks when the data can support it at all.

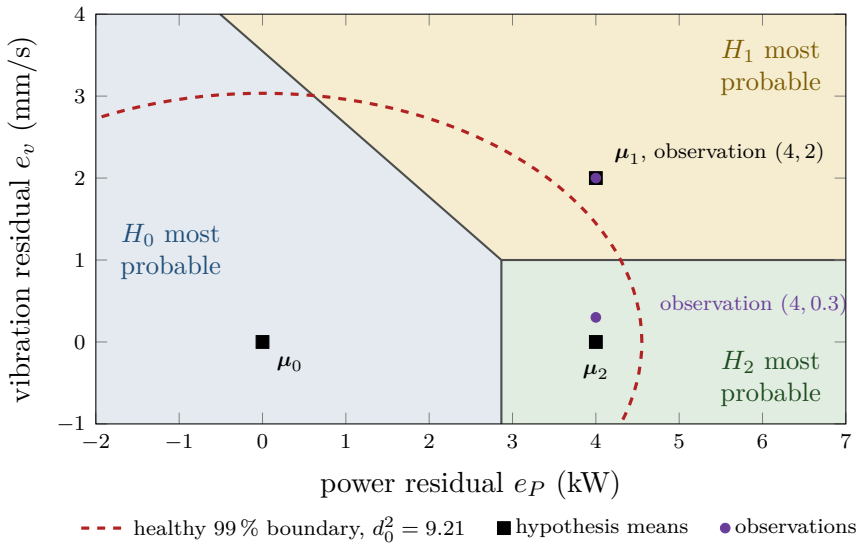


Figure 22.6: Most probable hypothesis over the plane of power and vibration residuals, for the priors 0.70, 0.15, 0.15 and the Gaussian model of equations (22.19) and (22.20). Boundaries: $e_P = 2.87\,\text{kW}$ (H_0 and H_2), $e_v = 1\,\text{mm/s}$ (H_1 and H_2), $e_P = 3.99 - 1.125\,e_v$ (H_0 and H_1). Dashed: the boundary of the healthy 99 % region, $e_P^2/1.5^2 + e_v^2 = 9.21$.

22.5 Sensor faults, observability and the next measurement

Whether a set of measurements can distinguish the hypotheses of interest is a property of the measurement system, not of the inference algorithm. This section shows how a sensor fault can be indistinguishable from a physical one, states the condition for distinguishing faults in terms of their signatures, shows what an independent measurement and a physical balance equation add, and derives the information-theoretic and decision-theoretic value of the next measurement.

22.5.1 A sensor fault can look like a physical fault

A data center cooling system reports an unusually high return-water temperature. Three hypotheses explain it:

$$H_{\text{cap}} : \text{actual shortfall of cooling capacity},$$
$$H_{\text{bias}} : \text{bias of the return-temperature sensor},$$
$$H_{\text{load}} : \text{unexpected increase of the IT heat load}.$$

Each produces a high reading. If only that one temperature sensor is available, the three hypotheses are observationally indistinguishable: every value the sensor can report is predicted equally well by each of them, for a suitable size of the fault, and the posterior can never move far from the prior. That is not a weakness of the inference algorithm; it is a lack of information in the measurement system. The case is the mirror image of the flow-sensor bias of Section 4.5.7 of Chapter 4 of Volume I: there a biased sensor hid a real deficiency, here a biased sensor imitates one, and in both a sensor failure mode must be a hypothesis of the diagnosis on equal terms with the physical ones.

22.5.2 Diagnostic signatures

Near an operating point, the effect of small faults on the residuals is approximately linear,

$$\mathbf{e} = \mathbf{S}_f\,\mathbf{f} + \boldsymbol{\varepsilon}, \tag{22.25}$$

where $\mathbf{e}$ is the residual vector of the m measurements, $\mathbf{f}$ the vector of the magnitudes of the q candidate faults (zero when absent), $\mathbf{S}_f$ the $m \times q$ fault-signature, or sensitivity, matrix, and $\boldsymbol{\varepsilon}$ the residual noise. Column j of $\mathbf{S}_f$, the *signature* $\mathbf{s}_j$ of fault j, states how each residual responds to a unit of that fault. For two faults, $\mathbf{e} = \mathbf{s}_1 f_1 + \mathbf{s}_2 f_2 + \boldsymbol{\varepsilon}$. If $\mathbf{s}_1 = \mathbf{s}_2$, or more generally if one signature is a multiple of the other, any residual explained by f_1 is explained equally well by a suitable f_2, and with unknown fault magnitudes the residuals cannot distinguish them without further information.

Table 22.8 lists qualitative signatures for the cooling example. Each row is a measurement, each column a hypothesis, and an entry states whether the residual of that measurement rises (+) or stays near zero (0) under that hypothesis. With the first row only, the three columns are identical. The second row, an independently calibrated return-temperature sensor, separates the bias from the two physical causes; the supply temperature, which a capacity shortfall raises because the plant cannot hold its set point, and the IT electrical power, which rises only with the load, separate the physical causes from each other. With rows one, two and five the three columns are distinct.

Table 22.8: Qualitative diagnostic signatures of three explanations of a high return-water temperature in a constant-flow loop (illustrative): +, the residual rises; 0, it stays near zero.

Measurement	H_{cap}	H_{bias}	H_{load}
Return temperature, sensor 1	+	+	+
Return temperature, independent sensor 2	+	0	+
Supply temperature	+	0	0
Chilled-water flow	0	0	0
IT electrical power	0	0	+

The quantitative version is a rank condition. Suppose the one temperature residual is

$$e_T = \delta T_{\text{real}} + b_T + \varepsilon,$$

where δT_{real} is the actual deviation of the water temperature, b_T the bias of the sensor and ε its noise. A single residual of $e_T = 5°\text{C}$ cannot be divided uniquely between the actual deviation and the bias: the equation has two unknowns and one measurement. Its signature matrix is

$$\mathbf{S}_f = \begin{bmatrix} 1 & 1 \end{bmatrix},$$

of rank one, and both unknowns cannot be identified from this observation alone. In general, the least-squares estimate of the fault magnitudes, $\hat{\mathbf{f}} = (\mathbf{S}_f^{\mathsf{T}} \boldsymbol{\Sigma}_e^{-1} \mathbf{S}_f)^{-1} \mathbf{S}_f^{\mathsf{T}} \boldsymbol{\Sigma}_e^{-1} \mathbf{e}$, exists only when $\mathbf{S}_f^{\mathsf{T}} \boldsymbol{\Sigma}_e^{-1} \mathbf{S}_f$ is invertible, that is, when $\mathbf{S}_f$ has full column rank q; its covariance is then $(\mathbf{S}_f^{\mathsf{T}} \boldsymbol{\Sigma}_e^{-1} \mathbf{S}_f)^{-1}$.

22.5.3 An independent measurement

A second, independently calibrated temperature sensor at the same location gives

$$e_{T,2} = \delta T_{\text{real}} + \varepsilon_2,$$

while the first gives $e_{T,1} = \delta T_{\text{real}} + b_T + \varepsilon_1$. Together,

$$\begin{bmatrix} e_{T,1} \\ e_{T,2} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \delta T_{\text{real}} \\ b_T \end{bmatrix} + \boldsymbol{\varepsilon}, \tag{22.26}$$

and the signature matrix has full column rank, $\operatorname{rank} \mathbf{S}_f = 2$. Inverting it gives $\delta\hat{T}_{\text{real}} = e_{T,2}$ and $\hat{b}_T = e_{T,1} - e_{T,2}$. With the illustrative residuals $e_{T,1} = 5.0\,^\circ\text{C}$ and $e_{T,2} = 1.2\,^\circ\text{C}$, the water is $1.2\,^\circ\text{C}$ warmer than expected and the first sensor reads $3.8\,^\circ\text{C}$ high; with independent noise of standard deviation $0.3\,^\circ\text{C}$ on each sensor, the bias estimate has standard deviation $0.3\sqrt{2} = 0.42\,^\circ\text{C}$. Under suitable noise and calibration assumptions the real deviation and the bias can now be distinguished. This is a simple case of *diagnostic observability*: the faults of interest can be inferred from the available measurements.

Full rank is a structural condition and is not enough by itself. If the second sensor shares part of the first one's calibration error, for example through a common reference used to calibrate both, its row becomes $(1, 0.9)$ instead of $(1, 0)$. The matrix is still of rank two, but its columns are nearly dependent: its inverse is $\left[\begin{smallmatrix} -9 & 10 \\ 10 & -10 \end{smallmatrix}\right]$, the bias estimate is $10(e_{T,1} - e_{T,2})$, and its standard deviation is $0.3\sqrt{200} = 4.2\,^\circ\text{C}$, ten times larger than with independent calibration and larger than the bias being estimated. The condition number of $\mathbf{S}_f$, the ratio of its largest to its smallest singular value, rises from 2.6 to 38. A diagnosis can be structurally possible and numerically useless; and two sensors that share a calibration are partly a common-cause group in the sense of Chapter 13 of Volume I.

Principle 22.2. *Faults can be distinguished only when their signatures are linearly independent and the noise is small compared with the separation the signatures provide. When two explanations produce the same measurements, no inference algorithm can separate them without further evidence or further assumptions.*

22.5.4 Physical balance equations add diagnostic information

A second sensor is not the only source of redundancy. A physical balance equation relates measurements that are otherwise independent, and so provides *analytical redundancy*. For the cooling-water loop, the energy balance of the water gives the thermal power removed from the load,

$$P_{\text{th}} = \dot{m}\, c_p\, (T_{\text{return}} - T_{\text{supply}}), \tag{22.27}$$

where $\dot{m}$ is the mass flow rate in kg/s, c_p the specific heat capacity of water, about $4.186\,\text{kJ/(kg K)}$ (not the cost of a preventive replacement), and T_{return} and T_{supply} the return and supply water temperatures; the units are $\text{kg/s} \cdot \text{kJ/(kg K)} \cdot \text{K} = \text{kW}$. If electrical and IT telemetry give an independent estimate of the heat load, equation (22.27) becomes a consistency check between temperature, flow and electrical measurements.

As an illustration, let the pump of Section 22.4.1 circulate $0.050\,\text{m}^3/\text{s}$, so $\dot{m} = 50\,\text{kg/s}$, with a temperature rise of 6 K. The balance gives

$$P_{\text{th}} = 50 \times 4.186 \times 6 = 1256\,\text{kW},$$

consistent with an IT load estimated at 1256 kW. A reported return temperature that implies a rise of 12 K would mean 2512 kW, twice the plausible load. More subtly, a return sensor reading 1 K high gives an apparent rise of 7 K and 1465 kW. The balance residual is $e_{\text{bal}} = 1256 - 1465 = -209\,\text{kW}$. With a flow uncertainty of 2 %, a standard deviation of 0.2 K on each temperature (so $0.2\sqrt{2} = 0.283\,\text{K}$ on their difference) and 2 % on the IT estimate, the relative uncertainty of the thermal power is $\sqrt{0.02^2 + (0.283/7)^2} = 0.045$, that is 66 kW, the IT estimate contributes 25 kW, and $\sigma_e = \sqrt{66^2 + 25^2} = 71\,\text{kW}$. The standardized residual is $z = -209/71 = -2.96$: the inconsistency is detected.

This is the check that power-system state estimation calls *bad-data detection*, and that process monitoring calls the first half of fault detection and isolation. The detection does not isolate the fault. The same residual is explained by an incorrect temperature measurement, an

incorrect flow measurement (a flow reading 16.7 % high gives the same imbalance), an incorrect heat-load estimate, heat exchanged through a path the balance omits, or an inadequate model. The balance equation narrows the possible explanations; it does not identify a single fault. Isolation needs further equations, or measurements that enter some balances and not others, so that each fault has a distinct signature in the sense of Section 22.5.2. The credibility of each added constraint depends on its own measurements, on the definition of its boundary and on their uncertainty; a balance with a poorly defined boundary adds residuals that look like faults.

22.5.5 Entropy and expected information gain

Not every further measurement is equally valuable. After the first observation of Section 22.4.4, the posterior assigns 0.097 to the healthy pump, 0.179 to deterioration and 0.725 to the power-sensor bias. Four further measurements are possible: a new vibration measurement, an independent electrical power reading with a portable analyzer, a bearing temperature measurement, and a calibration test of the flow meter. A measurement whose outcome has nearly the same distribution under every hypothesis carries little diagnostic information; one whose outcome the hypotheses predict differently can reduce the uncertainty substantially.

The uncertainty of a discrete diagnosis is measured by the Shannon entropy of its posterior,

$$\mathcal{H}(\mathbf{y}) = -\sum_i \Pr(H_i \mid \mathbf{y}) \, \log_2 \Pr(H_i \mid \mathbf{y}), \tag{22.28}$$

in bits. It is zero when one hypothesis has probability one, and at most $\log_2 3 = 1.585$ bits for three hypotheses, reached when they are equally probable. For the current posterior,

$$\mathcal{H} = 0.0966 \times 3.372 + 0.1787 \times 2.484 + 0.7247 \times 0.465 = 0.326 + 0.444 + 0.337 = 1.106 \text{ bits}.$$

Let Y be the outcome of a candidate measurement. Each outcome would produce its own posterior by equation (22.23), with $\mathbf{y}$ extended by Y, and its own entropy $\mathcal{H}(\mathbf{y}, Y)$. Before the measurement is made the outcome is unknown, but its predictive probability is known, $\Pr(Y \mid \mathbf{y}) = \sum_i \Pr(Y \mid H_i) \, \Pr(H_i \mid \mathbf{y})$, assuming that the new measurement is conditionally independent of the old ones given the hypothesis. The *expected information gain* is the expected reduction of entropy,

$$\mathrm{EIG}(Y) = \mathcal{H}(\mathbf{y}) - \mathbb{E}_Y\big[\mathcal{H}(\mathbf{y}, Y)\big]. \tag{22.29}$$

It equals the mutual information between the diagnosis and the outcome, which gives a simpler formula. Writing $\pi_i = \Pr(H_i \mid \mathbf{y})$ and $q_{yi} = \Pr(Y = y \mid H_i)$, the joint probability is $\pi_i q_{yi}$, and expanding both entropies,

$$\mathrm{EIG}(Y) = \sum_{i,y} \pi_i q_{yi} \, \log_2 \frac{q_{yi}}{\Pr(Y = y \mid \mathbf{y})} = \mathcal{H}_Y - \sum_i \pi_i \, \mathcal{H}_{Y \mid H_i},$$

the entropy of the predicted outcome minus the average entropy of the outcome under each hypothesis. A measurement is informative when its outcome is uncertain overall but nearly certain under each hypothesis. $\mathrm{EIG} \geq 0$, with equality when the outcome distribution is the same under every hypothesis.

For a test with outcomes Y^+ and Y^-, $q_{+,i} = \Pr(Y^+ \mid H_i)$, and the entropies are binary. Table 22.9 evaluates the four candidates and two further ones, each reduced to a pass/fail outcome. The independent power reading has standard deviation 1 kW and reads "high" (Y^+) when its residual exceeds 2 kW: under H_1 the true power is 4 kW above the model and $\Pr(Y^+) = \Pr(\mathcal{N}(4,1) > 2) = 0.977$, while under H_0 and H_2 the true power is as modeled and $\Pr(Y^+) = 0.023$. The new vibration measurement, standard deviation 1 mm/s, reads high above 1 mm/s: 0.841 under H_1, 0.159 otherwise. The bearing temperature rises only weakly with the assumed deterioration: $\Pr(Y^+) = 0.30$ under H_1 against 0.20 otherwise. A calibration

test of the power meter detects a bias with sensitivity 0.90 (the probability of Y^+ when the bias is present) and specificity 0.95 (the probability of Y^- when it is absent). The flow-meter calibration finds an error with probability 0.05 under every hypothesis, since none of the three involves the flow meter. A teardown inspection reveals the true hypothesis. The columns give the three conditional probabilities, the predictive probability of Y^+, the posterior of H_2 after each outcome, and the expected information gain. One row of the arithmetic, the independent power reading: $\Pr(Y^+) = 0.0966 \times 0.023 + 0.1787 \times 0.977 + 0.7247 \times 0.023 = 0.193$, whose binary entropy is 0.708 bits; the binary entropy of each conditional outcome is 0.157 bits; so $\mathrm{EIG} = 0.708 - 0.157 = 0.552$ bits.

Table 22.9: Expected information gain of six candidate measurements after the observation $\mathbf{y} = (4.0\,\mathrm{kW}, 0.3\,\mathrm{mm/s})$, posterior $(0.097, 0.179, 0.725)$ and entropy 1.106 bits. Outcome probabilities illustrative.

Measurement	$q_{+,0}$	$q_{+,1}$	$q_{+,2}$	$\Pr(Y^+)$	$\Pr(H_2 \mid Y^+)$	$\Pr(H_2 \mid Y^-)$	EIG (bits)
Independent power reading	0.023	0.977	0.023	0.193	0.085	0.878	0.552
Power-meter calibration	0.05	0.05	0.90	0.666	0.979	0.217	0.500
New vibration measurement	0.159	0.841	0.159	0.281	0.410	0.848	0.225
Bearing temperature	0.20	0.30	0.20	0.218	0.665	0.741	0.006
Flow-meter calibration	0.05	0.05	0.05	0.050	0.725	0.725	0.000
Teardown inspection	reveals the hypothesis			–	–	–	1.106

The ranking follows the separation of the conditional probabilities. The independent power reading separates deterioration from the other two almost perfectly, and the power-meter calibration separates the bias from the other two; each is expected to remove about half of the current 1.106 bits of uncertainty. The new vibration measurement repeats a test whose outcome is noisier relative to its separation and gains less. The bearing temperature barely discriminates, and the flow-meter calibration not at all within this hypothesis set; its value lies in testing a hypothesis the set omits, the subject of Section 22.6.3. The teardown removes all the uncertainty. A measurement with a larger expected information gain is expected to reduce the uncertainty of the diagnosis more; it is not necessarily the best measurement to make.

22.5.6 Value of information for maintenance decisions

The purpose of a diagnosis is a decision. Let the maintenance actions be

$$a \in \{\text{continue operation},\ \text{inspect},\ \text{repair},\ \text{calibrate sensor}\},$$

and let $C(a, H_i)$ be the loss, in reais, of action a if hypothesis H_i is true: the cost of the action plus the expected cost of what follows from it over the decision period. Table 22.10 gives illustrative values: continuing with a deteriorating pump risks an unplanned failure, R$ 30 000; continuing with a biased meter costs R$ 2 000 in wrong energy accounting and misguided control; an inspection costs R$ 1 500 and, under deterioration, leads to a planned repair of R$ 8 000; a repair costs R$ 8 000 whatever is true and leaves a biased meter uncorrected; a calibration costs R$ 500 and leaves a deteriorating pump in service. The last column is the expected loss under the current posterior, $\sum_i \Pr(H_i \mid \mathbf{y})\, C(a, H_i)$.

Without a new measurement, the best action is the one with the smallest expected loss,

$$C^*_{\text{now}} = \min_a \sum_i \Pr(H_i \mid \mathbf{y})\, C(a, H_i), \tag{22.30}$$

Table 22.10: Losses $C(a, H_i)$ of four maintenance actions under the three hypotheses (illustrative, R$), and their expected values under the posterior $(0.097, 0.179, 0.725)$.

Action a	H_0	H_1	H_2	Expected loss
Continue operation	0	30 000	2 000	6 811
Inspect	1 500	9 500	3 500	4 379
Repair	8 000	8 000	10 000	9 449
Calibrate sensor	500	30 500	500	5 861

here the inspection, at R$ 4 379. With a candidate measurement Y, the action is chosen after its outcome is seen, and the expected loss is averaged over the outcomes,

$$C^*_Y = \mathbb{E}_Y\Big[\min_a \sum_i \Pr(H_i \mid \mathbf{y}, Y)\, C(a, H_i)\Big]. \tag{22.31}$$

The *expected value of sample information* is the reduction,

$$\mathrm{EVSI}(Y) = C^*_{\mathrm{now}} - C^*_Y. \tag{22.32}$$

It is never negative before the cost of the measurement. Write $\bar{C}(a \mid Y) = \sum_i \Pr(H_i \mid \mathbf{y}, Y) C(a, H_i)$; averaging it over Y returns the current expected loss of a, because the posterior averaged over the predictive distribution of the outcome is the current posterior. For every action a, $\min_{a'} \bar{C}(a' \mid Y) \le \bar{C}(a \mid Y)$ for each outcome, hence $C^*_Y \le \mathbb{E}_Y[\bar{C}(a \mid Y)]$, the current expected loss of a; taking the best a gives $C^*_Y \le C^*_{\mathrm{now}}$. A decision maker free to ignore the outcome cannot be worse off for having it. The upper limit is the expected value of perfect information, obtained when the outcome reveals the hypothesis: $\sum_i \Pr(H_i \mid \mathbf{y}) \min_a C(a, H_i) = 0.179 \times 8\,000 + 0.725 \times 500 = \text{R\$ } 1\,792$, so that the value of perfect information is $4\,379 - 1\,792 = \text{R\$ } 2\,587$.

Table 22.11 evaluates the candidates of Table 22.9. The columns give the best action after each outcome, the expected loss with the measurement, the EVSI, an illustrative cost of the measurement, and the net value, EVSI minus cost. One row of the arithmetic, the independent power reading: after Y^+ (probability 0.193) the posterior is $(0.011, 0.903, 0.085)$ and repair is best at R$ 8 171; after Y^- the posterior is $(0.117, 0.005, 0.878)$ and calibration is best at R$ 651; so $C^*_Y = 0.1934 \times 8\,171 + 0.8066 \times 651 = \text{R\$ } 2\,105$ and $\mathrm{EVSI} = 4\,379 - 2\,105 = \text{R\$ } 2\,274$.

Table 22.11: Expected value of sample information of the candidate measurements, for the losses of Table 22.10 (illustrative costs, R$). Without a measurement the best action is to inspect, at R$ 4 379.

Measurement	After Y^+	After Y^-	C^*_Y	EVSI	Cost	Net
Independent power reading	repair	calibrate	2 105	2 274	400	1 874
Power-meter calibration	calibrate	inspect	2 605	1 774	600	1 174
New vibration measurement	inspect	calibrate	3 064	1 315	200	1 115
Bearing temperature	inspect	inspect	4 379	0	50	−50
Flow-meter calibration	inspect	inspect	4 379	0	1 200	−1 200
Teardown inspection	best action per hypothesis		1 792	2 587	6 000	−3 413

Two results separate information from value. The bearing temperature has a positive information gain, 0.006 bits, and zero value: neither of its outcomes changes the best action, so it cannot change the loss. The teardown inspection removes all the uncertainty and has the largest EVSI, but its cost, which includes a long shutdown of the pump, makes it the worst

choice; the quick independent power reading, less informative, has the largest net value. The measurement that best distinguishes the hypotheses is not always the one that most improves the engineering decision, and a measurement is economically worthwhile only when its expected decision benefit exceeds its total acquisition and operating cost. The choice of measurements returns, for a whole network, in the later chapter on the value of information.

22.6 Validation and application of diagnosis

A diagnostic system is itself a model and must be shown to work. This section lists the measures of diagnostic performance, works a validation experiment, treats the hypotheses a diagnostic model omits, and places diagnosis in the workflow and the reliability decisions it serves.

22.6.1 Measures of diagnostic performance

A diagnostic algorithm evaluated only on carefully chosen examples can look better than it is; a credible assessment tests it systematically. Table 22.12 lists the measures that matter, each with its meaning. The *detection probability* P_{D} is the probability of flagging a fault when the fault is present, and $1 - P_{\mathrm{D}}$ is the probability of a missed detection; the *false-alarm probability* P_{FA} is the probability of flagging a fault during healthy operation. They are the sensitivity of the detector and the complement of its specificity.

Table 22.12: Measures of the performance of a diagnostic system.

Measure	Meaning
Detection probability P_{D}	probability of flagging a fault when the relevant fault is present
False-alarm probability P_{FA}	probability of flagging a fault during healthy operation
Detection delay	time between the occurrence of a fault and its detection
Diagnostic accuracy	frequency of correct classifications in the test population
Calibration	agreement between predicted probabilities and observed frequencies
Robustness	performance under changing load, environment, sensors and equipment
Unknown-fault handling	ability to recognize cases no candidate hypothesis explains

The measures answer different questions and can disagree. A detector can have excellent sensitivity and generate excessive false alarms. A classifier can achieve high accuracy while assigning poorly calibrated probabilities: if the cases to which it assigns 0.7 turn out to be correct 95 % of the time, its probabilities are not usable in the expected-loss calculation of Section 22.5.6, even though its classifications are good. A model can perform well in laboratory tests and fail when deployed under operating regimes it has not seen.

22.6.2 A validation experiment

Consider the validation of the pump model of Section 22.4. A test dataset is assembled with independent cases of healthy operation, of controlled hydraulic efficiency degradation, of controlled vibration-related mechanical faults, of artificially introduced power-sensor bias, of combined faults, and of changing speed and load. Model development is separated from evaluation by holding out equipment, operating periods or test runs, so that no case used to set the priors, signatures or covariance is used to judge them. The evaluation then checks six things: whether the healthy residuals follow the assumed distribution; whether false

alarms remain acceptable across operating conditions; whether the posterior probabilities are calibrated; whether the fault signatures remain distinguishable as noise increases; whether the system recognizes observations inconsistent with every modeled hypothesis; and whether the recommended actions remain appropriate under the diagnostic uncertainty.

Table 22.13 gives an illustrative outcome. Each row is a class of test cases, with its number of cases, the number for which the most probable hypothesis was H_0 (no alarm), H_1 or H_2, the detection rate and the rate of correct classification. One row of the arithmetic, the healthy cases: 11 of 600 raised an alarm, so $\hat{P}_{\mathrm{FA}} = 11/600 = 0.018$, with a 95 % binomial (Wilson) interval from 0.0103 to 0.0325, which excludes the nominal 0.01 of the design.

Table 22.13: Illustrative validation of the three-hypothesis pump model on held-out test cases. Detection: the most probable hypothesis is not H_0. Correct: the most probable hypothesis is the true one.

True class	Cases	Class H_0	Class H_1	Class H_2	Detection rate	Correct rate
Healthy	600	589	4	7	0.018 (false alarm)	0.982
Hydraulic efficiency loss	150	21	52	77	0.860	0.347
Mechanical (vibration) fault	150	9	135	6	0.940	0.900
Power-sensor bias	150	12	9	129	0.920	0.860
Combined faults	50	1	31	18	0.980	–
Flow-meter bias	50	6	4	40	0.880	–

The table shows what an experiment can reveal that the worked example could not. The false-alarm rate is about twice its nominal value, as the warnings of Section 22.2.4 about repeated, correlated tests anticipate. Detection is good for every fault class, but classification is poor for hydraulic efficiency loss: 77 of its 150 cases are classified as sensor bias, because an efficiency loss from impeller wear raises the power without raising the vibration, and so has the signature of H_2 rather than the $(4, 2)$ assumed for H_1. The single deterioration hypothesis lumps two mechanisms with different signatures, and the validation shows that it should be split. The combined faults have no correct class in the model, and the flow-meter bias is assigned to the power-sensor bias in 40 of 50 cases, the problem of the next subsection. Over the four classes the model represents, the accuracy is $(589 + 52 + 135 + 129)/1050 = 0.862$, a figure that hides all of these findings.

22.6.3 Missing hypotheses

The Bayesian model of Section 22.4 assumed that the truth is one of H_0, H_1 and H_2. Suppose instead that the flow meter reads low, 0.040 instead of $0.050\,\mathrm{m^3/s}$, on a healthy pump. The predicted power is then $0.040 \times 300\,000/0.75 = 16\,\mathrm{kW}$ while the pump draws 20 kW: the power residual is +4 kW with no vibration excess, exactly the signature of the power-sensor bias. The model assigns the posterior of equation (22.23), which sums to one over its three hypotheses whatever the truth, and here selects H_2 with probability 0.725; the recommended calibration of the power meter would find nothing. This is *model-set incompleteness*: a posterior is a statement relative to the hypothesis set, and a high posterior probability among incomplete hypotheses does not make the selected one correct.

One remedy is to add an explicit hypothesis for other or unmodeled explanations,

$$H_{\mathrm{U}} : \text{other or unmodeled explanation},$$

but adding H_{U} is not enough by itself: its likelihood must be defined, for example as a broad distribution of residuals, and its prior defended, or it either never wins or always does. A second

strategy is a model-adequacy check, the discrete counterpart of a posterior predictive check: compare the observed residuals with the distribution each hypothesis predicts for them. Under H_i the discrepancy d_i^2 of equation (22.21) is χ^2_2, so an observation for which even the smallest d_i^2 exceeds 9.21 is inconsistent at the 1 % level with every modeled explanation. For the residuals $(-3\,\mathrm{kW}, 3\,\mathrm{mm/s})$, a power below expectation with high vibration, $d^2 = (13.0, 22.8, 30.8)$; the posterior gives H_0 probability 0.998, yet $\Pr(\chi^2_2 > 13.0) = 0.0015$ even for H_0. The appropriate result is then that none of the modeled explanations adequately reproduces the measurements. That is a scientifically valuable conclusion: it directs attention to the system model, the fault library, the interpretation of the sensors or the operating-condition data before a confident maintenance recommendation is made. The flow-meter case passes this check, $d_2^2 = 0.09$, because its signature coincides with a modeled one; it is caught only by a measurement that tests it, such as the flow-meter calibration whose information gain within the three-hypothesis set was zero.

Common mistake

Reading the largest posterior probability as the identified fault. The posterior distributes probability over the hypotheses supplied, and it sums to one whether or not any of them is true. A diagnosis is complete only when the adequacy of the hypothesis set has been checked and the alternatives with the same signature have been named.

22.6.4 The condition-monitoring workflow

The methods of the chapter assemble into one workflow that connects measurements, physics, inference and action. Figure 22.7 draws it as a pipeline of eight stages, read from top to bottom; the box on the left names each stage and the text on the right states what it does. The arrow on the left returns from the last stage to the first: when the selected action is a further test, its result re-enters the pipeline as a new measurement, which is the sequential use of the expected information gain and the value of information.

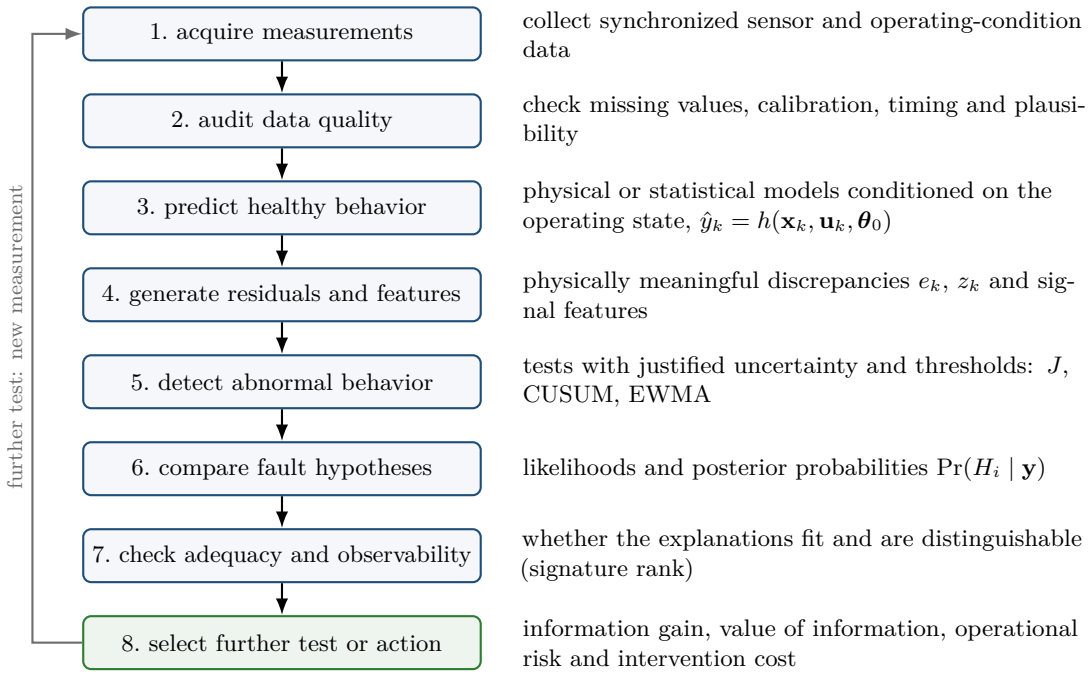


Figure 22.7: The condition-monitoring workflow, from measurements to a maintenance decision. A further test selected at stage 8 returns its outcome to stage 1.

Stages 1 to 5 are detection; stages 6 and 7 are diagnosis; stage 8 is the decision. Stage 7 is the one most often omitted: it asks whether the posterior of stage 6 can be trusted, by the adequacy check of Section 22.6.3, and whether the competing explanations are distinguishable at all, by the signature rank of Section 22.5.2.

22.6.5 Connection to system reliability

Fault diagnosis is not an isolated signal-processing exercise; its output changes the reliability of the system. Suppose further evidence has raised the posterior probability of significant pump deterioration to $\Pr(H_1 \mid \mathbf{y}) = 0.80$, the rest being on explanations in which the pump is mechanically healthy. A healthy pump has the constant failure rate of the running case, $\lambda_0 = 2 \times 10^{-4}\,\mathrm{h}^{-1}$ (Chapter 9 of Volume I); suppose a deteriorated one has the illustrative rate $\lambda_1 = 5 \times 10^{-3}\,\mathrm{h}^{-1}$, a value that would come from the degradation model of Chapter 21. The probability that the pump fails in the next 24 h is the mixture over the diagnosis,

$$\begin{aligned}\Pr(\text{pump fails in } 24\,\mathrm{h} \mid \mathbf{y}) &= 0.20\,(1 - \mathrm{e}^{-0.0048}) + 0.80\,(1 - \mathrm{e}^{-0.12}) \\ &= 0.20 \times 0.00479 + 0.80 \times 0.1131 = 0.0914.\end{aligned}$$

If a standby pump takes over with failure probability per demand $q_\mathrm{d} = 0.05$ (illustrative, as in Chapter 14 of Volume I), and ignoring failures of the standby while running, the probability of a loss of cooling in the next 24 h is

$$\Pr(\text{cooling loss in next } 24\,\mathrm{h} \mid \mathbf{y}) = 0.0914 \times 0.05 = 0.0046,$$

against $0.0210 \times 0.05 = 0.0011$ at the prior probability of deterioration, 0.15: the diagnosis has raised the risk more than fourfold. The operator can then compare actions: continue operating, reduce the pump's loading, switch to the redundant pump, inspect, schedule a repair, or replace the pump. Each changes the future state of the system, its failure exposure, its cost and its operational risk, and they are compared by the expected-loss calculation of Section 22.5.6 with losses drawn from the system model. The complete decision chain is

$$\text{measurements} \;\rightarrow\; \text{diagnosis} \;\rightarrow\; \text{updated reliability} \;\rightarrow\; \text{maintenance decision},$$

the bridge between diagnostics, prognostics and the availability optimization of Chapter 20 of Volume I. The diagnosis fixes which damage is present; the next chapter, Chapter 23, predicts how long the item will last with it.

22.7 What to retain

Definition 22.1 (Model-based fault diagnosis)**.** Fault detection tests the residuals $e = y - \hat{y}$ between measurements and a healthy model against their healthy uncertainty; fault diagnosis assigns posterior probabilities $\Pr(H_i \mid \mathbf{y}) \propto p(\mathbf{y} \mid H_i)\Pr(H_i)$ to competing physical and measurement explanations, which the evidence can separate only as far as their signatures differ, and whose credibility rests on a hypothesis set checked for adequacy.

The main conclusions are eight. Anomaly detection is not fault diagnosis: an unexpected measurement does not identify its cause. Operating conditions must be accounted for, since a healthy system shows different measurements at different loads, speeds and temperatures. Residual significance depends on uncertainty: 4 kW may be decisive or negligible. Bayesian diagnosis combines prior evidence with fault-specific predictions, and the most plausible explanation can change substantially when a new measurement arrives. Sensor faults are genuine competing hypotheses, since faulty instrumentation can imitate physical deterioration. Observability limits diagnosis: faults with indistinguishable signatures cannot be resolved by any algorithm

without further evidence or assumptions. Further measurements should be chosen by how they distinguish hypotheses and improve decisions. And model adequacy must be checked, because a high posterior among incomplete hypotheses does not make the selected explanation correct. Table 22.14 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 22.14: The terms of condition monitoring and fault diagnosis and where each is defined.

Term	Expression	Units	Meaning	Defined in
Residual	$e_k = y_k - \hat{y}_k$	as y	observed minus healthy prediction	22.2.2
Standardized residual	$z = e/\sigma_e$	–	residual in healthy standard deviations	22.2.3
Chi-square statistic	$J = \mathbf{e}^{\mathsf{T}}\boldsymbol{\Sigma}_e^{-1}\mathbf{e} \sim \chi^2_m$	–	multivariate detection test	22.2.4
Wire-to-water efficiency	$\eta = Q\,\Delta p / P_{\mathrm{e}}$	–	hydraulic over electrical power	22.2.1
RMS, crest factor, kurtosis	x_{RMS}, CF, Ku	mm/s, –, –	energy and impulsiveness of a signal	22.3.2
Bearing defect frequencies	f_{BPFO}, f_{BPFI}, f_{BSF}, f_{FTF}	Hz	kinematic frequencies of defects	22.3.3
EWMA	$s_k = \alpha y_k + (1-\alpha)s_{k-1}$	as y	smoothed level of an indicator	22.3.4
CUSUM	$S_k^+ = \max(0, S_{k-1}^+ + z_k - \kappa)$	–	accumulated evidence of a shift	22.3.5
Average run length	ARL_0, ARL_1	samples	mean time to alarm	22.3.5
Likelihood of a hypothesis	$p(\mathbf{y} \mid H_i) \propto \mathrm{e}^{-d_i^2/2}$	–	fit of the observation to H_i	22.4.3
Posterior probability	$\Pr(H_i \mid \mathbf{y})$	–	plausibility of H_i after the evidence	22.4.4
Fault signature	$\mathbf{e} = \mathbf{S}_f\mathbf{f} + \boldsymbol{\varepsilon}$	–	response of residuals to faults	22.5.2
Diagnostic observability	$\operatorname{rank}\mathbf{S}_f = q$	–	faults distinguishable from data	22.5.3
Bad-data detection	balance residual e_{bal}	kW	check by an energy balance	22.5.4
Entropy	$\mathcal{H} = -\sum_i \Pr(H_i \mid \mathbf{y}) \log_2 \Pr(H_i \mid \mathbf{y})$	bits	uncertainty of the diagnosis	22.5.5
Expected information gain	$\mathcal{H}(\mathbf{y}) - \mathbb{E}_Y[\mathcal{H}(\mathbf{y}, Y)]$	bits	expected entropy reduction	22.5.5
EVSI	$C^*_{\mathrm{now}} - C^*_Y$	R$	expected decision benefit	22.5.6
P_{D}, P_{FA}	detection, false-alarm probability	–	sensitivity; 1− specificity	22.6.1
Model-set incompleteness	truth outside $\{H_i\}$	–	posterior relative to the set	22.6.3

Further reading. Patton, Frank and Clark, *Fault Diagnosis in Dynamic Systems*: residual-based fault detection and model-based diagnosis. Gertler, *Fault Detection and Diagnosis in Engineering Systems*: analytical redundancy, fault signatures and diagnosability. NIST, *Process Monitoring and Statistical Control*: control charts, CUSUM, EWMA and multivariate monitoring. NIST, *EWMA Control Charts*: detecting sustained process shifts. NASA, *Prognostics Center of Excellence*: physics-based health monitoring, uncertainty and remaining-life prediction.

22.8 Exercises

Exercise 22.1. (a) State the distinction between fault detection and fault diagnosis. (Detection identifies abnormal behavior; diagnosis distinguishes its likely causes. A significant residual may establish abnormality without identifying the mechanism responsible.) (b) Explain why the electrical power of a pump should be conditioned on flow and pressure rise before it is used as a condition indicator. (Hydraulic power is $Q\,\Delta p$; a change of flow or pressure rise changes the power drawn without any change of efficiency.) (c) Define a standardized residual, and compute it for a residual of 4 kW with healthy standard deviations of 1, 1.5 and 5 kW. (The residual divided by its modeled standard deviation, which measures the discrepancy relative to the expected healthy variability; 4, 2.67 and 0.8.)

Exercise 22.2. (a) Explain why several measurements help diagnosis. (Different faults can produce different multivariable signatures; measurements help when the competing hypotheses predict distinguishable responses, and an extra sensor does not by itself guarantee it.) (b) State what a Bayesian posterior probability of a hypothesis depends on. (The prior probabilities and the likelihood of the measurements under each hypothesis, through Bayes' theorem, and therefore on the hypothesis set, the fault signatures and the noise model.) (c) State what should be done when every modeled hypothesis fits the measurements poorly. (Investigate missing faults, model inadequacy and measurement problems, rather than accept the highest-ranked hypothesis.) (d) A second return-temperature sensor is calibrated against the same reference as the first and shares its bias exactly, so that $e_{T,2} = \delta T_{\text{real}} + b_T + \varepsilon_2$. Write the signature matrix and its rank, and state what the second sensor adds. ($\mathbf{S}_f$ has two rows $(1, 1)$, rank one; it reduces the noise of the average but cannot separate the real deviation from the bias.)

Exercise 22.3. A pump runs at $Q = 0.040\,\text{m}^3/\text{s}$ against $\Delta p = 250\,\text{kPa}$, with a healthy wire-to-water efficiency of 0.72. The measured electrical power is 16.5 kW and the healthy residual standard deviation is 1.0 kW. (a) Compute the hydraulic power, the head, the expected electrical power and the apparent efficiency. (10.0 kW; 25.5 m; 13.89 kW; 0.606.) (b) Compute the residual, the standardized residual and the two-sided probability of a residual at least this extreme under the healthy model, and state the conclusion at the 1 % level. (2.61 kW; 2.61; 0.0090; the healthy model is rejected, since $2.61 > 2.576$, but the cause is not diagnosed.) (c) The standardized residuals of the following samples are 0.2, 1.1, 0.9, 1.4, 0.7, 1.3, 1.0, 1.6. Compute the one-sided CUSUM with $\kappa = 0.5$ and find when it crosses $h_{\text{C}} = 4$. (0, 0.6, 1.0, 1.9, 2.1, 2.9, 3.4, 4.5: an alarm at the eighth sample, although no sample exceeds 1.6.)

Exercise 22.4. For the three-hypothesis model of Section 22.4 (priors 0.70, 0.15, 0.15; means $(0, 0)$, $(4, 2)$, $(4, 0)$; standard deviations 1.5 kW and 1.0 mm/s), the residuals are $\mathbf{y} = (3.0\,\text{kW}, 1.2\,\text{mm/s})$. (a) Compute d_i^2 and the posterior probabilities. ($d^2 = (5.440, 1.084, 1.884)$; weights 0.0461, 0.0872, 0.0585; posterior $(0.240, 0.455, 0.305)$.) (b) Recompute the posterior with priors 0.50, 0.25, 0.25. ($(0.119, 0.527, 0.353)$.) (c) Compute the entropy of the posterior of (a) and the expected information gain of a power-meter calibration with sensitivity 0.90 and specificity 0.95. (1.534 bits; $\Pr(Y^+) = 0.309$, posteriors $(0.039, 0.074, 0.888)$ after Y^+ and $(0.331, 0.625, 0.044)$ after Y^-, EIG = 0.550 bits.) (d) Check the adequacy of the hypothesis set for this observation. (The smallest discrepancy is $d_1^2 = 1.084$, and $\Pr(\chi_2^2 > 1.084) = 0.58$: the observation is consistent with H_1, and the set is not contradicted.)

Chapter 23

Remaining Useful Life Estimation

23.1 Remaining useful life as a conditional random variable

Chapter 21 described how damage accumulates in a component, by deterministic laws and by stochastic processes, and defined failure as the crossing of a critical damage level. Chapter 22 described how the condition of a component is inferred from noisy measurements and how competing explanations of an abnormal reading are told apart. This chapter combines the two. Its question is

given what is known about a component's current condition, what is the probability that it fails within the next 10, 30 or 90 days, and how much confidence that prediction deserves.

The answer is a probability distribution, not a number. A prediction of the form "remaining useful life: 60 days" hides what a maintenance planner needs; a prediction of the form "a 10 % probability of reaching the critical degradation threshold within 32 days, 50 % within 58 days and 90 % within 91 days" states it. Those three numbers are illustrative (the bearing model worked in Section 23.2.3 gives 34.6, 55.9 and 90.7 days), but they show the point of the chapter: maintenance decisions depend more on the lower tail of the remaining life than on its mean.

The chapter proceeds in six steps. It defines the remaining useful life as a conditional first-passage time and the functions that describe it. It derives the remaining-life distribution of a Wiener degradation process from a known current state and works a cooling-pump bearing through it. It replaces the known state by a posterior distribution obtained by Bayesian filtering, and adds uncertainty in the model parameters. It treats monotone damage with the gamma process and nonlinear, non-Gaussian problems with the particle filter. It states how a predicted distribution is tested against outcomes. Finally, it turns the distribution into maintenance decisions and states what makes a prediction credible.

Jargon: RUL, prognostics, PHM, degradation state, failure threshold, first-passage time, Bayesian filter, Kalman filter, particle filter, calibration, sharpness, prognostic horizon, lead time

The *remaining useful life* (RUL) of a component at the present time t_0 is the time from t_0 until the component first fails to meet its required function, given everything known at t_0; it is a random variable, in days, hours or years. *Prognostics* is the prediction of the RUL; *prognostics and health management* (PHM) is the discipline that monitors condition, diagnoses faults, predicts the RUL and decides maintenance from the prediction. The *degradation state* $D(t)$ is a physical or normalized measure of accumulated damage; the *failure threshold* D_{crit} is the damage level at which the required function is lost. The *first-passage time* is the first time a stochastic process reaches a given level. A *Bayesian filter* updates the probability distribution of a hidden state as measurements arrive; the *Kalman filter* is its exact form for linear models with Gaussian noise, and the *particle filter* approximates it with weighted random samples. *Calibration* is the agreement between predicted probabilities and observed frequencies; *sharpness* is the concentration of the predicted distribution. The *prognostic horizon* is how long before failure a prediction first meets, and keeps meeting, an accuracy criterion. The *lead time* is the time an action needs before it takes effect, such as the delivery time of a spare part.

Key idea

The remaining useful life is a first-passage time of the degradation state to the failure threshold, conditioned on the evidence available now. Its distribution combines three separate uncertainties: what the current damage is, how fast damage will grow, and the randomness of the future damage path. Decisions are taken from its lower tail, not from its mean.

This section defines the remaining useful life as a random variable, names the functions and quantiles that describe it, and shows why its mean is a poor summary.

23.1.1 Failure as a threshold crossing

Let $D(t)$ be the degradation state of a component at time t and D_{crit} the failure threshold, the damage level at which the component no longer performs its required function. The time to failure is defined as the first time the state reaches the threshold,

$$T_{\text{f}} = \inf\{t \geq 0 : D(t) \geq D_{\text{crit}}\}, \tag{23.1}$$

the first-passage time of D to D_{crit}. The symbol T_{f} is the random time to failure T of Chapter 5 of Volume I when failure is produced by damage accumulation; the subscript marks that it is generated by a degradation model.

Remark 23.1. The letter D in this chapter is the degradation state of Chapter 21, in the units of the damage measure or normalized so that $D_{\text{crit}} = 1$. It is not the FMEA detection rating D of Chapter 4 of Volume I.

Suppose the present time is t_0 and the component has not failed. Its remaining useful life is

$$\text{RUL}(t_0) = T_{\text{f}} - t_0, \qquad \text{given } T_{\text{f}} > t_0, \tag{23.2}$$

and the object of prognostics is its conditional distribution

$$p(\text{RUL}(t_0) \mid \mathcal{Y}_{0:t_0}, T_{\text{f}} > t_0), \tag{23.3}$$

where $\mathcal{Y}_{0:t_0}$ denotes the measurements and other evidence collected from the start of service to t_0: inspections, condition indicators, operating history. The conditioning is essential. The

distribution changes when a measurement arrives, when the expected operating conditions change and when the degradation model is revised; an RUL prediction is always a prediction made at a stated time on stated evidence.

The definition extends a Volume I result. When the only evidence is that the component has survived to t_0, the remaining-life distribution is the conditional reliability of Chapter 5, $\Pr(\mathrm{RUL}(t_0) > \tau) = R(t_0 + \tau)/R(t_0)$ (equation (5.6)), and its mean is the mean residual life $\mathrm{MRL}(t_0) = \int_{t_0}^{\infty} R(u)\,\mathrm{d}u/R(t_0)$ (equation (5.20)). Both describe a population of components of age t_0. Condition evidence narrows the population to components in the measured state, and the remaining life of this particular component replaces the remaining life of an average one.

Figure 23.1 illustrates the definition with the Wiener degradation model worked in Section 23.2. The horizontal axis is time in days and the vertical axis the damage index, normalized so that $D_{\mathrm{crit}} = 1$. The black path is the history up to the present time $t_0 = 40$ days, where the damage is 0.40. The four dashed blue paths are possible futures simulated from that state; each stops where it first reaches the red threshold line, at 38, 52, 66 and 88 days after t_0. The shaded curve above the threshold is the density of the first-passage time measured from t_0, drawn on an arbitrary vertical scale. The spread of the crossing points is the spread of the remaining useful life.

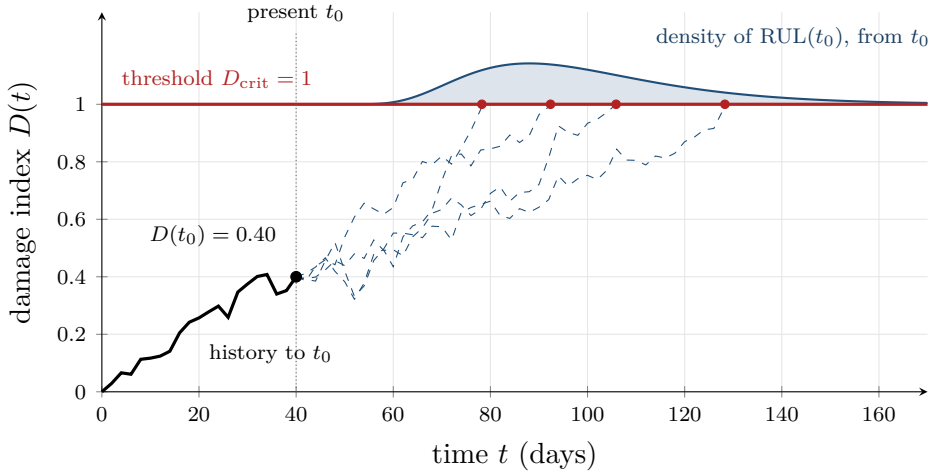


Figure 23.1: Remaining useful life as a first-passage time, illustrative. Black: a degradation history to $t_0 = 40$ days. Dashed: four future paths of the Wiener model $D(t_0 + \tau) = 0.40 + 0.01\tau + 0.03\,\tilde{B}(\tau)$, τ in days, simulated with daily steps; red dots: their first crossings of $D_{\mathrm{crit}} = 1$. Shaded: the inverse-Gaussian density of the first-passage time from t_0, equation (23.11), with $a = 0.6$, $\mu_D = 0.01\,\mathrm{day}^{-1}$, $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$, on an arbitrary vertical scale.

23.1.2 RUL reliability and probability of failure

The conditional distribution of equation (23.3) is described by the same functions as a lifetime in Chapter 5, with the age variable replaced by the time τ measured forward from t_0. The *RUL reliability* is the probability of surviving a further τ without reaching the threshold,

$$R_{\mathrm{RUL}}(\tau) = \Pr(\mathrm{RUL}(t_0) > \tau \mid \mathcal{Y}_{0:t_0}, T_{\mathrm{f}} > t_0), \tag{23.4}$$

and the conditional probability of failure within τ is its complement,

$$F_{\mathrm{RUL}}(\tau) = 1 - R_{\mathrm{RUL}}(\tau). \tag{23.5}$$

The dependence on t_0 and on the evidence is kept implicit in the notation. Where F_{RUL} is differentiable, the RUL density is $f_{\mathrm{RUL}}(\tau) = \mathrm{d}F_{\mathrm{RUL}}/\mathrm{d}\tau$, in inverse time, and the mean remaining life follows from the tail-integral identity of Chapter 5, $\mathbb{E}[\mathrm{RUL}(t_0)] = \int_0^{\infty} R_{\mathrm{RUL}}(\tau)\,\mathrm{d}\tau$.

A statement such as $F_{\mathrm{RUL}}(30\ \text{days}) = 0.15$ means that the modeled probability of failure within the next 30 days, given the evidence, is 15 %. It answers the question an operator actually asks about a horizon fixed by the plan (the next outage window, the next delivery of spares, the end of a peak season), and for that reason it is usually more useful than a mean remaining life.

23.1.3 RUL quantiles

The α-quantile of the remaining life, for a probability level $0 < \alpha < 1$, is the shortest horizon by which the failure probability reaches α,

$$\mathrm{RUL}_{\alpha}(t_0) = \inf\{\tau \geq 0 : F_{\mathrm{RUL}}(\tau) \geq \alpha\}. \tag{23.6}$$

Three quantiles are used routinely: $\mathrm{RUL}_{0.10}$, the horizon by which 10 % of the modeled failures have occurred; the median $\mathrm{RUL}_{0.50}$; and $\mathrm{RUL}_{0.90}$, by which 90 % have occurred. When F_{RUL} is continuous, the infimum is attained and $F_{\mathrm{RUL}}(\mathrm{RUL}_{\alpha}) = \alpha$ exactly. Thus if $\mathrm{RUL}_{0.10} = 25$ days, the modeled probability of failure within 25 days is 10 %.

The lower quantiles carry the decision. A risk-averse operator who accepts at most a 10 % chance of failure before a planned intervention schedules it no later than $\mathrm{RUL}_{0.10}$; one who accepts 5 % schedules it by $\mathrm{RUL}_{0.05}$. The median and the upper quantiles say how long the component is likely to last if it is left in service, which matters for ordering and for the value of the life that an early replacement discards.

23.1.4 Why the mean remaining life can be misleading

Two components can have the same mean remaining life and very different probabilities of early failure. Figure 23.2 shows two remaining-life densities, both with a mean of 60 days. The horizontal axis is the remaining life in days and the vertical axis the density per day. Both curves are inverse-Gaussian first-passage densities of the kind derived in Section 23.2.2, with the same remaining distance and mean degradation rate but different diffusion. The green curve is narrow (standard deviation 11.6 days); the gold curve is wide (standard deviation 46.5 days) and skewed to the right, with its peak well before 60 days. The dashed line marks the common mean and the dotted line a lead time of 30 days. The shaded area under the wide curve to the left of 30 days is 0.267; the corresponding area under the narrow curve is 0.000176. If a replacement part needs 30 days to arrive, the first component can wait for an order placed today and the second cannot. A narrow and a broad prediction with the same central value are not operationally equivalent.

Even a single distribution is poorly summarized by its mean when it is skewed. Figure 23.3 draws the remaining-life density of the bearing example of Section 23.2.3, with the same axes as Figure 23.2, and marks three summaries: the 10 % quantile at 34.6 days, the median at 55.9 days and the mean at 60 days. The density peaks earlier still, at its mode of 48 days. The shaded area to the left of the 10 % quantile is 0.10 by definition. The long right tail pulls the mean above the median, so that the probability of failure before the mean is 0.575, not 0.5, and the planning-relevant 10 % quantile lies 25 days before the mean.

Common mistake

Reporting the mean remaining life alone. Two predictions with the same mean can differ by three orders of magnitude in the probability of failure before a spare part arrives, and a skewed distribution puts more than half of its probability before its mean. A prediction that supports a decision states the probability of failure within the decision horizon, or the lower quantile, together with the evidence and the assumed operating conditions.

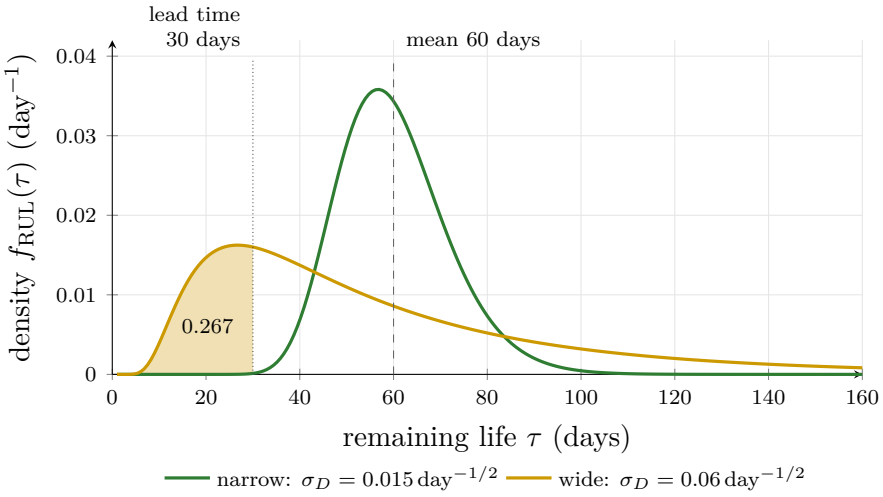


Figure 23.2: Two remaining-life distributions with the same mean of 60 days. Both are the inverse-Gaussian density of equation (23.11) with $a = 0.6$ and $\mu_D = 0.01\,\mathrm{day}^{-1}$; the diffusion differs. Shaded: the probability of failure within a 30-day lead time under the wide distribution, 0.267; under the narrow one it is 0.000176. Illustrative.

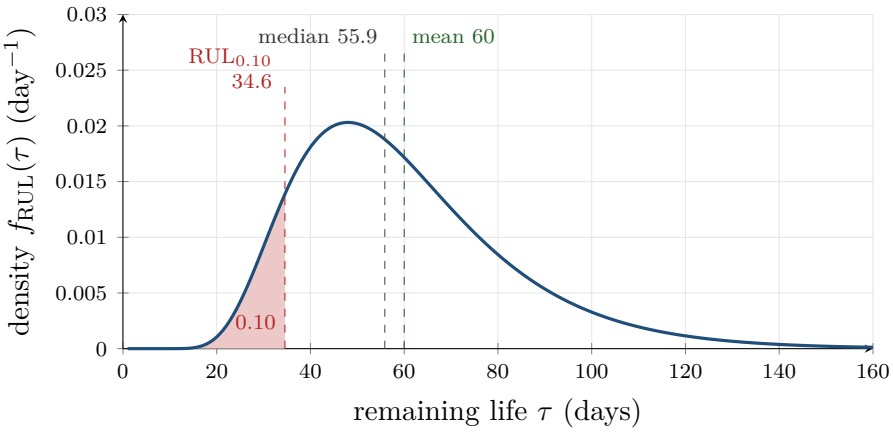


Figure 23.3: The right-skewed remaining-life density of the cooling-pump bearing, equation (23.11) with $a = 0.6$, $\mu_D = 0.01\,\mathrm{day}^{-1}$ and $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$. Shaded: the 10 % of the probability that lies before $\mathrm{RUL}_{0.10} = 34.6$ days. The median, 55.9 days, lies below the mean, 60 days; the mode is 48 days.

23.2 Wiener-process remaining useful life

When the degradation state is known exactly at t_0 and evolves as a Wiener process with constant parameters, the remaining-life distribution has a closed form. This section conditions the model on the current state, derives the first-passage distribution and works the cooling-pump bearing through it.

23.2.1 The degradation model conditioned on the current state

Chapter 21 introduced the Wiener (drift–diffusion) degradation model

$$D(t) = D_0 + \mu_D t + \sigma_D B(t), \tag{23.7}$$

where D_0 is the damage at the start of service, $\mu_D > 0$ the drift (the mean degradation rate, in damage units per unit time), $\sigma_D > 0$ the diffusion parameter (in damage units per square root of time), and $B(t)$ a standard Brownian motion: $B(0) = 0$, continuous paths, and independent increments $B(t+s) - B(t)$ distributed as $\mathcal{N}(0, s)$. Under this model $D(t) \sim \mathcal{N}(D_0 + \mu_D t,\, \sigma_D^2 t)$. The Brownian motion is written $B(t)$ because $W(t)$ is the cumulative intensity of a counting process in Chapter 17 of Volume I, and the drift μ_D is not the repair rate μ.

Suppose the damage at the present time is known exactly, $D(t_0) = d_0 < D_{\text{crit}}$, and define the *remaining threshold distance*

$$a = D_{\text{crit}} - d_0 > 0. \tag{23.8}$$

The future of the process follows in three steps. First, subtracting equation (23.7) at t_0 from the same equation at $t_0 + \tau$,

$$D(t_0 + \tau) - d_0 = \mu_D \tau + \sigma_D \left[B(t_0 + \tau) - B(t_0)\right] = \mu_D \tau + \sigma_D \tilde{B}(\tau),$$

where $\tilde{B}(\tau) = B(t_0 + \tau) - B(t_0)$. Second, $\tilde{B}$ is itself a standard Brownian motion, and by the independence of increments it is independent of the whole path of B up to t_0. Third, the conditioning event $\{T_{\text{f}} > t_0\}$, that the threshold was not reached before t_0, is determined by the path up to t_0; so is every past measurement. Conditioning on them therefore leaves the distribution of $\tilde{B}$ unchanged. The remaining life is the first time the future increment reaches the remaining distance,

$$\text{RUL}(t_0) = \inf\{\tau \geq 0 : \mu_D \tau + \sigma_D \tilde{B}(\tau) \geq a\}, \tag{23.9}$$

and its distribution depends on the history only through d_0. This is the Markov property of the Wiener process. The age t_0 does not appear: in a damage model the state carries the memory that the age carried in the aging hazard of Chapter 5 of Volume I.

23.2.2 The first-passage distribution

Chapter 21 gives the first-passage distribution of a Wiener process with positive drift to a level at distance a; with the conditioning of Section 23.2.1 the same result applies to the remaining life. Its cumulative distribution function is

$$F_{\text{RUL}}(\tau) = \Phi\left(\frac{\mu_D \tau - a}{\sigma_D \sqrt{\tau}}\right) + \exp\left(\frac{2\mu_D a}{\sigma_D^2}\right) \Phi\left(-\frac{\mu_D \tau + a}{\sigma_D \sqrt{\tau}}\right), \qquad \tau > 0, \tag{23.10}$$

where Φ is the standard normal cumulative distribution function. Because the formula is the working tool of this chapter, its origin is traced here in four steps, writing $X(\tau) = \mu_D \tau + \sigma_D \tilde{B}(\tau)$.

Step 1: split by the end point. The paths of X are continuous, so a path that ends at or above a has crossed a by τ. Hence $\Pr(\text{RUL} \le \tau) = \Pr\big(X(\tau) \ge a\big) + \Pr\big(\text{RUL} \le \tau,\, X(\tau) < a\big)$. The first term is a normal probability, $X(\tau) \sim \mathcal{N}(\mu_D\tau, \sigma_D^2\tau)$,

$$\Pr\big(X(\tau) \ge a\big) = 1 - \Phi\!\left(\frac{a - \mu_D\tau}{\sigma_D\sqrt{\tau}}\right) = \Phi\!\left(\frac{\mu_D\tau - a}{\sigma_D\sqrt{\tau}}\right),$$

which is the first term of equation (23.10).

Step 2: reflection without drift. For $\mu_D = 0$, the reflection principle, a standard property of Brownian motion, pairs each path that reaches a and ends at $x < a$ with the path reflected about a after its first passage, which ends at $2a - x$ and has the same probability. So $\Pr_0(\text{RUL} \le \tau,\, X(\tau) \in \mathrm{d}x) = \varphi_\tau(2a - x)\,\mathrm{d}x$, where φ_τ is the $\mathcal{N}(0, \sigma_D^2\tau)$ density and the subscript 0 marks zero drift.

Step 3: reweighting for the drift. Adding a constant drift changes the probability of a path by a factor that depends only on its end point, $\exp[\mu_D x/\sigma_D^2 - \mu_D^2\tau/(2\sigma_D^2)]$ for a path ending at x (the Cameron–Martin formula, also standard). Therefore

$$\Pr\big(\text{RUL} \le \tau,\, X(\tau) \in \mathrm{d}x\big) = \frac{1}{\sigma_D\sqrt{2\pi\tau}} \exp\!\left[-\frac{(2a - x)^2}{2\sigma_D^2\tau} + \frac{\mu_D x}{\sigma_D^2} - \frac{\mu_D^2\tau}{2\sigma_D^2}\right] \mathrm{d}x.$$

Step 4: complete the square. Expanding $(2a - x + \mu_D\tau)^2 = (2a - x)^2 - 2\mu_D\tau x + \mu_D^2\tau^2 + 4a\mu_D\tau$ shows that the exponent equals $-(x - 2a - \mu_D\tau)^2/(2\sigma_D^2\tau) + 2\mu_D a/\sigma_D^2$. The density in x is thus $\exp(2\mu_D a/\sigma_D^2)$ times a normal density with mean $2a + \mu_D\tau$ and variance $\sigma_D^2\tau$, and integrating over $x < a$ gives $\exp(2\mu_D a/\sigma_D^2)\,\Phi\big(-(a + \mu_D\tau)/(\sigma_D\sqrt{\tau})\big)$, the second term of equation (23.10). The second term is the probability of the paths that crossed the threshold and came back below it by τ; the probability that $D(t_0 + \tau)$ exceeds the threshold, the first term alone, misses them and understates the failure probability.

Differentiating equation (23.10) with respect to τ gives the first-passage density restated from Chapter 21,

$$f_{\text{RUL}}(\tau) = \frac{a}{\sigma_D\sqrt{2\pi\tau^3}} \exp\!\left[-\frac{(a - \mu_D\tau)^2}{2\sigma_D^2\tau}\right], \qquad \tau > 0, \tag{23.11}$$

the *inverse Gaussian distribution*, with the standard moments

$$\mathbb{E}[\text{RUL}] = \frac{a}{\mu_D}, \qquad \operatorname{Var}(\text{RUL}) = \frac{a\,\sigma_D^2}{\mu_D^3}. \tag{23.12}$$

As $\tau \to \infty$ the first term of equation (23.10) tends to one and the second to zero, so the threshold is reached with probability one when $\mu_D > 0$. (With $\mu_D < 0$ the same limit is $\exp(2\mu_D a/\sigma_D^2) < 1$, and there is a positive probability that the threshold is never reached.)

The two moments separate the roles of the parameters. The mean is the remaining distance divided by the mean rate, the deterministic estimate $(D_{\text{crit}} - d_0)/\mu_D$. The variance grows with σ_D^2: two components with the same estimated damage and the same mean rate have different remaining-life uncertainty if their damage paths fluctuate differently, which is the difference between the two curves of Figure 23.2. The distribution is skewed to the right, with skewness $3\sigma_D/\sqrt{a\mu_D}$. Its mode follows from setting the derivative of $\ln f_{\text{RUL}}$ to zero, $-3/(2\tau) + (a^2/\tau^2 - \mu_D^2)/(2\sigma_D^2) = 0$, which is the quadratic $\mu_D^2\tau^2 + 3\sigma_D^2\tau - a^2 = 0$ with positive root

$$\tau_{\text{mode}} = \frac{\sqrt{9\sigma_D^4 + 4\mu_D^2 a^2} - 3\sigma_D^2}{2\mu_D^2}, \tag{23.13}$$

always below the mean.

23.2.3 Worked example: the cooling-pump bearing

The bearing of the cooling-water pump carries a normalized degradation index with failure threshold $D_{\mathrm{crit}} = 1$. Condition monitoring has established the current damage as $d_0 = 0.40$, treated in this section as exact, so $a = 1 - 0.40 = 0.60$. The model parameters, the illustrative values of Chapter 21, are $\mu_D = 0.01\,\mathrm{day}^{-1}$ and $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$; time is in days in this example. Equation (23.12) gives

$$\mathbb{E}[\mathrm{RUL}] = \frac{0.60}{0.01} = 60 \text{ days}, \qquad \mathrm{Var}(\mathrm{RUL}) = \frac{0.60 \times 0.03^2}{0.01^3} = 540 \text{ days}^2,$$

so the standard deviation is $\mathrm{SD}(\mathrm{RUL}) = \sqrt{540} = 23.24$ days. Equation (23.13) gives the mode: $9\sigma_D^4 = 7.29 \times 10^{-6}$, $4\mu_D^2 a^2 = 1.44 \times 10^{-4}$, the square root of their sum is 0.0123, and $\tau_{\mathrm{mode}} = (0.0123 - 0.0027)/(2 \times 10^{-4}) = 48$ days.

Failure within 30 days. With $\tau = 30$, $\mu_D\tau = 0.30$ and $\sigma_D\sqrt{\tau} = 0.03\sqrt{30} = 0.16432$. The two arguments of equation (23.10) are

$$\frac{0.30 - 0.60}{0.16432} = -1.8257, \qquad -\frac{0.30 + 0.60}{0.16432} = -5.4772,$$

with $\Phi(-1.8257) = 0.033945$ and $\Phi(-5.4772) = 2.160 \times 10^{-8}$. The weight of the second term is $\exp(2 \times 0.01 \times 0.60/0.0009) = \mathrm{e}^{13.333} = 6.174 \times 10^5$, and the second term is $6.174 \times 10^5 \times 2.160 \times 10^{-8} = 0.013338$. Hence

$$F_{\mathrm{RUL}}(30) = 0.033945 + 0.013338 = 0.04728,$$

a 4.73 % probability of reaching the threshold within 30 days. The mean remaining life is 60 days, yet the probability of failure in half that time is far from zero. More than a quarter of it, 0.0133 out of 0.0473, comes from the reflected paths of Step 4: using only the probability that the damage exceeds the threshold at day 30 would have given 3.39 %.

Failure within 60 days. At $\tau = 60$ the mean increment equals the remaining distance, $\mu_D\tau = 0.60$, so the first argument is zero and $\Phi(0) = 0.5$. The second argument is $-1.20/0.23238 = -5.1640$, $\Phi(-5.1640) = 1.209 \times 10^{-7}$, and the second term is $6.174 \times 10^5 \times 1.209 \times 10^{-7} = 0.07463$. Hence $F_{\mathrm{RUL}}(60) = 0.5 + 0.0746 = 0.5746$. The probability of failure by the mean remaining life exceeds one half because the distribution is right-skewed: its median, 55.9 days, lies below its mean.

Table 23.1 lists the cumulative failure probability at six horizons, with the RUL reliability beside it, and Figure 23.4 plots the whole curve. In the figure the horizontal axis is the horizon τ in days and the vertical axis the probability of failure within τ; the dots are the rows of the table, the dashed horizontal line is a 5 % acceptance level and the dotted vertical line is the horizon at which the curve reaches it. The curve is flat for about two weeks, rises steeply between 30 and 70 days and approaches one slowly, the shape of the right-skewed density of Figure 23.3. If the largest acceptable probability of failure before maintenance is 5 %, the latest acceptable horizon is $\mathrm{RUL}_{0.05} = 30.3$ days, so a 30-day horizon is consistent with that criterion under this model. The calculation assumes that the degradation state and the model parameters are known exactly, which is rarely true; the next section removes both assumptions.

23.3 Bayesian degradation-state estimation and parameter uncertainty

The damage of a component in service is inferred from indirect measurements, and the parameters of its degradation model are estimated from short histories. This section propagates both

Table 23.1: Cumulative failure probability and RUL reliability of the cooling-pump bearing, equation (23.10) with $a = 0.6$, $\mu_D = 0.01\,\mathrm{day}^{-1}$, $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$, known current damage.

Horizon τ	$F_{\mathrm{RUL}}(\tau) = \Pr(\mathrm{RUL} \le \tau)$	$R_{\mathrm{RUL}}(\tau)$
15 days	0.000087	0.999913
30 days	0.04728	0.95272
45 days	0.28402	0.71598
60 days	0.57463	0.42537
90 days	0.89609	0.10391
120 days	0.97939	0.02061

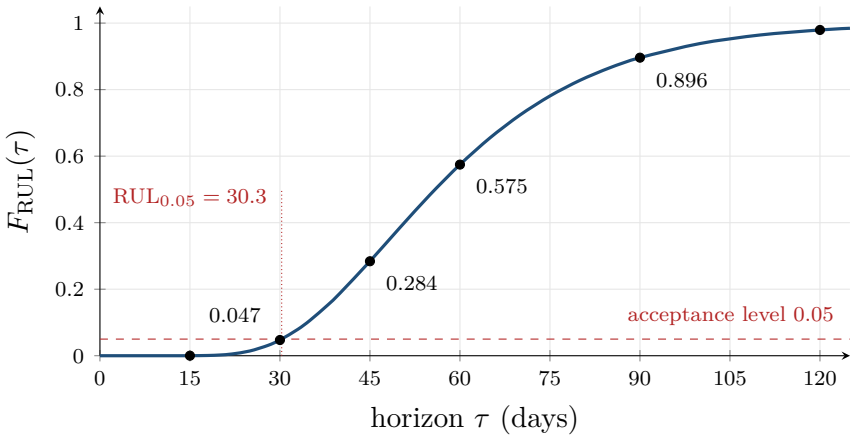


Figure 23.4: Probability that the cooling-pump bearing reaches its failure threshold within τ days, equation (23.10) with $a = 0.6$, $\mu_D = 0.01\,\mathrm{day}^{-1}$, $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$. Dots: the horizons of Table 23.1. Dashed: a 5 % acceptance level, reached at $\mathrm{RUL}_{0.05} = 30.3$ days.

uncertainties into the remaining life: the hidden state through a Bayesian filter, the parameters through a joint posterior.

23.3.1 The hidden damage state

Measurements such as vibration velocity, bearing temperature or motor current do not reveal the damage directly, as Chapter 22 showed. The model therefore distinguishes the latent damage $D_k = D(t_k)$ at the measurement time t_k from the measured condition indicator y_k. Sampling the Wiener model of equation (23.7) at the measurement times gives the discrete-time process model

$$D_{k+1} = D_k + \mu_D \, \Delta t_k + \sigma_D \sqrt{\Delta t_k} \, \xi_k, \qquad \xi_k \sim \mathcal{N}(0,1), \tag{23.14}$$

where $\Delta t_k = t_{k+1} - t_k$ and the ξ_k are independent; the increment $\sigma_D \sqrt{\Delta t_k}\, \xi_k$ has the variance $\sigma_D^2 \Delta t_k$ of a Brownian increment over Δt_k, so the discretization is exact. A linear observation model is

$$y_k = c \, D_k + b + \varepsilon_k, \qquad \varepsilon_k \sim \mathcal{N}(0, \sigma_y^2), \tag{23.15}$$

where c is the sensitivity of the measurement to damage, b the sensor offset (not the rate of a gamma prior) and ε_k the measurement error with standard deviation σ_y. The prognostic system must infer D_k from the observation history $y_{1:k} = (y_1, \ldots, y_k)$. When σ_y is large one measurement says little about the damage; when c or b is poorly known, different damage levels explain the same readings, and the damage is not identifiable from the indicator alone.

23.3.2 Bayesian filtering: prediction and update

Let $p(D_k \mid y_{1:k})$ denote the posterior distribution of the damage after the measurements up to t_k. A Bayesian filter advances it in two steps at each measurement time.

Prediction. The previous posterior is propagated through the process model,

$$p(D_k \mid y_{1:k-1}) = \int p(D_k \mid D_{k-1}) \, p(D_{k-1} \mid y_{1:k-1}) \, \mathrm{d}D_{k-1}, \tag{23.16}$$

which accounts for the expected degradation and for the process randomness between measurements.

Measurement update. The new measurement reweights the predicted distribution by Bayes' theorem (Chapter 8 of Volume I),

$$p(D_k \mid y_{1:k}) = \frac{p(y_k \mid D_k) \, p(D_k \mid y_{1:k-1})}{p(y_k \mid y_{1:k-1})}, \tag{23.17}$$

where the denominator $p(y_k \mid y_{1:k-1}) = \int p(y_k \mid D_k) \, p(D_k \mid y_{1:k-1}) \, \mathrm{d}D_k$ is the normalizing constant, the predictive probability of the measurement.

The scalar Kalman filter. When the process and observation models are linear with Gaussian noise, as in equations (23.14) and (23.15), both steps preserve normality and the filter reduces to recursions for a mean and a variance. Write the previous posterior as $\mathcal{N}(m_{k-1}, P_{k-1})$; here P is the variance of the damage estimate, the standard Kalman-filter symbol, not a path set or a state probability. The predicted damage is the sum of independent normal terms, so

$$m_k^- = m_{k-1} + \mu_D \, \Delta t_{k-1}, \qquad P_k^- = P_{k-1} + \sigma_D^2 \, \Delta t_{k-1} : \tag{23.18}$$

the mean advances by the drift and the variance grows by the process variance. For the update, the logarithm of the numerator of equation (23.17) is, up to a constant,

$$-\frac{(d-m_k^-)^2}{2P_k^-}-\frac{(y_k-c\,d-b)^2}{2\sigma_y^2},$$

a quadratic in the damage value d. Collecting the coefficient of d^2 and of d and completing the square shows that the posterior is normal with

$$\frac{1}{P_k}=\frac{1}{P_k^-}+\frac{c^2}{\sigma_y^2},\qquad m_k=P_k\left(\frac{m_k^-}{P_k^-}+\frac{c\,(y_k-b)}{\sigma_y^2}\right). \tag{23.19}$$

The posterior precision (inverse variance) is the prior precision plus the precision the measurement contributes, and the posterior mean is the precision-weighted average of the prediction and of the damage the measurement indicates. The same result in the gain form used in practice is

$$K_k=\frac{c\,P_k^-}{c^2P_k^-+\sigma_y^2},\qquad m_k=m_k^-+K_k\,(y_k-c\,m_k^--b),\qquad P_k=(1-K_kc)\,P_k^-,$$

where K_k is the Kalman gain and $y_k-c\,m_k^-b$ the innovation, the difference between the measurement and its prediction.

As an illustration of the prediction step, let the posterior after the previous inspection be $\mathcal{N}(0.30, 0.001)$, a standard deviation of 0.032, and let the next inspection come $\Delta t = 10$ days later. Equation (23.18) with the bearing parameters gives $m^- = 0.30 + 0.01 \times 10 = 0.40$ and $P^- = 0.001 + 0.0009 \times 10 = 0.010$, a predicted standard deviation of 0.10: ten days without measurements tripled the uncertainty of the damage.

23.3.3 Worked update: a noisy damage measurement

Take the predicted damage just derived, $D_k \mid y_{1:k-1} \sim \mathcal{N}(0.40, 0.10^2)$, so $P_k^- = 0.01$, and a calibrated measurement $y_k = D_k + \varepsilon_k$ ($c = 1$, $b = 0$) with $\sigma_y = 0.05$, so $\sigma_y^2 = 0.0025$. The inspection reads $y_k = 0.50$. Equation (23.19) gives

$$P_k=\left(\frac{1}{0.01}+\frac{1}{0.0025}\right)^{-1}=(100+400)^{-1}=0.002,\qquad \sqrt{P_k}=0.0447,$$

$$m_k=0.002\left(\frac{0.40}{0.01}+\frac{0.50}{0.0025}\right)=0.002\,(40+200)=0.48.$$

In the gain form, $K_k = 0.01/(0.01 + 0.0025) = 0.8$, $m_k = 0.40 + 0.8\,(0.50 - 0.40) = 0.48$ and $P_k = (1 - 0.8) \times 0.01 = 0.002$, the same values. The posterior is

$$D_k \mid y_{1:k} \sim \mathcal{N}(0.48,\, 0.0447^2).$$

The measurement, four times as precise as the prediction, receives four fifths of the weight. It moves the estimated damage from 0.40 to 0.48 and reduces its standard deviation from 0.10 to 0.045. The innovation, 0.10, is 0.89 standard deviations of its predictive distribution ($\sqrt{0.01 + 0.0025} = 0.112$), so the reading is consistent with the model and gives no sign of a fault.

Figure 23.5 shows the update. The horizontal axis is the damage value and the vertical axis a probability density. The gold curve is the predicted (prior) distribution, the dashed green curve the likelihood of the measurement as a function of the damage, normalized to unit area, and the blue curve the posterior. The posterior lies between the other two, nearer the narrower likelihood, and is narrower than either. The Gaussian form is a local approximation: a damage index is bounded, and a full model would impose $0 \le D_k < D_{\mathrm{crit}}$ together with the survival information. Here the posterior places a negligible probability outside that range (the threshold is 11.6 standard deviations away).

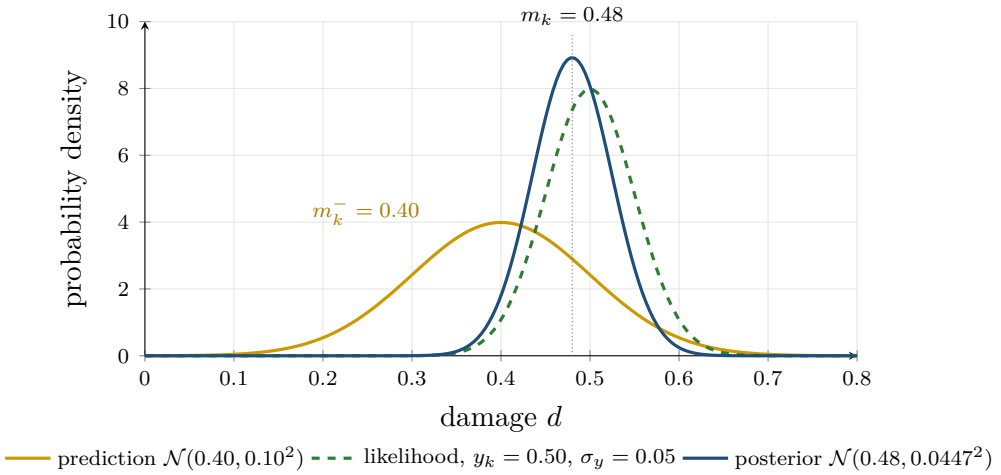


Figure 23.5: Scalar Kalman update of the bearing damage, equation (23.19) with $c = 1$, $b = 0$: prediction $\mathcal{N}(0.40, 0.01)$, measurement 0.50 with error variance 0.0025, posterior $\mathcal{N}(0.48, 0.002)$. Illustrative values.

23.3.4 How the state estimate changes the remaining life

After the update the remaining distance $a = 1 - D_k$ is uncertain, with posterior mean $\mathbb{E}[a \mid y_{1:k}] = 1 - 0.48 = 0.52$. A plug-in estimate substitutes this mean into equation (23.12): $\mathbb{E}[\text{RUL}] = 0.52/0.01 = 52$ days. The predictive distribution instead averages the conditional first-passage distribution over every plausible damage value,

$$\Pr(\text{RUL} \le \tau \mid y_{1:k}) = \int \Pr(\text{RUL} \le \tau \mid D_k = d)\, p(d \mid y_{1:k})\, \mathrm{d}d, \tag{23.20}$$

a mixture of the inverse-Gaussian distributions of Section 23.2.2, one for each d, weighted by the posterior. It contains both the randomness of the future path and the uncertainty of the present damage. Because $\Pr(\text{RUL} \le \tau \mid D_k = d)$ is a nonlinear function of d,

$$\mathbb{E}\big[F_{\text{RUL}}(\tau \mid D_k)\big] \neq F_{\text{RUL}}\big(\tau \mid \mathbb{E}[D_k]\big) \tag{23.21}$$

in general, and replacing the posterior by its mean before predicting the remaining life produces poorly calibrated failure probabilities. In the lower tail the conditional failure probability is a convex function of d, so the mixture is larger than the plug-in value: the damage values above the mean raise the early-failure probability more than the values below it lower it.

Table 23.2 quantifies the effect for the bearing. The four columns are the plug-in and the mixture predictions before the inspection (damage $\mathcal{N}(0.40, 0.10^2)$) and after it (damage $\mathcal{N}(0.48, 0.0447^2)$); the rows are three failure probabilities, the 10 % quantile and the mean. The mixtures were evaluated by numerical integration of equation (23.20) over the normal posterior. Three results stand out. The mean is the same for the plug-in and the mixture, because the conditional mean $(1 - d)/\mu_D$ is linear in d; the mean is blind to state uncertainty. The 30-day failure probability before the inspection is 4.73 % by plug-in but 7.91 % by the mixture, two thirds higher; the uncertainty of the damage, ignored by the plug-in, is what raises it. And the inspection itself raised the 30-day failure probability to 13.1 %, because it revealed more damage than predicted; the reduced uncertainty brought the plug-in and the mixture close together (12.2 % against 13.1 %). At 60 days the inequality reverses, as it must where the conditional failure probability is concave in d.

Table 23.2: Remaining-life predictions for the bearing before and after the inspection of Section 23.3.3. Plug-in: the inverse-Gaussian distribution at the mean damage. Mixture: equation (23.20) over the normal damage distribution. $\mu_D = 0.01\,\mathrm{day}^{-1}$, $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$, $D_{\mathrm{crit}} = 1$.

	before: $\mathcal{N}(0.40, 0.10^2)$		after: $\mathcal{N}(0.48, 0.0447^2)$	
Quantity	plug-in	mixture	plug-in	mixture
$\Pr(\mathrm{RUL} \le 20$ days)	0.0022	0.0121	0.0127	0.0174
$\Pr(\mathrm{RUL} \le 30$ days)	0.0473	0.0791	0.1218	0.1311
$\Pr(\mathrm{RUL} \le 60$ days)	0.5746	0.5689	0.7097	0.7061
$\mathrm{RUL}_{0.10}$ (days)	34.6	31.8	28.7	28.1
$\mathbb{E}[\mathrm{RUL}]$ (days)	60	60	52	52

23.3.5 Uncertain degradation rate and joint uncertainty

The drift $\mu_D = 0.01\,\mathrm{day}^{-1}$ has so far been treated as known. In practice it is estimated from a short and noisy degradation history and is described by a posterior $p(\mu_D \mid \mathcal{Y}_{0:t_0})$. The full predictive distribution of the remaining life integrates over the current damage d and over both parameters,

$$p(\mathrm{RUL} \mid \mathcal{Y}) = \iiint p(\mathrm{RUL} \mid d, \mu_D, \sigma_D)\; p(d, \mu_D, \sigma_D \mid \mathcal{Y})\, \mathrm{d}d\, \mathrm{d}\mu_D\, \mathrm{d}\sigma_D, \tag{23.22}$$

where $p(\mathrm{RUL} \mid d, \mu_D, \sigma_D)$ is the inverse-Gaussian density of equation (23.11). The same construction extends to an uncertain threshold, uncertain future loads and temperatures, or several competing physical models, each adding a dimension to the integral. This is the posterior predictive distribution of Chapter 8 applied to a first-passage time.

The integration must use the *joint* posterior. Damage and drift are estimated from the same measurements, and a history that rose steeply supports both a high current damage and a fast drift, so the two estimates tend to be positively correlated, $\mathrm{Cov}(D_k, \mu_D \mid \mathcal{Y}) > 0$. Their high values then occur together, and the combination of a short remaining distance and a fast approach is what produces early failures. Sampling the damage and the drift independently from their marginal distributions destroys the correlation and underestimates the lower tail of the remaining life.

Table 23.3 illustrates the effect with the posterior damage of Section 23.3.3 and an illustrative drift uncertainty, $\mu_D \sim \mathcal{N}(0.01, 0.002^2)\,\mathrm{day}^{-1}$, with σ_D fixed at $0.03\,\mathrm{day}^{-1/2}$. The three columns treat the drift as known, as uncertain and independent of the damage, and as uncertain with correlation $+0.7$ with the damage; the entries were computed by numerical integration of equation (23.22) over the bivariate normal distribution. The drift uncertainty raises the 20-day failure probability from 1.74 % to 2.10 %, and the correlation raises it further to 2.72 %, 56 % above the known-drift value; the 10 % quantile moves from 28.1 to 26.5 days. Near the center of the distribution the three columns almost agree, which is why a check of the mean or the median does not reveal the error.

23.3.6 Aleatory and epistemic uncertainty

The predictive distribution combines the two kinds of uncertainty defined in Volume I, and the distinction has a practical consequence. Table 23.4 gives one example of each. In the Wiener model of equation (23.7) the future increments of $B(t)$ are aleatory: they represent the variability of the physical evolution itself, which no measurement made today can remove. The uncertainty about d_0, μ_D and σ_D is epistemic: it comes from imperfect measurement and estimation, and more or better data can reduce it. The predictive distribution must contain both, but they are not the same thing and should be reported separately.

Table 23.3: Effect of drift uncertainty and of its correlation with the damage estimate on the bearing's remaining life. Damage $\mathcal{N}(0.48, 0.0447^2)$; drift $\mathcal{N}(0.01, 0.002^2)\,\mathrm{day}^{-1}$ in the last two columns; $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$. Illustrative drift uncertainty.

Quantity	drift known	drift uncertain, independent	drift uncertain, correlation 0.7
$\Pr(\mathrm{RUL} \le 20$ days)	0.0174	0.0210	0.0272
$\Pr(\mathrm{RUL} \le 30$ days)	0.1311	0.1447	0.1581
$\Pr(\mathrm{RUL} \le 45$ days)	0.4408	0.4457	0.4483
$\mathrm{RUL}_{0.10}$ (days)	28.1	27.3	26.5

Table 23.4: Aleatory and epistemic uncertainty in a remaining-life prediction.

Type	Example	Interpretation
Aleatory	random future increments of crack growth or of $B(t)$	variability of the future physical evolution
Epistemic	uncertain coefficient of a crack-growth law; uncertain d_0, μ_D, σ_D	incomplete knowledge of the model or of the system

The operational use of the distinction is that more instrumentation improves a prognosis only to the extent that the prediction uncertainty comes from what the instrumentation can reveal. For the bearing after the inspection, even with the damage known exactly at 0.48 and the drift known exactly, equation (23.12) leaves a standard deviation of $\sqrt{0.52 \times 0.0009/10^{-6}} = 21.6$ days. The predictive mixture over the damage posterior has variance $468 + 20 = 488$ days^2 (the mean conditional variance plus the variance of the conditional mean, $\mathrm{Var}(D_k)/\mu_D^2 = 0.002/10^{-4}$), a standard deviation of 22.1 days: almost all of the spread is aleatory. That floor is aleatory, set by σ_D; a better sensor narrows the damage posterior but cannot lower it. The question of how much a measurement is worth is taken up in a later chapter on the value of information.

23.4 Gamma-process prognostics and particle filtering

Two departures from the Wiener–Kalman setting are common: damage that can only grow, and models that are nonlinear or non-Gaussian. This section treats the first with the gamma process, whose remaining-life distribution is an incomplete gamma function, and the second with the particle filter.

23.4.1 The gamma degradation model

A Wiener path can decrease. That is acceptable for a health indicator that fluctuates, but not for damage that is physically irreversible: metal lost to corrosion, accumulated wear, some forms of irreversible capacity loss and other cumulative damage measures. For these a *gamma process* is the standard model (Chapter 21). Let $D(t_0) = d_0$ and let the future increments be independent with

$$D(t_0 + \tau) - D(t_0) \sim \mathrm{Gamma}(\alpha_D \tau,\, \beta_D), \tag{23.23}$$

a gamma distribution with shape $\alpha_D \tau$ and scale β_D. The shape grows linearly with the interval at the *shape rate* $\alpha_D > 0$, per unit time; the scale $\beta_D > 0$ is in damage units. The gamma distribution is written here with a scale, unlike the shape–rate prior $\mathrm{Gamma}(a, b)$ of Chapter 8 of Volume I; α_D is not an alpha factor and β_D is neither a beta factor nor a Weibull shape.

The density of an increment $x \ge 0$ with shape k and scale β_D is $x^{k-1}\mathrm{e}^{-x/\beta_D}/[\Gamma(k)\,\beta_D^k]$, with $\Gamma(\cdot)$ the gamma function, and its mean and variance are

$$\mathbb{E}[\Delta D(\tau)] = \alpha_D \beta_D\, \tau, \qquad \mathrm{Var}[\Delta D(\tau)] = \alpha_D \beta_D^2\, \tau. \tag{23.24}$$

Both grow linearly in τ, like those of the Wiener process, but every increment is nonnegative with probability one, so the damage never decreases.

23.4.2 The conditional RUL distribution

Let $a = D_{\mathrm{crit}} - d_0 > 0$ as before. Because the damage is nondecreasing, its largest value over $[t_0, t_0 + \tau]$ is its value at the end, $D(t_0 + \tau)$. The threshold has therefore been reached by $t_0 + \tau$ if and only if the damage is at or above it at $t_0 + \tau$:

$$\{\mathrm{RUL}(t_0) \le \tau\} = \{D(t_0 + \tau) \ge D_{\mathrm{crit}}\} = \{\Delta D(\tau) \ge a\}.$$

This identity is the convenience of a monotone model: no first-passage calculation is needed, because a path cannot cross the threshold and return below it. With $F_\Gamma(x; k, \beta_D)$ the gamma cumulative distribution function,

$$F_{\mathrm{RUL}}(\tau \mid d_0) = 1 - F_\Gamma(a;\, \alpha_D\tau,\, \beta_D), \qquad R_{\mathrm{RUL}}(\tau \mid d_0) = F_\Gamma(a;\, \alpha_D\tau,\, \beta_D). \tag{23.25}$$

The gamma cumulative distribution function is an incomplete gamma function. Substituting $u = x/\beta_D$ in the integral of the density,

$$F_\Gamma(a; k, \beta_D) = \frac{\gamma(k, a/\beta_D)}{\Gamma(k)}, \qquad 1 - F_\Gamma(a; k, \beta_D) = \frac{\Gamma(k, a/\beta_D)}{\Gamma(k)},$$

where $\gamma(k, x) = \int_0^x u^{k-1}\mathrm{e}^{-u}\,\mathrm{d}u$ and $\Gamma(k, x) = \int_x^\infty u^{k-1}\mathrm{e}^{-u}\,\mathrm{d}u$ are the lower and upper incomplete gamma functions. So

$$F_{\mathrm{RUL}}(\tau \mid d_0) = \frac{\Gamma(\alpha_D\tau,\, a/\beta_D)}{\Gamma(\alpha_D\tau)}, \tag{23.26}$$

the upper regularized incomplete gamma function with first argument $\alpha_D\tau$ and second argument a/β_D. As a function of the horizon τ this is not itself a gamma distribution: the horizon enters through the shape.

For an integer shape k the upper incomplete gamma function has a closed form. Integrating by parts, $\Gamma(k, x) = x^{k-1}\mathrm{e}^{-x} + (k-1)\,\Gamma(k-1, x)$, and with $\Gamma(1, x) = \mathrm{e}^{-x}$ the recursion gives $\Gamma(k, x) = (k-1)!\,\mathrm{e}^{-x}\sum_{j=0}^{k-1} x^j/j!$. Since $\Gamma(k) = (k-1)!$,

$$F_{\mathrm{RUL}}(\tau \mid d_0) = \mathrm{e}^{-x}\sum_{j=0}^{k-1}\frac{x^j}{j!}, \qquad k = \alpha_D\tau \in \{1, 2, \dots\}, \quad x = a/\beta_D. \tag{23.27}$$

The right side is the probability that a Poisson variable with mean x is at most $k-1$, an instance of the duality between gamma waiting times and Poisson counts.

23.4.3 Worked example: corrosion penetration

A pressure-containing component has a failure threshold of $D_{\mathrm{crit}} = 2.0\,\mathrm{mm}$ of corrosion penetration, and the current measured penetration is $d_0 = 1.2\,\mathrm{mm}$, so the remaining allowable penetration is $a = 0.8\,\mathrm{mm}$. The gamma process has $\alpha_D = 2\,\mathrm{yr}^{-1}$ and $\beta_D = 0.1\,\mathrm{mm}$, illustrative values; time is in years. The mean penetration rate is $\alpha_D\beta_D = 0.2\,\mathrm{mm}$ per year, and over one year the increment has mean $0.2\,\mathrm{mm}$ and standard deviation $\sqrt{\alpha_D\beta_D^2} = \sqrt{2 \times 0.1^2} = 0.1414\,\mathrm{mm}$. The second argument of equation (23.27) is $x = 0.8/0.1 = 8$ at every horizon, and $\mathrm{e}^{-8} = 3.3546 \times 10^{-4}$.

One year. The increment is Gamma$(2, 0.1)$, $k = 2$:

$$\Pr(\text{RUL} \le 1\,\text{yr}) = \mathrm{e}^{-8}\,(1 + 8) = 9 \times 3.3546 \times 10^{-4} = 0.003019,$$

a probability of 0.302 %.

Two years. The increment is Gamma$(4, 0.1)$, $k = 4$:

$$\Pr(\text{RUL} \le 2\,\text{yr}) = \mathrm{e}^{-8}\left(1 + 8 + \frac{8^2}{2!} + \frac{8^3}{3!}\right) = \mathrm{e}^{-8}\,(1 + 8 + 32 + 85.333) = 0.04238,$$

or 4.24 %. The expected penetration after two years is only 0.4 mm, half the allowance, yet some trajectories accumulate twice that.

Four years. The increment is Gamma$(8, 0.1)$, whose mean 0.8 mm equals the remaining allowance. With $k = 8$ the sum has eight terms, $1 + 8 + 32 + 85.333 + 170.667 + 273.067 + 364.089 + 416.102 = 1350.26$, and

$$\Pr(\text{RUL} \le 4\,\text{yr}) = 1350.26 \times 3.3546 \times 10^{-4} = 0.4530,$$

or 45.30 %. The probability is below one half because a gamma distribution is right-skewed, so its median lies below its mean: at the horizon where the expected damage reaches the threshold, the damage is more likely below it than above. The median remaining life is therefore later than four years, 4.17 years by solving $F_{\text{RUL}}(\tau) = 0.5$ numerically, and the mean remaining life, $\int_0^\infty R_{\text{RUL}}(\tau)\,\mathrm{d}\tau$, is 4.25 years. In this model, too, the time at which the expected damage reaches the threshold is neither the median nor the mean of the remaining life.

Table 23.5 collects the three horizons of the example with two more added to show the trend, and Figure 23.6 shows where the probabilities come from. In the figure the horizontal axis is the future penetration increment in millimeters and the vertical axis its density; the three curves are the increment densities after one, two and four years, and the vertical line is the remaining allowance $a = 0.8$ mm. The failure probability at each horizon is the area under that horizon's density to the right of the line, shaded for four years. As the horizon grows, the density moves right and spreads, and its tail beyond the allowance grows from 0.003 to 0.453. The 5 % quantile of the remaining life is 2.09 years and the 10 % quantile 2.50 years.

Table 23.5: Corrosion example: increment distribution and probability of reaching $D_{\text{crit}} = 2.0$ mm from $d_0 = 1.2$ mm, gamma process with $\alpha_D = 2\,\text{yr}^{-1}$ and $\beta_D = 0.1$ mm. Rows at 3 and 5 years added to show the trend.

Horizon τ	shape $\alpha_D\tau$	mean increment (mm)	SD of increment (mm)	$\Pr(\text{RUL} \le \tau)$
1 yr	2	0.2	0.141	0.00302
2 yr	4	0.4	0.200	0.04238
3 yr (added)	6	0.6	0.245	0.19124
4 yr	8	0.8	0.283	0.45296
5 yr (added)	10	1.0	0.316	0.71662

These are model-based probabilities under fixed parameters, a known current penetration and constant exposure. A change of temperature, fluid chemistry, flow velocity or operating regime changes the increment distribution and invalidates them, which is why the parameters of a corrosion model are tied to the operating conditions under which they were estimated. Noisy wall-thickness measurements are handled by the filtering of Section 23.3.2 with the gamma process as the process model; since its transition is not Gaussian, the Kalman filter no longer applies exactly.

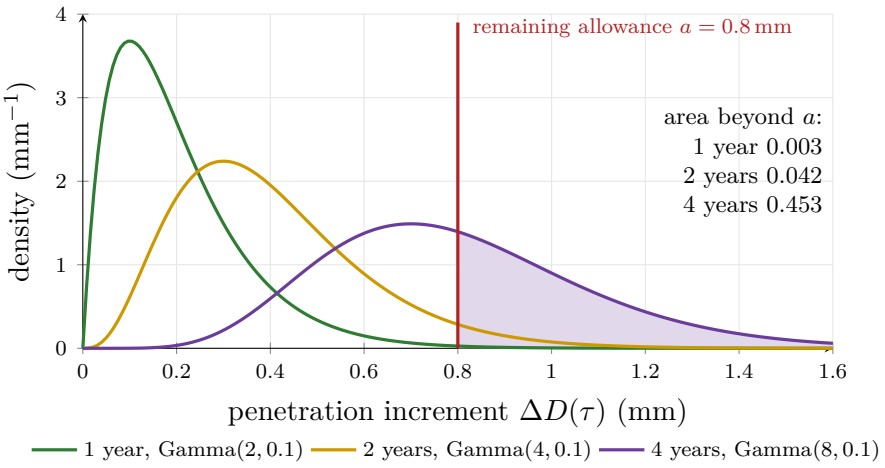


Figure 23.6: Densities of the future corrosion penetration after one, two and four years, gamma process with $\alpha_D = 2\,\mathrm{yr}^{-1}$ and $\beta_D = 0.1\,\mathrm{mm}$, density $x^{k-1}\mathrm{e}^{-x/\beta_D}/[(k-1)!\,\beta_D^k]$ with $k = \alpha_D\tau$. The probability of failure within each horizon is the area to the right of the remaining allowance (equation (23.25)); shaded for four years. Illustrative parameters.

23.4.4 When the Kalman filter is insufficient

The update of Section 23.3.3 was exact because the process and observation models were linear and the noise Gaussian. Realistic prognostic models often violate both conditions: nonlinear crack-growth equations, positive-only damage increments such as those of the gamma process, unknown fault modes, uncertain parameters estimated together with the state, non-Gaussian measurement errors, and changing operating conditions. The posterior is then not normal, and a mean and a variance no longer describe it.

A *particle filter* represents the posterior by a set of weighted samples. Let particle i carry the state vector

$$\mathbf{x}_k^{(i)} = \begin{bmatrix} D_k^{(i)} \\ \boldsymbol{\theta}_k^{(i)} \end{bmatrix},$$

a damage value together with values of the uncertain model parameters $\boldsymbol{\theta}$ (for the Wiener model, μ_D and σ_D), and a weight $w_k^{(i)} \geq 0$, the weights summing to one over the N_p particles. (The weight $w_k^{(i)}$ is not the rate of occurrence of failures $w(t)$ of Chapter 17 of Volume I.) The posterior is approximated by

$$p(\mathbf{x}_k \mid y_{1:k}) \approx \sum_{i=1}^{N_\mathrm{p}} w_k^{(i)}\, \delta(\mathbf{x}_k - \mathbf{x}_k^{(i)}), \tag{23.28}$$

where δ is the Dirac delta: a probability mass $w_k^{(i)}$ at each particle. Any expectation under the posterior, a failure probability included, becomes a weighted sum over the particles, and the approximation improves as N_p grows, whatever the shape of the posterior.

23.4.5 The particle filter algorithm

The basic algorithm, the bootstrap or sequential importance resampling filter, advances the particle set at each measurement time in six steps.

1. *Initialize.* Draw N_p particles from the prior distribution of the current damage and the model parameters, each with weight $1/N_\mathrm{p}$.

2. *Propagate.* Advance each particle through the stochastic degradation model, $\mathbf{x}_k^{(i)} \sim p(\mathbf{x}_k \mid \mathbf{x}_{k-1}^{(i)})$, for example by drawing a Wiener or gamma increment over Δt_{k-1} with the particle's own parameters. This is the sampled form of the prediction step, equation (23.16).

3. *Evaluate the measurement.* Compute the likelihood of the new observation under each particle, $p(y_k \mid \mathbf{x}_k^{(i)})$, from the observation model, for example a normal density centered on $c\,D_k^{(i)} + b$.

4. *Reweight and normalize.* Multiply each weight by its likelihood and rescale,

$$\tilde{w}_k^{(i)} = w_{k-1}^{(i)}\, p(y_k \mid \mathbf{x}_k^{(i)}), \qquad w_k^{(i)} = \frac{\tilde{w}_k^{(i)}}{\sum_{j=1}^{N_\mathrm{p}} \tilde{w}_k^{(j)}}.$$

   Particles consistent with the measurement gain weight; this is the sampled form of the update, equation (23.17).

5. *Resample when needed.* Monitor the effective sample size

$$N_\mathrm{eff} = \frac{1}{\sum_{i=1}^{N_\mathrm{p}} \big(w_k^{(i)}\big)^2},$$

   which equals N_p when the weights are equal and 1 when one particle carries all the weight. When it falls below a chosen level, a common choice being $N_\mathrm{p}/2$, draw a new set of N_p particles from the current set with probabilities $w_k^{(i)}$ and reset all weights to $1/N_\mathrm{p}$. This discards particles with negligible weight and duplicates those with large weight.

6. *Forecast.* Simulate future degradation trajectories from the posterior particles until each crosses the threshold, as described in Section 23.4.6.

Parameter learning. When a particle carries a static parameter such as the drift, the propagation step does not change it, and repeated resampling keeps only the parameter values of the particles that happened to survive. After a number of steps all particles may share a few drift values, and the filter has lost the parameter uncertainty it was meant to represent; this is the collapse of parameter diversity. Practical algorithms counter it by *rejuvenation* (small random perturbations of the parameters after resampling, or a Markov chain Monte Carlo move that leaves the posterior unchanged), by transition kernels designed for the parameters, by Rao–Blackwellization (integrating out analytically the part of the model that is linear and Gaussian, and sampling only the rest), or by a separate parameter-estimation method. A particle filter does not produce reliable remaining-life predictions merely because it uses many particles: the state and the parameters must still be identifiable from the measurements, and the model must still be calibrated against outcomes.

23.4.6 RUL from particles and what to report

Suppose the filtered posterior consists of the particles $\{D_k^{(i)}, \boldsymbol{\theta}^{(i)}, w_k^{(i)}\}$. For each particle, simulate N_f future degradation trajectories under an assumed future operating scenario, and let $\mathrm{RUL}^{(i,j)}$ be the threshold-crossing time of trajectory j started from particle i. The predictive failure probability within τ is estimated by

$$\widehat{\Pr}(\mathrm{RUL} \le \tau) = \sum_{i=1}^{N_\mathrm{p}} w_k^{(i)} \left[\frac{1}{N_\mathrm{f}} \sum_{j=1}^{N_\mathrm{f}} \mathbf{1}\{\mathrm{RUL}^{(i,j)} \le \tau\} \right]. \tag{23.29}$$

The bracket is a Monte Carlo estimate of the conditional failure probability $\Pr(\mathrm{RUL} \le \tau \mid D_k^{(i)}, \boldsymbol{\theta}^{(i)})$ of one particle, and the weighted sum over particles is the sampled form of the mixture

of equation (23.22). The estimate therefore combines the three uncertainties of the prediction: the current damage (the spread of the $D_k^{(i)}$), the model parameters (the spread of the $\boldsymbol{\theta}^{(i)}$) and the future degradation (the spread of the trajectories). When the conditional distribution has a closed form, as for the Wiener model, the bracket can be replaced by equation (23.10) evaluated at the particle's state and parameters, which removes the simulation noise of the inner average. From the same samples the quantiles RUL_α follow by sorting the weighted crossing times.

A prognostic report summarizes the distribution for a decision maker. Table 23.6 lists the quantities a useful report contains, with illustrative values that are not results of the bearing model: the current damage as a posterior median with an interval, the mean, median and 10 % quantile of the remaining life, the probability of failure within the decision horizon, and the main source of the uncertainty. A report must also state the assumed future operating scenario and the failure threshold, since the distribution is conditional on both.

Table 23.6: Contents of a prognostic report, with illustrative values.

Quantity	Example output
Current damage	posterior median 0.58
90 % interval of current damage	0.49 to 0.67
Mean RUL	48 days
Median RUL	44 days
10 % quantile of RUL	22 days
Probability of failure within 30 days	24 %
Main source of uncertainty	degradation rate and future loading
Stated assumptions	operating scenario; failure threshold

23.5 Calibration and validation of RUL predictions

A model can produce a smooth, plausible remaining-life distribution and still be systematically wrong. This section states what it means for a predicted distribution to agree with outcomes and gives the metrics that test it, including the treatment of components that have not yet failed.

23.5.1 A prediction distribution must be tested

A predicted failure probability is a statement about frequencies. Among comparable predictions that assign a 20 % probability of failure within 30 days, roughly 20 % of the components should experience the defined failure within 30 days, if the model is well calibrated and the cases are comparable enough to be grouped. A *calibration curve* (also called a reliability diagram) makes this check: the predictions are grouped by predicted probability, and the observed frequency of failure in each group is plotted against the predicted value; a calibrated model lies on the diagonal.

Calibration is necessary but not sufficient. *Sharpness* is the concentration of the predicted distributions: narrow intervals and failure probabilities near zero or one. A model that predicts very broad remaining-life intervals can achieve the stated coverage and still provide little decision value, since it never commits to an action. The aim is the sharpest prediction that remains calibrated; the two properties are evaluated together.

23.5.2 RUL error metrics

For equipment whose end of life has been observed, the predicted and realized remaining lives can be compared. For evaluation case $i = 1, \dots, N$, let RUL_i be the realized remaining life and $\widehat{\mathrm{RUL}}_i$ a point prediction made at the same time. The *mean absolute error* and the *root-mean-square error* are

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}\left|\widehat{\mathrm{RUL}}_i - \mathrm{RUL}_i\right|, \qquad \mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\widehat{\mathrm{RUL}}_i - \mathrm{RUL}_i\right)^2}, \tag{23.30}$$

both in units of time. The RMSE is never smaller than the MAE and weights large errors more heavily. These are the familiar forecasting metrics, and they ignore the predicted uncertainty entirely: a model can achieve a low MAE while being dangerously overconfident, stating narrow intervals that the outcomes frequently fall outside. Point metrics are therefore used alongside the probabilistic scores of the next subsection, not in place of them.

23.5.3 Evaluating probabilistic RUL forecasts

For a predicted cumulative distribution $F_i(\tau)$ of the remaining life and the realized value RUL_i, the *continuous ranked probability score* is

$$\mathrm{CRPS}(F_i, \mathrm{RUL}_i) = \int_{-\infty}^{\infty}\left[F_i(\tau) - \mathbf{1}\{\mathrm{RUL}_i \le \tau\}\right]^2 \mathrm{d}\tau, \tag{23.31}$$

in units of time; for a nonnegative remaining life the integrand vanishes for $\tau < 0$, and the integral starts at zero. The score compares the predicted distribution function with the step function of the outcome. Lower is better: it rewards distributions that place probability near the realized failure time without being broader than necessary. For a point prediction, whose distribution function is a step at $\widehat{\mathrm{RUL}}_i$, the integrand is one between $\widehat{\mathrm{RUL}}_i$ and RUL_i and zero elsewhere, so the CRPS reduces to the absolute error $|\widehat{\mathrm{RUL}}_i - \mathrm{RUL}_i|$; the CRPS is the generalization of the MAE to distributions. As an illustration, if the bearing of Section 23.2.3 failed 45 days after the prediction, the CRPS of its inverse-Gaussian forecast is 7.4 days (computed numerically), against an absolute error of 15 days for the point forecast at its mean.

For a fixed decision horizon τ, let $p_i = \Pr(\mathrm{RUL}_i \le \tau \mid \mathcal{Y}_i)$ be the predicted failure probability of case i and $o_i = \mathbf{1}\{\mathrm{RUL}_i \le \tau\}$ the observed outcome. The *Brier score* is

$$\mathrm{BS} = \frac{1}{N}\sum_{i=1}^{N}(p_i - o_i)^2, \tag{23.32}$$

a number between zero and one that evaluates the predicted failure probabilities at that horizon. Other useful metrics are the coverage of stated intervals (the fraction of outcomes inside the 80 % interval should be near 80 %), the interval width, the quantile (pinball) loss of predicted quantiles, and the calibration curves of Section 23.5.1. Prognostics-specific metrics also exist, among them the prognostic horizon of the next subsection and the alpha–lambda performance, which checks whether predictions made at specified fractions of the remaining life fall within a stated accuracy cone around the true remaining life.

23.5.4 Prognostic horizon

Let a component fail at time T_f. A prognostic model that predicts its remaining life repeatedly will, at some earlier time t_p, begin to meet a chosen accuracy criterion (for example, the true remaining life lies inside the predicted 90 % interval, and the median is within 20 % of the truth) and keep meeting it at every later prediction. The *prognostic horizon* is

$$\mathrm{PH} = T_\mathrm{f} - t_\mathrm{p}, \tag{23.33}$$

where the requirement that the criterion be sustained matters: a model that is accurate by chance at one early prediction and wrong afterwards does not earn that time.

The prognostic horizon must be compared with the lead time of the action the prediction supports. If a replacement component takes 45 days to arrive, an accurate prediction made five days before failure has little practical value, and a model with a prognostic horizon of 60 days supports the order while one with a horizon of 20 days does not. A useful prognostic model is one that becomes accurate early enough for the required action.

23.5.5 Censoring and incomplete failure observations

In most data sets many units have not reached the failure threshold when the analysis is done. A pump monitored for 200 days that is still operating has an unknown failure time, but it is known that $T_\mathrm{f} > 200$ days: a right-censored observation, as defined in Chapter 7 of Volume I. It is informative, and it enters the likelihood through the reliability function rather than the density. With $\delta_i = 1$ for an observed failure at t_i and $\delta_i = 0$ for a unit censored at t_i, the likelihood of the parameters $\boldsymbol{\theta}$ is the standard survival-analysis likelihood of Chapter 7,

$$L(\boldsymbol{\theta}) = \prod_i f(t_i;\boldsymbol{\theta})^{\delta_i}\, R(t_i;\boldsymbol{\theta})^{1-\delta_i}, \tag{23.34}$$

where for a degradation model f and R are the first-passage density and reliability, such as those of equations (23.11) and (23.10). Discarding the units that have not failed keeps only the short lives and can severely bias the estimated degradation rate and every remaining-life prediction built on it.

Censoring must also be handled when forecasts are scored. A unit censored before the end of a decision horizon has an unknown outcome o_i for that horizon: it cannot be counted as a non-failure in the Brier score of equation (23.32), nor can its CRPS be computed from an unknown RUL_i. Scores and calibration curves that silently treat censored units as survivors are biased in favor of optimistic models.

23.6 RUL-based maintenance decisions and the credibility of predictions

A remaining-life distribution is computed to support a decision. This section turns it into maintenance decisions that account for lead time, cost and risk, tabulates the failure probability for several damage levels and horizons, and states what makes a prediction credible and how it connects to system-level reliability.

23.6.1 Maintenance decisions depend on lead time

Suppose the bearing's remaining-life distribution is that of Section 23.2.3, with $F_\mathrm{RUL}(30) = 0.0473$, and a replacement assembly has a lead time of $\tau_\mathrm{L} = 30$ days. If the order is placed today, the probability that the bearing fails before the spare arrives is $F_\mathrm{RUL}(\tau_\mathrm{L}) = 4.73\,\%$, under the fixed operating scenario and the Wiener model with known state and parameters.

Whether that probability is acceptable depends on the consequences. For a pump that serves critical cooling without a standby, 4.73 % may be too high, and the plant would expedite the spare, reduce the load or prepare a temporary measure. For a noncritical component whose failure is inexpensive, it may be acceptable. A decision threshold such as 5 % is therefore a statement of risk tolerance, chosen from the consequences of failure, not a number chosen for convenience.

23.6.2 Expected maintenance loss

The available actions might be to continue operating, to inspect, to replace now, or to plan a replacement in 30 days. Consider the plan "replace at τ days unless the component fails first" over a single planning window, with the cost c_p of a planned (preventive) replacement and the cost c_f of a failure before replacement, which includes the corrective replacement and the consequences of the failure. The cost of the plan is a random variable: it is c_f if the remaining life ends before τ and c_p otherwise,

$$C_{\text{plan}}(\tau) = c_f\,\mathbf{1}\{\text{RUL} \le \tau\} + c_p\,\mathbf{1}\{\text{RUL} > \tau\}.$$

The expectation of an indicator is the probability of its event, so

$$\mathbb{E}[C_{\text{plan}}(\tau)] = c_f\,F_{\text{RUL}}(\tau) + c_p\,R_{\text{RUL}}(\tau) = c_p + (c_f - c_p)\,F_{\text{RUL}}(\tau), \tag{23.35}$$

using $R_{\text{RUL}} = 1 - F_{\text{RUL}}$. With $c_p = \text{R\$}\,2000$, $c_f = \text{R\$}\,50\,000$ (illustrative) and $F_{\text{RUL}}(30) = 0.0473$,

$$\mathbb{E}[C_{\text{plan}}(30)] = 2000 + 48\,000 \times 0.0473 = \text{R\$}\,4270,$$

the expected cost of the plan to wait 30 days. Of it, R\$ 2000 is the replacement that will be made in any case and R\$ 2270 the expected penalty of the 4.73 % chance of failing first.

Equation (23.35) describes the cost of one replacement event. It omits the operating costs, the useful life consumed, the effect of inspections and every cost after the replacement. Because F_{RUL} increases with τ, the expression is smallest at $\tau = 0$ and always favors replacing now; that conclusion is an artifact of the omission, since immediate replacement discards the remaining useful life that a later replacement would use. A complete comparison spreads the cost over the operating time each plan buys and continues over the future replacement cycles, as the age-replacement and maintenance-optimization models of Chapters 18 and 20 of Volume I do, here with the remaining-life distribution of the individual component in place of the population lifetime distribution.

23.6.3 Risk-constrained scheduling and load reduction

Instead of minimizing an expected cost, the operator may impose a limit on the probability of failure before planned maintenance,

$$\Pr(\text{RUL} \le \tau \mid \mathcal{Y}) \le p^*, \tag{23.36}$$

where p^* is the decision threshold, the acceptable probability of failure before the intervention. Since F_{RUL} is nondecreasing, the constraint holds for every τ up to the p^*-quantile, and the latest allowable maintenance horizon is

$$\tau_{\max} = \text{RUL}_{p^*}.$$

With $p^* = 0.05$ and the bearing of Section 23.2.3, $\tau_{\max} = \text{RUL}_{0.05} = 30.3$ days, the dotted line of Figure 23.4. This is a probabilistic maintenance deadline. It is not fixed: it is recomputed when a new measurement updates the damage posterior, as the inspection of Section 23.3.3 did (after it, the plug-in $\text{RUL}_{0.05}$ falls to 24.9 days), and when the future operating plan changes.

Load reduction as an alternative to replacement. Maintenance is not the only action. When the drift depends on the operating load u, $\mu_D = \mu_D(u)$, reducing the load slows damage accumulation and moves the remaining-life distribution to the right. For the bearing, an illustrative reduction of the drift from 0.010 to $0.008\,\text{day}^{-1}$ by unloading the pump raises the mean remaining life from 60 to 75 days, lowers the 30-day failure probability from 4.73 % to

2.11 %, and extends the 5 % deadline from 30.3 to 34.9 days, enough to cover a lead time that slightly exceeds 30 days. The decision then becomes a joint choice of the load profile and of the maintenance time,

$$\min_{u(\cdot),\,\tau} \; \mathbb{E}\big[\text{operating cost} + \text{maintenance cost} + \text{failure consequences}\big], \tag{23.37}$$

subject to the required functions of the system and to risk constraints of the form of equation (23.36). This is where reliability prognostics meets stochastic control and maintenance optimization.

23.6.4 Failure probability for several damage levels and horizons

Table 23.7 evaluates the Wiener model of Section 23.2.2 for several current damage levels and diffusion parameters, with $D_{\text{crit}} = 1$ and $\mu_D = 0.01\,\text{day}^{-1}$, treating the current damage and the parameters as exactly known so that the spread reflects only the future Brownian degradation. Each row gives the mean and standard deviation of the remaining life from equation (23.12), the probability of failure within 10, 30, 60 and 90 days from equation (23.10), and the 10 % quantile. One row of the arithmetic, for $d_0 = 0.60$, $\sigma_D = 0.03$ and $\tau = 30$: $a = 0.40$, $\mu_D\tau = 0.30$, $\sigma_D\sqrt{30} = 0.16432$; the first term is $\Phi(-0.10/0.16432) = \Phi(-0.6086) = 0.27140$; the weight is $\exp(2{\times}0.01{\times}0.40/0.0009) = \mathrm{e}^{8.889} = 7251$ and $\Phi(-0.70/0.16432) = \Phi(-4.2601) = 1.0218{\times}10^{-5}$, so the second term is 0.07409; the sum is 0.3455.

Table 23.7: Wiener-model remaining life for several current damage levels d_0 and diffusion parameters σ_D (in $\text{day}^{-1/2}$), with $D_{\text{crit}} = 1$ and $\mu_D = 0.01\,\text{day}^{-1}$, known state and parameters. The first block varies the damage; the second varies the diffusion at $d_0 = 0.40$. Entries below 10^{-4} are shown as $< 10^{-4}$.

d_0	σ_D	mean (days)	SD (days)	$F_{\text{RUL}}(10)$	$F_{\text{RUL}}(30)$	$F_{\text{RUL}}(60)$	$F_{\text{RUL}}(90)$	$\text{RUL}_{0.10}$ (days)
0.20	0.03	80	26.8	$< 10^{-4}$	0.0017	0.2393	0.6985	49.8
0.40	0.03	60	23.2	$< 10^{-4}$	0.0473	0.5746	0.8961	34.6
0.60	0.03	40	19.0	0.0013	0.3455	0.8663	0.9784	20.2
0.80	0.03	20	13.4	0.2126	0.8284	0.9819	0.9978	7.5
0.40	0.02	60	15.5	$< 10^{-4}$	0.0042	0.5507	0.9573	41.9
0.40	0.04	60	31.0	0.0001	0.1215	0.5972	0.8552	28.4
0.40	0.05	60	38.7	0.0014	0.1983	0.6182	0.8316	23.4

Two trends are visible. In the first block, raising the current damage from 0.40 to 0.60 shortens the mean remaining life by a third but multiplies the 30-day failure probability by more than seven, from 4.73 % to 34.6 %: near the threshold the failure probability is very sensitive to the damage estimate, which is why the state uncertainty of Section 23.3.4 matters most for components close to failure. In the second block the mean stays at 60 days while the diffusion grows from 0.02 to 0.05; the 30-day failure probability rises from 0.42 % to 19.8 % and the 10 % quantile falls from 41.9 to 23.4 days. Greater stochastic variability increases the probability of early failure even when the expected remaining life is unchanged, so damage-state uncertainty and future process variability must both be represented. The long-horizon column moves the other way: at 90 days a larger diffusion lowers the failure probability, because it also spreads probability into the far right tail.

23.6.5 What makes an RUL prediction unreliable

Table 23.8 lists the common causes of unreliable remaining-life predictions and their consequences. They fall in three groups: errors in what failure means (the threshold, unmodeled mechanisms),

errors in the present state (unobservable damage, sensor bias) and errors in the future (loading, the physical model, the parameters), with poor validation allowing any of them to pass unnoticed.

Table 23.8: Causes of unreliable remaining-life predictions.

Issue	Consequence
Incorrect failure threshold	the predicted end of life does not match the loss of the required function
Unobservable degradation state	different damage states explain the same measurements
Incorrect future loading	the forecast degradation rate differs from the one the actual operating conditions produce
Inadequate physical model	the extrapolated damage behavior is wrong
Uncertain model parameters, ignored	the predicted remaining life is overconfident
Sensor bias	the estimated damage is systematically wrong
Unmodeled failure mechanisms	the model misses an alternative route to failure
Poor validation	the apparent precision does not correspond to actual predictive performance

A credible prognostic system answers three questions separately, each with its own distribution and its own sources of uncertainty. What is known about the current damage, $p(D_{t_0} \mid \mathcal{Y}_{0:t_0})$: the state-estimation problem, limited by observability and sensor quality. How the damage is expected to evolve, $p(D_{t_0+\tau} \mid D_{t_0}, \boldsymbol{\theta}, \mathbf{u})$, with $\mathbf{u}$ the future operating conditions: the degradation model, limited by physics, parameters and the operating plan. And what the probability of crossing the functional failure threshold is, $\Pr(\mathrm{RUL}(t_0) \le \tau \mid \mathcal{Y}_{0:t_0})$: the first-passage calculation, which combines the first two and is limited by the definition of the threshold. Keeping the three steps distinct makes it possible to tell which one a wrong prediction came from.

23.6.6 Connection to system-level reliability

A remaining-life prediction is usually made for one component, while the operator is concerned with the availability of a function or the risk to the plant. A cooling pump with a 20 % probability of failing within 30 days does not give the cooling plant a 20 % probability of losing the required cooling: the plant may have redundant pumps, spares on site, thermal storage or the option to reduce the heat load. To obtain the system risk, the component remaining-life distributions are combined with the component dependencies, the repair times, the maintenance plans, the capacity constraints of the system and its physical operating requirements, by the block-diagram, fault-tree, Markov and availability models of Volume I (Chapters 9, 10, 19 and 20). In a multi-component system the future failure times are generally dependent, through shared operating conditions, the redistribution of load when one unit fails and common-cause mechanisms (Chapter 13 of Volume I), and the system model must preserve those dependencies rather than multiply marginal probabilities.

Figure 23.7 shows the resulting chain. Each box is a stage of the calculation and each arrow passes a probability distribution to the next stage; the note under each arrow names the model that the transition adds. Condition measurements become a damage posterior through the observation model and a filter; the posterior becomes component remaining-life distributions through the degradation model, the threshold and the operating scenario; these become system-state probabilities through the system logic, redundancy, repair and capacity constraints, with dependencies preserved; and the system-state probabilities become operational

risk through the consequences of each state. The chain links the prognostics of this chapter to the system reliability and risk models of Volume I. The next chapter uses it to decide when to intervene, accounting for the risks of acting too early, of acting too late, and of deciding from incomplete measurements.

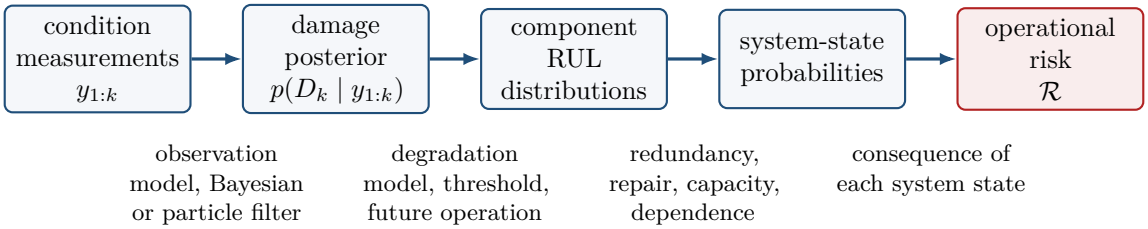


Figure 23.7: From condition measurements to operational risk. Each arrow passes a probability distribution; the notes below the arrows name the model that each transition adds. The first two transitions are the prognostics of this chapter; the last two are the system reliability and risk models of Volume I.

23.7 What to retain

Definition 23.1 (Remaining useful life)**.** The remaining useful life $\mathrm{RUL}(t_0) = T_\mathrm{f} - t_0$ is the first-passage time of the degradation state to the failure threshold, measured from the present and conditioned on the evidence available now; it is a random variable, and a prediction of it is a distribution that combines the uncertainty of the current damage, of the model parameters and of the future degradation path, reported through the probability of failure within the decision horizon and its lower quantiles rather than its mean.

Table 23.9 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 23.9: The terms of remaining-useful-life estimation and where each is defined.

Term	Expression	Units	Meaning	Defined in
First-passage time	$T_{\mathrm{f}} = \inf\{t : D(t) \geq D_{\mathrm{crit}}\}$	time	failure as threshold crossing	23.1.1
Remaining useful life	$\mathrm{RUL}(t_0) = T_{\mathrm{f}} - t_0 \mid \mathcal{Y}_{0:t_0}, T_{\mathrm{f}} > t_0$	time	conditional remaining life	23.1.1
RUL reliability	$R_{\mathrm{RUL}}(\tau) = \Pr(\mathrm{RUL} > \tau \mid \mathcal{Y})$	–	survives τ more	23.1.2
RUL failure probability	$F_{\mathrm{RUL}}(\tau) = 1 - R_{\mathrm{RUL}}(\tau)$	–	fails within τ	23.1.2
RUL quantile	$\mathrm{RUL}_\alpha = \inf\{\tau : F_{\mathrm{RUL}}(\tau) \geq \alpha\}$	time	horizon of probability α	23.1.3
Inverse Gaussian RUL	eqs. (23.10), (23.11); mean a/μ_D, variance $a\sigma_D^2/\mu_D^3$	time	Wiener first passage	23.2.2
Kalman prediction	$m^- = m + \mu_D \Delta t$, $P^- = P + \sigma_D^2 \Delta t$	–	propagate damage estimate	23.3.2
Kalman update	$1/P_k = 1/P_k^- + c^2/\sigma_y^2$	–	condition on a measurement	23.3.2
Predictive RUL	$\int \Pr(\mathrm{RUL} \leq \tau \mid d, \boldsymbol{\theta})\, p(d, \boldsymbol{\theta} \mid \mathcal{Y})$	–	mixture over the joint posterior	23.3.5
Gamma-process RUL	$\Gamma(\alpha_D \tau, a/\beta_D)/\Gamma(\alpha_D \tau)$	–	monotone damage	23.4.2
Particle approximation	$\sum_i w_k^{(i)} \delta(\mathbf{x}_k - \mathbf{x}_k^{(i)})$	–	weighted samples	23.4.4
Effective sample size	$N_{\mathrm{eff}} = 1/\sum_i (w_k^{(i)})^2$	–	resampling trigger	23.4.5
Calibration, sharpness	predicted against observed frequency; concentration	–	both required	23.5.1
MAE, RMSE	eq. (23.30)	time	point-prediction error	23.5.2
CRPS, Brier score	eqs. (23.31), (23.32)	time, –	probabilistic scores	23.5.3
Prognostic horizon	$\mathrm{PH} = T_{\mathrm{f}} - t_{\mathrm{p}}$	time	accurate early enough	23.5.4
Censored likelihood	$\prod f^{\delta_i} R^{1-\delta_i}$	–	keeps unfailed units	23.5.5
Expected maintenance loss	$c_p + (c_f - c_p) F_{\mathrm{RUL}}(\tau)$	R$	cost of one planned replacement	23.6.2
Risk-constrained deadline	$\tau_{\max} = \mathrm{RUL}_{p^*}$	time	latest acceptable intervention	23.6.3

Further reading. *Prognostics and Health Management of Engineering Systems*: physics-based and data-driven RUL estimation. *Reliability Analysis for a Wiener Degradation Process under Changing Failure Thresholds*: Wiener first-passage reliability. *Remaining Useful Lifetime Estimation and Noisy Gamma Deterioration Process*: gamma-process RUL with measurement noise and Bayesian inference. The NASA *Prognostics Metrics Library*: prognostic horizon, alpha–lambda performance and RUL evaluation metrics.

23.8 Exercises

Exercise 23.1. (a) Define the remaining useful life. (The time from the present until the defined failure threshold is first reached, conditional on the evidence available now; not the chronological age, and not a repair duration.) (b) Explain why the RUL distribution is more useful than its mean alone. (It describes the probability of early failure and the uncertainty, which decide the timing of maintenance; it does not remove the need for planning or guarantee a failure date.) (c) For Wiener degradation with fixed positive drift, name the distribution of the first-passage time to a fixed upper threshold. (The inverse Gaussian distribution, equation (23.11).) (d) Explain

why the gamma-process RUL is convenient for monotone damage. (Because the damage cannot decrease, crossing the threshold by τ is equivalent to the damage exceeding the threshold at τ, so the RUL distribution is a gamma tail probability, equation (23.25).)

Exercise 23.2. (a) State what a particle filter approximates. (The posterior distribution of hidden states, and possibly of parameters, by weighted samples.) (b) A model assigns a 20 % probability of failure within 30 days to many comparable components. What does calibration require? (That roughly 20 % of them fail within 30 days.) (c) Why must censored component histories be retained? (They show that failure has not occurred through the observation age, and discarding them biases the estimates toward short lives.) (d) After a measurement update, four particles have weights 0.5, 0.3, 0.1 and 0.1. Compute the effective sample size. ($1/(0.25 + 0.09 + 0.01 + 0.01) = 2.78$.) (e) Four predictions of failure within 30 days were 0.05, 0.20, 0.60 and 0.10, and the outcomes were no failure, failure, failure and no failure. Compute the Brier score. ($(0.0025 + 0.64 + 0.16 + 0.01)/4 = 0.203$.)

Exercise 23.3. A bearing follows the Wiener model with $D_{\text{crit}} = 1$, $\mu_D = 0.01\,\text{day}^{-1}$ and $\sigma_D = 0.03\,\text{day}^{-1/2}$. The predicted damage before an inspection is $\mathcal{N}(0.50, 0.08^2)$, and a calibrated measurement with $\sigma_y = 0.06$ reads 0.60. (a) Compute the Kalman gain and the posterior mean and standard deviation of the damage. ($K = 0.0064/0.01 = 0.64$; $m = 0.50 + 0.64 \times 0.10 = 0.564$; $P = 0.36 \times 0.0064 = 0.002304$, standard deviation 0.048.) (b) Compute the plug-in mean and standard deviation of the remaining life. ($a = 0.436$; 43.6 days; $\sqrt{0.436 \times 0.0009/10^{-6}} = 19.8$ days.) (c) Compute the plug-in probability of failure within 30 days with equation (23.10). ($\Phi(-0.828) + \text{e}^{9.689}\Phi(-4.479) = 0.2039 + 0.0605 = 0.264$.) (d) Would the predictive mixture over the posterior damage give a larger or a smaller 30-day probability, and why? (Larger, 0.274 by numerical integration: in the lower tail the failure probability is convex in the damage.)

Exercise 23.4. The corrosion component of Section 23.4.3 is found at a later inspection with $d_0 = 1.4\,\text{mm}$; the parameters are unchanged. (a) Compute the probability of reaching $D_{\text{crit}} = 2.0\,\text{mm}$ within one and within two years. ($a/\beta_D = 6$; $\text{e}^{-6}(1 + 6) = 0.0174$; $\text{e}^{-6}(1 + 6 + 18 + 36) = 0.1512$.) (b) With $c_p = \text{R\$}\,2000$ and $c_f = \text{R\$}\,50\,000$, compute the expected cost of the plans to replace in one year and in two years with equation (23.35). (R$ 2833 and R$ 9258.) (c) For a decision threshold $p^* = 0.05$, is the latest allowable replacement horizon shorter or longer than one and a half years? (Shorter: F_{RUL} already reaches 0.05 at $\text{RUL}_{0.05} = 1.40$ years, found numerically from equation (23.26).) (d) Explain why the gamma model, not the Wiener model, suits this damage. (Metal loss is irreversible; gamma increments are nonnegative, while a Wiener path can decrease.)

Chapter 24

Predictive Maintenance and PHM Decisions

24.1 Predictive maintenance as a decision problem

Chapter 21 modeled how physical damage accumulates, Chapter 22 showed how measurements support the detection and diagnosis of faults, and Chapter 23 turned the current condition of a component into a probability distribution over its remaining useful life (RUL). Those chapters produce predictions. This chapter turns predictions into actions. A diagnostic system may estimate a 12 % probability that a cooling-water pump reaches its failure threshold within the next 30 days. The number is useful, but it does not say whether to replace the pump today, inspect it tomorrow, reduce its load, or continue operating; answering that requires a decision model. The question of this chapter is

when to intervene in a degrading engineering system, given an uncertain remaining useful life, maintenance costs, operational consequences, and the possibility of collecting more information before acting.

Predicting a failure and deciding what to do about it are distinct steps. A defensible maintenance policy connects the condition of the component to the consequences of its failure, to the effectiveness of the available interventions, to the operational constraints of the plant, and to the economic value of each outcome. The chapter sets out the decision chain and the maintenance approaches, derives maintenance triggers and deadlines from the RUL distribution, works an integrated replacement decision for a degraded pump of the cooling plant, computes the expected value of an inspection, treats inspection scheduling, maintenance logistics and spare parts, states the sequential form of the problem, and ends with the uncertainty and credibility of the recommendation.

Jargon: PHM, predictive maintenance, belief state, maintenance policy, decision threshold, action lead time, value of information

Prognostics and health management (PHM) is the discipline that monitors the condition of equipment, diagnoses its faults, predicts its remaining useful life and uses those predictions to manage maintenance and operation. *Predictive maintenance* schedules an intervention from the predicted future condition and risk of the equipment, as distinct from its age or from a present reading alone. The *belief state* b_t is the posterior probability distribution, given all evidence up to time t, over the physical condition of the system and the uncertain parameters of its model. A *maintenance policy* (or decision rule) is a function δ that maps the information available at each time into an action. A *decision threshold* p^* is the value of a failure probability at which two actions have equal expected loss; it is dimensionless. The *action lead time* τ_{L}, in days, is the time from the decision to the completion of the intervention, during which the equipment remains exposed to failure. The *value of information* is the reduction in the expected loss of the best decision that a measurement or inspection is expected to bring, in reais, before its cost is paid.

Key idea

A prognostic forecast is an input to a decision, not the decision itself. The action follows from comparing the expected losses of the available actions, which combine the failure probability with the consequences, the intervention costs, the lead times and the possibility of learning more; a probability threshold derived this way is a ratio of costs, not an alarm setting.

This section describes the decision chain that links measurements to actions and compares the four approaches to maintenance by the information each uses.

24.1.1 The PHM decision chain

Figure 24.1 draws the decision chain. Each box is one step, with its name in bold and a short description below; the labels on the arrows are the objects that pass from one step to the next. The top row is inference: measurements of vibration, temperature, motor current, flow and operating history are collected (*observe*), the damage and the probabilities of competing fault mechanisms are estimated from them (*diagnose*, Chapter 22), and the conditional distribution of the future failure time is computed (*predict*, Chapter 23). The bottom row is decision: the consequences of the alternative operating and maintenance actions are estimated (*evaluate*), one of them is selected (*decide*), and the new measurements and the effect of the action are assimilated (*update*) before the next cycle. The arrow from the last box back to the first closes the loop.

The chain is a feedback process. After a maintenance action the system may have a new effective age, altered parameters or a different dominant failure mechanism; after a new measurement the posterior distribution of the degradation changes. A maintenance policy therefore maps information about the system into an action, not a fixed date into an action. Let $\mathbf{x}_t$ be the physical condition of the system at time t (the degradation states $D(t)$ of its components, written as in Chapter 21; D here is a degradation state, not the FMEA detection rating D), let $\boldsymbol{\theta}$ be the uncertain parameters of the model, and let $\mathcal{Y}_{0:t}$ be the measurements and other evidence collected up to t. The belief state is the posterior

$$b_t = p(\mathbf{x}_t, \boldsymbol{\theta} \mid \mathcal{Y}_{0:t}), \tag{24.1}$$

and a maintenance policy is a rule

$$a_t = \delta(b_t) \tag{24.2}$$

that selects the action a_t from the belief state. The action can be to continue operating, reduce the load, collect another measurement, perform an inspection, repair a component, replace it,

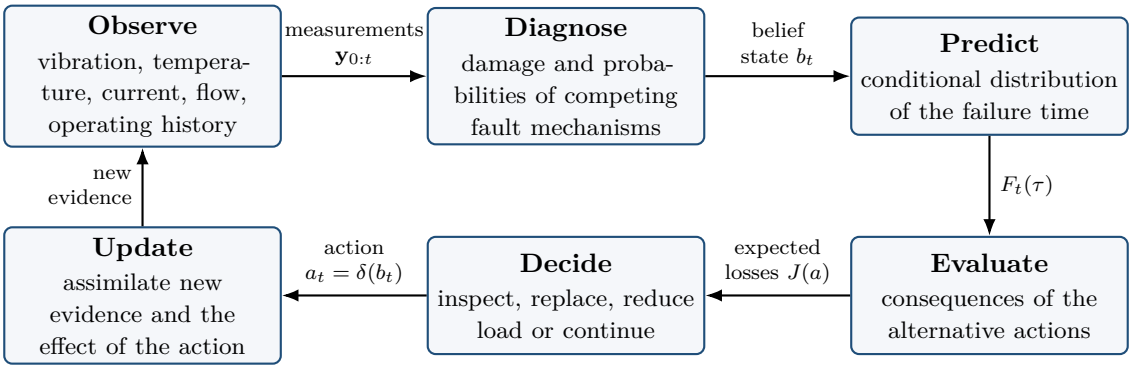


Figure 24.1: The PHM decision chain. The top row infers the condition and the remaining useful life from measurements; the bottom row evaluates and selects an action and assimilates its outcome. The process is a feedback loop, not a single calculation.

or switch to redundant equipment. The objective is a policy that produces acceptable reliability and operational performance at reasonable cost; the rest of the chapter constructs such policies, first for one decision and then for a sequence of them.

24.1.2 Four maintenance approaches

Maintenance approaches differ by the information their decision rule uses. Table 24.1 lists the four standard ones, with the decision rule of each and its main limitation. *Corrective maintenance*, defined in Chapter 20 of Volume I, restores the item after failure. *Time-based preventive maintenance* services the item at a predefined age or calendar interval, as the age-replacement policy of Chapter 18 of Volume I does. *Condition-based maintenance* (CBM) services the item when a measured condition indicator crosses a threshold. *Predictive maintenance* uses the estimated future condition and the risk it implies to schedule the action.

Table 24.1: Four maintenance approaches, ordered by the information their decision rule uses.

Approach	Information used	Main decision rule	Limitation
Corrective maintenance	failure event	repair or replace after failure	may incur costly unplanned outages
Time-based preventive maintenance	age or calendar time	service at a predefined age or interval	ignores differences in actual condition
Condition-based maintenance (CBM)	present measured condition	service when a measured condition crosses a threshold	depends on measurement quality and on the validity of the indicator
Predictive maintenance	predicted future condition and risk	schedule the action from the predicted failure probability and its consequences	requires calibrated prognostics and a decision model

The four approaches are not mutually exclusive. A facility may combine mandatory periodic inspection, continuous condition monitoring and predictive replacement for the same equipment. Condition-based practice monitors the actual equipment condition against a baseline and intervenes when deterioration or an impending failure is indicated; predictive maintenance adds the forecast and the decision model, and is the subject of the rest of the chapter.

24.2 Probability-based maintenance triggers

The simplest use of a prognosis is a trigger: a rule that starts maintenance planning when a predicted failure probability crosses a stated value. This section defines the trigger on the RUL distribution, derives a maintenance deadline from an RUL quantile, and adds the lead time of the action.

24.2.1 The RUL distribution as a trigger

Chapter 23 defined the remaining useful life at time t of a component that has not failed as $\mathrm{RUL}(t) = T_\mathrm{f} - t$, conditional on $T_\mathrm{f} > t$, where T_f is the first-passage time of the degradation state $D(t)$ to the failure threshold D_crit. Write its conditional cumulative distribution function, given the evidence, as

$$F_t(\tau) = \Pr\big(\mathrm{RUL}(t) \le \tau \mid \mathcal{Y}_{0:t}\big), \tag{24.3}$$

the probability, dimensionless, of reaching the failure threshold within the next τ days. A *probability-based maintenance trigger* is the rule

$$F_t(\tau) \ge p_\mathrm{trig}, \tag{24.4}$$

for a stated horizon τ and trigger probability p_trig. For example, $\Pr(\mathrm{RUL}(t) \le 30 \text{ days}) \ge 0.10$ could start maintenance planning. The value 0.10 is an assumed decision threshold, not an engineering standard: the appropriate value depends on the consequences of failure, the intervention costs, the redundancy and the action lead time, and Section 24.3.4 derives it from costs.

As an illustration, take the bearing of the cooling-water pump of Chapter 23, whose degradation index follows the Wiener model $D(t) = D_0 + \mu_D t + \sigma_D B(t)$ with drift $\mu_D = 0.01\,\mathrm{day}^{-1}$, diffusion $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$, standard Brownian motion $B(t)$ and threshold $D_\mathrm{crit} = 1$. With the current damage d_0 treated as known and the remaining distance $a = D_\mathrm{crit} - d_0$, the RUL has the inverse-Gaussian first-passage density and distribution function derived there,

$$f_t(\tau) = \frac{a}{\sigma_D\sqrt{2\pi\tau^3}} \exp\left[-\frac{(a-\mu_D\tau)^2}{2\sigma_D^2\tau}\right], \qquad F_t(\tau) = \Phi\left(\frac{\mu_D\tau - a}{\sigma_D\sqrt{\tau}}\right) + \mathrm{e}^{2\mu_D a/\sigma_D^2}\,\Phi\left(-\frac{\mu_D\tau + a}{\sigma_D\sqrt{\tau}}\right), \tag{24.5}$$

with mean a/μ_D, in days. At $d_0 = 0.40$, $a = 0.60$ and the mean RUL is 60 days, while $F_t(30) = 0.0473$: the 10 % trigger at 30 days does not fire. Solving $F_t(30) = 0.10$ for the current damage, with the drift and diffusion fixed, gives $d_0 = 0.462$, so that for this model the trigger is equivalent to a damage threshold of 0.462 on the estimated degradation index. A probability trigger is therefore a condition-based threshold whose value has been set through the degradation model, the horizon and the trigger probability, rather than by convention.

24.2.2 A maintenance deadline from an RUL quantile

Suppose the operator requires the probability of failure before the planned maintenance to be at most ε, the maximum acceptable probability, for example $\varepsilon = 0.05$:

$$F_t(\tau) \le \varepsilon. \tag{24.6}$$

Because F_t is nondecreasing in τ, the set of horizons that satisfy the requirement is an interval starting at zero, and its upper end is the largest allowable horizon. Defining the RUL quantile as in Chapter 23, $\mathrm{RUL}_\alpha(t) = \inf\{\tau : F_t(\tau) \ge \alpha\}$, the requirement holds for every τ up to

$$\tau_\mathrm{max} = \mathrm{RUL}_\varepsilon(t) = F_t^{-1}(\varepsilon), \tag{24.7}$$

the last equality holding when F_t is continuous and strictly increasing, so that $F_t(\tau_{\max}) = \varepsilon$ exactly. This produces a probabilistic maintenance deadline: maintenance should be completed by the ε-quantile of the RUL under the modeled operating scenario and criterion.

For the bearing with $d_0 = 0.40$, the 5 % quantile is $\mathrm{RUL}_{0.05}(t) = 30.3$ days; if the damage estimate were 0.55 instead, it would be 20.4 days, and at 0.50 it would be 23.6 days. Figure 24.2 draws the two distribution functions of equation (24.5). In both panels the horizontal axis is the horizon τ in days and the vertical axis the probability of reaching the threshold within it; the solid blue curve is the bearing at $d_0 = 0.40$ and the dashed red curve at $d_0 = 0.55$. Panel (a) shows the full distributions over 120 days; panel (b) magnifies the lower tail, where the dashed horizontal lines mark the trigger probability 0.10 and the tolerance $\varepsilon = 0.05$, the dotted vertical line marks the 30-day horizon, the dots on the tolerance line mark the 5 % quantiles, labeled in days, and the upper red dot marks the 30-day probability of the more degraded bearing. A damage increase of 0.15 shortens the deadline by ten days and multiplies the 30-day failure probability by five, from 4.73 % to 23.58 %, although the mean RUL falls only from 60 to 45 days: the deadline is governed by the lower tail of the distribution, not by its center.

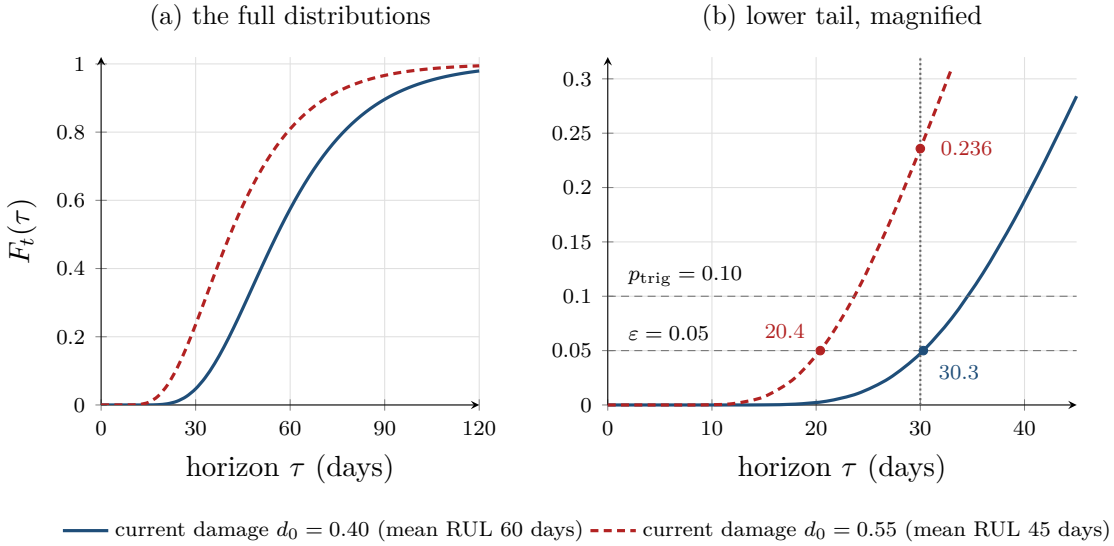


Figure 24.2: RUL distribution functions of the pump bearing, equation (24.5), Wiener degradation with $\mu_D = 0.01\,\mathrm{day}^{-1}$, $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$ and $D_{\mathrm{crit}} = 1$, for current damage $d_0 = 0.40$ and $d_0 = 0.55$ (illustrative values). Panel (b) magnifies the lower tail, with the trigger probability, the tolerance, the 30-day horizon, the 5 % quantiles $\mathrm{RUL}_{0.05}(t)$ in days, and $F_t(30) = 0.236$ for $d_0 = 0.55$.

Common mistake

Treating a maintenance deadline as an inspection deadline. The deadline of equation (24.7) is the latest time by which the intervention must be completed for the failure probability to stay within ε. An inspection before that time may provide new information and change the plan, but unless that possibility and the action that follows the inspection are explicitly modeled, deferring the maintenance in expectation of a future inspection exposes the system to the risk the deadline was meant to bound.

24.2.3 Action lead time

A replacement pump or bearing assembly must be ordered, delivered, installed and tested. The action lead time is the sum

$$\tau_{\mathrm{L}} = \tau_{\mathrm{proc}} + \tau_{\mathrm{mob}} + \tau_{\mathrm{inst}}, \tag{24.8}$$

in days, of the procurement, mobilization and installation times. During τ_{L} the component remains in service, so a decision taken now cannot prevent the failures that occur before the action is complete. If the lead time is deterministic, the relevant risk is

$$\Pr\bigl(\mathrm{RUL}(t) \le \tau_{\mathrm{L}} \mid \mathcal{Y}_{0:t}\bigr) = F_t(\tau_{\mathrm{L}}). \tag{24.9}$$

For the bearing with a total lead time of 21 days, $F_t(21) = 0.0035$ at $d_0 = 0.40$, but $F_t(21) = 0.0578$ at $d_0 = 0.55$. In the second case even an order placed at once cannot meet a tolerance of $\varepsilon = 0.05$: the deadline of 20.4 days falls before the action can be completed, and the remaining options are to shorten the lead time, to reduce the load and slow the degradation, or to accept the risk explicitly.

If the lead time is uncertain, with density $f_{\mathrm{L}}(\ell)$, and independent of the remaining life, the law of total probability over the value ℓ of the lead time gives

$$\Pr\bigl(\mathrm{RUL}(t) \le \tau_{\mathrm{L}}\bigr) = \int_0^{\infty} \Pr\bigl(\mathrm{RUL}(t) \le \ell\bigr)\, f_{\mathrm{L}}(\ell)\,\mathrm{d}\ell = \int_0^{\infty} F_t(\ell)\, f_{\mathrm{L}}(\ell)\,\mathrm{d}\ell, \tag{24.10}$$

the RUL distribution averaged over the lead-time distribution. As an illustration, let the lead time be 14, 21 or 35 days with equal probabilities, for the bearing at $d_0 = 0.40$. The sum in place of the integral gives $(0.00003 + 0.00346 + 0.10624)/3 = 0.0366$, against $F_t(23.3) = 0.0084$ at the mean lead time of 23.3 days. The lower tail of F_t is convex, so a variable lead time carries more risk than its mean suggests, here by a factor of four. If the degradation and the logistic delays share a cause, for example a regional heat wave that both accelerates wear and delays deliveries, the independence assumption fails and the joint distribution of $\mathrm{RUL}(t)$ and τ_{L} is required.

24.3 The integrated cooling-pump replacement decision

The trigger and the deadline answer whether a probability is high, not which action is best. This section works a complete decision for a degraded pump of the cooling plant: the evidence, the actions, their costs, the comparison of replacing now against waiting, and the probability threshold at which the preferred action changes.

24.3.1 The plant and the current evidence

The plant is the pumping branch of the data center cooling plant of Chapter 20 of Volume I: two identical cooling-water pumps in parallel, either of which can supply the required flow under the current operating conditions. That chapter modeled the pumps with a constant failure rate $\lambda = 0.001\,\mathrm{h}^{-1}$, a mean time to restore of 20 h and one repair crew, a three-state Markov model with steady-state availability $A = 0.99923136$. Here Pump A shows elevated vibration and signs of bearing degradation, while Pump B appears healthy, and the operator must decide what to do with Pump A. The numbers of the case are hypothetical, chosen to make the decision calculation transparent.

The prognostic model reports the probabilities that Pump A reaches its bearing failure threshold within four horizons, listed in Table 24.2. The second column is $F_t(\tau)$; the third is the conditional probability of reaching the threshold in each interval given survival to its start, $[F_t(\tau_j) - F_t(\tau_{j-1})]/[1 - F_t(\tau_{j-1})]$, computed from the second. The four values are nested cumulative probabilities from one forecast, and they are internally consistent because $F_t(7) \le F_t(30) \le F_t(60) \le F_t(90)$. The conditional probabilities per interval grow from 0.102 to 0.385 over intervals of 23 to 30 days, the signature of a wear-out mechanism. The forecast does not say that the pump must fail within 90 days: it assigns a probability of 0.40 to survival beyond that horizon. Linear interpolation in the table places the median RUL near 78 days and the 5 % quantile near 13.9 days.

Table 24.2: Prognosis of Pump A: cumulative probability of reaching the bearing failure threshold within each horizon, and the conditional probability for each interval given survival to its start (illustrative values).

Horizon τ (days)	$F_t(\tau)$	Interval probability given survival
7	0.02	0.020 (days 0–7)
30	0.12	0.102 (days 7–30)
60	0.35	0.261 (days 30–60)
90	0.60	0.385 (days 60–90)

With the trigger of Section 24.2.1, $F_t(30) = 0.12 \geq 0.10$, so maintenance planning starts. The trigger does not say whether the plan should be an immediate replacement, a deferred one, or an inspection first; the remaining subsections answer that.

Remark 24.1. The probabilities of Table 24.2 describe the bearing degradation mode of Pump A only. The constant rate $\lambda = 0.001\,\mathrm{h}^{-1}$ of Chapter 20 describes the random failure modes of a pump and was chosen deliberately high there; over 30 days it implies a probability $1 - \mathrm{e}^{-0.72} = 0.51$ of a random failure, larger than the bearing probability of the degraded pump. The two sets of illustrative numbers are therefore not mutually calibrated. The decision of this section is not affected, because the random modes have a constant hazard: replacing Pump A does not change their probability (the memoryless argument of Chapter 18 of Volume I), so they contribute equally to every action and cancel from the comparison. The rate λ is used below only for the few hours of a maintenance window.

24.3.2 Available actions and decision costs

Three immediate actions are available:

1. replace Pump A now;
2. continue operating for 30 days, then replace Pump A if it survives;
3. inspect Pump A now, update its prognosis, and then choose between immediate replacement and the 30-day deferral.

The first two are compared in this section and the third in Section 24.4. Table 24.3 lists the costs. A planned replacement costs $c_p = \mathrm{R\$}\,4000$. An unplanned failure requires a corrective replacement and causes additional disruption; its total cost, if Pump A fails before the planned replacement, is $c_f = \mathrm{R\$}\,60\,000$. The incremental failure penalty, relative to a planned replacement, is

$$\Delta c = c_f - c_p = \mathrm{R\$}\,60\,000 - \mathrm{R\$}\,4000 = \mathrm{R\$}\,56\,000. \tag{24.11}$$

Delaying the replacement preserves some useful pump life, avoids an immediate disruption and allows a better maintenance window to be chosen. The economic benefit of waiting 30 days, relative to replacing now, is represented by a single gain $G_\mathrm{w} = \mathrm{R\$}\,9000$. It is a modeled operational and asset-use value of the deferral, not necessarily a cash payment or a maintenance expense.

Two assumptions complete the model. If Pump A fails during the 30 days, the failure takes the place of the planned replacement and incurs the total corrective cost c_f. The comparison covers one decision: the costs and benefits that occur after the replacement, in the following life cycles of the new pump, are excluded, and the gain G_w is credited to the waiting action whether or not the failure occurs.

Table 24.3: Costs and characteristics of the pump decision (illustrative values). The last four rows are used in Sections 24.4 and 24.5.

Quantity	Symbol	Value
Planned replacement cost	c_p	R\$ 4000
Total cost of a failure before planned replacement	c_f	R\$ 60 000
Incremental failure penalty	$\Delta c = c_f - c_p$	R\$ 56 000
Gain from waiting 30 days	G_{w}	R\$ 9000
Probability of failure within 30 days	$p = F_t(30)$	0.12
Inspection cost	c_{I}	R\$ 500
Inspection sensitivity, specificity	–	0.85, 0.90
Duration of the planned replacement window	w	4 h
Consequence of losing pumping in the window	C_{sys}	R\$ 150 000

24.3.3 Replacing now against waiting

Write the expected cost of each action in full and then subtract a common reference. Replacing now costs c_p with certainty and gives up the gain G_{w}. Waiting 30 days yields the gain G_{w} and ends in a planned replacement with probability $1 - p$ or in a failure with probability $p = F_t(30)$, so its expected cost is

$$(1 - p)\, c_p + p\, c_f - G_{\mathrm{w}} = c_p + p\, \Delta c - G_{\mathrm{w}}.$$

Measuring both against the reference $c_p - G_{\mathrm{w}}$, the cost of a planned replacement after a successful deferral, gives the expected incremental decision losses

$$J_{\mathrm{rep}} = G_{\mathrm{w}}, \qquad J_{\mathrm{wait}} = p\, \Delta c, \tag{24.12}$$

in reais. The first is the gain given up by acting early; the second is the expected failure penalty of waiting. With $p = 0.12$,

$$J_{\mathrm{rep}} = \mathrm{R\$}\, 9000, \qquad J_{\mathrm{wait}} = 0.12 \times \mathrm{R\$}\, 56\,000 = \mathrm{R\$}\, 6720.$$

Under this objective, waiting has the lower expected loss, by R\$ 9000 − R\$ 6720 = R\$ 2280. The two numbers are losses relative to a common reference, not total maintenance expenditures; the planned replacement cost, common to both actions, has been removed from each.

24.3.4 The optimal probability threshold

The operator should replace immediately when $J_{\mathrm{rep}} < J_{\mathrm{wait}}$, that is, when $G_{\mathrm{w}} < p\, \Delta c$. Solving for the failure probability,

$$p > p^{*} = \frac{G_{\mathrm{w}}}{\Delta c} = \frac{G_{\mathrm{w}}}{c_f - c_p}. \tag{24.13}$$

For the case, $p^{*} = 9000/56\,000 = 0.1607$, or 16.1 %, and the decision rule is

$$\text{wait 30 days if } F_t(30) < 0.1607, \qquad \text{replace now if } F_t(30) > 0.1607,$$

with indifference at equality. At $p = 0.12$ the rule says wait.

Figure 24.3 draws the two losses of equation (24.12) against the failure probability. The horizontal axis is $p = F_t(30)$; the vertical axis is the expected incremental loss in reais. The flat blue line is J_{rep}, which does not depend on p; the rising red line is J_{wait}, with slope Δc.

They cross at p^*; to the left of the crossing waiting is cheaper, to the right replacement is. The vertical dotted line marks the present prognosis $p = 0.12$, at which the gap between the lines is the R$ 2280 advantage of waiting.

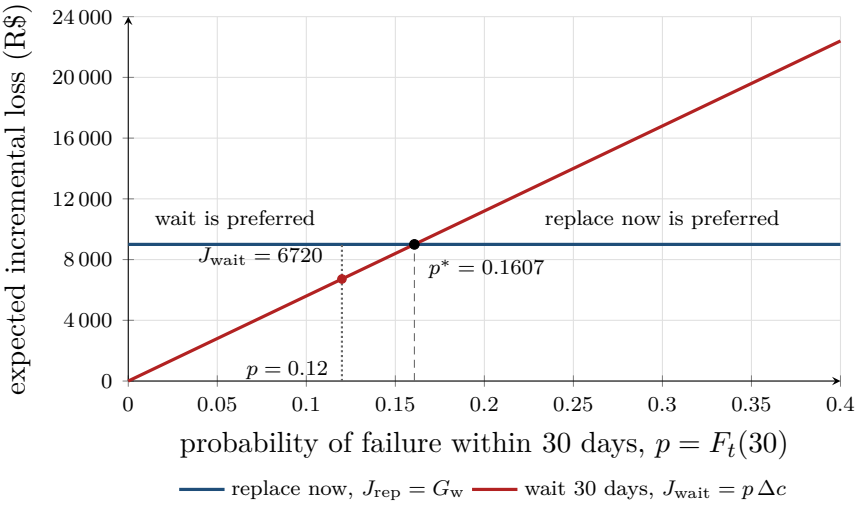


Figure 24.3: Expected incremental losses of replacing Pump A now and of waiting 30 days, equation (24.12), against the probability of failure within 30 days, with $G_{\text{w}} = $ R$ 9000 and $\Delta c = $ R$ 56 000 (illustrative values). The lines cross at the decision threshold $p^* = G_{\text{w}}/\Delta c$.

The threshold is not an arbitrary alarm setting; it is a ratio of costs. Equation (24.13) decreases as the failure penalty Δc grows, so a more consequential failure justifies acting at a lower probability, and increases as the gain from waiting G_{w} grows, so a more costly early replacement requires a higher probability before it is justified. The 10 % trigger of Section 24.2.1 fired at $p = 0.12$, while the cost threshold says wait: a trigger chosen by convention and a threshold derived from costs can disagree, and only the second is a statement about the decision.

Remark 24.2. The threshold depends on how the gain from waiting is credited. If the gain G_{w} is obtained only when Pump A survives the 30 days, for example because a failure forfeits the chosen maintenance window, the waiting loss becomes $p\,\Delta c + p\,G_{\text{w}}$ and the threshold becomes $G_{\text{w}}/(\Delta c + G_{\text{w}}) = 9000/65\,000 = 0.1385$. At $p = 0.12$ the decision is still to wait, but the margin is smaller. The model must state which assumption it uses.

Principle 24.1. *When the decision is between acting now and deferring, and the expected loss of deferring is linear in the failure probability, the optimal rule is a probability threshold, $p^* = G_{\text{w}}/(c_f - c_p)$: the ratio of the cost of acting early to the incremental penalty of failure.*

24.4 The expected value of inspection and inspection scheduling

The prognosis $p = 0.12$ lies below the threshold 0.1607, but not far below it, and the same probability could conceal two situations: a bearing that is close to failure, or a vibration caused by a benign imbalance or by a measurement artifact. This section computes whether an inspection that distinguishes them is worth its cost, and then treats the timing and the scheduling of inspections.

24.4.1 The inspection test and the probability of each outcome

An inspection is informative when its outcome is correlated with the event that matters for the decision. The question is not whether the inspection is informative, but whether its information is likely to change the decision enough to justify its cost. Define the event

$$F_{30} = \{\text{Pump A reaches its failure threshold within 30 days}\},$$

with complement F_{30}^{c}, so that the prior probability from the prognosis is $\Pr(F_{30}) = F_t(30) = 0.12$ and $\Pr(F_{30}^{\mathrm{c}}) = 0.88$. A specialized inspection returns one of two outcomes, positive Y^{+}, indicating elevated risk, or negative Y^{-}. Its *sensitivity* is the probability of a positive outcome when the event will occur, and its *specificity* the probability of a negative outcome when it will not; assume

$$\Pr(Y^{+} \mid F_{30}) = 0.85, \qquad \Pr(Y^{-} \mid F_{30}^{\mathrm{c}}) = 0.90,$$

so that the false-positive probability is $\Pr(Y^{+} \mid F_{30}^{\mathrm{c}}) = 0.10$ and the false-negative probability $\Pr(Y^{-} \mid F_{30}) = 0.15$. These are hypothetical test characteristics, defined against the 30-day failure event, not against the visibility of bearing wear. Assume also that the inspection gives its result at once, does not change the degradation state of the pump, and costs $c_{\mathrm{I}} = \mathrm{R\$}\,500$.

The probability of each outcome follows from the law of total probability over the two states of the event:

$$\begin{aligned} \Pr(Y^{+}) &= \Pr(Y^{+} \mid F_{30})\,\Pr(F_{30}) + \Pr(Y^{+} \mid F_{30}^{\mathrm{c}})\,\Pr(F_{30}^{\mathrm{c}}) \\ &= 0.85 \times 0.12 + 0.10 \times 0.88 = 0.102 + 0.088 = 0.19, \end{aligned} \tag{24.14}$$

and $\Pr(Y^{-}) = 1 - 0.19 = 0.81$. Of the 19 % of positive outcomes, slightly more than half (0.102) come from pumps that will fail and the rest (0.088) are false alarms from the much larger group of pumps that will not.

24.4.2 Bayes updates after inspection

Bayes' theorem (Chapter 8 of Volume I) gives the posterior failure probability after each outcome. After a positive outcome,

$$\Pr(F_{30} \mid Y^{+}) = \frac{\Pr(Y^{+} \mid F_{30})\,\Pr(F_{30})}{\Pr(Y^{+})} = \frac{0.85 \times 0.12}{0.19} = \frac{0.102}{0.19} = 0.5368, \tag{24.15}$$

which exceeds the threshold $p^{*} = 0.1607$: after a positive inspection the pump should be replaced at once. After a negative outcome,

$$\Pr(F_{30} \mid Y^{-}) = \frac{\Pr(Y^{-} \mid F_{30})\,\Pr(F_{30})}{\Pr(Y^{-})} = \frac{0.15 \times 0.12}{0.81} = \frac{0.018}{0.81} = 0.0222, \tag{24.16}$$

well below p^{*}: after a negative inspection the replacement should be deferred 30 days. The same updates read more simply in odds. The prior odds of failure are $0.12/0.88 = 0.1364$; a positive outcome multiplies them by the likelihood ratio $0.85/0.10 = 8.5$, to 1.159, a probability of $1.159/2.159 = 0.5368$, and a negative outcome by $0.15/0.90 = 0.1667$, to 0.0227, a probability of 0.0222.

The inspection therefore creates a contingent decision rule: replace if positive, wait if negative. That the two outcomes lead to different actions is the source of all the value the inspection has in this model.

24.4.3 Expected cost after inspection and the value of inspection

Without the inspection, the best action is to wait, with expected loss

$$J^{*}_{\text{without}} = \min\{G_{\text{w}},\ p\,\Delta c\} = \min\{9000,\ 6720\} = \text{R\$}\,6720.$$

With the inspection, the best action is chosen after the outcome is known, so the expected loss is the average over the outcomes of the minimum loss given each outcome:

$$J^{*}_{\text{with}} = \Pr(Y^{+})\,\min\{G_{\text{w}},\ \Pr(F_{30} \mid Y^{+})\,\Delta c\} + \Pr(Y^{-})\,\min\{G_{\text{w}},\ \Pr(F_{30} \mid Y^{-})\,\Delta c\}, \quad (24.17)$$

before the inspection cost. A positive outcome, with probability 0.19, leads to replacement and a loss of R\$ 9000; a negative outcome, with probability 0.81, leads to waiting and an expected loss of $0.0222 \times 56\,000 = \text{R\$}\,1244$. Since $\Pr(Y^{-})\,\Pr(F_{30} \mid Y^{-}) = 0.15 \times 0.12 = 0.018$,

$$J^{*}_{\text{with}} = 0.19 \times 9000 + 0.018 \times 56\,000 = 1710 + 1008 = \text{R\$}\,2718.$$

The *expected value of sample information* (EVSI) is the reduction in the expected loss of the best decision that the inspection brings, before its cost,

$$\text{EVSI} = J^{*}_{\text{without}} - J^{*}_{\text{with}} = \text{R\$}\,6720 - \text{R\$}\,2718 = \text{R\$}\,4002, \quad (24.18)$$

which is the value of this inspection. After its cost, the expected loss of inspecting and then deciding is $J_{\text{insp}} = \text{R\$}\,2718 + \text{R\$}\,500 = \text{R\$}\,3218$, and the net expected benefit relative to the best action without inspection is $\text{R\$}\,6720 - \text{R\$}\,3218 = \text{R\$}\,3502$, equal to $\text{EVSI} - c_{\text{I}}$. The inspection is the preferred action of the three.

Figure 24.4 draws the decision as a tree. Squares are decision nodes, where the operator chooses a branch; circles are chance nodes, where the branches carry probabilities; the numbers at the right are the losses at the end of each path, in reais, and the numbers in boxes above the nodes are the expected losses obtained by rolling the tree back from the right, taking the probability-weighted average at a chance node and the minimum at a decision node. Thick branches are the ones chosen. The tree shows the two outcomes of the inspection leading to different choices, which is why the inspection branch has the lowest expected loss.

Figure 24.5 compares the three immediate policies. The horizontal axis lists the policies; the height of each bar is the expected incremental loss in reais, with its value above it, and the bar of the inspection policy is split into the expected loss after the outcome-dependent decision (lower part) and the inspection cost (upper part). The inspection policy costs about a third of replacing now and half of waiting. The losses are one-decision losses measured from the same reference; the common replacement expenses and the later operating cycles are excluded from all three.

An upper bound on the value of any inspection is the value of *perfect* information about the event, an inspection with sensitivity and specificity equal to one. It would reveal whether the failure will occur, so the operator would replace exactly when it will, with expected loss $\Pr(F_{30})\,G_{\text{w}} = 0.12 \times 9000 = \text{R\$}\,1080$, and the *expected value of perfect information* is

$$\text{EVPI} = J^{*}_{\text{without}} - \Pr(F_{30})\,G_{\text{w}} = \text{R\$}\,6720 - \text{R\$}\,1080 = \text{R\$}\,5640. \quad (24.19)$$

No inspection can be worth more than R\$ 5640 for this decision, and the real inspection captures $4002/5640 = 71\,\%$ of that bound.

The value arises because different outcomes lead to different actions. If the operator would take the same action whatever the result, the inspection would have zero decision value, apart from any benefit outside the model. Two consequences follow. A test can greatly reduce the uncertainty without changing the preferred action, and then it is worth nothing to this decision. Conversely, a modestly informative test can be valuable when the system is close to a decision boundary and the consequences of choosing wrongly are high. This is the basis of inspection and maintenance scheduling driven by the value of information.

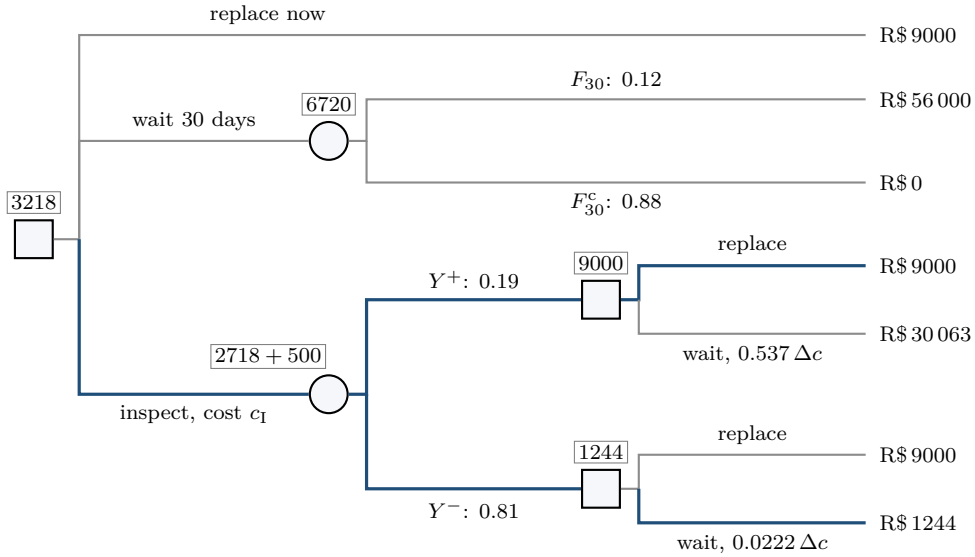


Figure 24.4: Decision tree of the pump decision. Squares are decisions, circles are chance outcomes with their probabilities; losses in reais at the right, relative to the common reference of Section 24.3.3; boxed values are the rolled-back expected losses. Thick branches are chosen: inspect, then replace after Y^{+} and wait after Y^{-} (illustrative values).

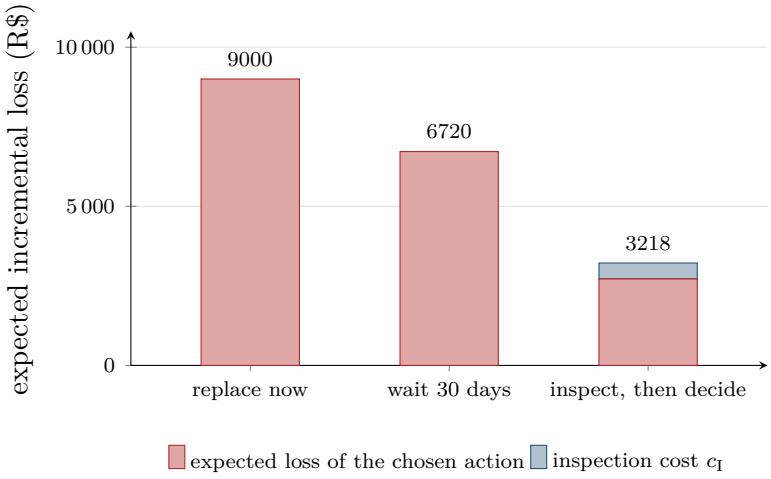


Figure 24.5: Expected incremental losses of the three immediate policies for Pump A: replace now (G_w), wait 30 days ($p\,\Delta c$), and inspect then decide ($J^{*}_{\text{with}} + c_I$), in reais, measured from the same reference (illustrative values).

Principle 24.2. *An inspection has decision value only when its possible outcomes lead to different actions. Its value is the expected value of sample information, $J^{*}_{\text{without}} - J^{*}_{\text{with}}$, bounded by the expected value of perfect information; it is worth performing when that value exceeds its cost.*

24.4.4 The periodic inspection model

The preceding analysis concerned one inspection. A degrading component is inspected repeatedly, and the interval between inspections must be chosen. Frequent inspections detect degradation early but cost money and may require planned outages; infrequent inspections save resources but increase the probability that undetected degradation reaches the failure threshold. Let the equipment be inspected every Δ days, each inspection cost c_{I}, let the component be restored, replaced or otherwise managed at each inspection according to a fixed policy, let $p_{\mathrm{F}}(\Delta)$ be the probability of an unplanned failure within one inspection cycle under that policy, and let Δc be the incremental cost of a failure. By the renewal-reward argument of Chapter 18 of Volume I, the long-run cost rate is the expected cost of one cycle divided by its expected length,

$$C(\Delta) = \frac{c_{\mathrm{I}} + \Delta c \, p_{\mathrm{F}}(\Delta)}{\mathbb{E}[\text{cycle length}]}, \tag{24.20}$$

in reais per day. If the cycle length is effectively Δ, the simplified form is

$$C(\Delta) \approx \frac{c_{\mathrm{I}} + \Delta c \, p_{\mathrm{F}}(\Delta)}{\Delta}. \tag{24.21}$$

The approximation needs care: a failure ends a cycle early, so the expected cycle length is shorter than Δ, and the exact renewal-reward form must be used when $p_{\mathrm{F}}(\Delta)$ is not small.

To make the model concrete, assume each inspection returns the component to a condition equivalent to new, because the degraded part found is corrected at the inspection, so that inspection epochs are regeneration points; and assume the inspection-and-restoration cost c_{I} is paid in every cycle, at the scheduled time or as part of the corrective work after a failure. Then, with T the time to failure from the restored condition, $p_{\mathrm{F}}(\Delta) = F(\Delta)$, the expected cycle length is $\mathbb{E}[\min(T, \Delta)] = \int_0^{\Delta} R(t)\,\mathrm{d}t$, and equation (24.20) is exactly the age-replacement cost rate of Chapter 18,

$$C(T_p) = \frac{c_p \, R(T_p) + c_f \, F(T_p)}{\int_0^{T_p} R(t)\,\mathrm{d}t},$$

with $T_p = \Delta$, $c_p = c_{\mathrm{I}}$ and $c_f - c_p = \Delta c$. Its optimum satisfies the condition derived there, $(c_f - c_p)\,\lambda(\Delta^{*}) = C(\Delta^{*})$: at the best interval, the marginal expected failure penalty per day of extending the interval equals the average cost rate. A model that represents the detection of a defect separately from its development into a failure, the delay-time model, is the standard refinement when an inspection detects defects without restoring the component.

As an illustration, let the time to failure from the restored condition be Weibull with shape $\beta = 2$ and scale $\eta = 400$ days, so that $F(\Delta) = 1 - \exp[-(\Delta/400)^2]$ and $\lambda(\Delta) = 2\Delta/400^2$, with $c_{\mathrm{I}} = \mathrm{R\$}\,500$ and $\Delta c = \mathrm{R\$}\,56\,000$. For small Δ, $F(\Delta) \approx (\Delta/\eta)^2$ and the cycle length is close to Δ, so equation (24.21) becomes

$$C(\Delta) \approx \frac{c_{\mathrm{I}}}{\Delta} + \frac{\Delta c \, \Delta}{\eta^2}.$$

The first term, the inspection cost spread over the interval, falls with Δ; the second, the failure penalty per day, rises. Setting the derivative $-c_{\mathrm{I}}/\Delta^2 + \Delta c/\eta^2$ to zero gives

$$\Delta^{*} \approx \eta\sqrt{\frac{c_{\mathrm{I}}}{\Delta c}} = 400\sqrt{\frac{500}{56\,000}} = 37.8 \text{ days}, \qquad C(\Delta^{*}) \approx \frac{2\sqrt{c_{\mathrm{I}}\,\Delta c}}{\eta} = \mathrm{R\$}\,26.46 \text{ per day}. \tag{24.22}$$

Minimizing the exact rate (24.20) numerically gives $\Delta^* = 37.8$ days and $C(\Delta^*) = \text{R\$}\,26.48$ per day, with $F(37.8) = 0.0089$ and an expected cycle length of 37.7 days; the optimality condition checks, since $\Delta c\,\lambda(37.82) = 56\,000 \times 2 \times 37.82/160\,000 = \text{R\$}\,26.48$ per day. Without inspection every cycle ends in a failure, and the rate is $(c_\text{I} + \Delta c)/\text{MTTF} = 56\,500/354.5 = \text{R\$}\,159.4$ per day, six times the optimum.

Figure 24.6 draws the cost rate against the inspection interval. The horizontal axis is Δ in days and the vertical axis is the cost rate in reais per day. The solid blue curve is the exact rate (24.20); the dashed gray curve is the approximation (24.21); the dash-dotted red line is the marginal failure penalty $\Delta c\,\lambda(\Delta)$, which crosses the exact curve at its minimum; and the dotted vertical line is the interval permitted by a 5 % risk constraint, discussed in Section 24.4.5. The curve falls steeply for short intervals, where inspections are wasted, and rises slowly for long ones, where failures accumulate; its minimum is flat, with $C(30) = \text{R\$}\,27.19$ and $C(50) = \text{R\$}\,27.51$ per day. The two curves agree closely while $p_\text{F}(\Delta)$ is small and separate beyond about 100 days, where the approximation understates the rate by neglecting the cycles shortened by failure.

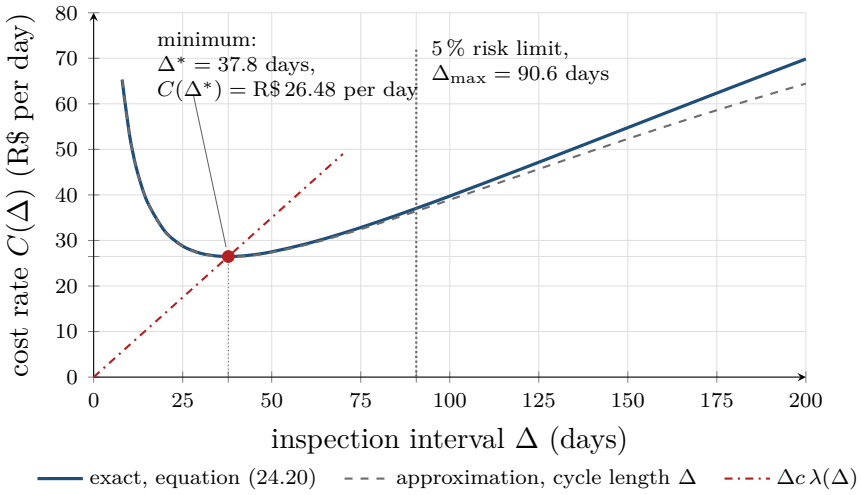


Figure 24.6: Long-run cost rate of periodic inspection with restoration against the inspection interval: Weibull time to failure from the restored condition with $\beta = 2$ and $\eta = 400$ days, $c_\text{I} = \text{R\$}\,500$, $\Delta c = \text{R\$}\,56\,000$ (illustrative values). The marginal failure penalty $\Delta c\,\lambda(\Delta) = 0.7\,\Delta$ crosses the exact rate at its minimum.

24.4.5 Inspection timing under changing risk

A cost rate presumes a stationary renewal cycle. When the condition of the component is known and changing, the interval can instead be set from today's RUL distribution. If the probability of failure before the next inspection must remain below 5 %,

$$F_t(\Delta) \le 0.05, \qquad \Delta_\text{max} = F_t^{-1}(0.05) = \text{RUL}_{0.05}(t), \tag{24.23}$$

which is the deadline of equation (24.7) applied to the next inspection: a simple risk-constrained inspection rule. For the Weibull component of Section 24.4.4, freshly restored, $\Delta_\text{max} = \eta\sqrt{-\ln 0.95} = 400 \times 0.2265 = 90.6$ days, longer than the cost-optimal 37.8 days, so the risk constraint is not binding there (Figure 24.6). For the bearing of Section 24.2.1 at $d_0 = 0.40$, $\Delta_\text{max} = 30.3$ days.

The rule rests on two assumptions that must be checked. It assumes that the inspection occurs before the failure, and that the actions following the inspection can manage the condition it reveals. If an inspection produces no intervention and the equipment continues to operate in the same state, collecting the measurement alone does not prevent the threshold crossing: the

inspection changes what is known, not what happens. The complete decision problem therefore includes the maintenance action that follows each inspection, as the decision tree of Figure 24.4 does for one inspection and Section 24.6 does in general.

24.4.6 Fixed and adaptive inspection schedules

Two policies can be compared. A *fixed schedule* inspects every 30 days whatever the condition. An *adaptive schedule* chooses the next inspection time from the current damage estimate, its uncertainty, the predicted failure probability and the expected value of information: sooner when the damage or its uncertainty rises or the operating conditions worsen, later when the component is demonstrably healthy and the predicted risk is low.

Table 24.4 shows what the difference amounts to for the bearing model of Section 24.2.1. Each row is a current damage estimate d_0, treated as known; the columns give the mean RUL a/μ_D, the probability $F_t(30)$ of failure within one fixed 30-day interval, and the adaptive interval $\mathrm{RUL}_{0.05}(t)$ that holds the probability of failure before the next inspection at 5 %. One row of the arithmetic, $d_0 = 0.50$: $a = 0.50$, the mean RUL is $0.50/0.01 = 50$ days, and equation (24.5) at $\tau = 30$ gives $\Phi(-0.2/0.1643) + \mathrm{e}^{11.11}\,\Phi(-0.8/0.1643) = 0.1118 + 0.0376 = 0.1494$.

Table 24.4: Fixed and adaptive inspection intervals for the pump bearing, Wiener degradation with $\mu_D = 0.01\,\mathrm{day}^{-1}$, $\sigma_D = 0.03\,\mathrm{day}^{-1/2}$ and $D_{\mathrm{crit}} = 1$, for several current damage estimates (illustrative values).

Current damage d_0	Mean RUL (days)	Fixed 30 days: $F_t(30)$	Adaptive: $\mathrm{RUL}_{0.05}(t)$ (days)
0.30	70	0.0108	37.2
0.40	60	0.0473	30.3
0.45	55	0.0875	26.9
0.50	50	0.1494	23.6
0.55	45	0.2358	20.4
0.60	40	0.3455	17.3
0.70	30	0.6024	11.4

Under the fixed schedule the probability of failure between inspections ranges from 1.1 % to 60 % across the rows: the schedule is cautious for a healthy bearing and reckless for a degraded one. The adaptive schedule holds the between-inspection risk at 5 % by shortening the interval from 37 to 11 days as the damage grows; for $d_0 \geq 0.45$ the adaptive interval is shorter than the fixed one, and for the most degraded rows it is so short that the inspection becomes a countdown to the maintenance deadline, at which point the decision is no longer when to inspect but when to replace. Comparisons of fixed and adaptive condition-based policies weigh their cost, downtime and between-inspection risk jointly; the adaptive policy is worth its added complexity when the condition evolves and the prognostic model is credible. The best inspection schedule depends on the evolving distribution of the condition, not on calendar age alone.

24.5 Maintenance logistics and system availability

A component prognosis is not a system risk, and an intervention is not free of risk. This section connects the prognosis of Pump A to the availability of the pumping function, prices the risk of the planned maintenance window and its reduction, and adds the lead time of spare parts and the inventory decision.

24.5.1 Component failure and system failure

Either pump of the plant can meet the cooling demand. A 12 % probability that Pump A reaches its failure threshold within 30 days therefore does not imply that the probability of losing the cooling service within 30 days is also 0.12, because Pump B may preserve the service after Pump A fails. The probability of a complete loss depends on the condition and failure process of Pump B; on whether the two pumps share a common cause (Chapter 13 of Volume I); on the performance of the automatic switchover; on the repair duration; on the cooling load; on any thermal buffer or storage; and on whether the remaining pump can actually deliver the required hydraulic duty alone.

An order of magnitude follows from the pump model of Chapter 20 of Volume I, with independent pumps and a mean repair time of 20 h. When Pump A fails, the pumping function is lost if Pump B is already under repair, with probability about $\lambda/(\lambda+\mu) = 0.001/0.051 = 0.0196$, or if Pump B, available, fails before the repair of Pump A is complete, again with probability about 0.0196 under exponential repair times. The conditional probability of a loss given the failure of Pump A is therefore about $0.0196+0.9804\times0.0196 = 0.039$, and the probability that the bearing failure of Pump A causes a loss of pumping within 30 days is about $0.12\times0.039 = 0.0047$, twenty-five times smaller than the component probability. A common cause, a switchover that fails on demand or a single pump unable to carry the summer load would raise it, possibly by orders of magnitude. A system-level decision must integrate the component prognosis with these mechanisms.

Common mistake

Reading a component failure probability as a system failure probability. The prognosis of one pump gives the probability of a component event; the probability that the required function is lost follows only from a system model that adds redundancy, restoration, dependencies and the physical requirement, and it may be much smaller or, with common causes, much closer to the component value.

24.5.2 Risk during the planned maintenance window

Replacing Pump A while Pump B stays in service requires a maintenance window of $w = 4\,\mathrm{h}$, during which the cooling system depends entirely on Pump B. Chapter 20 computed this risk for the same plant: if Pump B is operating at the start of the window, has the constant failure rate $\lambda = 0.001\,\mathrm{h}^{-1}$ and cannot be repaired within the window, it survives the window with probability $R_B(w) = \mathrm{e}^{-\lambda w}$, so

$$\Pr(\text{Pump B fails during the window}) = 1 - \mathrm{e}^{-\lambda w} = 1 - \mathrm{e}^{-0.004} = 0.003992, \tag{24.24}$$

about 0.399 %. If the state of Pump B at the start is uncertain, the probability of continuous service is multiplied by the probability p_{start} that it is operating, as equation (20.25) of that chapter states; confirming its condition before the window opens sets $p_{\mathrm{start}} = 1$.

The 0.399 % is much smaller than the 12 % of Pump A, but the two are not directly comparable: they concern different events over different horizons, a random failure of Pump B within four hours against a degradation failure of Pump A within thirty days. What makes the small number important is its consequence. A failure of Pump B during the window removes all cooling-water pumping capacity at once, without the redundancy that protects the plant the rest of the time.

24.5.3 Consequences of the maintenance window and their reduction

Suppose a loss of the pumping function during the window has a modeled consequence of $C_{\mathrm{sys}} = \mathrm{R\$}\,150\,000$, fixed per occurrence and with no further mitigation available. The expected

consequence of the window is

$$\mathbb{E}[C_{\text{window}}] = \left(1 - \mathrm{e}^{-\lambda w}\right) C_{\text{sys}} = 0.003992 \times \text{R\$}\,150\,000 = \text{R\$}\,599, \tag{24.25}$$

and the expected cost of the immediate intervention, including this risk, is R\$ 4000 + R\$ 599 = R\$ 4599. The consequence is a fixed amount per occurrence; it is much larger than the downtime cost model of Chapter 20 would give, $c_{\text{D}} = \text{R\$}\,10\,000$ per hour times the expected downtime of the window, $w - (1 - \mathrm{e}^{-\lambda w})/\lambda = 0.007989$ h, or about R\$ 80, because it represents consequences beyond the hours of downtime, such as a thermal excursion in the data hall or contractual penalties.

The window changes the comparison of Section 24.3.3, but both actions contain a window. Replacing now incurs it with certainty. Waiting incurs a planned window at day 30 with probability $1 - p$, and with probability p a corrective replacement instead, during which Pump B is alone for the mean repair time of 20 h, a loss probability of $1 - \mathrm{e}^{-0.02} = 0.0198$ and an expected consequence of R\$ 2970. With $q_w = 0.003992$ and $q_f = 0.0198$, the losses of equation (24.12) become

$$\begin{aligned} J_{\text{rep}} &= G_{\text{w}} + q_w\, C_{\text{sys}} = 9000 + 599 = \text{R\$}\,9599, \\ J_{\text{wait}} &= p\,(\Delta c + q_f\, C_{\text{sys}}) + (1 - p)\, q_w\, C_{\text{sys}} = 0.12 \times 58\,970 + 0.88 \times 599 = \text{R\$}\,7603, \end{aligned}$$

and the threshold becomes $p^* = G_{\text{w}}/[\Delta c + (q_f - q_w)\, C_{\text{sys}}] = 9000/58\,371 = 0.1542$. The decision is still to wait; including the window risk lowers the threshold slightly, because a failure exposes the plant to a longer single-pump period than a planned replacement does. A rigorous analysis models the window risks, costs and operating consequences of every candidate intervention in the same way.

The window risk can itself be reduced. The operator may verify the condition of Pump B before the window; schedule the work during a period of reduced cooling demand; arrange temporary pumping capacity; prepare the tools and spare components before isolating Pump A; test the switchover controls; shorten the maintenance duration; or provide an additional independent cooling path. Table 24.5 shows the effect of the duration alone. Each row is a window length; the columns are the probability $1 - \mathrm{e}^{-\lambda w}$ of losing Pump B within it and the expected consequence at $C_{\text{sys}} = \text{R\$}\,150\,000$. Preparing the work so that the window falls from four to two hours halves the expected consequence, from R\$ 599 to R\$ 300. The decision concerns not only when the maintenance is done but how.

Table 24.5: Risk of the maintenance window against its length, with Pump B operating at the start, $\lambda = 0.001\,\text{h}^{-1}$ and $C_{\text{sys}} = \text{R\$}\,150\,000$ per occurrence (illustrative values). The last row is a corrective replacement of mean duration 20 h.

Window length w	$1 - \mathrm{e}^{-\lambda w}$	Expected consequence (R\$)
2 h, prepared work	0.001998	300
4 h, base case	0.003992	599
8 h	0.007968	1195
20 h, corrective replacement	0.019801	2970

24.5.4 Spare-part lead time

Suppose the bearing assembly of Pump A is not in inventory, and let τ_{L} be its procurement lead time. The replacement cannot be completed until the part arrives, so the immediate replacement of Section 24.3.3 is not available: the actual choice is when to order and how fast to

have the part delivered. If procurement takes 30 days, the probability that Pump A reaches its failure threshold before delivery is, by equation (24.9), $F_t(30) = 0.12$. If an emergency spare can be bought from another supplier with a 7-day lead time, the probability falls to $F_t(7) = 0.02$. Let c_{E} be the additional cost of the emergency procurement. Its net expected cost, relative to the normal order, is

$$\Delta J = c_{\mathrm{E}} - \big[F_t(30) - F_t(7)\big]\,\Delta c = c_{\mathrm{E}} - (0.12 - 0.02) \times \mathrm{R\$}\,56\,000 = c_{\mathrm{E}} - \mathrm{R\$}\,5600, \qquad (24.26)$$

so the emergency procurement is worthwhile when it costs less than R$ 5600. The comparison assumes that both strategies deliver the same replacement outcome and differ only in lead time and procurement expense. A complete model would add the scheduling of the replacement, the remaining service life, the effect on the inventory and the risk of the installation window.

24.5.5 Spare inventory as a strategic decision

The value of stocking a critical part depends on more than the current prognosis of one component. For a fleet of pumps, generators, chillers or inverters, the demand for spares is random, and the inventory decision depends on the expected demand, the uncertainty of the lead time, the consequence of a stockout and the cost of holding stock; these are modeling ingredients, not terms of a sum in one unit. A full cost model for a stock level s is

$$J(s) = C_{\mathrm{hold}}(s) + \mathbb{E}\big[C_{\mathrm{exp}}(s)\big] + \mathbb{E}\big[C_{\mathrm{down}}(s)\big], \qquad (24.27)$$

in reais per year, the holding cost plus the expected costs of expediting and of downtime. Chapter 20 modeled the demand during a lead time τ as Poisson, $K(\tau) \sim \mathrm{Poisson}(\nu\tau)$ with demand rate ν, and chose s from the stockout probability $\Pr[K(\tau) > s]$, a model based on the historical failure frequency of the fleet.

Condition monitoring changes the inventory problem, because the demand in the coming lead time becomes partly predictable. If unit i of the fleet has RUL distribution F_i, and the units fail independently, the number of spares needed within a lead time τ_{L} is the sum of independent indicators with probabilities $F_i(\tau_{\mathrm{L}})$, a Poisson-binomial count with mean $\sum_i F_i(\tau_{\mathrm{L}})$. As an illustration, four pumps with 30-day probabilities 0.12, 0.05, 0.02 and 0.01 give an expected demand of 0.20; the probability of no demand is $0.88 \times 0.95 \times 0.98 \times 0.99 = 0.8111$, of exactly one demand 0.1780, and of two or more 0.0109. One spare in stock covers the next 30 days with probability 0.989, and the forecast also says which pump will most probably need it. A Poisson model with the same mean would give 0.0175 for two or more demands, because it ignores that each pump can fail only once within the lead time. Joint condition-based maintenance, inspection scheduling and spare-parts ordering, including components whose degradations are correlated, is an active area of maintenance optimization; the correlation matters because it widens the distribution of the demand.

24.6 Sequential decisions under uncertainty

The calculations so far compared immediate choices over fixed horizons. This section states the sequential form of the problem, gives its Bellman equation, evaluates the inspection decision over a range of inputs, and treats the uncertainty of the predicted probability, the credibility of the model and the competing objectives of a real plant.

24.6.1 State, actions and observations

Maintenance decisions are sequential, and one-step optimization is often insufficient. If Pump A is not replaced today, new vibration measurements will arrive in seven days and revise the posterior of its degradation; the operator may then replace the pump, inspect it further,

reduce its load or continue operating. The value of waiting today includes the possibility of a better-informed decision later, which a one-step comparison over a fixed horizon does not credit.

The sequential problem is stated on decision epochs $k = 0, 1, \ldots, N$. The hidden physical state collects the degradation of both pumps and the discrete equipment and repair states,

$$\mathbf{x}_k = (D_A(t_k),\ D_B(t_k),\ \text{equipment and repair states}). \tag{24.28}$$

The action is chosen from a finite set, $a_k \in \{\text{operate}, \text{inspect}, \text{reduce load}, \text{repair}, \text{replace}\}$. The measurement is $\mathbf{y}_k = h(\mathbf{x}_k) + \mathbf{v}_k$, with h the measurement function and $\mathbf{v}_k$ the measurement noise. The system dynamics and the observation model are the conditional distributions

$$\mathbf{x}_{k+1} \sim p(\mathbf{x}_{k+1} \mid \mathbf{x}_k, a_k), \qquad \mathbf{y}_{k+1} \sim p(\mathbf{y}_{k+1} \mid \mathbf{x}_{k+1}, a_k). \tag{24.29}$$

The action enters both. It changes the future physical state: a load reduction may decrease the drift of the degradation, and a replacement resets the damage state of a component. It also changes what is observed: an inspection improves the estimate of the condition without directly changing the physical state, which is why the observation model is conditioned on a_k.

24.6.2 The Bellman equation

Because the state is hidden, the decision can depend only on the belief state b_k, the posterior distribution of $\mathbf{x}_k$ given the history, equation (24.1). After action a and observation $\mathbf{y}$, the belief is updated by the Bayes filter of Chapter 23, a prediction through the dynamics followed by a correction by the likelihood,

$$b'(\mathbf{x}') = \mathcal{B}(b, a, \mathbf{y})(\mathbf{x}') = \frac{p(\mathbf{y} \mid \mathbf{x}', a) \int p(\mathbf{x}' \mid \mathbf{x}, a)\, b(\mathbf{x})\, \mathrm{d}\mathbf{x}}{\Pr(\mathbf{y} \mid b, a)}, \tag{24.30}$$

where the denominator, the probability of observing $\mathbf{y}$, is the numerator integrated over $\mathbf{x}'$. Let $c(b, a)$ be the expected immediate cost of action a under belief b, the cost averaged over the belief, and let $V_k(b)$ be the minimum expected future cost from epoch k with belief b. For a finite planning horizon, with a terminal cost V_N, the principle of optimality gives the backward recursion

$$V_k(b) = \min_a \Big[\, c(b, a) + \gamma \sum_{\mathbf{y}} \Pr(\mathbf{y} \mid b, a)\, V_{k+1}(\mathcal{B}(b, a, \mathbf{y}))\Big], \qquad k = N-1, \ldots, 0, \tag{24.31}$$

with the sum replaced by an integral for continuous measurements and $\gamma \in (0, 1]$ a discount factor, equal to one over short horizons and below one when costs occur over long periods. The minimizing action at each b defines the optimal policy δ of equation (24.2). Equation (24.31) is the *Bellman equation*, and the decision problem it solves, a Markov decision process whose state is observed only through noisy measurements, is a *partially observable Markov decision process* (POMDP), also called a partially observed Markov decision process. The belief state is a sufficient statistic of the history for this problem, so a POMDP is a Markov decision process on the space of beliefs.

The decision tree of Figure 24.4 is the two-epoch case of equation (24.31). At the last epoch the only actions are to replace or to wait, and $V_1(b) = \min\{G_{\mathrm{w}},\ p(b)\,\Delta c\}$, where $p(b)$ is the 30-day failure probability under belief b. At the first epoch the inspection adds a third action, whose immediate cost is c_{I} and whose observation updates the belief by equations (24.15) and (24.16):

$$\begin{aligned} V_0(b) &= \min\Big\{G_{\mathrm{w}},\ p\,\Delta c,\ c_{\mathrm{I}} + \Pr(Y^+)\, V_1(b^+) + \Pr(Y^-)\, V_1(b^-)\Big\} \\ &= \min\{9000,\ 6720,\ 500 + 0.19 \times 9000 + 0.81 \times 1244\} = \text{R\$}\,3218, \end{aligned}$$

where b^+ and b^- are the beliefs after a positive and a negative outcome. The equation captures a central principle: a maintenance action affects both the immediate consequences and the future opportunities. An inspection creates information that improves later decisions. A replacement incurs an immediate cost but changes the future degradation state. A load reduction may preserve useful life while sacrificing operating efficiency or capacity. Solving the recursion exactly is rarely possible for continuous beliefs; approximate methods, such as discretized beliefs, point-based value iteration and online Monte Carlo tree search, have been used to prioritize maintenance when resources are limited.

Principle 24.3. *The optimal maintenance policy is in general sequential: today's action changes tomorrow's equipment condition, the information available and the maintenance opportunities. The Bellman equation of the partially observable Markov decision process values each action by its immediate cost plus the expected optimal cost from the belief it leads to.*

24.6.3 Expected cost with and without inspection

The value of the inspection of Section 24.4.3 depends on the prognosis, the costs and the quality of the test. The model is the one-period comparison of equations (24.12), (24.17) and (24.18): the outcome probabilities from equation (24.14), the posteriors from Bayes' theorem, the best action after each outcome, and the EVSI as the difference of the expected losses without and with the test; the inspection pays when EVSI $> c_{\mathrm{I}}$, so the EVSI column is also the largest inspection cost worth paying. Table 24.6 evaluates it for ten failure probabilities at the base costs and test. One row of the arithmetic, $p = 0.05$: $\Pr(Y^+) = 0.85 \times 0.05 + 0.10 \times 0.95 = 0.1375$; $\Pr(F_{30} \mid Y^+) = 0.0425/0.1375 = 0.309 > p^*$, so replace, contributing $0.1375 \times 9000 = 1237.5$; $\Pr(F_{30} \mid Y^-) = 0.0075/0.8625 = 0.0087 < p^*$, so wait, contributing $0.0075 \times 56\,000 = 420$; hence $J^*_{\text{with}} = 1657.5$, against $J^*_{\text{without}} = 0.05 \times 56\,000 = 2800$, and EVSI = R$ 1142.5.

Table 24.6: Expected losses of the pump decision with and without the inspection, for several 30-day failure probabilities p, with $G_{\mathrm{w}} =$ R$ 9000, $\Delta c =$ R$ 56 000 ($p^* = 0.1607$), sensitivity 0.85 and specificity 0.90 (illustrative values; losses in reais, before the inspection cost). R: replace now; W: wait 30 days.

p	$\Pr(Y^+)$	$\Pr(F_{30} \mid Y^+)$	$\Pr(F_{30} \mid Y^-)$	Best without	J^*_{without}	Actions after Y^+/Y^-	J^*_{with}	EVSI
0.02	0.1150	0.1478	0.0034	W	1120	W/W	1120	0
0.05	0.1375	0.3091	0.0087	W	2800	R/W	1657.5	1142.5
0.08	0.1600	0.4250	0.0143	W	4480	R/W	2112	2368
0.10	0.1750	0.4857	0.0182	W	5600	R/W	2415	3185
0.12	0.1900	0.5368	0.0222	W	6720	R/W	2718	4002
0.16	0.2200	0.6182	0.0308	W	8960	R/W	3324	5636
0.20	0.2500	0.6800	0.0400	R	9000	R/W	3930	5070
0.30	0.3250	0.7846	0.0667	R	9000	R/W	5445	3555
0.50	0.4750	0.8947	0.1429	R	9000	R/W	8475	525
0.60	0.5500	0.9273	0.2000	R	9000	R/R	9000	0

Figure 24.7 draws the EVSI column as a function of p. The horizontal axis is the 30-day failure probability and the vertical axis the EVSI in reais; the dashed horizontal lines are three inspection costs, R$ 500, R$ 2000 and R$ 4000, and the inspection pays wherever the curve lies above the line of its cost. The curve is piecewise linear, because each expected loss is linear in p between the points where a best action changes. It is zero below $p = 0.0220$, where even a positive outcome leaves the posterior below p^*, so the operator waits whatever the result; it

rises to its maximum of R$ 5665 at $p = p^*$, where the operator is indifferent without the test and any information is most useful; and it falls back to zero at $p = 0.5347$, beyond which even a negative outcome leaves the posterior above p^* and the operator replaces whatever the result.

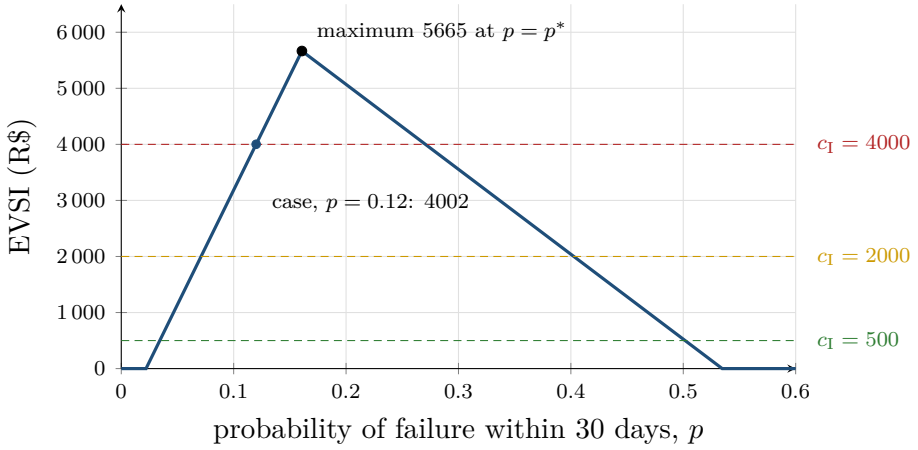


Figure 24.7: Expected value of sample information of the inspection against the 30-day failure probability, from Table 24.6's model: sensitivity 0.85, specificity 0.90, G_w = R$ 9000, Δc = R$ 56 000 (illustrative values). The inspection pays where the curve exceeds its cost.

Three trends follow. First, raising the inspection cost narrows the range of prognoses in which the inspection pays, from 0.034–0.502 at R$ 500 through 0.071–0.403 at R$ 2000 to 0.120–0.271 at R$ 4000; at the case's $p = 0.12$ a R$ 4000 inspection is barely worth its cost, and an inspection costing more than R$ 5665 never pays, whatever the prognosis: the information may still be valuable but obtaining it becomes uneconomic. Second, the inspection is most valuable near the decision threshold, and worthless far from it on either side. Third, the value depends on the quality of the test and on the costs. Table 24.7 varies them one at a time at $p = 0.12$. The rows reduce the sensitivity and the specificity toward 0.5, raise them toward one, and change the failure penalty and the gain from waiting; the columns are the resulting threshold, outcome probabilities, posteriors, expected losses and EVSI.

Table 24.7: The inspection decision at $p = 0.12$ for other test qualities and costs, changed one at a time from the base case (sensitivity 0.85, specificity 0.90, G_w = R$ 9000, Δc = R$ 56 000); illustrative values, losses in reais before the inspection cost.

Case	p^*	$\Pr(Y^+)$	$\Pr(F_{30} \mid Y^+)$	$\Pr(F_{30} \mid Y^-)$	J^*_{without}	J^*_{with}	EVSI
Sensitivity, specificity 0.50, 0.50	0.1607	0.500	0.1200	0.1200	6720	6720	0
0.70, 0.70	0.1607	0.348	0.2414	0.0552	6720	5148	1572
0.85, 0.90 (base)	0.1607	0.190	0.5368	0.0222	6720	2718	4002
0.95, 0.95	0.1607	0.158	0.7215	0.0071	6720	1758	4962
1, 1 (perfect)	0.1607	0.120	1	0	6720	1080	5640
Δc = R$ 30 000	0.3000	0.190	0.5368	0.0222	3600	2250	1350
Δc = R$ 120 000	0.0750	0.190	0.5368	0.0222	9000	3870	5130
G_w = R$ 4500	0.0804	0.190	0.5368	0.0222	4500	1863	2637
G_w = R$ 18 000	0.3214	0.190	0.5368	0.0222	6720	4428	2292

A test with sensitivity and specificity of 0.5 is uninformative: its outcome is independent of the failure, the posteriors equal the prior, and its value is zero. As the test improves, the

EVSI rises toward the EVPI of R$ 5640, reached by the perfect test. The costs act through the threshold: with a failure penalty of R$ 30 000, $p^* = 0.30$ lies far above the prognosis and the EVSI falls to R$ 1350; with a penalty of R$ 120 000, $p^* = 0.075$ lies below it, the preferred action without the test becomes immediate replacement, and the inspection, worth R$ 5130, would mainly prevent unnecessary replacements. The value of the inspection is a property of the decision it informs, not of the test alone.

24.6.4 Uncertainty in the predicted failure probability

The case used $F_t(30) = 0.12$ as though it were known exactly. In practice it is a posterior predictive quantity, derived from uncertain damage states, degradation parameters, operating conditions and failure thresholds. Collecting these in a vector $\mathbf{z}$, the probability used in the decision is

$$\Pr\big(\mathrm{RUL}(t) \le 30 \mid \mathcal{Y}_{0:t}\big) = \int \Pr\big(\mathrm{RUL}(t) \le 30 \mid \mathbf{z}\big)\, p(\mathbf{z} \mid \mathcal{Y}_{0:t})\, \mathrm{d}\mathbf{z}, \tag{24.32}$$

which must be computed from one coherent joint distribution of $\mathbf{z}$. Averaging predictions from incompatible operating scenarios, without stating the probability of each, produces a number that describes no scenario.

For the one-step decision the predictive probability is all that matters for the choice of action. The loss of waiting, $\Delta c\, \mathbf{1}\{\mathrm{RUL}(t) \le 30\}$, is linear in the failure indicator, so its expectation over both the future and the uncertain $\mathbf{z}$ is Δc times the probability of equation (24.32). The spread of the uncertain probability matters in two other places. It sets the value of learning more about $\mathbf{z}$: if the probability were 0.05 or 0.19 with equal chances, with the same mean 0.12, waiting would still be the best immediate action with expected loss R$ 6720, but knowing which value holds would allow waiting in the first case (R$ 2800) and replacing in the second (R$ 9000), an expected loss of R$ 5900, so the information is worth R$ 820. It also governs any chance constraint of the form "the failure probability must be below ε with confidence 0.95", which depends on the distribution of the probability and not only on its mean, as the availability chance constraint of Chapter 20 of Volume I did.

24.6.5 Model credibility

A numerical maintenance recommendation is only as credible as the assumptions that support it. Table 24.8 lists the questions to settle before trusting the resulting action, each with the reason it matters. They follow the three separate inference steps of the prognosis, what is known about the current damage, how the damage is expected to evolve, and how likely the failure threshold is to be crossed, and add the two steps that the decision introduces: the effectiveness of the intervention and the consequences at the system level.

When these questions cannot be answered, the defensible outcome may be to collect more evidence rather than to issue a definite replacement date. In high-consequence systems this is often the correct recommendation, and the value-of-information calculation of Section 24.4.3 is the tool for deciding which evidence to collect.

24.6.6 Competing objectives and constraints

A facility cares about several quantities at once: the maintenance cost C_{maint}, the downtime cost C_{down}, the system availability A_{system} and the probability of a critical failure. A general formulation minimizes the expected total cost over the policies δ and a horizon of N epochs, subject to a risk constraint and to the operating constraints of the plant,

$$\min_{\delta}\ \mathbb{E}\Big[\sum_{k=0}^{N} c(\mathbf{x}_k, a_k)\Big] \quad \text{subject to} \quad \Pr(\text{critical system failure during the horizon}) \le \varepsilon, \tag{24.33}$$

Table 24.8: Questions that establish the credibility of a maintenance recommendation.

Question	Why it matters
Is the degradation mechanism correct?	The wrong physical model can predict a misleading RUL.
Is the damage identifiable from the measurements?	Sensor signals may support several explanations of the condition.
Is the RUL forecast calibrated?	A reported 12 % failure probability must correspond to an observed frequency of failures among comparable cases.
Are the future operating conditions represented?	Changes of load and temperature may accelerate the damage.
Is the intervention effective?	Replacement, repair and load reduction change the future failure behavior differently.
Are the dependencies included?	Common-cause events can defeat the redundant protection.
Are the consequences credible?	A component fault may have little or catastrophic system-level impact, depending on the operating state.

and, at every epoch, $C_{\text{available}}(t_k) \geq L(t_k)$ (available cooling capacity at least the required load), the number of concurrent maintenance tasks at most the number of available crews, and the consumption of spares at most the available inventory. Problem (24.33) joins the prognostics of this part of the book with the availability and maintenance models of Chapter 20: the repair crews, spare parts and scheduling constraints of that chapter become constraints of the sequential problem, and its cost-of-downtime objective becomes one term of $c(\mathbf{x}_k, a_k)$. Stochastic repair times, capacity constraints and shared maintenance resources make the optimal schedule substantially harder to find than a component replacement rule, which is why practical methods combine the thresholds and value-of-information calculations of this chapter with approximate solutions of the sequential problem. The next chapter, Chapter 25, moves from the degradation and maintenance of individual equipment to the adequacy of a power system's generation, where the same probabilistic reasoning is applied to the capacity of many units against demand.

24.7 What to retain

Definition 24.1 (Predictive maintenance decision)**.** A predictive maintenance decision selects the action, among continuing, inspecting, reducing load, repairing and replacing, that minimizes the expected loss given the current belief about the condition of the system, combining the predicted failure probability with the consequences, the intervention costs, the lead times and the value of the information the action produces; a prognostic forecast is an input to that decision, not the decision itself.

A probability-based trigger and an RUL-quantile deadline turn a prognosis into a planning rule, but only a decision model says which action to take. For a choice between acting now and deferring, the optimal rule is a probability threshold equal to the ratio of the gain from waiting to the incremental penalty of failure. An inspection has value only when its outcomes lead to different actions; its value is the expected value of sample information, bounded by the value of perfect information, and it is greatest near the decision threshold. Inspection intervals follow from a renewal-reward cost rate or from a quantile of the current RUL distribution, and adaptive schedules hold the between-inspection risk where fixed ones cannot. A component prognosis is not a system risk: redundancy, restoration, common causes, maintenance windows

and spare-part lead times connect the two. The optimal policy is in general sequential and is stated by the Bellman equation of a partially observable Markov decision process. Table 24.9 lists the chapter's terms, their meaning, expression and units, and the subsection that defines each.

Table 24.9: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Belief state	posterior over condition and parameters	$b_t = p(\mathbf{x}_t, \boldsymbol{\theta} \mid \mathcal{Y}_{0:t})$	–	24.1.1
Maintenance policy	maps information to an action	$a_t = \delta(b_t)$	–	24.1.1
Condition-based, predictive maintenance	act on present, on predicted condition	Table 24.1	–	24.1.2
Probability-based trigger	plan when failure probability is high	$F_t(\tau) \geq p_{\text{trig}}$	–	24.2.1
Maintenance deadline	latest time meeting a risk tolerance	$\tau_{\max} = \text{RUL}_{\varepsilon}(t)$	days	24.2.2
Action lead time	decision to completed action	τ_{L}; risk $F_t(\tau_{\text{L}})$	days	24.2.3
Incremental failure penalty	extra cost of a failure	$\Delta c = c_f - c_p$	R$	24.3.2
Gain from waiting	value of deferring the action	G_{w}	R$	24.3.2
Decision threshold	failure probability of indifference	$p^* = G_{\text{w}}/\Delta c$	–	24.3.4
Sensitivity, specificity	detection and correct-rejection probabilities	$\Pr(Y^+ \mid F_{30})$, $\Pr(Y^- \mid F_{30}^{\text{c}})$	–	24.4.1
Posterior failure probability	updated by the inspection outcome	equations (24.15), (24.16)	–	24.4.2
Expected value of sample information	loss reduction from a test	$J^*_{\text{without}} - J^*_{\text{with}}$	R$	24.4.3
Expected value of perfect information	upper bound on EVSI	equation (24.19)	R$	24.4.3
Periodic inspection cost rate	renewal-reward rate of inspection	$C(\Delta)$, equation (24.20)	R$/day	24.4.4
Risk-constrained interval	next inspection within tolerance	$\Delta_{\max} = \text{RUL}_{0.05}(t)$	days	24.4.5
Adaptive inspection	interval set by current condition	Table 24.4	days	24.4.6
Maintenance-window risk	loss of redundancy while servicing	$(1 - \text{e}^{-\lambda w})\, C_{\text{sys}}$	R$	24.5.3
Emergency procurement value	benefit of a shorter lead time	$[F_t(30) - F_t(7)]\, \Delta c$	R$	24.5.4
Inventory cost	holding, expediting and downtime	$J(s)$, equation (24.27)	R$/yr	24.5.5
Belief update	Bayes filter after action and observation	$\mathcal{B}(b, a, \mathbf{y})$	–	24.6.2
Bellman equation, POMDP	sequential optimal policy	equation (24.31)	R$	24.6.2
Posterior predictive failure probability	averaged over uncertain quantities	equation (24.32)	–	24.6.4

Further reading. NIST, *Condition-Based Maintenance Policy Optimization*, for policy optimization under stochastic repair times and shared resources; Fauriat and Zio, *Inspection and Condition-Based Maintenance Driven by Value of Information*, for value-of-information inspection scheduling; NIST, *Online Improvement of Condition-Based Maintenance Policy*, for sequential maintenance decisions by approximate online optimization; Jardine and Tsang, *Maintenance, Replacement, and Reliability*, for maintenance policies, replacement decisions and life-cycle economics.

24.8 Exercises

Exercise 24.1. (a) State what a probability-based maintenance trigger is. (b) In the replace-now against wait decision, state what determines the probability threshold. (c) State why an inspection can have positive economic value. (d) State what the expected value of sample information measures. (e) State why the RUL distribution of a component is insufficient to establish the risk of a system outage. (f) State an advantage of adaptive inspection. (g) State what information a sequential PHM decision model preserves. (Answers: (a) a rule based on the probability of failure within a specified horizon, $F_t(\tau) \geq p_{\text{trig}}$; (b) the gain from deferral divided by the incremental failure penalty, $p^* = G_{\text{w}}/\Delta c$; (c) its outcomes can lead to different, better-informed actions; (d) the expected reduction in the loss of the optimal decision before the measurement is paid for; (e) redundancy, common causes, restoration and operating constraints also determine whether the function is lost; (f) the inspection timing responds to the current risk, its uncertainty and the expected decision benefit; (g) the current beliefs, the actions, the observations and their influence on future decisions.)

Exercise 24.2. In the pump case, a revised estimate of the consequences raises the total cost of a failure before planned replacement to $c_f = \text{R\$}\,100\,000$; $c_p = \text{R\$}\,4000$, $G_{\text{w}} = \text{R\$}\,9000$ and $F_t(30) = 0.12$ are unchanged. (a) Compute the incremental failure penalty and the decision threshold. (b) Compute the expected losses of replacing now and of waiting, and state the decision. (c) Compute the threshold if the gain from waiting is forfeited when the failure occurs. (d) Compute the gain from waiting at which the operator would be indifferent at $p = 0.12$. (Answers: (a) $\Delta c = \text{R\$}\,96\,000$, $p^* = 9000/96\,000 = 0.0938$; (b) $J_{\text{rep}} = \text{R\$}\,9000$, $J_{\text{wait}} = 0.12 \times 96\,000 = \text{R\$}\,11\,520$: replace now; (c) $9000/105\,000 = 0.0857$, and the decision is unchanged; (d) $G_{\text{w}} = 0.12 \times 96\,000 = \text{R\$}\,11\,520$.)

Exercise 24.3. A later prognosis gives $F_t(30) = 0.20$, with $G_{\text{w}} = \text{R\$}\,9000$ and $\Delta c = \text{R\$}\,56\,000$. A different inspection has sensitivity 0.90, specificity 0.80 and costs R\$ 1000. (a) Compute the probabilities of a positive and a negative outcome. (b) Compute the posterior failure probabilities and the best action after each outcome. (c) Compute the expected losses without and with the inspection, the EVSI and the expected loss of inspecting, and state whether to inspect. (d) Compute the expected value of perfect information. (Answers: (a) $\Pr(Y^+) = 0.18 + 0.16 = 0.34$, $\Pr(Y^-) = 0.66$; (b) $\Pr(F_{30} \mid Y^+) = 0.18/0.34 = 0.5294$, replace; $\Pr(F_{30} \mid Y^-) = 0.02/0.66 = 0.0303$, wait; (c) $J^*_{\text{without}} = \min\{9000, 11\,200\} = \text{R\$}\,9000$ (replace now); $J^*_{\text{with}} = 0.34 \times 9000 + 0.02 \times 56\,000 = 3060 + 1120 = \text{R\$}\,4180$; EVSI = R\$ 4820; $J_{\text{insp}} = \text{R\$}\,5180 < \text{R\$}\,9000$, so inspect; (d) EVPI $= 9000 - 0.20 \times 9000 = \text{R\$}\,7200$.)

Exercise 24.4. (a) The replacement of Pump A is prepared poorly and the window lasts $w = 6\,\text{h}$; with Pump B operating at the start, $\lambda = 0.001\,\text{h}^{-1}$ and $C_{\text{sys}} = \text{R\$}\,150\,000$, compute the probability of losing pumping during the window and its expected consequence. (b) For the bearing of Section 24.2.1 at current damage $d_0 = 0.55$, compute $F_t(14)$ and $F_t(30)$ from equation (24.5), and the largest additional cost worth paying to shorten the spare-part lead time from 30 to 14 days, with $\Delta c = \text{R\$}\,56\,000$. (c) For periodic inspection with restoration, with a Weibull time to failure of shape 2 and scale 600 days, $c_{\text{I}} = \text{R\$}\,500$ and

$\Delta c = \text{R\$}\,56\,000$, compute the approximate optimal interval and cost rate of equation (24.22), and the interval that holds the probability of failure before the next inspection at $5\,\%$. (Answers: (a) $1 - \mathrm{e}^{-0.006} = 0.005982$, expected consequence R\$ 897; (b) $F_t(14) = 0.0045$ and $F_t(30) = 0.2358$, so up to $(0.2358 - 0.0045) \times 56\,000 = \text{R\$}\,12\,955$; (c) $\Delta^* \approx 600\sqrt{500/56\,000} = 56.7$ days, $C(\Delta^*) \approx 2\sqrt{500 \times 56\,000}/600 = \text{R\$}\,17.64$ per day (exact 56.7 days and R\$ 17.65 per day), and $\Delta_{\max} = 600\sqrt{-\ln 0.95} = 135.9$ days.)

Part VII

Power-System Reliability and Adequacy

Chapter 25

Generation Adequacy and Reliability Indices

25.1 Generation adequacy and the adequacy event

An electric power system is built to supply its customers at every hour of the year, and the first reliability question it raises concerns the generating fleet alone: whether the installed generating units are enough to meet the demand, given that each of them is from time to time unavailable because of a forced outage. A fleet whose installed capacity exceeds the annual peak demand by a quarter can still fail to meet that demand in a sizable fraction of the hours, because the units are not all available at once. The question of this chapter is

whether the installed generating resources suffice to meet the demand at every hour of a study period when each unit is subject to random forced outages, and how the probability, the expected duration and the expected energy of the resulting shortfalls are computed.

The chapter defines the adequacy event and the three principal indices that measure it, models each generator as a two-state random capacity, builds the probability distribution of the capacity on outage for a small fleet by complete enumeration, and carries the indices from one hour to a year. It then adds uncertain demand and dependent outages, derives the recursive method by which the distribution is built for large fleets, and closes with the limits of the method: what adequacy does not cover, which definitions must accompany every reported index, and what a static capacity model leaves out.

Jargon: generation adequacy, forced outage, forced outage rate (FOR), loss of load, LOLP, LOLE, LOLH, EUE (EENS), COPT, reserve margin

Generation adequacy (also *resource adequacy* when storage, demand response and imports count as resources) is the ability of the installed resources to meet the demand over a study period, with their planned and unplanned outages accounted for. A *forced outage* is an unplanned removal of a unit from service because of a failure; its probability in a randomly chosen hour is the unit's *forced outage rate* (FOR), which despite its name is a dimensionless probability, the unit's unavailability. *Loss of load* is the event that the available capacity falls short of the demand. The *loss-of-load probability* (LOLP) is the probability of that event in one interval; the *loss-of-load expectation* (LOLE) is the expected number of intervals with loss of load in a study period, called *loss-of-load hours* (LOLH) when the intervals are hours; the *expected unserved energy* (EUE), also *expected energy not supplied* (EENS), is the expected energy of demand not met, in MWh. A *capacity outage probability table* (COPT) lists the probability of each total amount of generating capacity on outage. The *reserve margin* is the excess of installed capacity over peak demand as a fraction of the peak demand.

Key idea

Because units fail at random, the available capacity of a fleet is a random variable, and adequacy is a property of its probability distribution measured against the demand: the probability of a shortfall (LOLP), its expected duration (LOLE) and its expected energy (EUE). Installed capacity and reserve margin do not determine any of the three.

This section states the adequacy question, defines the adequacy event and the capacity shortfall, and defines the three indices.

25.1.1 Installed resources against demand under random outages

Chapter 3 of Volume I separated two aspects of power-system reliability (Section 3.6.1): *adequacy*, the existence of enough generation and transmission capacity to satisfy the demand with outages accounted for, and *security*, the ability to withstand sudden disturbances. Adequacy is a static question about capacity against demand over a period; security is a dynamic question about the response to an event. That chapter also computed the simplest adequacy index: for two independent 100 MW generators, each unavailable with probability 0.02, serving a demand of 180 MW, the probability of insufficient generation at a random instant is $1 - 0.98^2 = 0.0396$, and if the same demand held for all 8760 h of a year the loss-of-load expectation would be 346.9 h per year and the expected unserved energy 28 102 MWh per year. This chapter develops those two indices in full.

The question is posed hour by hour. In each hour t of a study period the system must supply a demand L_t; the generating units that are available in that hour supply a total capacity C_t, which is the installed capacity minus the capacity of the units on forced outage, planned maintenance or other unavailability. Generation adequacy asks how often, by how much and for how long C_t falls short of L_t. In its classical form the question treats the transmission network as a single bus that can carry any generation to any load (a *copper-plate* model); in the terminology of power-system reliability this is the first hierarchical level, generation only. Adding the transmission network gives a composite generation-and-transmission assessment, which a later chapter on composite-system adequacy develops.

The distinction between the physical system, the model and the evidence applies directly. The physical system is a fleet of machines that trip, are repaired and return, and a demand that follows the weather and the economy. The model represents each unit by two states with a fixed probability of being out and the demand by a number or a distribution in each hour.

The evidence is the record of each unit's forced outages and of past demand, from which the outage probabilities and the demand distribution are estimated. Every index in this chapter is a property of the model, conditional on that evidence.

25.1.2 The adequacy event and the capacity shortfall

Let L_t be the electrical demand during interval t, in MW, and C_t the available generation and other qualifying resources in the same interval, in MW; C_t is the quantity written $C_{\text{available}}(t)$ in Chapter 3. The basic adequacy event, *loss of load*, is

$$C_t < L_t. \tag{25.1}$$

The *capacity shortfall* is the amount by which demand exceeds the available capacity, and zero otherwise,

$$S_t = (L_t - C_t)^+ = \max(0,\ L_t - C_t), \tag{25.2}$$

in MW. The positive part $(\cdot)^+$ keeps only deficits: a surplus of capacity contributes nothing. The event (25.1) and the condition $S_t > 0$ are the same event. The symbol S_t denotes the shortfall throughout this chapter; it is not the FMEA severity rating S of Chapter 4 of Volume I. As an illustration, if $C_t = 700\,\text{MW}$ and $L_t = 750\,\text{MW}$, then $S_t = (750 - 700)^+ = 50\,\text{MW}$: the system lacks 50 MW of the resources needed to supply its modeled demand.

The inequality in equation (25.1) is strict. If exactly 80 MW is available and the demand is 80 MW, the simplified model records no shortfall; whether an operator would in practice carry some operating reserve above the demand, and count a state without that reserve as deficient, is a separate modeling choice, and a study that makes it must state it.

A modeled shortfall need not mean that customers experienced that much load shedding. Emergency imports from neighboring systems, reserve-sharing arrangements, demand response, voltage reduction or other actions not included in the model may prevent the outcome. The shortfall is therefore a property of the model relative to the resources and the emergency actions it includes, and a study must state which of them it represents before its indices can be interpreted.

25.1.3 Three principal adequacy indices

The *loss-of-load probability* of interval t is the probability of the adequacy event in that interval,

$$\text{LOLP}_t = \Pr(C_t < L_t) = \Pr(S_t > 0). \tag{25.3}$$

It is a dimensionless probability. A value $\text{LOLP}_t = 0.02$ states that the model assigns a 2 % probability to insufficient resources in interval t. It is estimated, like every quantity of this chapter, by evaluating the model: from the probability distribution of C_t, which is built from the outage probabilities of the units, and the demand of the interval.

The *loss-of-load expectation* over a study period of N intervals of length Δt is the sum

$$\text{LOLE} = \sum_{t=1}^{N} \text{LOLP}_t\, \Delta t, \tag{25.4}$$

which is the time-discretized form stated in Chapter 3 (Section 3.6.3). Its meaning depends on the interval. For hourly intervals, $\Delta t = 1\,\text{h}$, it is the expected number of hours with a shortfall, commonly reported as *loss-of-load hours* (LOLH), in hours per study period (h/yr for an annual study). When each interval is a day and LOLP_t is the probability of a shortfall at the day's peak demand, it is the expected number of days with a shortfall, in days per year.

The *expected unserved energy* weights each shortfall by its size and duration,

$$\mathrm{EUE} = \sum_{t=1}^{N} \mathbb{E}\big[(L_t - C_t)^+\big]\,\Delta t = \sum_{t=1}^{N} \mathbb{E}[S_t]\,\Delta t. \tag{25.5}$$

With L_t and C_t in MW and Δt in hours, the EUE is in MWh over the study period (MWh/yr for an annual study).

The three indices are complementary. LOLP describes the probability of a shortfall in one interval; LOLH describes the expected duration of shortfalls over a period; EUE also captures their magnitude. Two systems with the same expected shortage hours can have very different expected unserved energy, as the illustrative pair of Chapter 3 showed: two one-hour shortfalls of 5 MW per year and one eight-hour shortfall of 500 MW every four years both give LOLE = 2 h per year, with EUEs of 10 and 1000 MWh per year. Adequacy studies therefore report the indices together.

25.2 Forced outages and the capacity outage probability table

The indices of the previous section need the probability distribution of the available capacity in each interval. This section models each unit as a two-state random capacity, enumerates the joint outage states of a three-generator fleet, aggregates them into a capacity outage probability table, and computes the loss-of-load probability and the expected shortfall from it.

25.2.1 The forced outage rate

Consider generator i with rated capacity K_i, in MW. (The rating is written K_i in this chapter so that C_t keeps its meaning of available capacity.) In the simplest model the unit is either fully available or fully unavailable, and its state is described by the *outage indicator*

$$B_i = \begin{cases} 1, & \text{generator } i \text{ is unavailable (on forced outage)},\\ 0, & \text{generator } i \text{ is available.} \end{cases} \tag{25.6}$$

The probability that the unit is unavailable in the modeled assessment state is its forced outage rate,

$$\mathrm{FOR}_i = U_i = \Pr(B_i = 1). \tag{25.7}$$

It is the per-unit unavailability U_i, the complement of the unit's availability, and both symbols are used below for the same number.

If the unit alternates between operation and forced outage with exponential up and down durations, a constant failure rate λ_i while operating and a constant repair rate μ_i, it is the two-state repairable item of Chapter 3 of Volume I, whose steady-state availability is $A_i = \mu_i/(\lambda_i + \mu_i) = \mathrm{MTTF}_i/(\mathrm{MTTF}_i + \mathrm{MTTR}_i)$. The unavailability is the complement,

$$U_i = 1 - A_i = \frac{\lambda_i}{\lambda_i + \mu_i} = \frac{\mathrm{MTTR}_i}{\mathrm{MTTF}_i + \mathrm{MTTR}_i}, \tag{25.8}$$

and the second form holds for any up and down distributions with these means, as the steady-state result of an alternating renewal process. A unit with $\mathrm{MTTF}_i = 900$ h and $\mathrm{MTTR}_i = 100$ h (illustrative) has $U_i = 100/1000 = 0.10$: it spends one hour in ten on forced outage, in the long run. The quantity is a probability, not a rate per unit time; the name "forced outage rate" is historical.

The forced outage rate is estimated from operating records as the forced outage hours divided by the sum of the in-service hours and the forced outage hours over the period observed, which is the long-run fraction of the exposure spent on forced outage. The reported industry

statistics are not all interchangeable with U_i. An observed FOR, an *equivalent forced outage rate on demand* (EFORd), which counts partial deratings as equivalent full outages and restricts the exposure to the periods when the unit is called on to run, and other outage factors have different denominators and treat deratings, demand periods and exposure differently. An adequacy study must use the outage quantity that matches its generator-state model: in the two-state model of this chapter, the probability that the unit is fully unavailable in a randomly selected assessment hour.

25.2.2 Generator availability as a random capacity

A generator with rating $K = 50\,\text{MW}$ and forced outage rate $U = 0.10$ has an available capacity that is a random variable with two values,

$$C = \begin{cases} 50\,\text{MW}, & \text{with probability } 0.90, \\ 0\,\text{MW}, & \text{with probability } 0.10. \end{cases}$$

Its expected available capacity is

$$\mathbb{E}[C] = 50 \times 0.90 + 0 \times 0.10 = 45\,\text{MW}.$$

The random generator cannot be replaced by a firm 45 MW resource when shortage probabilities are computed, because $\Pr(C < L)$ is not determined by $\mathbb{E}[C]$ alone. For a demand of 40 MW, the firm 45 MW resource never falls short, while the random unit falls short with probability 0.10, and then by the whole 40 MW, so that its expected shortfall is $0.10 \times 40 = 4\,\text{MW}$. For a demand of 48 MW, the firm resource falls short with probability 1 by 3 MW, while the random unit falls short with probability 0.10 by 48 MW: the same expected available capacity gives a certain small shortfall in one case and a rare total one in the other. A resource that delivers either 50 MW or nothing has different adequacy consequences from a perfectly reliable 45 MW resource. This is the reason adequacy analysis works with probability distributions of capacity rather than with average availabilities.

Common mistake

Replacing each unit by its expected available capacity $(1 - U_i)K_i$ and comparing the sum with the demand. The shortfall depends nonlinearly on the capacity state, and a fleet whose expected capacity exceeds the demand can still fall short in a large fraction of the hours (Section 25.2.6).

25.2.3 A three-generator fleet

A small system with three dispatchable generators illustrates the calculation. Table 25.1 defines the fleet: each row is one unit with its rated capacity and its forced outage rate; the last row is the installed capacity. The values are illustrative and deliberately pessimistic, so that the indices are large enough to read.

The model rests on six assumptions: (1) each generator is either fully available or unavailable; (2) the outage states of the three units are statistically independent; (3) the forced outage rates apply to a randomly selected assessment hour; (4) there are no imports, no demand response and no energy storage; (5) transmission constraints are ignored; (6) all available generation can be dispatched to meet the demand. The demand is first taken constant,

$$L = 80\,\text{MW}.$$

Figure 25.1 draws the model. The three boxes on the left are the units, each with its rating and forced outage rate; the thick vertical line is the single bus to which every unit delivers its

Table 25.1: The three-generator fleet (illustrative). Each unit is either fully available or on forced outage.

Generator	Rated capacity K_i	Forced outage rate $\mathrm{FOR}_i = U_i$
G1	50 MW	0.10
G2	30 MW	0.15
G3	20 MW	0.20
Total installed capacity	100 MW	

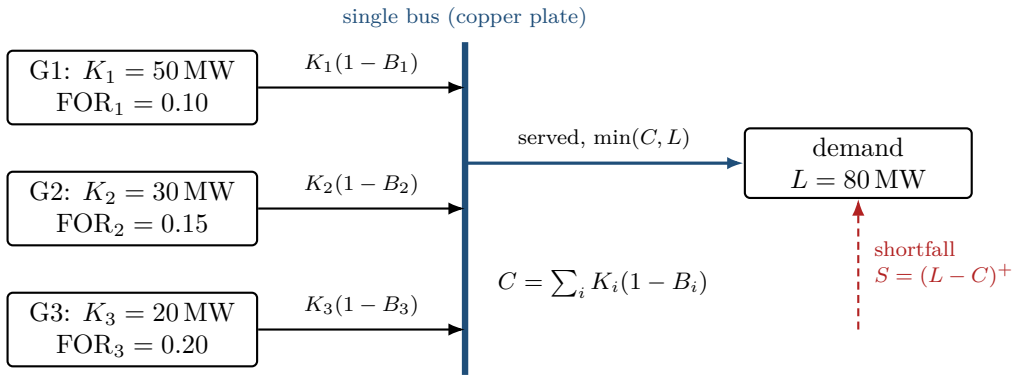


Figure 25.1: The three-generator fleet as a single-bus adequacy model. Each unit delivers K_i when available ($B_i = 0$) and nothing when on forced outage ($B_i = 1$); the shortfall is the part of the demand that the available capacity C does not cover (illustrative values; no imports, storage or network constraints).

available capacity $K_i(1 - B_i)$; the box on the right is the demand. The dashed red arrow is the shortfall, the part of the demand that the available capacity does not cover.

The installed reserve margin, in planning often called the planning reserve margin (PRM), is the excess of installed capacity over the demand as a fraction of the demand,

$$\text{PRM} = \frac{100 - 80}{80} = 0.25 = 25\,\%.$$

A quarter of spare capacity appears substantial. Whether it implies a low probability of insufficient generation is a question that only the distribution of C can answer, and the next three subsections compute it.

25.2.4 Enumeration of the outage states

Each generator has two conditions, so the three generators have $2^3 = 8$ joint states. A state is the vector of outage indicators $\mathbf{b} = (b_1, b_2, b_3)$ with each $b_i \in \{0, 1\}$. The probability of one unit's indicator taking the value b_i can be written in a single expression, $\Pr(B_i = b_i) = U_i^{b_i}(1-U_i)^{1-b_i}$, which is U_i when $b_i = 1$ and $1 - U_i$ when $b_i = 0$. Under assumption (2), the joint probability of the three indicators is the product of the marginal ones,

$$\Pr(\mathbf{B} = \mathbf{b}) = \prod_{i=1}^{3} U_i^{b_i}(1 - U_i)^{1-b_i}. \tag{25.9}$$

This product requires independence; Section 25.4 replaces it when outages share a cause. The total capacity on outage in state $\mathbf{b}$ and the available capacity are

$$X = \sum_{i=1}^{3} K_i B_i, \qquad C = 100\,\text{MW} - X. \tag{25.10}$$

The capacity outage X, in MW, is the random variable that the capacity outage probability table describes; it is unrelated to the interfailure times X_j and the Markov state $X(t)$ of earlier chapters.

Table 25.2 lists the eight states. Each row is one state: the units on outage, the capacity outage X, the available capacity C, and the probability from equation (25.9), with its three factors written out.

Table 25.2: The eight outage states of the three-generator fleet, with the probability of each from equation (25.9) (independent units).

Units on outage	Outage X	Available C	Product	Probability
none	0 MW	100 MW	$0.90 \times 0.85 \times 0.80$	0.612
G3	20 MW	80 MW	$0.90 \times 0.85 \times 0.20$	0.153
G2	30 MW	70 MW	$0.90 \times 0.15 \times 0.80$	0.108
G2, G3	50 MW	50 MW	$0.90 \times 0.15 \times 0.20$	0.027
G1	50 MW	50 MW	$0.10 \times 0.85 \times 0.80$	0.068
G1, G3	70 MW	30 MW	$0.10 \times 0.85 \times 0.20$	0.017
G1, G2	80 MW	20 MW	$0.10 \times 0.15 \times 0.80$	0.012
G1, G2, G3	100 MW	0 MW	$0.10 \times 0.15 \times 0.20$	0.003
Total				1.000

Three rows show the arithmetic. With no unit on outage, $\Pr = (1-0.10)(1-0.15)(1-0.20) = 0.90 \times 0.85 \times 0.80 = 0.612$. With only G2 on outage, $\Pr = 0.90 \times 0.15 \times 0.80 = 0.108$. With

G1 and G3 on outage, Pr $= 0.10 \times 0.85 \times 0.20 = 0.017$. The eight probabilities sum to $0.612 + 0.153 + 0.108 + 0.027 + 0.068 + 0.017 + 0.012 + 0.003 = 1.000$, as they must, since the states are mutually exclusive and exhaust the possibilities. The table is therefore a complete probability model of the available generating capacity. It is exact, not rounded: every product of factors with two decimals gives three decimals here.

Figure 25.2 arranges the same eight states by the number of units on outage, from none at the top to all three at the bottom. Each box gives the units out, the available capacity and the probability; a line joins two states that differ in the status of one unit, so the lines are the single outages and restorations that move the fleet between states. Red boxes are the states whose available capacity is below the demand of 80 MW. Six of the eight states are red, including all three states with two units out and two of the three with one unit out; the only one-unit outage that the fleet survives is the loss of the smallest unit, G3.

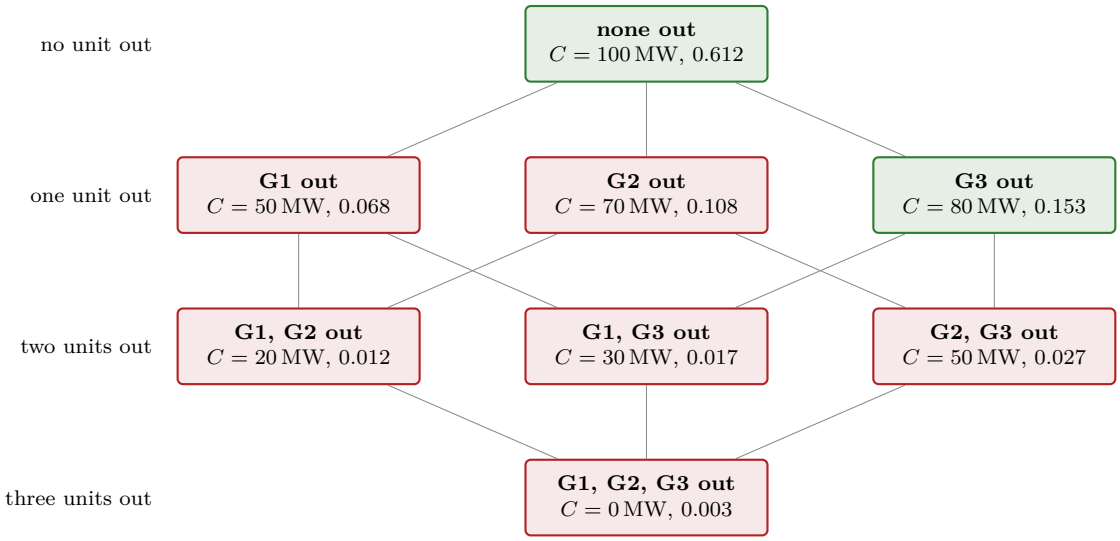


Figure 25.2: The eight outage states of the three-generator fleet, arranged by the number of units on outage. Lines join states that differ in one unit. Green: available capacity at least 80 MW; red: available capacity below the demand of 80 MW (illustrative fleet of Table 25.1).

25.2.5 The capacity outage probability table

Several states can have the same capacity on outage. Two states of Table 25.2 have $X = 50$ MW: G1 alone on outage, with probability 0.068, and G2 and G3 together, with probability 0.027. Since they are mutually exclusive,

$$\Pr(X = 50\,\text{MW}) = 0.068 + 0.027 = 0.095.$$

Aggregating every state by its outage level gives the *capacity outage probability table* (COPT), the probability distribution of X. Table 25.3 lists it. Each row is one outage level x; the columns are the available capacity $100 - x$, the probability $\Pr(X = x)$, and the cumulative probability $\Pr(X \geq x)$, the probability that at least x is on outage. The cumulative column is obtained by summing the probability column from the bottom up: for example $\Pr(X \geq 70) = 0.017 + 0.012 + 0.003 = 0.032$ and $\Pr(X \geq 50) = 0.095 + 0.032 = 0.127$.

Figure 25.3 draws the distribution. The horizontal axis is the capacity on outage in MW; the vertical axis is the probability of each outage level. Green bars are outage levels the fleet survives at a demand of 80 MW; red bars are levels that leave less than 80 MW available. The dashed line between 20 and 30 MW separates the two groups. The installed capacity is 100 MW, but the capacity actually available is uncertain: with probability 0.388 at least 20 MW is out, and

Table 25.3: Capacity outage probability table of the three-generator fleet (independent units). The cumulative column is the probability that the capacity on outage is at least x.

Outage x	Available $100 - x$	$\Pr(X = x)$	$\Pr(X \geq x)$	States
0 MW	100 MW	0.612	1.000	none out
20 MW	80 MW	0.153	0.388	G3
30 MW	70 MW	0.108	0.235	G2
50 MW	50 MW	0.095	0.127	G1; G2 and G3
70 MW	30 MW	0.017	0.032	G1 and G3
80 MW	20 MW	0.012	0.015	G1 and G2
100 MW	0 MW	0.003	0.003	all three

the expected capacity on outage is $\mathbb{E}[X] = 50 \times 0.10 + 30 \times 0.15 + 20 \times 0.20 = 13.5$ MW, so the expected available capacity is 86.5 MW. This distribution is the central input of conventional generation adequacy calculations.

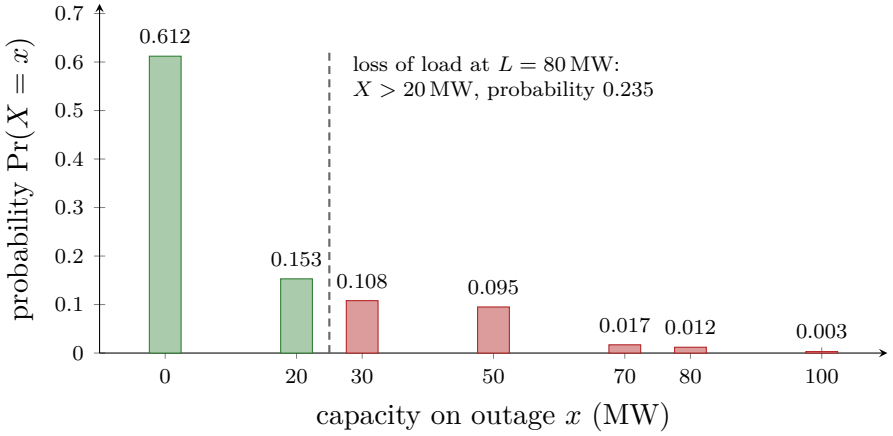


Figure 25.3: Capacity outage distribution of three independent generators of 50, 30 and 20 MW with forced outage rates 0.10, 0.15 and 0.20 (Table 25.3; illustrative). Red bars are the outage levels that leave less than 80 MW available; their total is the loss-of-load probability at that demand.

25.2.6 Loss-of-load probability and expected shortfall at 80 MW

Loss of load occurs when $C < L$. With $C = 100 - X$ and $L = 80$ MW, the condition $100 - X < 80$ is

$$X > 20 \,\text{MW}.$$

The inequality is strict: the state with G3 alone on outage leaves exactly 80 MW and, under the strict definition of Section 25.1.2, no shortfall. The deficient states are the outage levels 30, 50, 70, 80 and 100 MW, and

$$\text{LOLP}(80) = 0.108 + 0.095 + 0.017 + 0.012 + 0.003 = 0.235,$$

which is the cumulative entry $\Pr(X \geq 30)$ of Table 25.3. Despite a nominal reserve margin of 25 %, the model assigns a probability of 23.5 % to insufficient generation in a randomly chosen hour. The value is deliberately large, because the example uses only three relatively unreliable units, but the conclusion is general: installed capacity and reserve margin do

not describe resource adequacy. In this fleet the result has a simple form. A shortfall at 80 MW occurs exactly when G1 or G2 is out, whatever the state of G3, since G1 and G2 together supply 80 MW and any state without one of them has at most 70 MW; therefore $\mathrm{LOLP}(80) = 1 - (1 - U_1)(1 - U_2) = 1 - 0.90 \times 0.85 = 0.235$, and the forced outage rate of G3 does not enter it.

Principle 25.1. *Installed capacity and reserve margin do not determine adequacy. A fleet with* 25 % *installed reserve can fall short of its demand in almost a quarter of the hours when its units are few and their forced outage rates are high; the probability distribution of the available capacity decides.*

The expected shortfall is $\mathbb{E}[S] = \mathbb{E}[(L - C)^+]$. Table 25.4 computes it at $L = 80$ MW. Each row is one available-capacity level of the COPT, with its probability, the shortfall $(80 - C)^+$ in that state, and the product of the two; the expected shortfall is the sum of the last column.

Table 25.4: Expected shortfall at a demand of 80 MW from the COPT of Table 25.3.

Available C	Probability	Shortfall $(80 - C)^+$	Probability × shortfall
100 MW	0.612	0 MW	0 MW
80 MW	0.153	0 MW	0 MW
70 MW	0.108	10 MW	1.08 MW
50 MW	0.095	30 MW	2.85 MW
30 MW	0.017	50 MW	0.85 MW
20 MW	0.012	60 MW	0.72 MW
0 MW	0.003	80 MW	0.24 MW
Total	1.000		5.74 MW

The expected capacity shortfall in the modeled interval is therefore

$$\mathbb{E}[(80 - C)^+] = 1.08 + 2.85 + 0.85 + 0.72 + 0.24 = 5.74\,\mathrm{MW}.$$

If the interval lasts one hour and the capacity state is constant throughout it, equation (25.5) with $N = 1$ and $\Delta t = 1$ h gives an expected unserved energy of $\mathrm{EUE}_{1\,\mathrm{h}} = 5.74\,\mathrm{MW} \times 1\,\mathrm{h} = 5.74\,\mathrm{MWh}$. The number is an expectation over the capacity states. The shortfall realized in a particular hour is one of the listed values, 0, 10, 30, 50, 60 or 80 MW, and is never 5.74 MW.

The expected shortfall mixes how often a deficit occurs with how large it is. The two are separated by the expected shortfall conditional on a deficit, $\mathbb{E}[S \mid S > 0]$. By the law of total expectation, $\mathbb{E}[S] = \Pr(S > 0)\,\mathbb{E}[S \mid S > 0] + \Pr(S = 0)\,\mathbb{E}[S \mid S = 0]$, and the second term vanishes because $S = 0$ whenever there is no deficit. Hence

$$\mathbb{E}[S] = \mathrm{LOLP} \cdot \mathbb{E}[S \mid S > 0], \qquad \mathbb{E}[S \mid S > 0] = \frac{\mathbb{E}[S]}{\mathrm{LOLP}}, \tag{25.11}$$

and at 80 MW,

$$\mathbb{E}[S \mid S > 0] = \frac{5.74}{0.235} = 24.43\,\mathrm{MW}.$$

When the capacity is insufficient, the deficit averages about 24.4 MW, almost a third of the demand. This number is not contained in the LOLP, and it is what a planner needs to size emergency imports or demand response.

25.3 From one hour to annual adequacy, with uncertain demand

The calculation so far concerned one interval with a known demand. This section carries it over the hours of a year, shows why the hourly probabilities can be summed whatever their dependence, works a daily load profile, and then lets the demand itself be uncertain.

25.3.1 Hourly loss-of-load probability

In a real power system the demand changes substantially through the day and the year, and the available capacity changes with it: planned maintenance removes units in some weeks, the weather changes the capability of thermal units and the output of wind and solar plants, and fuel or water limits bind in some seasons. The loss-of-load probability is therefore defined for each hour,

$$\mathrm{LOLP}_t = \Pr(C_t < L_t), \qquad t = 1, \dots, 8760,$$

for a year of 365 days (8760 h), each hour with its own demand and its own capacity distribution.

The hourly values need not be independent across time. Generator forced outages commonly last many hours, so a unit out at hour t is very likely still out at hour $t+1$, and weather conditions that raise the demand or reduce the capacity can persist for days. The capacity states of consecutive hours are strongly correlated, and the shortfall events of consecutive hours with them.

25.3.2 Expected loss-of-load hours

Define the hourly shortage indicator

$$I_t = \mathbf{1}_{\{C_t < L_t\}},$$

equal to 1 if hour t has a shortfall and 0 otherwise. The number of shortage hours in one realized year is the random variable

$$N_{\mathrm{short}} = \sum_{t=1}^{8760} I_t.$$

Its expectation, by linearity of expectation, is

$$\mathrm{LOLH} = \mathbb{E}[N_{\mathrm{short}}] = \mathbb{E}\Big[\sum_{t=1}^{8760} I_t\Big] = \sum_{t=1}^{8760} \mathbb{E}[I_t] = \sum_{t=1}^{8760} \Pr(C_t < L_t), \tag{25.12}$$

since the expectation of an indicator is the probability of its event. Linearity of expectation holds for any random variables, dependent or not, so equation (25.12) needs no assumption of independence between hours. The hourly shortage probabilities may be summed to obtain the expected shortage hours even when outages are strongly correlated in time; this is equation (25.4) with $\Delta t = 1\,\mathrm{h}$, in hours per year.

Dependence does matter for every other property of the shortage hours: how they cluster into events and days, the distribution of N_{short} around its mean, and the probability of a year with no shortage at all. Two models with the same hourly probabilities give the same LOLH and can differ greatly in the number of shortage days, as Section 25.6.1 shows.

> **Principle 25.2.** *Expected shortage hours are the sum of the hourly loss-of-load probabilities, whatever the dependence between hours. The frequency, duration and clustering of shortage events require chronological information that the hourly probabilities do not contain.*

25.3.3 Expected unserved energy over a year

With the hourly shortfall $S_t = (L_t - C_t)^+$ of equation (25.2), the expected unserved energy over a year of one-hour intervals is

$$\mathrm{EUE} = \sum_{t=1}^{8760} \mathbb{E}[S_t] \times 1\,\mathrm{h},$$

in MWh per year, again by linearity and without any independence assumption. In a general time discretization, the uncertainty in each interval is often represented by a finite set of scenarios s, each a joint value of demand and available capacity $(L_{t,s}, C_{t,s})$ with probability $p_{t,s}$, $\sum_s p_{t,s} = 1$. The expectation in equation (25.5) then becomes a probability-weighted sum of shortfall energies,

$$\mathrm{EUE} = \sum_t \sum_s p_{t,s}\,(L_{t,s} - C_{t,s})^+\,\Delta t. \tag{25.13}$$

The scenarios can be the states of a COPT, samples of a Monte Carlo simulation or weather-driven combinations of demand and renewable output. In every form the EUE captures both how likely a deficit is and how large it is.

25.3.4 A daily load profile

Let the three-generator fleet serve a simplified daily profile: four hours at 80 MW and twenty hours at 60 MW. The fleet's outage distribution is the same in every hour, the COPT of Table 25.3. At 80 MW the previous section gave $\mathrm{LOLP}(80) = 0.235$ and $\mathbb{E}[S \mid L = 80] = 5.74$ MW, where the conditioning denotes the shortfall evaluated at that demand.

At 60 MW, loss of load requires $100 - X < 60$, that is $X > 40$ MW, which the COPT reaches first at $X = 50$ MW:

$$\mathrm{LOLP}(60) = \Pr(X \geq 50) = 0.127.$$

The expected shortfall sums $(60 - C)^+$ over the deficient states,

$$\begin{aligned}\mathbb{E}[S \mid L = 60] &= 0.095 \times 10 + 0.017 \times 30 + 0.012 \times 40 + 0.003 \times 60 \\ &= 0.95 + 0.51 + 0.48 + 0.18 = 2.12\,\mathrm{MW}.\end{aligned}$$

Table 25.5 collects the two load periods. Each row is a period, with its number of hours, its LOLP and its expected shortfall.

Table 25.5: Indices of the two load periods of the illustrative daily profile, from the COPT of Table 25.3.

Load period	Hours per day	LOLP	Expected shortfall
80 MW (peak)	4	0.235	5.74 MW
60 MW (off-peak)	20	0.127	2.12 MW

By equation (25.12), the expected shortage hours per day are

$$\mathrm{LOLH}_{\mathrm{day}} = 4 \times 0.235 + 20 \times 0.127 = 0.94 + 2.54 = 3.48\,\mathrm{h},$$

and the expected unserved energy per day is

$$\mathrm{EUE}_{\mathrm{day}} = 4\,\mathrm{h} \times 5.74\,\mathrm{MW} + 20\,\mathrm{h} \times 2.12\,\mathrm{MW} = 22.96 + 42.40 = 65.36\,\mathrm{MWh}.$$

These are the expected shortage hours and shortfall energy of the stylized 24 h profile. They are high because the fleet is small, its forced outage rates are large and it has little redundancy. If the same profile and the same outage distribution held every day of a 365-day year, the annual expectations would be

$$\begin{aligned}\mathrm{LOLH}_{\mathrm{year}} &= 365 \times 3.48 = 1270.2\,\mathrm{h\ per\ year}, \\ \mathrm{EUE}_{\mathrm{year}} &= 365 \times 65.36 = 23\,856.4\,\mathrm{MWh\ per\ year}.\end{aligned}$$

The model expects a shortfall in 14.5 % of the hours of the year. These numbers illustrate the aggregation and are not a realistic adequacy benchmark: real systems have larger fleets, imports, scheduled maintenance, variable generation and far more varied demand.

Figure 25.4 draws the profile against the capacity states. The horizontal axis is the hour of the day; the vertical axis is power in MW. The thick black line is the demand, 60 MW except in the four peak hours, placed for illustration from 17:00 to 21:00. The dashed horizontal lines are the seven available-capacity levels of the COPT, each labeled at the right with its probability. Green levels (100 and 80 MW) never fall below the demand; the gold level (70 MW) falls below it only in the peak hours; the red levels (50 MW and below) fall below it in every hour. The LOLP of an hour is the total probability of the levels that lie below the demand line in that hour: $0.108 + 0.127 = 0.235$ at the peak and 0.127 off the peak.

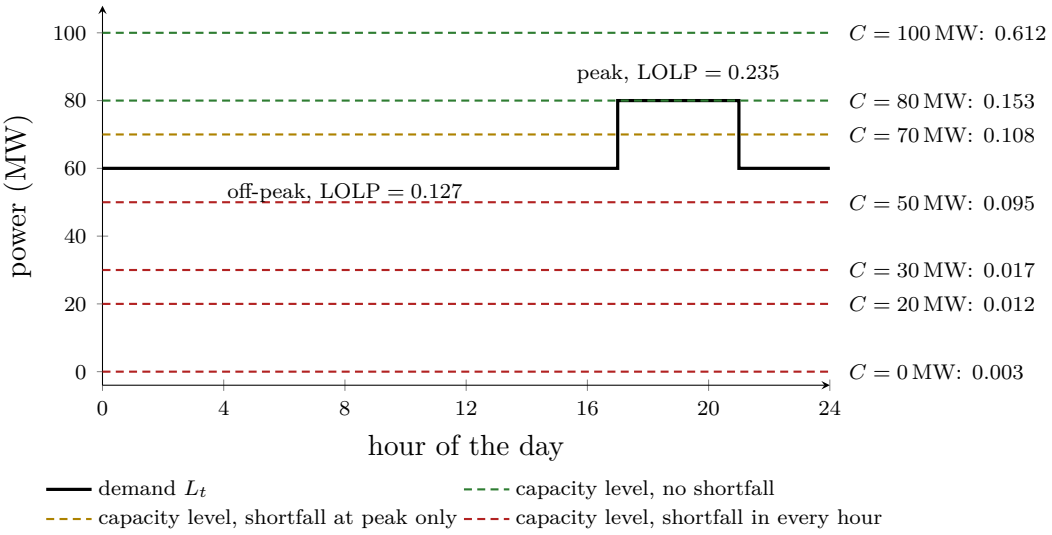


Figure 25.4: Illustrative daily profile, four hours at 80 MW and twenty at 60 MW, against the available-capacity levels of the three-generator COPT (Table 25.3) with their probabilities. Expected shortage hours: 3.48 h per day; expected unserved energy: 65.36 MWh per day.

25.3.5 Demand uncertainty

So far the demand of each interval was a known number. Demand a day ahead, and still more a year ahead, is uncertain. A common representation is

$$L_t = \hat{L}_t + \varepsilon_t,$$

where $\hat{L}_t$ is the forecast demand and ε_t the forecast error, both in MW. The distribution of L_t depends on temperature, time of day, economic activity, electrification and other variables. Adequacy then becomes a joint probability problem: the loss-of-load probability is still $\mathrm{LOLP}_t = \Pr(C_t < L_t)$, but the probability now runs over the joint distribution of C_t and L_t.

Let $p_{L_t}(\ell)$ be the probability density of the demand. By the law of total probability, conditioning on the demand,

$$\mathrm{LOLP}_t = \int \Pr(C_t < \ell \mid L_t = \ell)\, p_{L_t}(\ell)\, \mathrm{d}\ell.$$

If the available capacity and the demand are independent, the conditional probability equals the unconditional one, and

$$\mathrm{LOLP}_t = \int \Pr(C_t < \ell)\, p_{L_t}(\ell)\, \mathrm{d}\ell, \tag{25.14}$$

the LOLP of each demand level averaged over the demand distribution. Without independence the conditional capacity distribution must be used,

$$\mathrm{LOLP}_t = \int \Pr(C_t < \ell \mid L_t = \ell)\, p_{L_t}(\ell)\, \mathrm{d}\ell. \tag{25.15}$$

The distinction matters most under extreme weather, which can simultaneously raise the demand (heating or air conditioning) and reduce the available supply (forced outages of units exposed to cold or heat, derated thermal plants, low wind). Equation (25.15) then weights the high-demand levels with a capacity distribution shifted toward larger outages, and equation (25.14) underestimates the risk. The expected shortfall follows the same pattern, with $\Pr(C_t < \ell)$ replaced by $\mathbb{E}[(\ell - C_t)^+]$.

25.3.6 Three demand levels

Let the demand of a selected hour take three values, Table 25.6, independent of the generator outage states of the COPT. Each row is one demand level with its probability, its LOLP and its expected shortfall from the COPT, and the contributions of that level to the averages. The values at 60 and 80 MW were computed above; at 90 MW, loss of load requires $X > 10$ MW, so $\mathrm{LOLP}(90) = \Pr(X \geq 20) = 0.388$, and $\mathbb{E}[S \mid L = 90] = 0.153 \times 10 + 0.108 \times 20 + 0.095 \times 40 + 0.017 \times 60 + 0.012 \times 70 + 0.003 \times 90 = 1.53 + 2.16 + 3.80 + 1.02 + 0.84 + 0.27 = 9.62$ MW.

Table 25.6: Indices at three demand levels (illustrative demand distribution, independent of the outages) and their probability-weighted contributions.

Demand	Probability	LOLP	Expected shortfall	Contribution to LOLP	Contribution to $\mathbb{E}[S]$
60 MW	0.30	0.127	2.12 MW	0.0381	0.636 MW
80 MW	0.50	0.235	5.74 MW	0.1175	2.870 MW
90 MW	0.20	0.388	9.62 MW	0.0776	1.924 MW
Total	1.00			0.2332	5.430 MW

By the discrete form of equation (25.14),

$$\mathrm{LOLP} = 0.30 \times 0.127 + 0.50 \times 0.235 + 0.20 \times 0.388 = 0.2332,$$

or 23.32 %, and by the same averaging of the expected shortfall,

$$\mathbb{E}[S] = 0.30 \times 2.12 + 0.50 \times 5.74 + 0.20 \times 9.62 = 0.636 + 2.870 + 1.924 = 5.43 \text{ MW},$$

which for a one-hour interval is an expected unserved energy of 5.43 MWh.

The expected demand is $\mathbb{E}[L] = 0.30 \times 60 + 0.50 \times 80 + 0.20 \times 90 = 76$ MW. Evaluating the shortfall at this single demand is not the same as averaging it over the uncertain demand. At 76 MW the COPT gives $\mathbb{E}[(76-C)^+] = 0.108\times6+0.095\times26+0.017\times46+0.012\times56+0.003\times76 = 4.80$ MW, below the 5.43 MW of the full calculation. Replacing the capacity by its expectation as well gives $(\mathbb{E}[L] - \mathbb{E}[C])^+ = (76 - 86.5)^+ = 0$: a model built on expected demand and expected capacity finds no shortfall at all. The general statement is an inequality. The function $z \mapsto z^+$ is convex, and by Jensen's inequality the expectation of a convex function of a random variable is at least the function of its expectation; applied to $z = L - C$,

$$\mathbb{E}\big[(L - C)^+\big] \geq \big(\mathbb{E}[L] - \mathbb{E}[C]\big)^+. \tag{25.16}$$

The two sides are equal only in degenerate cases, for instance when $L - C$ never changes sign. A model that substitutes expected demand and expected generation for their distributions therefore underestimates the expected shortfall, and can do so by the whole of it.

The intermediate value 4.80 MW follows from the same argument applied to the demand alone. For a fixed capacity distribution, let $g(\ell) = \mathbb{E}[(\ell - C)^+]$ be the expected shortfall at

demand ℓ. Each term $(\ell - c)^+$ has slope 0 for $\ell < c$ and 1 for $\ell > c$, so the slope of g is the probability of the states with $c < \ell$:

$$\frac{\mathrm{d}g}{\mathrm{d}\ell} = \Pr(C < \ell) = \mathrm{LOLP}(\ell), \qquad g(\ell) = \int_0^{\ell} \mathrm{LOLP}(y)\,\mathrm{d}y, \tag{25.17}$$

at every ℓ that is not a capacity level, with $g(0) = 0$ for a nonnegative capacity. Since $\mathrm{LOLP}(\ell)$ does not decrease with ℓ, g is convex, and Jensen's inequality over the demand distribution gives $\mathbb{E}[g(L)] \geq g(\mathbb{E}[L])$: here $5.43 \geq 4.80$ MW. The integral form checks the earlier numbers: at 80 MW it is $0.003 \times 20 + 0.015 \times 10 + 0.032 \times 20 + 0.127 \times 20 + 0.235 \times 10 = 5.74$ MW.

Figure 25.5 draws both functions of the demand for the three-generator COPT. In the left panel the horizontal axis is the demand ℓ and the vertical axis the LOLP; it is a step function that jumps just above each available-capacity level and is flat in between. In the right panel, drawn for demands from 50 to 100 MW so that the points near 76 MW can be told apart, the vertical axis is the expected shortfall $g(\ell)$; it is piecewise linear and convex, with slope equal to the left panel. The dots on the right mark the three demand levels of Table 25.6; the open square at $\ell = 76$ MW is their probability-weighted average, 5.43 MW, which lies above the curve value 4.80 MW at the expected demand. The left panel shows why no such inequality holds for the LOLP: the step function is not convex, and here the averaged LOLP, 0.2332, is slightly below the LOLP at the expected demand, $\mathrm{LOLP}(76) = 0.235$. The underestimation by expected values is a property of the expected shortfall and the EUE, not of every index.

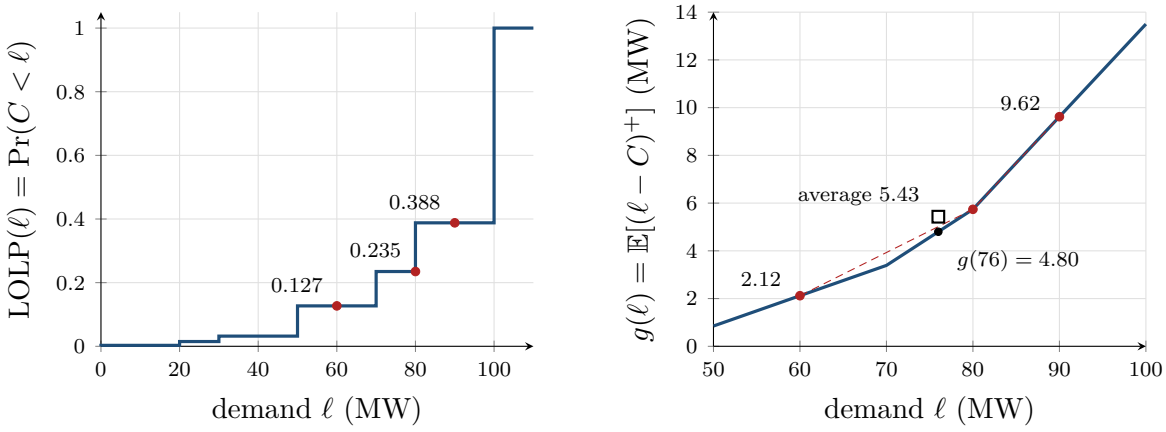


Figure 25.5: Loss-of-load probability (left) and expected shortfall (right) as functions of the demand, for the three-generator COPT of Table 25.3 (illustrative); the right panel spans 50 to 100 MW. The expected shortfall is the integral of the LOLP, equation (25.17), and is convex. Dots: the demand levels 60, 80 and 90 MW of Table 25.6; square: their probability-weighted average expected shortfall, above the value at the expected demand of 76 MW.

Principle 25.3. *The expected shortfall is a convex function of the demand and of the capacity deficit, so a calculation at expected demand and expected capacity underestimates the expected unserved energy. Uncertain demand must be averaged over, not replaced by its mean.*

25.4 Dependent outages and the construction of the capacity outage probability table

The state probabilities of Section 25.2.4 were products of marginal probabilities, which assumes independent units, and they were obtained by listing every state, which is possible only for a

few units. This section relaxes the first assumption with a common-shock model and replaces the listing by the recursive construction used for real fleets.

25.4.1 Sources of dependence between generator outages

Real generating fleets share exposure to common causes: extreme heat or cold, which forces units off line and limits the output of others; interruptions of a shared fuel supply, such as a gas pipeline serving several plants; shortages of cooling water during a drought or a heat wave; common electrical or control systems; common transmission facilities, whose loss disconnects several units at once; and floods, wildfires and storms. These are the coupling factors of common-cause failure, treated in Chapter 13 of Volume I, in generating-fleet form.

The same chapter (Section 13.4) showed the mechanism by which a shared condition creates dependence. If the failure events A and B of two units are independent given an environmental condition Z, so that $\Pr(A \cap B \mid Z) = \Pr(A \mid Z)\,\Pr(B \mid Z)$, they are still dependent once Z is averaged out, because both conditional probabilities move with the same Z; the joint probability is $\Pr(A \cap B) = \sum_z \Pr(A \mid z)\,\Pr(B \mid z)\,\Pr(Z = z)$, which exceeds $\Pr(A)\,\Pr(B)$ when the conditional probabilities rise together. Weather is the common stochastic driver of power systems, and it acts on both sides of the balance: during severe weather the demand and the generator unavailabilities can increase together. Such dependence concentrates probability in the joint outage of many units at the moment of high demand, and it can dominate the upper tail of the shortfall distribution.

25.4.2 A common-shock model

As a hypothetical illustration, let a common shock occur during the assessment interval with probability $q_{\mathrm{cs}} = 0.05$. When the shock occurs, all three generators are unavailable; when it does not, the generator states follow the independent COPT of Table 25.3. The model is a mixture of two distributions, and by the law of total probability any probability or expectation is the shock-weighted average of its two conditional values. At a demand of 80 MW, the shock gives a shortfall with certainty, and of the whole 80 MW:

$$\mathrm{LOLP} = q_{\mathrm{cs}} \times 1 + (1 - q_{\mathrm{cs}})\,\mathrm{LOLP}_{\mathrm{ind}}, \qquad \mathbb{E}[S] = q_{\mathrm{cs}} \times L + (1 - q_{\mathrm{cs}})\,\mathbb{E}[S]_{\mathrm{ind}}, \tag{25.18}$$

where the subscript ind denotes the value under independent outages. With $\mathrm{LOLP}_{\mathrm{ind}} = 0.235$ and $\mathbb{E}[S]_{\mathrm{ind}} = 5.74\,\mathrm{MW}$,

$$\mathrm{LOLP} = 0.05 + 0.95 \times 0.235 = 0.05 + 0.22325 = 0.27325,$$

or 27.325 %, and

$$\mathbb{E}[S] = 0.05 \times 80 + 0.95 \times 5.74 = 4 + 5.453 = 9.453\,\mathrm{MW}.$$

Table 25.7 compares the two models, with a third row discussed in the next subsection. The shock raises the LOLP by 16.3 % ($0.27325/0.235 = 1.163$) and the expected shortfall by 64.7 % ($9.453/5.74 = 1.647$). The proportional effect on the expected shortfall is about four times that on the LOLP because the common event produces the most extreme outcome, in which all 80 MW of demand is unserved. The conditional shortfall of equation (25.11) rises accordingly, from 24.43 MW to $9.453/0.27325 = 34.59$ MW. This is why the EUE can reveal consequences that a frequency-oriented index such as the LOLP does not characterize.

25.4.3 Dependence at fixed marginal forced outage rates

The comparison of the first two rows of Table 25.7 is between two distinct models, not a controlled comparison: the common shock also raises the marginal forced outage rate of every

Table 25.7: Loss-of-load probability and expected shortfall at $L = 80\,\text{MW}$ under three outage models of the three-generator fleet (illustrative).

Model	LOLP	$\mathbb{E}[S]$	$\mathbb{E}[S \mid S > 0]$	$\Pr(S = L)$
Independent, FOR 0.10/0.15/0.20	0.2350	5.740 MW	24.43 MW	0.0030
Common shock $q_{\text{cs}} = 0.05$ added	0.2733	9.453 MW	34.59 MW	0.0529
Independent, FOR 0.145/0.1925/0.24	0.3096	8.319 MW	26.87 MW	0.0067

unit. Under the shock model the probability that generator i is out is $q_{\text{cs}} + (1 - q_{\text{cs}})U_i$, which is $0.05 + 0.95 \times 0.10 = 0.145$ for G1, $0.05 + 0.95 \times 0.15 = 0.1925$ for G2 and $0.05 + 0.95 \times 0.20 = 0.24$ for G3. Part of the increase in the indices comes from these higher marginal rates and part from the dependence. The common shock plays the role of the common-cause event of the beta-factor model of Chapter 14 of Volume I, which also adds a shared failure to the independent ones.

The third row of Table 25.7 separates the two effects: it applies the inflated marginal rates 0.145, 0.1925 and 0.24 to independent units. With the same marginals, independence gives a *larger* LOLP than the shock model, 0.3096 against 0.2733, and a *smaller* expected shortfall, 8.319 against 9.453 MW. The probability that all three units are out is $0.145 \times 0.1925 \times 0.24 = 0.0067$ under independence and $0.05 + 0.95 \times 0.003 = 0.0529$ under the shock, about eight times larger. At fixed marginals, positive dependence moves probability out of the moderate deficits and into the total loss: fewer hours are short, and those that are short are much shorter of capacity.

Figure 25.6 shows this shift. The horizontal axis lists the possible shortfalls at $L = 80\,\text{MW}$, from 0 to 80 MW; the vertical axis, on a logarithmic scale, is the probability of each. The three bars of each group are the three models of Table 25.7. The shock model (red) is close to the original independent model (blue) for every shortfall except the largest, where it is about eighteen times higher (0.0529 against 0.0030); the independent model with inflated marginals (gold) is higher than both for the moderate shortfalls and far below the shock model at 80 MW.

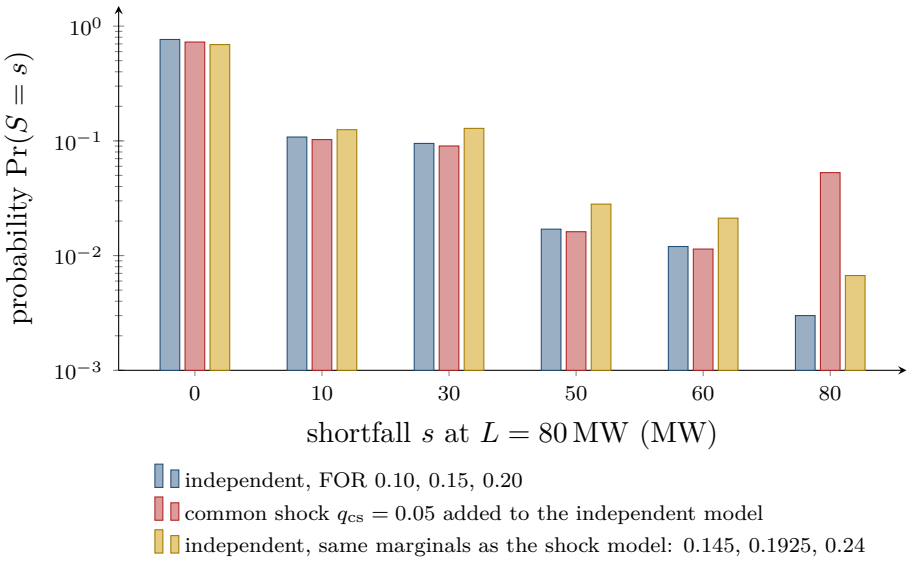


Figure 25.6: Probability of each shortfall at a demand of 80 MW under the three outage models of Table 25.7 (illustrative), on a logarithmic scale. The common shock adds probability almost entirely to the total loss of 80 MW.

Principle 25.4. *Dependence between generator outages matters most in the tail. At fixed marginal forced outage rates, a shared cause can lower the loss-of-load probability while raising the expected unserved energy, because it replaces many moderate deficits with fewer total ones.*

25.4.4 The need for an efficient construction

The three-generator system had eight joint states, and all of them could be listed. A fleet of n two-state units has 2^n joint states; for $n = 100$ that is $2^{100} \approx 1.27 \times 10^{30}$, and direct enumeration is impossible. The number of distinct outage levels is far smaller. When every rating is a multiple of a common capacity step δ (for example 1 or 5 MW), the capacity on outage can take at most $K_{\text{tot}}/\delta + 1$ values, where K_{tot} is the installed capacity. A fleet of 100 units totaling 10 000 MW on a 1 MW grid has at most 10 001 outage levels. The conventional method builds the distribution over these levels directly, adding one unit at a time, so that the work grows with the number of units times the number of levels rather than with 2^n.

25.4.5 Recursive convolution

Let X_n be the capacity on outage of a fleet of n units and $p_n(x) = \Pr(X_n = x)$ its distribution, the COPT of that fleet. Add unit $n+1$, with rating K and forced outage rate U, independent of the first n. The new capacity on outage is $X_{n+1} = X_n + K B_{n+1}$: it is unchanged if the new unit is available and increases by K if it is out. Conditioning on the state of the new unit and using the law of total probability,

$$\begin{aligned} p_{n+1}(x) &= \Pr(B_{n+1} = 0)\ \Pr(X_n = x \mid B_{n+1} = 0) + \Pr(B_{n+1} = 1)\ \Pr(X_n + K = x \mid B_{n+1} = 1) \\ &= (1 - U)\ \Pr(X_n = x) + U\ \Pr(X_n = x - K), \end{aligned}$$

where independence removed the conditioning. Hence

$$p_{n+1}(x) = (1 - U)\, p_n(x) + U\, p_n(x - K), \tag{25.19}$$

with $p_n(x) = 0$ for outage levels x that the first n units cannot produce (in particular every $x < 0$). The recursion starts from an empty fleet, which has nothing on outage with certainty, $p_0(0) = 1$, and adds the units one at a time. Equation (25.19) is the convolution of p_n with the two-point distribution of the new unit's outage. Summing it over all levels at or above x gives the same recursion for the cumulative column of the COPT,

$$\Pr(X_{n+1} \geq x) = (1 - U)\ \Pr(X_n \geq x) + U\ \Pr(X_n \geq x - K), \tag{25.20}$$

with $\Pr(X_n \geq y) = 1$ for $y \leq 0$, which is the form in which the table is often built.

For the fleet of Table 25.1, start with G1, $K_1 = 50$ MW and $U_1 = 0.10$. From $p_0(0) = 1$, equation (25.19) gives $p_1(0) = 0.90 \times 1 = 0.90$ and $p_1(50) = 0.10 \times p_0(0) = 0.10$, the two-row COPT of one unit. Add G2, $K_2 = 30$ MW and $U_2 = 0.15$:

$$\begin{aligned} p_2(0) &= 0.85\, p_1(0) + 0.15\, p_1(-30) = 0.85 \times 0.90 = 0.765, \\ p_2(30) &= 0.85\, p_1(30) + 0.15\, p_1(0) = 0 + 0.15 \times 0.90 = 0.135, \\ p_2(50) &= 0.85\, p_1(50) + 0.15\, p_1(20) = 0.85 \times 0.10 + 0 = 0.085, \\ p_2(80) &= 0.85\, p_1(80) + 0.15\, p_1(50) = 0 + 0.15 \times 0.10 = 0.015, \end{aligned}$$

which sum to 1. Add G3, $K_3 = 20$ MW and $U_3 = 0.20$, so that $p_3(x) = 0.80\, p_2(x) + 0.20\, p_2(x - 20)$:

$$\begin{aligned} p_3(0) &= 0.80 \times 0.765 = 0.612, & p_3(20) &= 0.20 \times 0.765 = 0.153, \\ p_3(30) &= 0.80 \times 0.135 = 0.108, & p_3(50) &= 0.80 \times 0.085 + 0.20 \times 0.135 = 0.095, \\ p_3(70) &= 0.20 \times 0.085 = 0.017, & p_3(80) &= 0.80 \times 0.015 = 0.012, \\ p_3(100) &= 0.20 \times 0.015 = 0.003. \end{aligned}$$

This is the COPT of Table 25.3, obtained without listing the eight states. The entry at 50 MW shows the aggregation at work: it receives 0.068 from the level 50 MW of p_2 (G1 out, G3 available) and 0.027 from the level 30 MW (G2 out, G3 out).

Figure 25.7 traces the construction. Each column is the COPT after one more unit, from the empty fleet on the left to the full fleet on the right; each box is one outage level with its probability, placed vertically by the outage level. A solid arrow carries probability to the same level with weight $1-U$ (the new unit available); a dashed arrow carries it to the level K higher with weight U (the new unit out). A box with two incoming arrows, such as $p_3(50)$, aggregates two states.

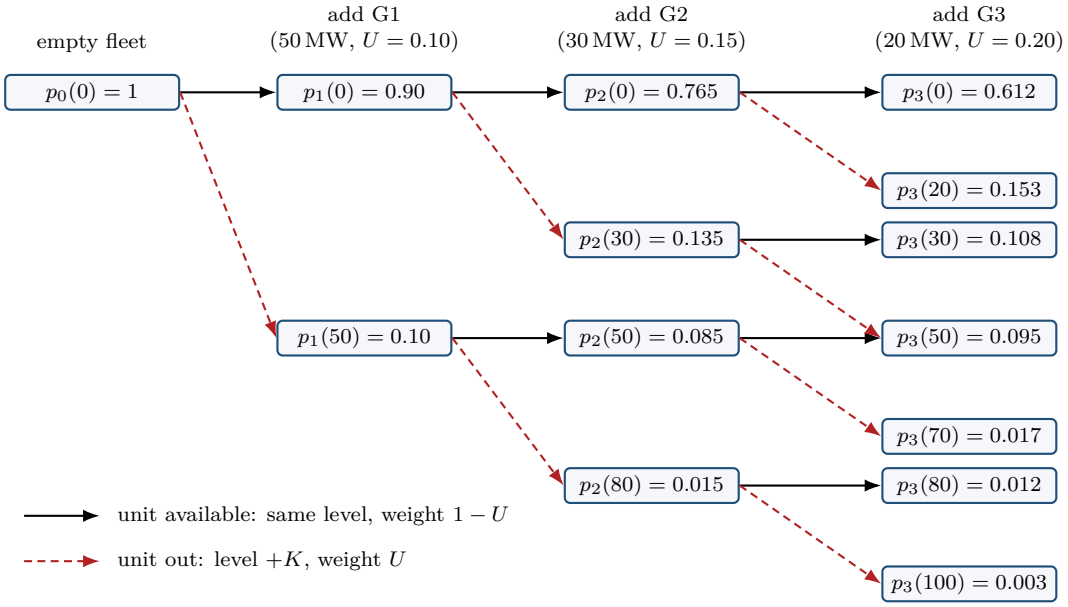


Figure 25.7: Recursive construction of the COPT of the three-generator fleet by equation (25.19), $p_{n+1}(x) = (1-U)\,p_n(x) + U\,p_n(x-K)$, starting from $p_0(0) = 1$. Boxes are placed vertically by the outage level x (illustrative fleet of Table 25.1).

The technique is efficient when the ratings lie on a discrete capacity grid. For large fleets with partial deratings, which give each unit more than two states, or with ratings that share no convenient common step, the levels multiply, and rounding to a grid, approximation of the distribution or other numerical methods are needed. The recursion also inherits the independence assumption of its derivation: each unit is added independently of the fleet already built, so a common shock must be added separately, as a mixture in the manner of equation (25.18).

25.5 Resource adequacy, operating security and the sensitivity of the indices

The indices computed so far answer a specific question under specific assumptions. This section separates adequacy from operating security, adds network and energy constraints to the shortfall, and tabulates the sensitivity of the indices to the demand, to the forced outage rates and to the length of the assessment interval.

25.5.1 Adequacy and contingency security

Generation adequacy asks whether sufficient resources will be available to meet the demand over the assessment period. Operating security asks whether the system can withstand a specified

contingency, such as the sudden loss of a large unit or a line, without violating operating limits or causing unacceptable cascading consequences. Chapter 3 of Volume I introduced the two as related but different aspects of power-system reliability, and the indices of this chapter belong to the first.

The difference has consequences in both directions. A generation adequacy study may conclude that $C_t \geq L_t$ in an hour while the network is unable to deliver enough power to a particular load area, because of transmission constraints; the copper-plate model cannot see that deficit. Conversely, a system may show a shortfall in a simplified adequacy model and still avoid any actual load shedding through emergency imports or operating actions that the model omits. An adequacy index is therefore conditional on the model's boundary in both senses: what it leaves out can create deficits it does not report and remove deficits it does report.

25.5.2 Transmission and operating constraints

A more complete adequacy model includes the feasibility of delivering the available generation through the network. For each interval t and scenario s, define the minimum load shedding needed to satisfy the physical operating constraints,

$$S^*_{t,s} = \min_{\mathbf{x},\,\mathbf{u},\,\Delta \boldsymbol{L}} \sum_n \Delta L_n \tag{25.21}$$

subject to the power balance at every bus, the operating limits of the available generators, the transmission constraints (line and transformer ratings, and voltage and stability limits where modeled), any other relevant network constraints, and $0 \leq \Delta L_n \leq L_n$. Here ΔL_n is the load curtailed at bus n, L_n the demand at that bus, $\mathbf{x}$ the network state (bus voltage angles and magnitudes) and $\mathbf{u}$ the dispatch. The expected unserved energy is then

$$\mathrm{EUE} = \sum_{t,s} p_{t,s}\, S^*_{t,s}\, \Delta t, \tag{25.22}$$

which reduces to equation (25.13) when the network is a single bus and $S^*_{t,s} = (L_{t,s} - C_{t,s})^+$. This formulation connects probabilistic adequacy assessment to power flow, security-constrained dispatch and network optimization, and it matters whenever geographically distributed resources cannot always supply geographically distributed demand.

Figure 25.8 gives an illustrative case. The fleet of Table 25.1 is split between two areas joined by one tie line: G1 and G2 in area A on the left, G3 and all 80 MW of demand in area B on the right. The tie line, drawn between the two buses, can carry at most $K_{AB} = 60$ MW from A to B. With only two areas and one line, problem (25.21) has the closed-form solution $S^* = \big(80 - \min(C_A, K_{AB}) - C_B\big)^+$, where C_A and C_B are the available capacities in the two areas. With all three units available, area A has 80 MW but can export only 60 MW, which with G3's 20 MW just meets the demand. With only G3 out, the copper-plate model has 80 MW and no shortfall; the two-area model delivers 60 MW and sheds 20 MW. Summing over the eight states of Table 25.2, the LOLP rises from 0.235 to $1 - 0.612 = 0.388$ and the expected shortfall from 5.74 MW to $0.153 \times 20 + 0.108 \times 10 + 0.027 \times 30 + 0.068 \times 30 + 0.017 \times 50 + 0.012 \times 60 + 0.003 \times 80 = 8.80$ MW. The generators and their outage probabilities are unchanged; only the deliverability is. A tie-line rating is itself weather-dependent, as Chapter 19 of Volume I showed for two lines whose configuration probabilities are combined with a network feasibility function, $A_{\text{network}} = \sum_{\mathbf{x}} \Pr(\text{configuration } \mathbf{x})\, \phi(\mathbf{x}, \mathbf{u})$: a hot, still day that lowers K_{AB} raises the shortfall in exactly this way.

25.5.3 Energy-limited resources: storage and renewables

A battery rated at 50 MW cannot necessarily deliver 50 MW indefinitely. With a maximum discharge power $P_{\max} = 50$ MW and a usable stored energy $E_0 = 100$ MWh, and ignoring

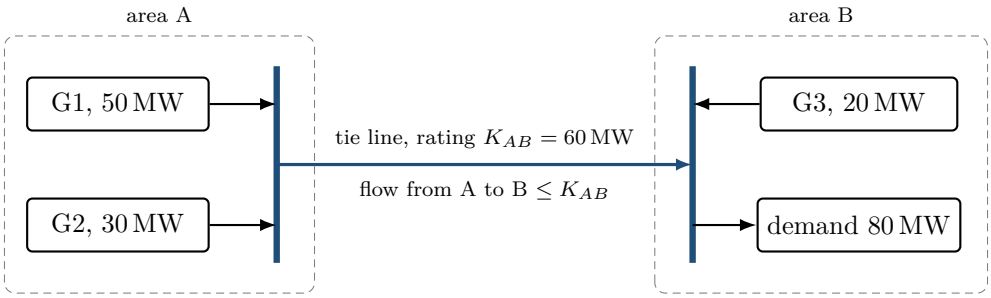


Figure 25.8: Illustrative two-area version of the three-generator system. The generators and outage probabilities are those of Table 25.1; the tie-line rating limits what area A can deliver to the demand in area B. At $L = 80$ MW the LOLP rises from 0.235 (single bus) to 0.388 and the expected shortfall from 5.74 to 8.80 MW.

efficiency losses, it can deliver its full power for at most

$$\frac{E_0}{P_{\max}} = \frac{100\,\text{MWh}}{50\,\text{MW}} = 2\,\text{h}.$$

Its adequacy contribution therefore depends on the duration and the timing of the shortfalls, not only on its rating. As an illustration, if the three-generator fleet is in the state with G1 out ($C = 50$ MW) through the four peak hours of the daily profile, the shortfall is 30 MW for 4 h, 120 MWh in all; a full battery covers it for $100/30 = 3.33$ h and leaves 20 MWh unserved, while a firm 50 MW generator would have covered all of it. A COPT that counted the battery as a 50 MW unit would overstate its contribution. Solar and wind generation have weather-dependent availability of a different kind: their output is not a two-state random capacity with a fixed outage probability, but a time series driven by the weather, and correlated with the demand.

A credible resource adequacy assessment therefore often needs a chronological simulation that steps through the hours of the year with the hourly demand, the weather-dependent renewable generation, the state of charge of each battery, the outages and repairs of the units, the imports and transmission limits, and the demand response and dispatch controls. The peak reserve margin and the static COPT are important introductory methods, but they are not always sufficient for modern power systems; energy-limited resources and a changing resource mix require energy-oriented adequacy measures, such as the EUE and its distribution, alongside the traditional capacity assessment.

25.5.4 Loss-of-load probability, expected shortfall and unserved energy for several demands and forced outage rates

The indices of the three-generator fleet respond differently to the demand, to each unit's forced outage rate and to the length of the interval. The model is the exact eight-state model of Section 25.2.4: generators of 50, 30 and 20 MW with forced outage rates U_1, U_2, U_3, independent states from equation (25.9), a constant demand L, and a capacity state that is constant over an assessment interval of length τ. For each case the table gives the LOLP from equation (25.3), the expected shortfall $\mathbb{E}[S]$ from equation (25.2), the expected unserved energy over the interval $\text{EUE}_\tau = \mathbb{E}[S]\,\tau$, and the conditional shortfall of equation (25.11).

Table 25.8 lists the results. The upper block varies the demand at the base forced outage rates with $\tau = 4$ h; the lower block varies the forced outage rates at $L = 80$ MW and gives the EUE for intervals of 4 and 12 h.

The row with $U_1 = 0.02$ shows the arithmetic. With $(U_1, U_2, U_3) = (0.02, 0.15, 0.20)$, the states that leave less than 80 MW have probabilities $0.98 \times 0.15 \times 0.80 = 0.1176$ (G2 out, $C = 70$), $0.02 \times 0.85 \times 0.80 + 0.98 \times 0.15 \times 0.20 = 0.0136 + 0.0294 = 0.0430$ ($C =$

Table 25.8: Indices of the three-generator fleet (50, 30, 20 MW, independent two-state units, capacity state constant over the interval) for several demands and forced outage rates (illustrative).

L (MW)	U_1	U_2	U_3	LOLP	$\mathbb{E}[S]$ (MW)	$\mathrm{EUE}_{4\,\mathrm{h}}$ (MWh)	$\mathrm{EUE}_{12\,\mathrm{h}}$ (MWh)	$\mathbb{E}[S \mid S>0]$ (MW)
50	0.10	0.15	0.20	0.0320	0.850	3.40		26.56
60	0.10	0.15	0.20	0.1270	2.120	8.48		16.69
70	0.10	0.15	0.20	0.1270	3.390	13.56		26.69
80	0.10	0.15	0.20	0.2350	5.740	22.96		24.43
90	0.10	0.15	0.20	0.3880	9.620	38.48		24.79
100	0.10	0.15	0.20	0.3880	13.500	54.00		34.79
80	0.10	0.15	0.20	0.2350	5.740	22.96	68.88	24.43
80	0.05	0.15	0.20	0.1925	3.920	15.68	47.04	20.36
80	0.02	0.15	0.20	0.1670	2.828	11.31	33.94	16.93
80	0.10	0.05	0.20	0.1450	4.180	16.72	50.16	28.83
80	0.10	0.15	0.05	0.2350	5.035	20.14	60.42	21.43
80	0.05	0.075	0.10	0.1213	2.568	10.27	30.81	21.18
80	0.20	0.30	0.40	0.4400	13.720	54.88	164.64	31.18

50), $0.02 \times 0.85 \times 0.20 = 0.0034$ ($C = 30$), $0.02 \times 0.15 \times 0.80 = 0.0024$ ($C = 20$) and $0.02 \times 0.15 \times 0.20 = 0.0006$ ($C = 0$). Their sum is the LOLP, 0.1670, equal to $1 - 0.98 \times 0.85$. The expected shortfall is $0.1176 \times 10 + 0.0430 \times 30 + 0.0034 \times 50 + 0.0024 \times 60 + 0.0006 \times 80 = 1.176 + 1.290 + 0.170 + 0.144 + 0.048 = 2.828$ MW; over 4 h the EUE is 11.31 MWh, and the conditional shortfall is $2.828/0.1670 = 16.93$ MW.

Four trends follow from the table. First, raising the demand from 80 to 90 MW raises the LOLP from 0.235 to 0.388 and the expected shortfall by 68 %, from 5.74 to 9.62 MW, although the installed capacity is still 100 MW: adequacy risk rises sharply as the demand approaches the capacity. The LOLP is a step function of the demand, equal at 60 and 70 MW and at 90 and 100 MW, while the expected shortfall rises in every row, as Figure 25.5 showed.

Second, the forced outage rate of the largest unit has a strong effect on the expected shortfall: lowering U_1 from 0.10 to 0.02 halves it, from 5.74 to 2.83 MW, because the loss of G1 removes half the installed capacity. Its effect on the LOLP is more modest, from 0.235 to 0.167, and per unit change of the forced outage rate it is not the largest. Since every index here is a linear function of each U_i with the others fixed, its derivative with respect to U_i is the difference between its values with unit i out and with unit i available, which is the Birnbaum importance of Chapter 12 of Volume I. For the LOLP at 80 MW, $\mathrm{LOLP} = 1 - (1 - U_1)(1 - U_2)$ gives $\partial\,\mathrm{LOLP}/\partial U_1 = 1 - U_2 = 0.85$, $\partial\,\mathrm{LOLP}/\partial U_2 = 1 - U_1 = 0.90$ and $\partial\,\mathrm{LOLP}/\partial U_3 = 0$; for the expected shortfall the three derivatives are 36.4, 15.6 and 4.7 MW per unit of forced outage rate. G1 dominates the expected shortfall, G2 the LOLP by a small margin, and G3 does not affect the LOLP at 80 MW at all, although lowering U_3 from 0.20 to 0.05 still lowers the expected shortfall from 5.74 to 5.04 MW.

Third, the interval length scales the EUE and nothing else: going from 4 to 12 h triples the EUE, from 22.96 to 68.88 MWh in the base case, while the LOLP and the expected shortfall are unchanged. This is a consequence of the fixed-state assumption of the table, under which the capacity state drawn at the start of the interval persists through it; it is the distinction between the probability of a capacity deficiency and the energy exposed to it. Fourth, halving all three forced outage rates roughly halves both the LOLP (0.1213) and the expected shortfall (2.568 MW), and doubling them more than doubles the expected shortfall

(13.72 MW), which grows faster than the LOLP because multiple outages, the large deficits, become disproportionately likely.

25.6 Limits and interpretation of adequacy indices

An adequacy index is a number attached to an event definition, a time discretization and a model boundary. This section states the definitions that must accompany a reported index, the units and readings of the one-day-in-ten-years criterion, and what a static capacity outage probability table leaves out.

25.6.1 Definitions that must accompany an index

A study can be mathematically correct and still be misunderstood if its indices are reported without their event definitions. Expected shortage hours, expected shortage days and expected shortage events are three different quantities. Table 25.9 gives two illustrative realized years. Each row is one year; the columns count its loss-of-load hours, loss-of-load days (days with at least one shortage hour) and loss-of-load events (maximal runs of consecutive shortage hours, often written LOLEV in their expected form).

Table 25.9: Hour, day and event counts of two illustrative realized years with the same number of shortage hours.

Realized year	Shortage hours	Shortage days	Shortage events
One continuous four-hour shortfall within one day	4	1	1
Four isolated one-hour shortfalls on four different days	4	4	4

The same number of shortage hours does not imply the same number of days or events, and the same expected number of hours does not imply the same expected number of days. The daily profile of Section 25.3.4 shows the size of the difference. Its expected shortage hours are 1270.2 h per year under any assumption about time dependence, by equation (25.12). The expected number of shortage days is not determined by them. If the capacity state is drawn once a day and persists through the day, which approximates outages lasting many hours, a day has a shortfall exactly when the capacity is below the peak demand, with probability $\mathrm{LOLP}(80) = 0.235$, so the daily LOLE is $365 \times 0.235 = 85.8$ days per year. If instead the capacity state were drawn independently every hour, a day would escape a shortfall only if all twenty-four hours did, with probability $0.765^4 \times 0.873^{20} = 0.0226$, and the daily LOLE would be $365 \times (1 - 0.0226) = 356.7$ days per year. Temporal dependence and event duration decide the day and event counts, and only a chronological model represents them.

25.6.2 The one-day-in-ten-years criterion

A frequently cited North American resource adequacy benchmark is

$$\mathrm{LOLE} = 0.1 \text{ day per year},$$

traditionally described as “one day in ten years”. Its unit is days per year: it is a daily LOLE, an expected number of days with a shortfall per year, not an hourly one. Two readings of it are common. In the first, the classical one, it bounds the expected number of days per year on which the available capacity falls short of the daily peak demand, computed as $\sum_d \Pr(C_d < L_{d,\mathrm{peak}})$

over the days d of the year, so that a day counts once however many of its hours are short. In the second, it is read as a frequency of events: on average one loss-of-load event in ten years, 0.1 events per year. The two coincide only when every event falls within a single day and no day has two events.

Neither reading means one 24 h outage every ten years, and the criterion does not directly imply LOLH $= 0.1 \times 24 = 2.4$ h per year: the conversion depends on how a shortage day is defined, on the duration of events and on the chronological model, as the two daily LOLEs of the previous subsection, with the same LOLH, showed. The criterion is also not a universal enforceable standard for every balancing area or resource-adequacy program; some apply it, some use an hourly or an energy-based target, and its continued suitability is debated as the generation portfolio changes, which is a reason to report hourly and energy-based indices alongside it.

Common mistake

Converting "one day in ten years" into 2.4 h of shortage per year, or into one full-day outage per decade. The criterion counts days with at least one shortfall, or events, per year; the number of hours it allows depends on the event definition and the chronological model.

25.6.3 What the static COPT does not represent

A conventional static COPT built from independent two-state generators does not by itself represent the features listed in Table 25.10. Each row names a feature and the kind of extension that brings it into the model: some can be incorporated by extending the state model while keeping an analytical calculation, others require chronological simulation, correlated scenario generation or integrated optimization. The appropriate fidelity depends on the engineering question: a screening comparison of two expansion plans may need only the COPT, while a study of a system with much storage and renewable generation needs the chronological model.

Table 25.10: Features a static COPT of independent two-state units does not represent, and the kind of model extension that represents each.

Feature	Kind of extension
Outage durations; repairs and time correlation	state-transition (Markov) models or chronological simulation of up and down times
Planned maintenance schedules	a separate COPT for each period, with the units on maintenance removed
Common-cause failures	dependent state model, for example the common-shock mixture of Section 25.4.2
Generator deratings	multistate unit models in the recursion
Fuel or water limitations	energy constraints across hours; chronological simulation
Transmission deliverability	network model, minimum load shedding of equation (25.21)
Weather dependence of renewables	correlated weather, demand and output scenarios
Demand response activation	dispatch rules in a chronological simulation
Battery charging and discharge limits	state-of-charge tracking in a chronological simulation

The next chapter takes up the state models themselves: multistate generators with partial deratings, the recursive convolution for large fleets, correlated outages, and the choice between analytical enumeration and Monte Carlo simulation.

25.7 What to retain

Definition 25.1 (Generation adequacy and its indices)**.** Generation adequacy is the probability, magnitude and duration of the shortfall $S_t = (L_t - C_t)^+$ of random available capacity against demand, not whether installed capacity exceeds the peak. It is measured by the loss-of-load probability $\mathrm{LOLP}_t = \Pr(C_t < L_t)$, the loss-of-load expectation $\mathrm{LOLE} = \sum_t \mathrm{LOLP}_t \, \Delta t$ (LOLH when hourly) and the expected unserved energy $\mathrm{EUE} = \sum_t \mathbb{E}[S_t] \, \Delta t$, computed from the capacity outage probability table under stated assumptions on independence, demand and the resources modeled, and reported together.

Forced outages make available capacity a random variable, and its expected value is no substitute for its distribution. The COPT is the fundamental model of outage combinations, and its independence assumption must be explicit. Expected shortage hours follow from the hourly probabilities without temporal independence; event frequency and duration do not. Uncertain demand and dependent outages, above all common weather exposure, change the risk materially, and energy-limited resources and network constraints require more detailed models than a static COPT. Table 25.11 lists the chapter's terms.

Table 25.11: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Adequacy event (loss of load)	capacity below demand	$C_t < L_t$	–	25.1.2
Capacity shortfall	unmet demand in an interval	$S_t = (L_t - C_t)^+$	MW	25.1.2
LOLP	probability of loss of load in one interval	$\Pr(C_t < L_t)$	–	25.1.3
LOLE, LOLH	expected shortage intervals; hours when hourly	$\sum_t \mathrm{LOLP}_t\,\Delta t$	h/yr, d/yr	25.1.3
EUE (EENS)	expected unserved energy	$\sum_t \mathbb{E}[S_t]\,\Delta t$	MWh/yr	25.1.3
Forced outage rate	per-unit unavailability	$\mathrm{FOR}_i = U_i = \mathrm{MTTR}_i/(\mathrm{MTTF}_i + \mathrm{MTTR}_i)$	–	25.2.1
Random available capacity	two-state unit capacity	$K_i(1 - B_i)$	MW	25.2.2
Installed reserve margin	spare installed capacity	$(K_{\mathrm{tot}} - L)/L$	–	25.2.3
Outage state probability	independent units	$\prod_i U_i^{b_i}(1 - U_i)^{1-b_i}$	–	25.2.4
Capacity outage	capacity of units out	$X = \sum_i K_i B_i$	MW	25.2.4
COPT	distribution of capacity on outage	$\Pr(X = x)$, $\Pr(X \geq x)$	–	25.2.5
Conditional shortfall	mean deficit given a deficit	$\mathbb{E}[S]/\mathrm{LOLP}$	MW	25.2.6
Expected shortfall curve	expected shortfall against demand	$g(\ell) = \int_0^\ell \mathrm{LOLP}(y)\,\mathrm{d}y$	MW	25.3.6
Common-shock model	shared outage of all units	equation (25.18)	–	25.4.2
Recursive convolution	COPT built unit by unit	equation (25.19)	–	25.4.5
Minimum load shedding	shortfall with network constraints	equation (25.21)	MW	25.5.2
One day in ten years	daily LOLE criterion	$\mathrm{LOLE} = 0.1$	d/yr	25.6.2

Further reading. EPRI, *Resource Adequacy Metrics Explained*, for the definitions and distinctions among LOLP, LOLE, LOLH, EUE and event-count metrics; NERC, *Probabilistic Adequacy and Measures*, for resource adequacy metrics, modeling practices and interpretation; Billinton and Allan, *Reliability Evaluation of Power Systems*, for generation adequacy, capacity outage probability tables and reliability indices; EPRI, *Resource Adequacy Program*, for resource adequacy modeling and research.

25.8 Exercises

Exercise 25.1. (a) State the event that defines insufficient generation adequacy in the basic capacity model. (b) State the units of an hourly LOLP. (c) State what the EUE measures.

(d) Explain why the expected available capacity cannot replace its probability distribution in an adequacy study. (e) State what a capacity outage probability table describes. (f) State whether summing hourly LOLPs to obtain the expected shortage hours requires the hours to be independent. (g) Explain why common weather conditions can increase adequacy risk. (Answers: (a) the available capacity is below the demand, $C_t < L_t$; (b) none, it is a dimensionless probability; (c) the expected energy of the shortfall, which reflects its magnitude as well as its probability; (d) the shortfall depends nonlinearly on the capacity state, so the distribution, not its mean, determines the indices; (e) the probability distribution of the generating capacity on outage; (f) no, linearity of expectation suffices; (g) they may reduce generation availability and increase demand at the same time.)

Exercise 25.2. (a) A 10 MW unit G4 has MTTF = 1900 h and MTTR = 100 h. Compute its forced outage rate from equation (25.8). (b) Add G4 to the COPT of Table 25.3 with the recursion of equation (25.19), and list the new COPT with its cumulative column. (c) Compute the LOLP, the expected shortfall and the conditional shortfall at $L = 80$ MW, and explain why the LOLP falls so much. (Answers: (a) $100/2000 = 0.05$. (b) $x = 0$: 0.5814; 10: 0.0306; 20: 0.14535; 30: 0.11025; 40: 0.0054; 50: 0.09025; 60: 0.00475; 70: 0.01615; 80: 0.01225; 90: 0.0006; 100: 0.00285; 110 MW: 0.00015; cumulative from $x = 110$ down: 0.00015, 0.003, 0.0036, 0.01585, 0.032, 0.03675, 0.127, 0.1324, 0.24265, 0.388, 0.4186, 1. (c) With 110 MW installed, loss of load requires $X > 30$ MW, so LOLP $= \Pr(X \geq 40) = 0.1324$; $\mathbb{E}[S] = 3.5075$ MW; $\mathbb{E}[S \mid S > 0] = 26.49$ MW. The state with G2 alone out now leaves 80 MW when G4 is available, so it no longer causes a shortfall.)

Exercise 25.3. The demand of one hour is 70 MW with probability 0.6 and 90 MW with probability 0.4, independent of the outages of the three-generator fleet. (a) Compute the LOLP and the expected shortfall, using Table 25.8. (b) Compute the expected demand and the expected shortfall at that single demand, and the deterministic estimate $(\mathbb{E}[L] - \mathbb{E}[C])^+$. (c) Check the expected shortfall at 70 MW by integrating the LOLP as in equation (25.17). (d) State which inequality the results illustrate. (Answers: (a) $0.6 \times 0.127 + 0.4 \times 0.388 = 0.2314$; $0.6 \times 3.39 + 0.4 \times 9.62 = 5.882$ MW. (b) $\mathbb{E}[L] = 78$ MW; $\mathbb{E}[(78 - C)^+] = 5.27$ MW; $(78 - 86.5)^+ = 0$. (c) $0.003 \times 20 + 0.015 \times 10 + 0.032 \times 20 + 0.127 \times 20 = 3.39$ MW. (d) Jensen's inequality for the convex expected shortfall, $5.882 \geq 5.27 \geq 0$; the LOLP at the expected demand, 0.235, exceeds the averaged LOLP, so no such inequality holds for the LOLP.)

Exercise 25.4. In every hour of the daily profile of Section 25.3.4 (four hours at 80 MW, twenty at 60 MW), a common shock that takes all three units out occurs with probability $q_{\mathrm{cs}} = 0.02$; otherwise the independent COPT applies. (a) Compute the LOLP and the expected shortfall at 80 and at 60 MW. (b) Compute the expected shortage hours and the expected unserved energy per day and per 365-day year. (c) Compare the proportional increases of the two indices over the independent model and explain the difference. (d) State the marginal forced outage rate of G1 under this model. (Answers: (a) 80 MW: $0.02 + 0.98 \times 0.235 = 0.2503$ and $0.02 \times 80 + 0.98 \times 5.74 = 7.2252$ MW; 60 MW: $0.02 + 0.98 \times 0.127 = 0.14446$ and $0.02 \times 60 + 0.98 \times 2.12 = 3.2776$ MW. (b) $4 \times 0.2503 + 20 \times 0.14446 = 3.8904$ h and $4 \times 7.2252 + 20 \times 3.2776 = 94.45$ MWh per day; 1420.0 h and 34 475 MWh per year. (c) The LOLH rises by 11.8 % (from 3.48 h) and the EUE by 44.5 % (from 65.36 MWh), because the shock adds rare total losses of the whole demand, which weigh far more in the energy than in the count of hours. (d) $0.02 + 0.98 \times 0.10 = 0.118$.)

Chapter 26

State Enumeration and Capacity Outage Models

26.1 Multistate generating units and the capacity outage model

Chapter 25 computed the loss-of-load probability, the loss-of-load hours and the expected unserved energy of a small fleet whose generators were either fully available or fully unavailable,

$$\mathrm{LOLP}_t = \Pr(C_t < L_t), \qquad \mathrm{LOLH} = \sum_t \mathrm{LOLP}_t \times 1\,\mathrm{h}, \qquad \mathrm{EUE} = \sum_t \mathbb{E}\big[(L_t - C_t)^+\big]\,\Delta t,$$

from a capacity outage probability table (COPT) built by listing the joint outage states or by recursive convolution. Real resources have more than two operating states. A generator may deliver its full rating, a fraction of it or nothing; a thermal unit may be derated when its cooling water is warm; a battery may have the power to serve the demand but not the stored energy; several units may depend on the same fuel supply, cooling system or weather. The question of this chapter is

how the probability distribution of the available capacity of a fleet is computed when each resource can be fully available, partially derated or unavailable, and when the outages may be dependent; and which computational method, exact enumeration, recursive convolution or Monte Carlo simulation, fits which model.

The chapter builds multistate models of generating units, enumerates the joint states of a small multistate fleet exactly and computes its adequacy indices, derives the recursion and the generating function that build the same distribution for large fleets, represents dependent outages through a shared latent state, develops nonsequential and chronological Monte Carlo simulation with their sampling error, and ends with the choice of method, the numerical approximations that make large problems tractable, and the checks that a capacity model must pass before it is trusted.

Jargon: multistate unit, derated state, capacity outage, unit addition algorithm, convolution, probability generating function, state aggregation, latent state, nonsequential and sequential Monte Carlo, variance reduction

A *multistate unit* is a resource modeled with more than two capacity states; a *derated state* (a partial outage, or *derating*) is a state in which the unit is in service but can deliver only part of its rated capacity. The *capacity outage* of a unit is its rated capacity minus its available capacity; summed over a fleet it is the variable that a COPT describes. The *unit addition algorithm* is the recursion that builds the COPT one unit at a time; each step is a *convolution*, the operation that gives the distribution of a sum of independent random variables from their individual distributions. A *probability generating function* encodes a discrete distribution as a polynomial whose coefficients are the probabilities, so that convolution becomes multiplication. *State aggregation* merges joint states that share a quantity of interest, here the total capacity. A *latent state* is a shared, often unobserved, variable such as the weather that conditions several outage probabilities at once. *Nonsequential* (state-sampling) Monte Carlo draws independent snapshots of the system state; *sequential* (chronological) Monte Carlo simulates the system hour by hour through a study period. *Variance reduction* is the family of sampling designs that lower the statistical error of a Monte Carlo estimate for a given number of samples.

Key idea

Adequacy indices are functionals of the probability distribution of the available capacity. That distribution can be computed exactly by enumerating the joint states, efficiently by convolving the units' distributions one at a time when the units are independent, as a mixture of such convolutions when the dependence enters through a shared state, and approximately by sampling when chronology, energy limits or complex dependence make the state space too rich. The method is chosen by the structure of the model: what is independent, which states can be aggregated, and which dependencies must be retained.

This section replaces the two-state generator by a multistate one, relates the available-capacity and the capacity-outage descriptions of the same unit, and states where derated states come from and how their probabilities are estimated.

26.1.1 Beyond binary availability

The conventional two-state generator of Chapter 25 has available capacity

$$C_i = \begin{cases} K_i, & \text{available,} \\ 0, & \text{unavailable,} \end{cases}$$

where K_i is its rated capacity in MW. A *multistate* generator instead has an available capacity C_i that takes one of m_i values,

$$C_i \in \{c_{i,1}, c_{i,2}, \ldots, c_{i,m_i}\}, \qquad \Pr(C_i = c_{i,j}) = p_{i,j}, \qquad \sum_{j=1}^{m_i} p_{i,j} = 1, \tag{26.1}$$

where $c_{i,j}$ is the capacity of state j, in MW, and $p_{i,j}$ its probability in a randomly selected assessment interval. The probabilities sum to one because the states are mutually exclusive and exhaust the conditions of the unit. The states should represent distinct levels of capacity that the resource can credibly deliver during the modeled interval: full output, one or more derated levels, and the forced outage. The two-state model is the case $m_i = 2$ with $c_{i,1} = K_i$, $c_{i,2} = 0$ and $p_{i,2} = U_i$, the forced outage rate. The subscript i on C_i denotes a unit, while the subscript

t on C_t in Chapter 25 denotes an hour; the fleet capacity in one interval is written C without subscript.

Figure 26.1 shows an illustrative 100 MW generator with three states. Each row is one state; the gray outline is the rated capacity of 100 MW, the colored bar the capacity available in that state, and the right column its probability. The unit delivers its full 100 MW with probability 0.80, a derated 60 MW with probability 0.15, and nothing with probability 0.05.

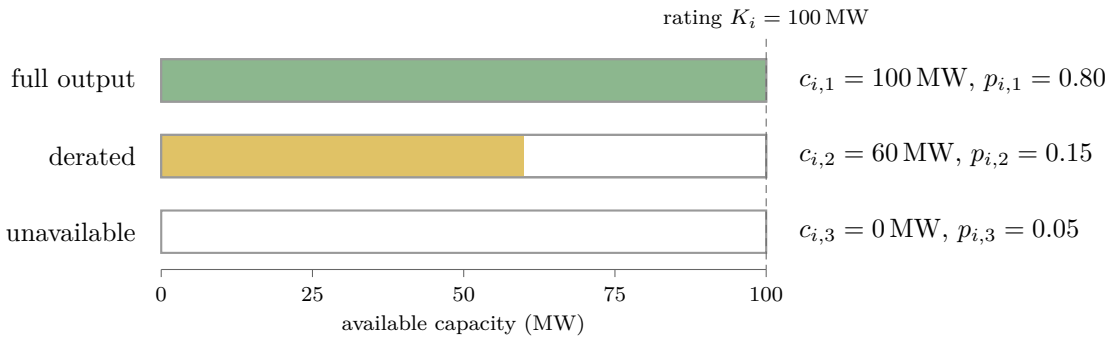


Figure 26.1: Operating states of an illustrative three-state 100 MW generator. The colored bars are the available capacity in each state, the gray outlines the rated capacity. Expected available capacity: $0.80 \times 100 + 0.15 \times 60 + 0.05 \times 0 = 89$ MW.

The expected available capacity of a multistate unit is the probability-weighted sum of its state capacities,

$$\mathbb{E}[C_i] = \sum_{j=1}^{m_i} p_{i,j}\, c_{i,j}, \tag{26.2}$$

which for the generator of Figure 26.1 is

$$\mathbb{E}[C_i] = 0.80 \times 100 + 0.15 \times 60 + 0.05 \times 0 = 80 + 9 + 0 = 89\,\text{MW}.$$

As Section 25.2.2 established for the two-state unit, a random resource with mean capacity 89 MW is not equivalent to a guaranteed 89 MW resource when shortage probabilities are computed. Against a demand of 80 MW, the guaranteed resource never falls short; the multistate unit falls short with probability $0.15 + 0.05 = 0.20$, by 20 MW in the derated state and by 80 MW in the outage state, an expected shortfall of $0.15 \times 20 + 0.05 \times 80 = 7$ MW.

Common mistake

Collapsing a multistate unit to its expected capacity, or to a two-state unit whose outage probability is the probability of any departure from full output. The first discards the distribution altogether; the second treats every derating as a full outage and overstates both the probability and the size of the shortfalls (Section 26.2.4).

26.1.2 Available capacity and capacity outage

Two equivalent descriptions of a unit's state are in common use. The *available capacity* C_i is the capacity the resource can supply; the *capacity outage*

$$D_i = K_i - C_i \tag{26.3}$$

is the part of the rating it cannot supply, with values $d_{i,j} = K_i - c_{i,j}$ and the same probabilities $p_{i,j}$. Table 26.1 lists both for the 100 MW generator of Figure 26.1: each row is one operating state, with its available capacity and its capacity outage. The symbol D_i denotes a capacity

Table 26.1: Available capacity and capacity outage of the illustrative 100 MW generator in its three states.

Operating state	Available capacity C_i	Capacity outage D_i	Probability
Full output	100 MW	0 MW	0.80
Derated	60 MW	40 MW	0.15
Unavailable	0 MW	100 MW	0.05

outage in this chapter; it is not the degradation state $D(t)$ of Chapter 21 nor the FMEA detection rating D.

When the rated capacity is fixed, the two descriptions contain the same information: each determines the other through equation (26.3). For a fleet of n units the totals are

$$C = \sum_{i=1}^{n} C_i, \qquad X = \sum_{i=1}^{n} D_i, \qquad C = K_{\text{tot}} - X, \tag{26.4}$$

where $K_{\text{tot}} = \sum_i K_i$ is the installed capacity and X the total capacity on outage, the variable of the COPT. The worked example of the next section uses the available-capacity form, because the shortfall $(L - C)^+$ is written in it; the recursion of Section 26.3 uses the outage form, because the empty fleet has a fixed outage of zero whatever the ratings of the units that will be added.

26.1.3 Sources and estimation of derated states

Derated states arise from physical limits that remove part of a unit's capability without removing the unit from service. A thermal generator whose condenser is cooled by river or lake water loses output when the water is warm, because the condenser pressure rises and the cycle efficiency falls; a gas turbine loses output in hot air, whose lower density reduces the mass flow; a boiler feed pump or a mill out of service may limit a unit to part load; a hydro unit is limited by its head and its water allowance. A battery is limited in a different way: it may have the power rating to serve the demand at an instant but not the stored energy to sustain it, which is a constraint across hours rather than a state of one hour (Section 26.5.7). Several units can share the cause of a derating or an outage, a fuel supply, a cooling-water source, a heat wave, which makes their states dependent (Section 26.4).

The distinction between the physical system, the model and the evidence applies to each derated state. The physical unit has a continuum of achievable outputs that changes with the ambient conditions; the model replaces it by a few discrete capacity levels and their probabilities; the evidence is the unit's operating record, in which each hour is classified as full output, one of the derated levels, or forced outage. The probability $p_{i,j}$ is estimated as the fraction of the exposure hours that the unit spent in state j, in the same way as the forced outage rate of Section 25.2.1. A derated level that is estimated from few hours carries parameter uncertainty, which Section 26.6.5 separates from the randomness of the operating state.

Industry statistics sometimes fold deratings into a single number, the *equivalent forced outage rate*, in which a derating counts as a fraction of a full outage in proportion to the capacity lost. Such a number preserves the expected available capacity of the unit but not its distribution, and Section 26.2.4 shows that the adequacy indices computed from it differ from those of the multistate model.

26.2 Exact enumeration of a multistate fleet

The three-generator fleet of Chapter 25 is given derated states for two of its units. This section defines the fleet, enumerates its eighteen joint states, aggregates them into the exact distribution of the available capacity, and computes the adequacy indices at a demand of 80 MW and at other demands.

26.2.1 The multistate fleet

Table 26.2 defines the fleet. Each row is one generator, with its rating, its available-capacity states and their probabilities. G1 keeps the two states of Chapter 25. G2 and G3 have a derated state at half their rating. The probability that G2 is not at full output is $0.10 + 0.05 = 0.15$, and that G3 is not at full output is $0.15 + 0.05 = 0.20$: these are the forced outage rates of Table 25.1, which is therefore the same fleet with every departure from full output counted as a full outage. The values are illustrative.

Table 26.2: The multistate fleet (illustrative). The installed capacity is 100 MW.

Generator	Rating K_i	Available-capacity states $c_{i,j}$	State probabilities $p_{i,j}$	$\mathbb{E}[C_i]$
G1	50 MW	50, 0 MW	0.90, 0.10	45.0 MW
G2	30 MW	30, 15, 0 MW	0.85, 0.10, 0.05	27.0 MW
G3	20 MW	20, 10, 0 MW	0.80, 0.15, 0.05	17.5 MW
Fleet	100 MW			89.5 MW

The operating states of the three generators are taken as independent at first; Section 26.4 relaxes this. The other assumptions of Section 25.2.3 are kept: no imports, no demand response or storage, no transmission constraints, and all available capacity dispatchable. A joint state is the vector of the units' state indices, $\mathbf{j} = (j_1, j_2, j_3)$, and the number of joint states is the product of the numbers of unit states,

$$N_{\text{states}} = \prod_{i=1}^{3} m_i = 2 \times 3 \times 3 = 18.$$

Under independence the probability of a joint state is the product of the unit probabilities, and its available capacity is the sum of the unit capacities,

$$\Pr(\mathbf{j}) = \prod_{i=1}^{n} p_{i,j_i}, \qquad C(\mathbf{j}) = \sum_{i=1}^{n} c_{i,j_i}, \tag{26.5}$$

with $n = 3$ here; the first expression is equation (25.9) of the two-state fleet, generalized to any number of states per unit.

26.2.2 Representative states and their aggregation

Four joint states show the arithmetic. With every generator at full output,

$$C = 50 + 30 + 20 = 100\,\text{MW}, \qquad \Pr = 0.90 \times 0.85 \times 0.80 = 0.612.$$

With G2 derated and the others at full output,

$$C = 50 + 15 + 20 = 85\,\text{MW}, \qquad \Pr = 0.90 \times 0.10 \times 0.80 = 0.072.$$

With G1 unavailable and the others at full output,

$$C = 0 + 30 + 20 = 50\,\text{MW}, \qquad \Pr = 0.10 \times 0.85 \times 0.80 = 0.068.$$

A capacity of 50 MW is also produced by a second combination, G1 at full output with G2 and G3 both unavailable,

$$C = 50 + 0 + 0 = 50\,\text{MW}, \qquad \Pr = 0.90 \times 0.05 \times 0.05 = 0.00225.$$

The two joint states are mutually exclusive, so the probability of the capacity level is their sum,

$$\Pr(C = 50\,\text{MW}) = 0.068 + 0.00225 = 0.07025.$$

Adequacy at a single bus depends only on the total available capacity, so joint states with the same total are aggregated in this way.

Table 26.3 lists all eighteen joint states. The left half holds the nine states with G1 available, the right half the nine with G1 unavailable; in each half the columns are the capacities of the three units in that state, the total C, and the probability from equation (26.5). The only pair with equal totals is the one just computed, at 50 MW, so the eighteen joint states give seventeen distinct capacity levels.

Table 26.3: The eighteen joint states of the multistate fleet, with unit capacities in MW and probabilities from equation (26.5) (independent units).

c_1	c_2	c_3	C	Probability	c_1	c_2	c_3	C	Probability
50	30	20	100	0.61200	0	30	20	50	0.06800
50	30	10	90	0.11475	0	30	10	40	0.01275
50	30	0	80	0.03825	0	30	0	30	0.00425
50	15	20	85	0.07200	0	15	20	35	0.00800
50	15	10	75	0.01350	0	15	10	25	0.00150
50	15	0	65	0.00450	0	15	0	15	0.00050
50	0	20	70	0.03600	0	0	20	20	0.00400
50	0	10	60	0.00675	0	0	10	10	0.00075
50	0	0	50	0.00225	0	0	0	0	0.00025
Subtotal, G1 available				0.90000	Subtotal, G1 unavailable				0.10000

26.2.3 The exact available-capacity distribution

Aggregating the eighteen joint states by their total capacity gives the exact distribution of the available capacity, which in outage form is the COPT of the fleet. Table 26.4 lists it. Each row is one capacity level: the available capacity c, the capacity on outage $x = 100 - c$, the probability $\Pr(C = c) = \Pr(X = x)$, the cumulative probability $\Pr(X \geq x)$ obtained by summing the probability column from the bottom up, and the number of joint states aggregated into the level. The probabilities sum to 1.00000.

Figure 26.2 draws the distribution. The horizontal axis is the available capacity in MW; the vertical axis is the probability of each level, on a logarithmic scale so that the small probabilities of the low-capacity levels remain visible. Green bars are the levels at or above 80 MW, red bars the levels below it. The probability mass at 100 MW dominates, 0.612, and the next largest masses sit at 90 and 85 MW, the deratings of G3 and G2 alone. The adequacy risk, however, is determined by the lower-capacity tail: the red bars, whose total is the loss-of-load probability at 80 MW, and in particular the mass of 0.07025 at 50 MW, where G1 is out.

Two checks confirm the table. The probabilities sum to one. The mean of the distribution is

$$\mathbb{E}[C] = \sum_{c} c\,\Pr(C = c) = 89.5\,\text{MW},$$

Table 26.4: Exact available-capacity distribution and capacity outage probability table of the multistate fleet (independent units). The cumulative column is the probability that the capacity on outage is at least x.

Available c	Outage x	$\Pr(C = c) = \Pr(X = x)$	$\Pr(X \geq x)$	Joint states
100 MW	0 MW	0.61200	1.00000	1
90 MW	10 MW	0.11475	0.38800	1
85 MW	15 MW	0.07200	0.27325	1
80 MW	20 MW	0.03825	0.20125	1
75 MW	25 MW	0.01350	0.16300	1
70 MW	30 MW	0.03600	0.14950	1
65 MW	35 MW	0.00450	0.11350	1
60 MW	40 MW	0.00675	0.10900	1
50 MW	50 MW	0.07025	0.10225	2
40 MW	60 MW	0.01275	0.03200	1
35 MW	65 MW	0.00800	0.01925	1
30 MW	70 MW	0.00425	0.01125	1
25 MW	75 MW	0.00150	0.00700	1
20 MW	80 MW	0.00400	0.00550	1
15 MW	85 MW	0.00050	0.00150	1
10 MW	90 MW	0.00075	0.00100	1
0 MW	100 MW	0.00025	0.00025	1
Total		1.00000		18

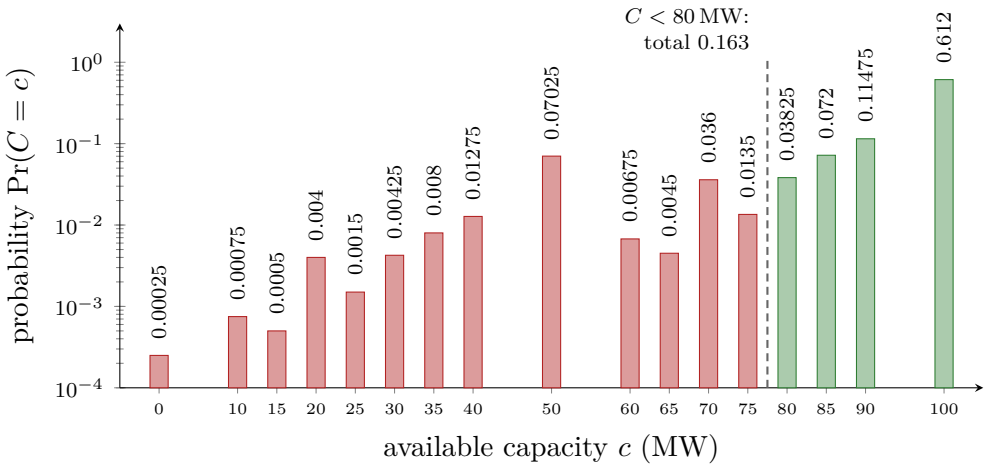


Figure 26.2: Exact available-capacity distribution of the multistate fleet of Table 26.2 (Table 26.4; illustrative), logarithmic vertical scale. Red bars: levels below a demand of 80 MW, whose total is the loss-of-load probability 0.163; green bars: levels at or above it.

which equals the sum of the unit means of Table 26.2, $45.0 + 27.0 + 17.5 = 89.5$ MW, as linearity of expectation requires for any dependence between the units. The two-state fleet of Chapter 25 had $\mathbb{E}[C] = 86.5$ MW: counting the deratings as full outages removes 3 MW of expected capacity, 0.10×15 from G2 and 0.15×10 from G3.

26.2.4 Loss-of-load probability and expected shortfall at 80 MW

Let the demand be constant, $L = 80$ MW. The fleet is inadequate when $C < 80$ MW, and the loss-of-load probability is the sum of the probabilities of the levels below 80 MW,

$$\begin{aligned}\mathrm{LOLP}(80) = {} & 0.01350 + 0.03600 + 0.00450 + 0.00675 + 0.07025 + 0.01275 \\ & + 0.00800 + 0.00425 + 0.00150 + 0.00400 + 0.00050 + 0.00075 + 0.00025 = 0.163,\end{aligned}$$

or 16.3 %. It is the cumulative entry $\Pr(X \geq 25) = 0.16300$ of Table 26.4, since $C < 80$ is $X > 20$ and the next outage level above 20 MW is 25 MW. The expected shortfall is

$$\mathbb{E}[S] = \sum_{c} \Pr(C = c)\,(80 - c)^{+}, \tag{26.6}$$

which, term by term from $c = 75$ MW down to $c = 0$, is

$$\begin{aligned}\mathbb{E}[S] = {} & 0.01350 \times 5 + 0.03600 \times 10 + 0.00450 \times 15 + 0.00675 \times 20 + 0.07025 \times 30 \\ & + 0.01275 \times 40 + 0.00800 \times 45 + 0.00425 \times 50 + 0.00150 \times 55 + 0.00400 \times 60 \\ & + 0.00050 \times 65 + 0.00075 \times 70 + 0.00025 \times 80 \\ = {} & 0.0675 + 0.36 + 0.0675 + 0.135 + 2.1075 + 0.51 + 0.36 \\ & + 0.2125 + 0.0825 + 0.24 + 0.0325 + 0.0525 + 0.02 = 4.2475 \text{ MW}.\end{aligned}$$

For a one-hour assessment with the capacity state constant through the hour, EUE $= 4.2475$ MW $\times$ 1 h $= 4.2475$ MWh. The available resources are insufficient with probability 16.3 %, and the probability-weighted expected shortage is 4.2475 MW; the conditional shortfall of equation (25.11) is $\mathbb{E}[S \mid S > 0] = 4.2475/0.163 = 26.06$ MW. The largest single contribution, 2.1075 MW of the 4.2475, comes from the level of 50 MW, the outage of G1.

Table 26.5 compares three models of the same fleet at this demand. Each row is one model, with the expected available capacity, the LOLP and the expected shortfall. The first row is the exact multistate model. The second is the two-state fleet of Chapter 25, in which every derating is a full outage. The third is a two-state model whose outage probabilities are the equivalent forced outage rates of Section 26.1.3: each derating to half capacity counts as half an outage, so that G2 has $0.05 + 0.10 \times 0.5 = 0.10$ and G3 has $0.05 + 0.15 \times 0.5 = 0.125$. The equivalent model reproduces the expected capacity of 89.5 MW but still overstates both indices, the LOLP by 17 % (0.19 against 0.163) and the expected shortfall by 10 % (4.675 against 4.2475 MW), because it replaces each derating of G2 by half its probability of a full outage, and a full outage of G2 takes the fleet below 80 MW while its derating alone does not. The full-outage model overstates them by 44 and 35 %.

26.2.5 Dependence on demand

Table 26.6 repeats the calculation at other demands. Each row is one demand level; the columns are the LOLP, the expected shortfall and the conditional shortfall of the multistate fleet, and, for comparison, the LOLP and expected shortfall of the two-state fleet of Chapter 25. The rows at 60, 75, 80 and 90 MW are the four demands of the original example; the others show the trend.

The row at 60 MW shows the arithmetic: the levels below 60 MW are those from 50 MW down, whose probabilities sum to 0.10225, and the expected shortfall is $0.07025 \times 10 + 0.01275 \times$

Table 26.5: Three models of the same fleet at a demand of 80 MW (illustrative).

Model	$\mathbb{E}[C]$	LOLP	$\mathbb{E}[S]$
Multistate, Table 26.2 (exact)	89.5 MW	0.163	4.2475 MW
Two-state, derating counted as full outage (FOR 0.10, 0.15, 0.20)	86.5 MW	0.235	5.7400 MW
Two-state, equivalent outage rates (0.10, 0.10, 0.125)	89.5 MW	0.190	4.6750 MW

Table 26.6: Adequacy indices of the multistate fleet and of the two-state fleet of Chapter 25 at several demands (independent units, illustrative).

	Multistate fleet			Two-state fleet	
Demand	LOLP	$\mathbb{E}[S]$	$\mathbb{E}[S \mid S>0]$	LOLP	$\mathbb{E}[S]$
50 MW	0.03200	0.5500 MW	17.19 MW	0.0320	0.850 MW
60 MW	0.10225	1.5725 MW	15.38 MW	0.1270	2.120 MW
70 MW	0.11350	2.6850 MW	23.66 MW	0.1270	3.390 MW
75 MW	0.14950	3.4325 MW	22.96 MW	0.2350	4.565 MW
80 MW	0.16300	4.2475 MW	26.06 MW	0.2350	5.740 MW
85 MW	0.20125	5.2538 MW	26.11 MW	0.3880	7.680 MW
90 MW	0.27325	6.6200 MW	24.23 MW	0.3880	9.620 MW
100 MW	0.38800	10.5000 MW	27.06 MW	0.3880	13.500 MW

$20+0.008\times25+0.00425\times30+0.0015\times35+0.004\times40+0.0005\times45+0.00075\times50+0.00025\times60 = 1.5725$ MW. Equation (25.17) gives the same numbers more quickly: the expected shortfall $g(\ell) = \mathbb{E}[(\ell - C)^+]$ at demand ℓ is the integral of the LOLP from zero to ℓ, so between consecutive capacity levels it rises with slope equal to the LOLP. From 75 to 80 MW the LOLP is 0.163, and $g(80) = g(75) + 5 \times 0.163 = 3.4325 + 0.815 = 4.2475$ MW.

Figure 26.3 draws both indices against the demand for the two fleets. In the left panel the horizontal axis is the demand ℓ from 40 to 100 MW and the vertical axis the LOLP; both curves are step functions that jump just above each capacity level. In the right panel the vertical axis is the expected shortfall $g(\ell)$, a continuous, piecewise linear and convex curve whose slope is the left panel. The dots mark the four demands of the original example on the multistate curves.

The LOLP changes sharply because the capacity distribution is discrete: it is constant between consecutive capacity levels and jumps by the probability of a level as the demand crosses it. Between 80 and 90 MW it rises from 0.163 to 0.27325, since the levels 80 and 85 MW, with probabilities 0.03825 and 0.072, become deficient. The expected shortfall rises continuously with the demand, because even within a fixed set of deficient states the size of each shortfall grows as the demand rises; it is the integral of the LOLP. The multistate fleet has more capacity levels than the two-state fleet, so its LOLP rises in more and smaller steps, and it lies below the two-state curves at every demand shown: a model that counts deratings as full outages is conservative for this fleet. The conditional shortfall does not move monotonically with the demand, as the third column of Table 26.6 shows, because each new deficient level enters with a small shortfall and lowers the average deficit.

26.3 Recursive convolution and probability generating functions

Listing eighteen joint states is easy; listing the joint states of a fleet of a hundred multistate units is impossible. This section derives the multistate form of the recursion of Section 25.4.5,

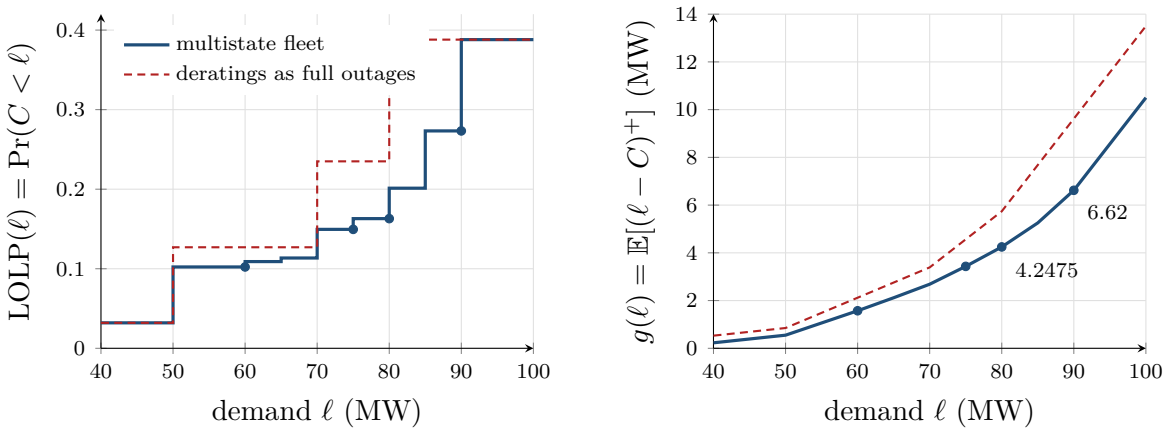


Figure 26.3: Loss-of-load probability (left) and expected shortfall (right) against demand for the multistate fleet of Table 26.2 (solid) and the same fleet with every derating counted as a full outage (dashed), from Table 26.6 (illustrative). Dots: the demands 60, 75, 80 and 90 MW. The expected shortfall is the integral of the LOLP over the demand.

applies it to the fleet, states it as an algorithm, recasts it as a product of generating functions, and states its computational cost and its limits.

26.3.1 The multistate recursion

Let X_n be the total capacity on outage of the first n units of a fleet and

$$p_n(x) = \Pr(X_n = x)$$

its distribution, the COPT of those n units. Unit $n+1$ has outage states $d_{n+1,1}, \dots, d_{n+1,m}$, in MW, with probabilities $p_{n+1,1}, \dots, p_{n+1,m}$, and is independent of the first n units. (The fleet distribution is a function, $p_n(\cdot)$; a unit's state probability carries two indices, $p_{i,j}$.) The COPT of the enlarged fleet is

$$p_{n+1}(x) = \sum_{j=1}^{m} p_{n+1,j}\; p_n(x - d_{n+1,j}), \tag{26.7}$$

with $p_n(y) = 0$ for every outage level y that the first n units cannot produce, negative y in particular. The recursion starts from the empty fleet, which has no capacity on outage with certainty,

$$p_0(0) = 1,$$

and adds the units one at a time; after the last unit, p_n is the COPT of the fleet. Each new unit contributes its outage distribution to the existing fleet distribution: the probability at each existing level is copied to that level plus each of the unit's outage values, weighted by the probability of that value. In the literature on generation adequacy the procedure is called the *unit addition algorithm*.

For a two-state unit with rating K and forced outage rate U, the outage takes the values $d \in \{0, K\}$ with probabilities $1 - U$ and U, and equation (26.7) reduces to

$$p_{n+1}(x) = (1 - U)\, p_n(x) + U\, p_n(x - K),$$

which is equation (25.19) of Chapter 25. Summing equation (26.7) over all levels at or above x gives the same recursion for the cumulative column,

$$\Pr(X_{n+1} \geq x) = \sum_{j=1}^{m} p_{n+1,j}\; \Pr(X_n \geq x - d_{n+1,j}), \tag{26.8}$$

with $\Pr(X_n \geq y) = 1$ for $y \leq 0$.

26.3.2 Convolution from first principles

Equation (26.7) follows from two facts. The capacity on outage of the enlarged fleet is the sum of that of the first n units and that of the new one,

$$X_{n+1} = X_n + D_{n+1},$$

where D_{n+1} is the new unit's capacity outage. The events $\{D_{n+1} = d_{n+1,j}\}$, $j = 1, \ldots, m$, are mutually exclusive and exhaustive, so by the law of total probability

$$\Pr(X_{n+1} = x) = \sum_{j=1}^{m} \Pr(X_{n+1} = x,\ D_{n+1} = d_{n+1,j}) = \sum_{j=1}^{m} \Pr(X_n = x - d_{n+1,j},\ D_{n+1} = d_{n+1,j}),$$

where the second step uses $X_n = X_{n+1} - D_{n+1}$. Under independence of X_n and D_{n+1} the joint probability factors,

$$\Pr(X_n = x - d_{n+1,j},\ D_{n+1} = d_{n+1,j}) = p_n(x - d_{n+1,j})\, p_{n+1,j},$$

which gives equation (26.7). A sum of the form $\sum_y a(y)\, b(x - y)$ is the *discrete convolution* of the sequences a and b, and the distribution of a sum of independent discrete random variables is the convolution of their distributions.

The independence assumption is essential: it is the only step at which the joint probability is replaced by a product. If the new unit's state is correlated with the states of the units already in the fleet, the joint probability does not factor, and convolving the marginal distributions gives a wrong distribution, in general one with too little probability in the joint outages that a shared cause produces. Section 26.4 restores the convolution under a weaker assumption, independence given a shared state.

26.3.3 Adding G2 to the fleet

Table 26.7 gives the states of G2 in both descriptions. Each row is one state, with its available capacity, its capacity outage and its probability.

Table 26.7: The states of generator G2 in available-capacity and outage form.

Available capacity	Capacity outage $d_{2,j}$	Probability $p_{2,j}$
30 MW	0 MW	0.85
15 MW	15 MW	0.10
0 MW	30 MW	0.05

For G2, equation (26.7) reads

$$p_{n+1}(x) = 0.85\, p_n(x) + 0.10\, p_n(x - 15) + 0.05\, p_n(x - 30). \tag{26.9}$$

If the existing fleet has $p_n(50) = 0.10$ and no probability at the outage levels 35 and 20 MW, then after G2 is added

$$p_{n+1}(50) = 0.85 \times 0.10 + 0.10 \times p_n(35) + 0.05 \times p_n(20) = 0.085 + 0 + 0 = 0.085.$$

The other two states of G2 carry the probability at 50 MW to the levels 65 MW, with $0.10 \times 0.10 = 0.010$, and 80 MW, with $0.05 \times 0.10 = 0.005$.

That situation is exactly the fleet consisting of G1 alone. From $p_0(0) = 1$, adding G1 gives $p_1(0) = 0.90$ and $p_1(50) = 0.10$. Adding G2 with equation (26.9) gives the six levels

$$p_2(0) = 0.85 \times 0.90 = 0.765, \quad p_2(15) = 0.10 \times 0.90 = 0.090, \quad p_2(30) = 0.05 \times 0.90 = 0.045,$$
$$p_2(50) = 0.85 \times 0.10 = 0.085, \quad p_2(65) = 0.10 \times 0.10 = 0.010, \quad p_2(80) = 0.05 \times 0.10 = 0.005,$$

which sum to one. Figure 26.4 traces this step. The left column is the COPT of G1 alone, the right column the COPT after G2 is added; each box is one outage level with its probability, placed vertically by the outage level. Each box on the left sends three arrows, one per state of G2: a solid arrow to the same level with weight 0.85, a dashed arrow to the outage level 15 MW higher (drawn lower) with weight 0.10, and a dotted arrow to the level 30 MW higher with weight 0.05.

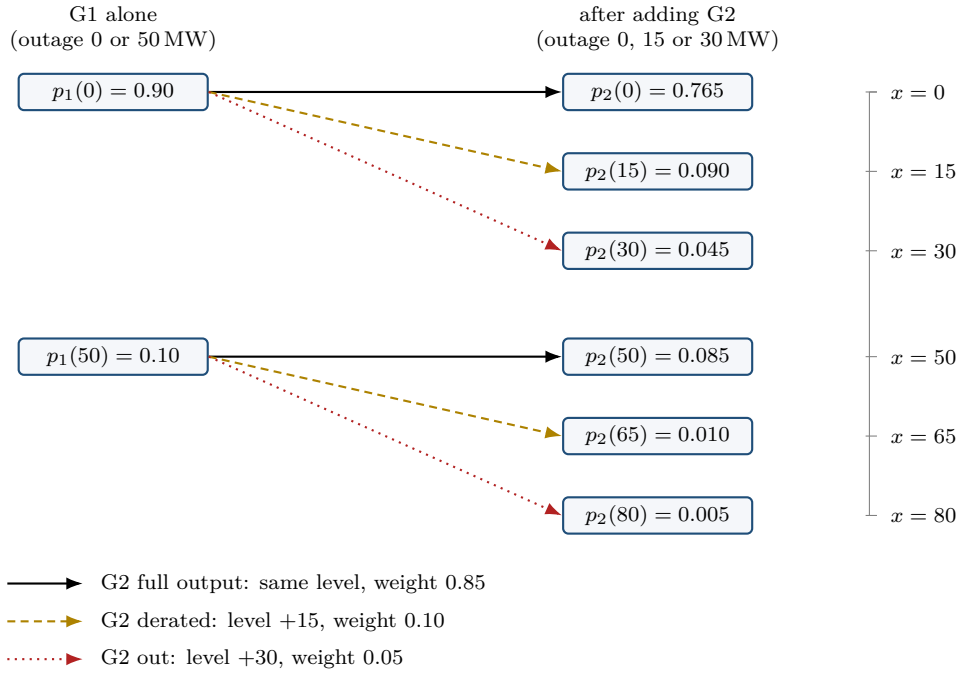


Figure 26.4: One step of the unit addition algorithm, equation (26.9): the three-state generator G2 is added to the COPT of G1. Boxes are placed vertically by the outage level x (MW), marked on the right (illustrative fleet of Table 26.2).

In this step no two arrows meet, because the outage levels of G1, 0 and 50 MW, are farther apart than the largest outage of G2. Adding G3, with outages 0, 10 and 20 MW and probabilities 0.80, 0.15 and 0.05, uses $p_3(x) = 0.80\,p_2(x) + 0.15\,p_2(x - 10) + 0.05\,p_2(x - 20)$, and here two contributions meet at the level of 50 MW:

$$\begin{aligned} p_3(50) &= 0.80\,p_2(50) + 0.15\,p_2(40) + 0.05\,p_2(30) \\ &= 0.80 \times 0.085 + 0 + 0.05 \times 0.045 = 0.068 + 0.00225 = 0.07025, \end{aligned}$$

the two joint states of Section 26.2.2, G1 out alone and G2 and G3 both out. Two further levels show the pattern: $p_3(0) = 0.80 \times 0.765 = 0.612$, and $p_3(25) = 0.15\,p_2(15) = 0.15 \times 0.090 = 0.0135$, G2 and G3 both derated. The six levels of p_2 times the three states of G3 give eighteen products, which aggregate into the seventeen levels of Table 26.4.

26.3.4 The algorithm

A direct implementation stores the COPT as a list of pairs (outage level in MW, probability), keyed by the outage level, so that only the levels that occur are kept. The procedure is:

1. Initialize the table with the single entry $(0, 1)$: the empty fleet has no capacity on outage.
2. For each unit $i = 1, \dots, n$, given as a list of its outage states $(d_{i,j}, p_{i,j})$, create an empty table.
3. For every entry (x, p) of the current table and every state $(d_{i,j}, p_{i,j})$ of the unit, add $p \times p_{i,j}$ to the entry of the new table at the level $x + d_{i,j}$, creating the entry with probability zero if it does not exist.
4. Replace the current table by the new one and continue with the next unit.
5. After the last unit, check that the probabilities sum to one within round-off, and form the cumulative column by summing from the largest outage level downward.

For the fleet of Table 26.2, with G1 given as the states $(0, 0.90)$ and $(50, 0.10)$, G2 as $(0, 0.85)$, $(15, 0.10)$ and $(30, 0.05)$, and G3 as $(0, 0.80)$, $(10, 0.15)$ and $(20, 0.05)$, the procedure returns the seventeen entries of Table 26.4, whose sum is one to machine precision. Step 3 is where states with the same total are merged: the table never holds more entries than there are distinct outage levels, however many joint states lead to them.

26.3.5 Probability generating functions

Suppose every capacity outage is an integer multiple of a capacity step ΔK (written δ in Section 25.4.4). The *probability generating function* of unit i is the polynomial

$$G_i(\zeta) = \mathbb{E}\big[\zeta^{D_i/\Delta K}\big] = \sum_{j=1}^{m_i} p_{i,j}\, \zeta^{d_{i,j}/\Delta K}, \tag{26.10}$$

in which the coefficient of ζ^k is the probability that the unit's outage is $k\,\Delta K$. The variable ζ is a formal (dummy) variable with no physical meaning, written ζ in this chapter so that z remains free for the values of the weather state of Section 26.4; G_i is a function, not the name of generator Gi. For a fleet of independent units, the generating function of the total outage $X = \sum_i D_i$ is the product of the unit functions,

$$G_{\text{fleet}}(\zeta) = \mathbb{E}\big[\zeta^{X/\Delta K}\big] = \mathbb{E}\Big[\prod_i \zeta^{D_i/\Delta K}\Big] = \prod_i \mathbb{E}\big[\zeta^{D_i/\Delta K}\big] = \prod_i G_i(\zeta), \tag{26.11}$$

where the third equality uses the fact that the expectation of a product of independent random variables is the product of their expectations. The coefficient of ζ^k in G_{fleet} is the probability that the total outage equals $k\,\Delta K$. Multiplying two polynomials collects, for each power k, the products of the coefficients whose powers add to k, which is the convolution of the two coefficient sequences: equation (26.11) is the recursion of equation (26.7) written as polynomial multiplication.

As an example, take two two-state 10 MW generators, each unavailable with probability 0.1, and $\Delta K = 10$ MW. Each has $G_i(\zeta) = 0.9 + 0.1\,\zeta$, and the fleet has

$$G(\zeta) = (0.9 + 0.1\,\zeta)^2 = 0.81 + 0.18\,\zeta + 0.01\,\zeta^2.$$

The total outage is therefore 0 MW with probability 0.81, 10 MW with probability 0.18 and 20 MW with probability 0.01, the distribution obtained by enumerating the four joint states (0.9^2; two states with one unit out, $2 \times 0.9 \times 0.1$; 0.1^2). For the multistate fleet, with $\Delta K = 5$ MW, $G_{\text{G1}}(\zeta) = 0.90 + 0.10\,\zeta^{10}$ and $G_{\text{G2}}(\zeta) = 0.85 + 0.10\,\zeta^3 + 0.05\,\zeta^6$; their product is $0.765 + 0.09\,\zeta^3 + 0.045\,\zeta^6 + 0.085\,\zeta^{10} + 0.01\,\zeta^{13} + 0.005\,\zeta^{16}$, whose coefficients are the six levels of p_2 at 0, 15, 30, 50, 65 and 80 MW. Differentiating equation (26.11) at $\zeta = 1$ gives the mean, $\mathbb{E}[X] = \Delta K\, G'_{\text{fleet}}(1)$, a quick check of a computed table.

26.3.6 Computational complexity

Direct enumeration of a fleet of n units with m states each lists m^n joint states, so its cost grows as $O(m^n)$. For $n = 100$ three-state units that is $3^{100} \approx 5.2 \times 10^{47}$ states, against $2^{100} \approx 1.3 \times 10^{30}$ for two-state units; neither can be listed. Convolution instead works on the distinct values of the total capacity. If the capacities lie on a grid of step ΔK, the outage of the fleet takes at most $N_{\text{bin}} = K_{\text{tot}}/\Delta K + 1$ values, and each unit addition multiplies every entry of the table by each of the unit's m states, so a direct dynamic-programming implementation needs about

$$n \cdot m \cdot N_{\text{bin}} \tag{26.12}$$

multiplications and additions, which grows linearly in the number of units. A fleet of 100 three-state units with 10 000 MW installed on a 1 MW grid needs about $100 \times 3 \times 10\,001 \approx 3 \times 10^6$ operations, which is immediate on a computer. The exact cost depends on the sparsity of the table (only the levels that occur need be stored), the implementation, the capacity grid, and the number of distinct states.

For very large grids, the repeated convolutions can be performed as polynomial multiplications with the fast Fourier transform, whose cost for two sequences of length N_{bin} grows as $N_{\text{bin}} \log N_{\text{bin}}$ rather than N_{bin}^2. The method has numerical costs of its own: floating-point round-off can produce small negative or spurious probabilities in the far tail, which is where the adequacy risk lies, and the transform computes a circular convolution, so the sequences must be padded with zeros to the length of the result or the high outage levels wrap around onto the low ones. The central computational advantage, by either method, is that states are aggregated by total capacity instead of retaining every combination of component states.

26.3.7 State aggregation and lost information

Aggregation discards the identity of the units behind a capacity level. Suppose two distinct fleet configurations both leave $C = 70$ MW; in the multistate fleet the level of 50 MW is such a case, reached with G1 out alone or with G2 and G3 both out. For copper-plate adequacy against a demand the two configurations are equivalent: the LOLP and the expected shortfall depend only on the total. They need not be equivalent for transmission congestion, reactive power and voltage support, ramping capability, fuel constraints, restoration time, geographic deliverability or contingency security. Section 25.5.2 gave a two-area example in which the location of the unit on outage decided whether a state was deficient, and in such a model total MW is not a sufficient description of the state.

> **Principle 26.1.** *Two states can be merged only when they are equivalent for the quantities that must be calculated. For generation-only LOLP and EUE under a copper-plate model, the total available capacity is a sufficient description of the state; for network-constrained reliability it is not.*

26.4 Correlated generator outages

The recursion, the generating function and the enumeration of Section 26.2 all multiply marginal probabilities. This section states why that fails for real fleets, restores an exact calculation when the dependence acts through a shared state, recomputes the indices of the multistate fleet under a common shock, and shows that dependence can move the probability and the size of shortfalls in opposite directions.

26.4.1 Why independence fails

Generating units share exposure to common causes: interruptions of a shared fuel supply; cooling-water restrictions in a drought or a heat wave; extreme heat or cold; floods and storms; common electrical infrastructure such as a substation or a transmission connection; and common control systems. These are the coupling factors of common-cause failure of Chapter 13 of Volume I, and Section 25.4.1 described them for generating fleets. When they act, the joint distribution of the unit capacities is not the product of the marginal ones,

$$\Pr(C_1 = c_1, C_2 = c_2, \ldots, C_n = c_n) \neq \prod_{i=1}^{n} \Pr(C_i = c_i),$$

and the convolution of the marginal distributions is no longer valid. The dependence structure must be represented in the model; the marginal probabilities alone do not determine the joint distribution, as the Fréchet bounds of Chapter 13 (Section 13.2.3) showed for two components.

26.4.2 Conditional independence through a latent state

A shared operating or environmental state Z, for example

$$Z \in \{\text{normal},\ \text{extreme weather}\},$$

conditions the outage probabilities of several units at once. Suppose the units are independent given Z, so that within each value z of the shared state their joint distribution factors. By the law of total probability, conditioning on Z,

$$\Pr(\boldsymbol{C} = \mathbf{c}) = \sum_{z} \Pr(Z = z)\ \Pr(\boldsymbol{C} = \mathbf{c} \mid Z = z) = \sum_{z} \Pr(Z = z) \prod_{i=1}^{n} \Pr(C_i = c_i \mid Z = z), \tag{26.13}$$

where $\boldsymbol{C} = (C_1, \ldots, C_n)$ is the vector of unit capacities and $\mathbf{c}$ a value of it. This is a *mixture* of conditionally independent models, the construction of Chapter 13 (Section 13.4.1) applied to a whole fleet, with Z the weather state of that chapter. Within each value of Z the units are independent, so the conditional COPT can be built by the recursion; the dependence enters only through the average over Z. The workflow is:

1. define the possible common environmental states z;
2. assign their probabilities $\Pr(Z = z)$;
3. specify the state distribution of each unit conditional on each environmental state, $\Pr(C_i = c \mid Z = z)$;
4. build a COPT within each environmental state with the recursion of equation (26.7);
5. average the conditional COPTs with the environmental-state probabilities, which is equation (26.13) summed over the joint states with the same total.

Since the LOLP and the expected shortfall are expectations of functions of the total capacity, the law of total expectation gives them in the same mixed form,

$$\text{LOLP} = \sum_{z} \Pr(Z = z)\, \text{LOLP}_z, \qquad \mathbb{E}[S] = \sum_{z} \Pr(Z = z)\, \mathbb{E}[S \mid Z = z], \tag{26.14}$$

where LOLP_z is the loss-of-load probability computed from the conditional COPT of state z.

The mechanism is the one of Chapter 13. Extreme weather can raise the outage probability of many units together; within each weather state the units may still have approximately independent residual failure mechanisms; but mixing over the weather states creates positive dependence in the combined distribution, because the conditional outage probabilities of all units move up together in the extreme state. The units are dependent marginally although they are independent given the shared environment.

Figure 26.5 draws the construction for the common shock of the next subsection, the simplest latent state. The box at the top is the shared state with its two values and their probabilities; the two boxes in the middle are the conditional COPTs and their indices at 80 MW; the box at the bottom is the mixture, whose indices are the weighted averages of equation (26.14).

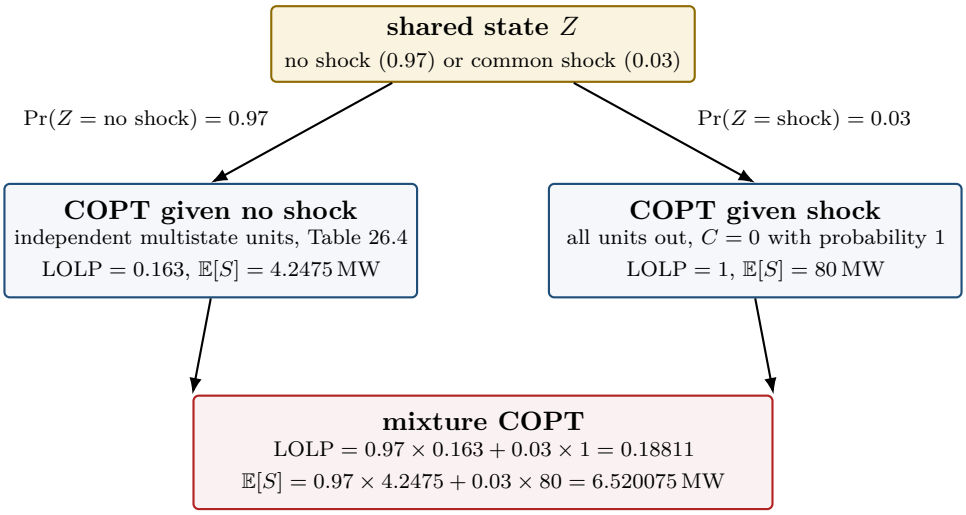


Figure 26.5: A common shock as a latent state of the multistate fleet at a demand of 80 MW. Given the state, the units are independent and each conditional COPT is built by convolution; the unconditional indices are the probability-weighted averages of the conditional ones, equation (26.14) (illustrative).

26.4.3 A common shock in the multistate fleet

As a hypothetical illustration, let a common shock occur during the assessment interval with probability $q_{\text{cs}} = 0.03$ and disable all three generators; when it does not occur, the independent multistate model of Table 26.4 applies. The shared state has two values, and equation (26.14) becomes equation (25.18) of Chapter 25. At $L = 80$ MW the independent model gives $\text{LOLP}_{\text{ind}} = 0.163$ and $\mathbb{E}[S]_{\text{ind}} = 4.2475$ MW. Under the shock $C = 0$, so a shortfall is certain and equals the whole demand. Hence

$$\text{LOLP} = q_{\text{cs}} \times 1 + (1 - q_{\text{cs}})\,\text{LOLP}_{\text{ind}} = 0.03 + 0.97 \times 0.163 = 0.03 + 0.15811 = 0.18811,$$

or 18.811 %, and

$$\mathbb{E}[S] = q_{\text{cs}} \times 80 + (1 - q_{\text{cs}})\,\mathbb{E}[S]_{\text{ind}} = 0.03 \times 80 + 0.97 \times 4.2475 = 2.4 + 4.120075 = 6.520075\,\text{MW},$$

which over a one-hour interval is an expected unserved energy of 6.520075 MWh.

Table 26.8 compares the models; its third row is discussed below. Each row is one outage model, with the LOLP, the expected shortfall, the conditional shortfall and the probability that no capacity is available. The shock raises the LOLP by 15.4 % (0.18811/0.163) and the expected shortfall by 53.5 % (6.520075/4.2475), because it produces the most extreme outcome, the loss of all 80 MW of demand; the conditional shortfall rises from 26.06 to 6.520075/0.18811 = 34.66 MW.

This comparison adds a shock on top of the existing unit-state distributions, so the marginal unavailability of each unit also rises, and it is not a controlled comparison at fixed marginal availability. Under the shock model the full-outage probabilities become $0.03 + 0.97 \times 0.10 = 0.127$ for G1, $0.03 + 0.97 \times 0.05 = 0.0785$ for G2 and 0.0785 for G3, and the other state probabilities are multiplied by 0.97 (G2 derated 0.097, G3 derated 0.1455). The third row of Table 26.8 applies these marginals to independent units. At the same marginals, independence gives a *larger* LOLP

Table 26.8: Loss-of-load probability and expected shortfall of the multistate fleet at $L = 80\,\mathrm{MW}$ under three outage models (illustrative).

Model	LOLP	$\mathbb{E}[S]$	$\mathbb{E}[S \mid S > 0]$	$\Pr(C = 0)$
Independent multistate units	0.16300	4.2475 MW	26.06 MW	0.00025
Common shock $q_{\mathrm{cs}} = 0.03$ added	0.18811	6.5201 MW	34.66 MW	0.03024
Independent, marginals of the shock model	0.21450	5.7320 MW	26.72 MW	0.00078

than the shock model, 0.2145 against 0.18811, and a *smaller* expected shortfall, 5.732 against 6.520 MW; the probability of a total loss of capacity is $0.127 \times 0.0785 \times 0.0785 = 0.00078$ under independence and $0.03 + 0.97 \times 0.00025 = 0.03024$ under the shock, almost forty times larger. Both models have the same expected available capacity, 86.815 MW, which no dependence structure can change.

26.4.4 Dependence at fixed marginal probabilities

A two-unit example isolates the effect of dependence completely. Two 50 MW generators each have unavailability $q = 0.10$, and the demand is 80 MW, so both units must be available for the demand to be met. Under independent outages,

$$\Pr(\text{both out}) = 0.10^2 = 0.01, \qquad \Pr(\text{exactly one out}) = 2 \times 0.10 \times 0.90 = 0.18,$$

the shortage probability is $1 - 0.90^2 = 0.19$, and the expected shortfall is $0.18 \times 30 + 0.01 \times 80 = 5.4 + 0.8 = 6.2$ MW: one unit out leaves 50 MW and a 30 MW shortfall, both out leave nothing. Now let the two units share a perfectly common failure state: they are either both available or both unavailable, each still with marginal unavailability 0.10. Then $\Pr(\text{both out}) = 0.10$, the shortage probability is 0.10, and the expected shortfall is $0.10 \times 80 = 8.0$ MW. Positive dependence has reduced the probability of any shortage, from 19 to 10 %, because failures occur together instead of separately, and increased the expected shortfall, from 6.2 to 8.0 MW, because the total loss of 80 MW has become ten times more likely.

The general relation for this pair follows from the marginals. Let $P_1 = \Pr(\text{exactly one out})$ and $P_2 = \Pr(\text{both out})$. The expected number of units out is $q+q = 2q$ whatever the dependence, by linearity of expectation, and it is also $P_1 + 2P_2$. Hence $P_1 = 2q - 2P_2$, and

$$\mathrm{LOLP} = P_1 + P_2 = 2q - P_2, \qquad \mathbb{E}[S] = 30\,P_1 + 80\,P_2 = 60\,q + 20\,P_2, \tag{26.15}$$

in MW for the second expression. At fixed marginals both indices are linear in the joint outage probability P_2, with opposite signs: every unit of probability moved into the joint outage lowers the LOLP by one unit and raises the expected shortfall by 20 MW. The Fréchet bounds confine P_2 to $[\max(0, 2q-1),\, q] = [0, 0.10]$, so the LOLP lies between 0.10 and 0.20 and the expected shortfall between 6.0 and 8.0 MW for every dependence structure with these marginals.

Table 26.9 traces the interval with a common-shock model of the pair: a shared event of probability c' takes both units out, and otherwise each fails independently with probability $q_{\mathrm{I}} = (q - c')/(1 - c')$, chosen so that the marginal stays $c' + (1 - c')q_{\mathrm{I}} = 0.10$, the beta-factor construction of Chapter 14 of Volume I. Each row is one value of c', with the independent rate, the probabilities of one and of both units out, the correlation of the two outage indicators, $\rho = (P_2 - q^2)/[q(1-q)]$, and the two indices from equation (26.15). The first row is independence, the last perfect dependence.

Table 26.9: Two 50 MW units with fixed marginal unavailability 0.10 at a demand of 80 MW, under increasing dependence (illustrative).

Shared c'	q_{I}	P_1	P_2	ρ	LOLP	$\mathbb{E}[S]$
0	0.1000	0.1800	0.0100	0	0.1900	6.200 MW
0.02	0.0816	0.1469	0.0265	0.184	0.1735	6.531 MW
0.05	0.0526	0.0947	0.0526	0.474	0.1474	7.053 MW
0.10	0	0	0.1000	1	0.1000	8.000 MW

Principle 26.2. *Dependence can change the frequency and the severity of shortfalls in different directions. At fixed marginal outage probabilities, positive dependence moves probability from partial outages into joint ones; it can lower the loss-of-load probability while raising the expected unserved energy, so correlation must not be assumed to increase every risk index.*

Whether positive dependence raises or lowers the LOLP depends on the demand threshold and on the dependence structure. In the pair above the threshold requires both units, so joint outages, which concentrate the failures, lower the LOLP. If the demand were 40 MW, one unit would suffice, the shortage event would be the joint outage itself, and positive dependence would raise the LOLP from 0.01 to as much as 0.10.

26.5 Monte Carlo adequacy simulation

Exact methods need discrete, independent or conditionally independent states that can be aggregated by total capacity. This section develops sampling as the alternative: nonsequential sampling of capacity states, the sampling error of its estimates, the number of samples that rare shortfalls require, the designs that reduce that number, and chronological simulation for the time-coupled features that snapshots cannot represent.

26.5.1 Sampling of capacity states

The exact COPT methods work well when resource states are discrete and independent, or independent given a few shared states, and can be aggregated by total capacity. Modern adequacy models may also include wind and solar generation driven by the weather, battery energy constraints, chronological unit commitment, outage and repair durations, transmission constraints, correlated demand and supply, demand response and complex operating actions. Monte Carlo simulation represents these mechanisms through sampled system states or trajectories, at the price of a statistical error, the general method of Chapter 11 of Volume I (Section 11.5.5) applied to adequacy.

A basic *nonsequential* simulation samples one independent system state per trial. For each trial $k = 1, \dots, N$:

1. sample the generator states from their joint distribution;
2. compute the total available capacity $C^{(k)}$;
3. sample the demand $L^{(k)}$, if it is uncertain;
4. compute the shortage indicator and the shortfall,

$$I^{(k)} = \mathbf{1}_{\{C^{(k)} < L^{(k)}\}}, \qquad S^{(k)} = \left(L^{(k)} - C^{(k)}\right)^{+}.$$

The estimates are the sample means,

$$\widehat{\mathrm{LOLP}} = \frac{1}{N}\sum_{k=1}^{N} I^{(k)}, \qquad \widehat{\mathbb{E}[S]} = \frac{1}{N}\sum_{k=1}^{N} S^{(k)}, \qquad \widehat{\mathrm{EUE}} = \widehat{\mathbb{E}[S]}\,\Delta t \tag{26.16}$$

for an interval of duration Δt. Each estimate is unbiased: the expectation of a sample mean of independent draws is the expectation of one draw, $\mathbb{E}[I^{(k)}] = \text{LOLP}$ and $\mathbb{E}[S^{(k)}] = \mathbb{E}[S]$. The procedure is straightforward and flexible: step 1 can draw from any joint distribution, including the mixture of equation (26.13), and step 4 can be replaced by any model of the shortfall, including the minimum load shedding of equation (25.21). Rare adequacy failures, however, may require a very large number of samples.

26.5.2 Standard error and coefficient of variation

The shortage indicator of one trial is a Bernoulli variable, $I^{(k)} \sim \text{Bernoulli}(\text{LOLP})$, with mean LOLP and variance $\text{LOLP}(1 - \text{LOLP})$. For independent trials the variance of a sum is the sum of the variances, so the variance of the mean of N trials is

$$\operatorname{Var}(\widehat{\text{LOLP}}) = \frac{1}{N^2} \sum_{k=1}^{N} \operatorname{Var}(I^{(k)}) = \frac{\text{LOLP}\,(1 - \text{LOLP})}{N}, \tag{26.17}$$

and the standard error, estimated by replacing the unknown LOLP by its estimate, is

$$\text{SE}(\widehat{\text{LOLP}}) \approx \sqrt{\frac{\widehat{\text{LOLP}}\,(1 - \widehat{\text{LOLP}})}{N}}. \tag{26.18}$$

The *coefficient of variation* of the estimate, its standard error relative to the quantity estimated, is

$$\text{CV}(\widehat{\text{LOLP}}) = \frac{\text{SE}}{\text{LOLP}} = \sqrt{\frac{1 - \text{LOLP}}{N\,\text{LOLP}}}. \tag{26.19}$$

For the multistate fleet at 80 MW, $\text{LOLP} = 0.163$, and with $N = 10\,000$ trials

$$\text{SE} = \sqrt{\frac{0.163 \times 0.837}{10\,000}} = \sqrt{1.364 \times 10^{-5}} = 0.00369,$$

a coefficient of variation of 2.27 %. By the central limit theorem the estimate is approximately normal for large N, and a rough 95 % confidence interval has half-width $1.96 \times 0.00369 = 0.00724$, or 0.724 percentage points: the estimate would typically lie between about 15.6 and 17.0 %. Even with ten thousand samples, Monte Carlo has a noticeable statistical error. The exact COPT gives the probability without sampling error, conditional on the component-state probabilities and the model structure being known.

The same argument applies to the expected shortfall, with the variance of the shortfall in place of the Bernoulli variance: $\operatorname{Var}(\widehat{\mathbb{E}[S]}) = \operatorname{Var}(S)/N$. For the multistate fleet at 80 MW, $\mathbb{E}[S^2] = \sum_c \Pr(C = c)\,(80 - c)^{+2} = 144.425\,\text{MW}^2$, so $\operatorname{Var}(S) = 144.425 - 4.2475^2 = 126.38\,\text{MW}^2$ and the standard deviation of one trial's shortfall is 11.24 MW, 2.65 times its mean. With $N = 10\,000$ the standard error of the expected shortfall is $11.24/100 = 0.112$ MW, a coefficient of variation of 2.65 %, slightly larger than that of the LOLP: most trials have no shortfall at all, and the few that do vary in size.

Figure 26.6 shows the behavior of one simulation. The horizontal axis is the number of trials N, on a logarithmic scale from 100 to 100 000; the vertical axis is the LOLP. The horizontal line is the exact value 0.163; the two dashed curves are the band $0.163 \pm 1.96\,\text{SE}$ of equation (26.18), which narrows as $1/\sqrt{N}$; the dots are the running estimate $\widehat{\text{LOLP}}$ of one sequence of pseudo-random samples of the fleet, recorded at ten values of N. The estimate wanders by several percentage points for the first few thousand trials and stays within the band, as about 95 % of such sequences would.

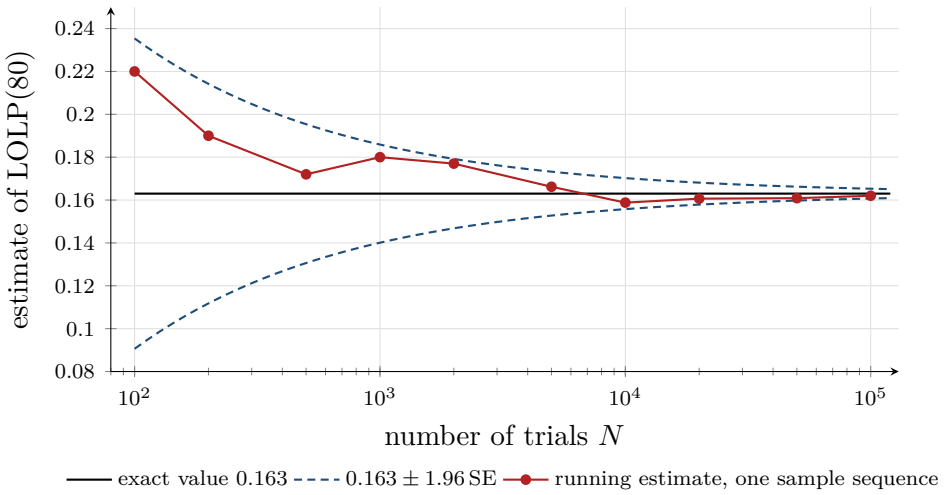


Figure 26.6: Convergence of a nonsequential Monte Carlo estimate of the loss-of-load probability of the multistate fleet at 80 MW. Band: $0.163 \pm 1.96\sqrt{0.163 \times 0.837/N}$. Dots: running estimate of one illustrative pseudo-random sample sequence.

26.5.3 Rare events and the required number of samples

Suppose a large power system has LOLP $= 10^{-4}$ in an hour. With $N = 10\,000$ trials the expected number of shortage observations is $N\,\mathrm{LOLP} = 1$. The simulation may observe zero, one or several shortage cases purely by chance: the number observed is approximately Poisson with mean one, and the probability of observing none is $(1 - 10^{-4})^{10\,000} \approx \mathrm{e}^{-1} = 0.368$. More than a third of such simulations would report a LOLP of zero, and the estimate cannot be precise with that number of trials.

For a target coefficient of variation ϵ, equation (26.19) solved for N gives

$$N = \frac{1 - \mathrm{LOLP}}{\epsilon^2\,\mathrm{LOLP}} \approx \frac{1}{\epsilon^2\,\mathrm{LOLP}} \quad \text{for small LOLP,} \tag{26.20}$$

the approximation using $\mathrm{SE}/\mathrm{LOLP} \approx 1/\sqrt{N\,\mathrm{LOLP}}$. For LOLP $= 10^{-4}$ and $\epsilon = 0.10$,

$$N \approx \frac{1}{0.01 \times 10^{-4}} = 1\,000\,000$$

independent samples are needed to obtain a relative standard error of about 10 % under the idealized Bernoulli model (the exact form gives 999 900): the simulation must observe about a hundred shortfalls, since $1/\epsilon^2 = 100$. Table 26.10 extends the calculation. Each row is one LOLP; the columns are the standard error and coefficient of variation at $N = 10\,000$, the probability that such a simulation observes no shortage at all, and the number of trials for a coefficient of variation of 10 % from the exact form of equation (26.20).

The required number of trials grows in inverse proportion to the probability: each factor of ten in rarity costs a factor of ten in samples for the same relative precision. A system that meets the one-day-in-ten-years criterion of Section 25.6.2, with one to ten shortage hours per shortage day, has an expected 0.1 to 1 h of shortage per year, an average hourly loss-of-load probability of the order of 10^{-5} to 10^{-4}; this is why rare-event methods matter in reliability engineering.

26.5.4 Variance reduction

Several sampling designs lower the variance of the estimate for a given number of trials. Table 26.11 lists the common ones. Each row is one method, with its basic idea and the condition on which its correctness or its benefit depends.

Table 26.10: Sampling error of a nonsequential Monte Carlo estimate of the loss-of-load probability (independent Bernoulli trials).

LOLP	SE at $N = 10^4$	CV at $N = 10^4$	Probability of no shortage in 10^4 trials	N for CV = 10 %
0.163	3.69×10^{-3}	0.023	≈ 0	514
10^{-2}	9.95×10^{-4}	0.099	2.2×10^{-44}	9 900
10^{-3}	3.16×10^{-4}	0.316	4.5×10^{-5}	99 900
10^{-4}	1.00×10^{-4}	1.000	0.368	999 900
10^{-5}	3.16×10^{-5}	3.162	0.905	9 999 900

Table 26.11: Variance-reduction techniques for Monte Carlo adequacy simulation.

Method	Basic idea	Important caution
Importance sampling	Sample high-risk states more frequently and reweight each sample	The weights must correctly account for the changed probabilities
Stratified sampling	Divide the state space into strata and sample each deliberately	The stratum probabilities must be known
Conditional Monte Carlo	Integrate some random variables analytically and sample the rest	The conditional model must be correct
Control variates	Use a correlated quantity with known expectation to correct the estimate	The benefit depends on the correlation
Quasi-Monte Carlo	Use low-discrepancy point sets instead of pseudo-random samples	The error behavior depends on smoothness and dimension
Splitting methods	Progressively replicate trajectories that approach the rare event	The choice of intermediate thresholds affects efficiency

Importance sampling is particularly relevant for rare adequacy events. Let $P(\mathbf{x})$ be the probability of system state $\mathbf{x}$ under the model, $\psi(\mathbf{x})$ the quantity whose expectation is wanted (the shortage indicator or the shortfall), and $\tilde{P}(\mathbf{x})$ a sampling distribution chosen to make the important states more frequent. Multiplying and dividing by $\tilde{P}$,

$$\mathbb{E}_P[\psi(\mathbf{X})] = \sum_{\mathbf{x}} \psi(\mathbf{x})\, P(\mathbf{x}) = \sum_{\mathbf{x}} \psi(\mathbf{x})\, \frac{P(\mathbf{x})}{\tilde{P}(\mathbf{x})}\, \tilde{P}(\mathbf{x}) = \mathbb{E}_{\tilde{P}}\Big[\psi(\mathbf{X})\, \frac{P(\mathbf{X})}{\tilde{P}(\mathbf{X})}\Big], \tag{26.21}$$

provided $\tilde{P}(\mathbf{x}) > 0$ wherever $P(\mathbf{x})\,\psi(\mathbf{x}) \neq 0$, so that no term is lost in the division. The estimator draws the states from $\tilde{P}$ and averages $\psi\, P/\tilde{P}$; the ratio $P/\tilde{P}$ is the likelihood ratio, or weight, of the sample. It is unbiased when properly implemented, and its variance is small when $\tilde{P}$ is roughly proportional to ψP.

As an illustration, the probability that the multistate fleet has no capacity at all is $\Pr(C = 0) = 0.10 \times 0.05 \times 0.05 = 0.00025$. Crude sampling needs $N = (1 - 0.00025)/(0.01 \times 0.00025) \approx 400\,000$ trials for a coefficient of variation of 10 %. If instead each unit's full-outage probability is raised to 0.5 in the sampling distribution (its other states scaled down in proportion), the event has sampling probability $\tilde{P} = 0.5^3 = 0.125$ and weight $0.00025/0.125 = 0.002$. Each sample's weighted indicator then has mean 0.00025 and second moment $0.125 \times 0.002^2 = 5 \times 10^{-7}$, so its variance relative to the squared mean is $5 \times 10^{-7}/0.00025^2 - 1 = 7$, and $N = 7/0.01 = 700$ trials give the same 10 % coefficient of variation, a reduction by a factor of about 570. For an event made of a single state, the ideal sampling distribution puts all its mass on that state and has zero variance; the example shows the mechanism, not a realistic design, which must weight many states at once and is easily made worse than crude sampling by a poor choice of $\tilde{P}$.

26.5.5 Nonsequential sampling and chronological simulation

Nonsequential Monte Carlo samples the marginal distribution of the resource states in one interval. It is suitable for many capacity adequacy calculations, those whose indices are sums of interval expectations, such as the LOLH and the EUE of equation (25.12). It does not preserve the timing and duration of outages. A generator that fails and remains unavailable for 40 h produces forty strongly dependent hours: its state in one hour predicts its state in the next. A chronological (sequential) simulation must model the state transitions.

For a two-state generator with constant failure rate λ and repair rate μ, the transitions are

$$\text{available} \xrightarrow{\ \lambda\ } \text{failed}, \qquad \text{failed} \xrightarrow{\ \mu\ } \text{available},$$

which produce alternating up and down periods with exponential durations of means $1/\lambda$ and $1/\mu$: the continuous-time Markov chain of Chapter 19 of Volume I, whose steady-state unavailability is $\lambda/(\lambda + \mu)$. A unit with mean up time 360 h and mean repair time 40 h (illustrative) has $\lambda = 1/360\,\mathrm{h}^{-1}$, $\mu = 1/40\,\mathrm{h}^{-1}$ and unavailability $40/400 = 0.10$, the forced outage rate of G1. With other failure and repair time distributions the up and down periods are simulated as an alternating renewal process or a semi-Markov process, as in Chapter 18 of Volume I.

Figure 26.7 compares the two kinds of sampling for such a unit over 400 h. The horizontal axis is time in hours. The upper row is one chronological realization: the unit is up (green) except for a single forced outage (red) from hour 150 to hour 190. The lower row samples the state of each hour independently with the same unavailability 0.10, as a nonsequential simulation of hourly snapshots would. Both rows have forty hours down, 10 % of the period; the chronological row has them in one outage event, the independent row in thirty-two separate events, most one hour long.

The two rows give the same expected shortage hours, because the LOLH is a sum of hourly probabilities that does not depend on the dependence between hours (Section 25.3.2). They

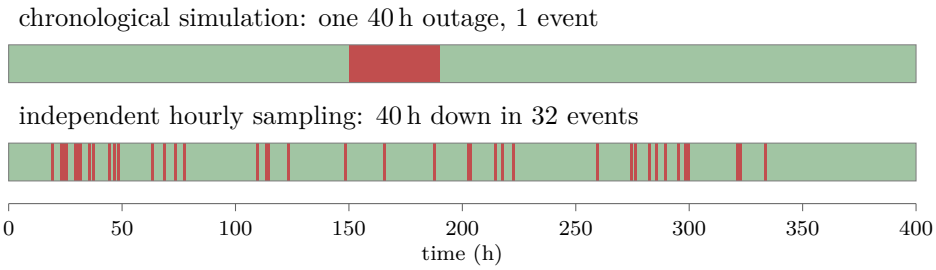


Figure 26.7: The same unavailability, 0.10, sampled chronologically (upper row; mean up time 360 h, mean repair time 40 h) and hour by hour independently (lower row) over 400 h. Green: available; red: on forced outage. Both rows have 40 h on outage; their numbers of outage events differ (illustrative realizations).

give very different numbers of shortage events and shortage days, different event durations, and different demands on any resource that must cover a shortfall for its whole duration. Section 25.6.1 showed the same contrast analytically for the daily LOLE.

26.5.6 Chronological adequacy indices

A chronological simulation steps through each simulated year k hour by hour and records the hourly shortfall $S_t^{(k)}$. The annual shortage hours and the annual unserved energy of year k are

$$N_{\text{short}}^{(k)} = \sum_t \mathbf{1}_{\{S_t^{(k)}>0\}}\,\Delta t, \qquad \text{UE}^{(k)} = \sum_t S_t^{(k)}\,\Delta t, \tag{26.22}$$

in hours and MWh, and across N simulated years the estimates are the sample means,

$$\widehat{\text{LOLH}} = \frac{1}{N}\sum_{k=1}^{N} N_{\text{short}}^{(k)}, \qquad \widehat{\text{EUE}} = \frac{1}{N}\sum_{k=1}^{N} \text{UE}^{(k)}, \tag{26.23}$$

in hours per year and MWh per year. The simulated years are independent replications, so the standard error of each estimate is the sample standard deviation of the annual values divided by $\sqrt{N}$, by the same argument as equation (26.17); the hours within a year are dependent, which is why the year, not the hour, is the unit of replication.

The simulation also yields the indices that only chronology defines. A day is a shortage day if any shortage occurs during it, and the annual count of shortage days averaged over the years estimates the daily LOLE. An event is a maximal run of contiguous shortage intervals, and the annual count of events averaged over the years estimates the loss-of-load event frequency; the run lengths give the distribution of event durations. The event definition must be explicit, particularly when a brief restoration splits what an operator would treat as one disturbance into two events: a rule that merges runs separated by fewer than a stated number of hours changes the count. The whole distribution of the annual values is available, not only its mean, so the probability of a year with more than a given unserved energy can be reported too.

26.5.7 Energy-limited resources

A battery adds a dynamic energy state. With stored energy E_t in MWh at the start of hour t, charging power $P_{\text{c},t}$ and discharging power $P_{\text{d},t}$ in MW, and charging and discharging efficiencies η_{c} and η_{d}, the energy balance of the store is

$$E_{t+1} = E_t + \eta_{\text{c}}\,P_{\text{c},t}\,\Delta t - \frac{P_{\text{d},t}}{\eta_{\text{d}}}\,\Delta t, \tag{26.24}$$

subject to

$$0 \le E_t \le E_{\max}, \qquad 0 \le P_{\mathrm{d},t} \le P_{\max}, \qquad 0 \le P_{\mathrm{c},t} \le P_{\max}.$$

Charging stores less energy than it draws, by the factor η_{c}, and discharging draws more from the store than it delivers, by the factor $1/\eta_{\mathrm{d}}$. A battery with $P_{\max} = 50\,\mathrm{MW}$, $E_{\max} = 100\,\mathrm{MWh}$ and $\eta_{\mathrm{d}} = 0.95$ (illustrative), starting full, delivers 50 MW for $100 \times 0.95/50 = 1.9\,\mathrm{h}$, and covers a 30 MW shortfall for $100 \times 0.95/30 = 3.17\,\mathrm{h}$.

The ability of the battery to serve demand in an hour depends on the preceding charging and discharging decisions, through E_t, so a static capacity model cannot determine its adequacy contribution. In the chronological row of Figure 26.7, with G1's outage leaving a 30 MW shortfall, a full battery covers the first 3.17 h of the 40 h event and nothing afterwards. A snapshot model that treated the battery as available in every shortage hour would cover all forty hours of the independent row and overstate its contribution more than tenfold. Chronological simulation, dynamic programming or dispatch optimization over the hours is necessary when energy constraints materially affect reliability; Section 25.5.3 reached the same conclusion for a single shortfall.

26.6 Choice of computational method and validation of a capacity model

The methods of the previous sections compute the same indices under different assumptions and at different costs. This section compares them, quantifies the error introduced by placing capacities on a coarse grid, and lists the checks a capacity outage model must pass and the sources of uncertainty it must keep apart.

26.6.1 Comparison of the methods

Table 26.12 compares six methods. Each row is one method, with its strength, its weakness and the problems it suits best. No single method is best for every reliability problem; the useful strategy is to match the method to the structure of the model, as Section 11.5.6 of Chapter 11 of Volume I concluded for fault trees.

If the system consists of independent discrete generator states and the demand is known, COPT convolution is often highly efficient. If it contains dependence through a few shared states, the mixture of Section 26.4.2 keeps the calculation exact. If it contains complex operating constraints, storage or event-duration questions, chronological simulation may be necessary. If the estimation of rare shortfall probabilities dominates the cost, hybrid methods, analytical convolution inside each sampled shared state, or variance-reduced Monte Carlo, may be attractive. Under the same model the methods must agree, within the sampling error for Monte Carlo, which makes the exact methods the natural check of a simulation on a reduced version of the system.

26.6.2 State-space reduction and conservative rounding

For large systems even the convolution may carry too many capacity levels, and nearby capacities are aggregated onto a coarser grid, for example $\Delta K = 10\,\mathrm{MW}$. This reduces the computation, but it changes the probability distribution, and the change matters most near the demand threshold. With $L = 80\,\mathrm{MW}$, a state with $C = 79\,\mathrm{MW}$ is inadequate; rounded to the nearest 10 MW it becomes 80 MW and is classified as adequate. An apparently small change of one megawatt switches the shortage indicator from one to zero and moves the whole probability of the state out of the LOLP. Discretization can therefore introduce larger errors in the LOLP than in the average capacity, which it barely changes.

Table 26.12: Computational methods for capacity outage models and adequacy indices.

Method	Strength	Weakness	Best suited for
Direct enumeration	Exact and transparent for small models	Exponential growth in the number of joint states	Small fleets, verification of other methods
COPT convolution	Efficient for independent discrete capacities	Loses unit identity; requires independence	Large copper-plate adequacy models
Generating functions and FFT	Fast convolution on large capacity grids	Numerical round-off and discretization issues	Large discrete capacity distributions
Nonsequential Monte Carlo	Flexible; easy scenario sampling	Sampling error for rare events; no chronology	Adequacy estimates from interval marginals
Sequential Monte Carlo	Preserves the chronology of outages, repairs, demand and storage	Computationally expensive	Chronological adequacy of modern systems
Conditional and hybrid methods	Combine exact calculation with sampling	More complex implementation	Rare events and structured dependence

One remedy is to round available capacities downward,

$$\tilde{C} = \Delta K \left\lfloor \frac{C}{\Delta K} \right\rfloor, \tag{26.25}$$

so that $\tilde{C} \le C$ in every state. For a fixed demand, $L - \tilde{C} \ge L - C$, and since the positive part is nondecreasing,

$$(L - \tilde{C})^+ \ge (L - C)^+, \qquad \mathbf{1}_{\{\tilde{C}<L\}} \ge \mathbf{1}_{\{C<L\}},$$

the second because $C < L$ implies $\tilde{C} < L$. Taking expectations over the same scenario distribution, downward rounding gives an upper bound on the expected shortfall, hence on the EUE, and on the LOLP: it is *conservative*. Rounding upward, $\Delta K \lceil C/\Delta K \rceil$, gives lower bounds by the same argument, an *optimistic* estimate. The pair of bounds quantifies the discretization error instead of assuming that it is harmless. When the rounding is applied to each unit's states before the convolution, which is what reduces the computation, the bound still holds, since the sum of rounded-down unit capacities is at most the rounded-down total, which is at most the exact total.

26.6.3 Error against computational cost

A practical numerical study evaluates the same model on a sequence of grids, for example $\Delta K = 20$, 10, 5 and 1 MW, and compares the number of retained capacity states, the computing time and memory, the LOLP, the EUE, and the tail quantiles and probabilities of severe shortfalls. Table 26.13 carries out such a study for the multistate fleet at 80 MW, with the available-capacity states of each unit rounded down and up to the grid before the convolution. Each row is one grid step; the columns are the number of points of the grid up to the rounded-down installed capacity, the number of distinct levels actually retained, and the LOLP and expected shortfall under downward (upper bound) and upward (lower bound) rounding.

Table 26.13: Grid step against accuracy for the multistate fleet at $L = 80\,\mathrm{MW}$, unit capacities rounded to the grid before convolution. Exact values: LOLP $= 0.163$, $\mathbb{E}[S] = 4.2475\,\mathrm{MW}$.

ΔK	Grid points	Retained levels	LOLP, down	LOLP, up	$\mathbb{E}[S]$, down	$\mathbb{E}[S]$, up
20 MW	5	5	0.3880	0.10225	11.000 MW	2.5450 MW
10 MW	11	11	0.1630	0.14950	4.3875 MW	4.1075 MW
5 MW	21	17	0.1630	0.16300	4.2475 MW	4.2475 MW
1 MW	101	17	0.1630	0.16300	4.2475 MW	4.2475 MW

Every capacity of the fleet is a multiple of 5 MW, so the grids of 5 and 1 MW are exact, and a sparse table retains only the seventeen levels that occur, while a dense array of 1 MW bins holds 101. At $\Delta K = 10\,\mathrm{MW}$ only G2's derated state moves, to 10 or 20 MW; the LOLP bounds are $[0.1495, 0.163]$, the expected-shortfall bounds $[4.1075, 4.3875]\,\mathrm{MW}$, and the exact values lie inside, the LOLP at the upper bound because the only state that crosses the threshold, G2 derated alone at 85 MW, falls to 80 MW, which is still adequate. At $\Delta K = 20\,\mathrm{MW}$ the grid does not divide the rating of G1, which becomes 40 or 60 MW; the bounds widen to $[0.102, 0.388]$ for the LOLP and $[2.5, 11.0]\,\mathrm{MW}$ for the expected shortfall, too wide to be useful. Grids are therefore chosen to divide the ratings, with the rounding confined to the derated levels, and convergence is established for the risk indices that matter, not for the mean capacity: an aggregation can approximate the expected available capacity well while misrepresenting a rare, high-consequence shortage state, which would be unacceptable for some planning applications.

26.6.4 Checks before a capacity outage model is trusted

A model whose probabilities sum to one is not necessarily physically credible. The following checks apply before a COPT, or a simulation built on the same unit models, is used for a decision; the numbers in parentheses are those of the multistate fleet.

1. **Probability normalization.** The state probabilities of every unit and the probabilities of the fleet's capacity levels sum to one (Tables 26.2 and 26.4).
2. **Expected capacity.** The mean of the fleet distribution equals the sum of the expected unit capacities ($89.5 = 45.0 + 27.0 + 17.5\,\mathrm{MW}$), whatever the dependence model.
3. **State completeness.** Partial deratings, full outages and important operating restrictions are represented, not folded into a single outage probability (Table 26.5).
4. **Dependencies.** Shared fuel, weather, cooling and electrical infrastructure are not treated as independent when they are not (Table 26.8).
5. **Temporal consistency.** Snapshot outage probabilities are not confused with probabilities of experiencing at least one failure during an interval. A unit with unavailability 0.10 and failure rate $1/360\,\mathrm{h}^{-1}$ is out at a random instant with probability 0.10, but an available unit fails within the next 24 h with probability $1 - \mathrm{e}^{-24/360} = 0.0645$; the two numbers answer different questions.
6. **Operational feasibility.** Nameplate capacity that is available is not assumed deliverable when network or energy constraints matter (Sections 26.3.7 and 26.5.7).
7. **Numerical convergence.** Aggregation error and sampling error are small enough for the decision being made (Tables 26.10 and 26.13).
8. **Parameter uncertainty.** Uncertain forced outage and derating probabilities are propagated when their uncertainty materially affects the conclusions (Section 26.6.5).

26.6.5 Three sources of uncertainty

Three different sources of uncertainty enter an adequacy assessment, and they require different treatments. Table 26.14 lists them. Each row is one source, with an example from the multistate fleet and the treatment it requires.

Table 26.14: Three sources of uncertainty in a capacity outage model.

Source	Example	Treatment
Random operating state	Whether G1 is unavailable in an assessment hour	State enumeration, convolution or scenario sampling
Parameter uncertainty	Uncertainty in G1's forced outage probability, estimated from a finite outage record	Bayesian or statistical propagation of the parameter distribution
Model uncertainty	Whether extreme weather creates common outages; how many derated levels a unit has	Alternative model structures, validation against records, model averaging where justified

The first source is aleatory: the state of the fleet in a given hour is random even when every parameter is known, and the COPT is its distribution. The second and third are epistemic: they express what is not known about the parameters and the structure, in the sense of Chapter 8 of Volume I. Monte Carlo sampling of generator states addresses the first source only; it does not by itself address the other two, and a larger number of samples does not reduce them. The effect of parameter uncertainty can be seen in the multistate fleet. At 80 MW a shortfall occurs whenever G1 is out, and otherwise only when G2 and G3 together supply less than 30 MW, which has probability 0.07; hence

$$\mathrm{LOLP}(80) = U_1 + (1 - U_1) \times 0.07 = 0.07 + 0.93\,U_1,$$

which gives 0.163 at $U_1 = 0.10$. If U_1 is uncertain with mean 0.10 and standard deviation 0.03 (illustrative), the expected LOLP is still 0.163, because the index is linear in U_1, but the LOLP itself has a standard deviation of $0.93 \times 0.03 = 0.028$, more than seven times the sampling standard error, 0.00369, of a 10 000-trial simulation. A comprehensive adequacy assessment keeps the three sources distinct and reports the effect of each.

The capacity outage model of this chapter treats the network as a single bus. Chapter 27 connects component outages to the flows of the transmission network, where the location of each outage, and not only its size, decides whether the demand can be served.

26.7 What to retain

Definition 26.1 (Capacity outage models)**.** The available capacity of a fleet is a discrete random variable whose distribution, the capacity outage probability table, is computed from multistate unit models by enumerating the joint states, by recursive convolution $p_{n+1}(x) = \sum_j p_{n+1,j}\, p_n(x - d_{n+1,j})$ when the units are independent, by a mixture of such convolutions when the dependence enters through a shared state, or by Monte Carlo sampling with a stated standard error when chronology, energy limits or complex dependence make the exact methods impractical.

Multistate models capture partial deratings, not only full outages, and collapsing them to two states, even with an equivalent outage rate, changes the indices. Convolution avoids enumerating every joint configuration when the units are independent and the total capacity is a sufficient description of the state; the generating function is the same operation as polynomial

multiplication. Correlation changes the distribution of shortfalls, and at fixed marginals it can lower the probability of a shortage while raising its expected size. Exact methods eliminate sampling error but require factorization and aggregation; Monte Carlo accommodates richer dependence and operating constraints but needs about $1/(\epsilon^2\,\mathrm{LOLP})$ samples for a relative error ϵ, which makes rare events expensive. Chronological models are needed for storage, repairs and event durations. State reduction must preserve the consequences that the decision depends on, and rounding capacities down gives conservative bounds. The governing principle is to exploit the probabilistic structure of the model before enumerating every state: a reliability model becomes scalable less through faster computation than through recognizing which variables are independent, which states can be aggregated, and which dependencies must be retained. Table 26.15 lists the chapter's terms.

Table 26.15: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Multistate unit	unit with several capacity states	$\Pr(C_i = c_{i,j}) = p_{i,j}$	MW, –	26.1.1
Expected unit capacity	probability-weighted capacity	$\sum_j p_{i,j}\, c_{i,j}$	MW	26.1.1
Capacity outage	rating minus available capacity	$D_i = K_i - C_i$; $X = \sum_i D_i$	MW	26.1.2
Joint-state count	number of joint states	$\prod_i m_i$	–	26.2.1
Joint-state probability	independent units	$\prod_i p_{i,j_i}$	–	26.2.1
Unit addition algorithm	multistate COPT recursion	$p_{n+1}(x) = \sum_j p_{n+1,j}\, p_n(x - d_{n+1,j})$	–	26.3.1
Probability generating function	distribution as a polynomial	$G_{\text{fleet}}(\zeta) = \prod_i G_i(\zeta)$	–	26.3.5
Convolution cost	operations of the recursion	$n\, m\, N_{\text{bin}}$	–	26.3.6
Latent-state mixture	conditional independence	$\sum_z \Pr(z) \prod_i \Pr(C_i \mid z)$	–	26.4.2
Fixed-marginal relation	two units, demand needs both	$\mathrm{LOLP} = 2q - P_2$	–	26.4.4
Monte Carlo LOLP estimate	fraction of trials short	$N^{-1} \sum_k I^{(k)}$	–	26.5.1
Standard error, CV	sampling error of the estimate	$\sqrt{p(1-p)/N}$; $\sqrt{(1-p)/(Np)}$	–	26.5.2
Required samples	trials for relative error ϵ	$N \approx 1/(\epsilon^2\,\mathrm{LOLP})$	–	26.5.3
Importance sampling	reweighted sampling	$\mathbb{E}_{\tilde{P}}[\psi\, P/\tilde{P}]$	–	26.5.4
Chronological indices	annual hours and energy	$N^{(k)}_{\text{short}}$, $\mathrm{UE}^{(k)}$	h, MWh	26.5.6
Battery energy balance	stored-energy state	equation (26.24)	MWh	26.5.7
Conservative rounding	capacity rounded down	$\Delta K \lfloor C/\Delta K \rfloor$	MW	26.6.2

In the table, p stands for the LOLP being estimated.

Further reading. Billinton and Allan, *Reliability Evaluation of Power Systems*, for capacity outage probability tables, generation adequacy and reliability evaluation; Billinton and Allan, *Reliability Assessment of Large Electric Power Systems*, for analytical and simulation methods for large systems; *Monte Carlo Methods in Reliability Engineering*, for simulation, sampling uncertainty and rare-event estimation; EPRI's resource adequacy research and NERC's probabilistic adequacy reports, for practical definitions of adequacy metrics and modern resource-modeling challenges.

26.8 Exercises

Exercise 26.1. (a) State what distinguishes a multistate generator from a binary one. (b) State when ordinary COPT convolution is valid. (c) Explain why the eighteen joint states of the multistate fleet produce fewer than eighteen fleet capacity outcomes. (d) State whether positive outage dependence always increases the LOLP. (e) State the main statistical difficulty of simple Monte Carlo estimation of very rare shortages. (f) Explain why a battery generally requires a chronological adequacy model. (g) Explain why rounding available capacity downward gives a conservative adequacy estimate. (Answers: (a) it can occupy partial-capacity states between full availability and complete unavailability; (b) when the outage states being combined are independent of the fleet already built; (c) different combinations can have the same total capacity and are aggregated; (d) no, the effect depends on the threshold and the dependence structure, and correlated outages can lower the chance of any shortfall while raising the chance and size of severe ones; (e) few shortage observations occur, so the relative sampling error is large, one expected observation in 10 000 trials at a probability of 10^{-4}; (f) its deliverable energy depends on the preceding charging and discharging, through the state of charge; (g) for a fixed demand it can only increase the modeled shortfall, so it cannot reduce the LOLP or the EUE.)

Exercise 26.2. A small fleet has a 40 MW unit with states 40, 20 and 0 MW of probabilities 0.92, 0.05 and 0.03, and a two-state 20 MW unit with forced outage rate 0.10, independent. (a) Write the probability generating functions of the two units' capacity outages with $\Delta K = 20\,\mathrm{MW}$, multiply them, and list the COPT. (b) Check the expected available capacity against the sum of the unit means. (c) Compute the LOLP and the expected shortfall at a demand of 45 MW, and the conditional shortfall. (Answers: (a) $0.92 + 0.05\,\zeta + 0.03\,\zeta^2$ and $0.9 + 0.1\,\zeta$; product $0.828 + 0.137\,\zeta + 0.032\,\zeta^2 + 0.003\,\zeta^3$, so the outage is 0, 20, 40, 60 MW with probabilities 0.828, 0.137, 0.032, 0.003. (b) $60 \times 0.828 + 40 \times 0.137 + 20 \times 0.032 = 55.8\,\mathrm{MW} = 37.8 + 18.0$. (c) $\mathrm{LOLP} = 0.137 + 0.032 + 0.003 = 0.172$; $\mathbb{E}[S] = 0.137 \times 5 + 0.032 \times 25 + 0.003 \times 45 = 1.62\,\mathrm{MW}$; $\mathbb{E}[S \mid S > 0] = 9.42\,\mathrm{MW}$.)

Exercise 26.3. Two 50 MW units serve a demand of 80 MW. The weather is normal with probability 0.9, in which each unit is out with probability 0.05, and extreme with probability 0.1, in which each is out with probability 0.55; given the weather the units are independent. (a) Compute the marginal unavailability of each unit. (b) Build the conditional COPT in each weather state and the mixture COPT. (c) Compute the LOLP and the expected shortfall, and check them with equation (26.15). (d) Compute the correlation of the two outage indicators and compare the indices with those of independent units of the same marginal unavailability. (Answers: (a) $0.9 \times 0.05 + 0.1 \times 0.55 = 0.10$. (b) Normal: outage 0, 50, 100 MW with 0.9025, 0.095, 0.0025; extreme: 0.2025, 0.495, 0.3025; mixture: 0.8325, 0.135, 0.0325. (c) $\mathrm{LOLP} = 0.135 + 0.0325 = 0.1675 = 0.2 - 0.0325$; $\mathbb{E}[S] = 0.135 \times 30 + 0.0325 \times 80 = 6.65\,\mathrm{MW} = 6 + 20 \times 0.0325$. (d) $\rho = (0.0325 - 0.01)/0.09 = 0.25$; independence gives 0.19 and 6.2 MW, so the shared weather lowers the LOLP and raises the expected shortfall.)

Exercise 26.4. For the multistate fleet at 80 MW (LOLP = 0.163, $\mathbb{E}[S] = 4.2475\,\mathrm{MW}$, $\mathrm{Var}(S) = 126.38\,\mathrm{MW}^2$): (a) compute the number of nonsequential trials needed for a coefficient

of variation of 5 % in the LOLP estimate, and the number needed if the LOLP were 10^{-3}; (b) compute the number needed for a 5 % coefficient of variation in the expected-shortfall estimate; (c) compute the probability that a simulation of 10 000 trials observes no shortage when the LOLP is 10^{-4}; (d) round the derated state of G2 down to a 10 MW grid, recompute the LOLP and the expected shortfall, and explain why the LOLP does not change. (Answers: (a) $0.837/(0.0025 \times 0.163) \approx 2\,054$ and $0.999/(0.0025 \times 0.001) = 399\,600$. (b) $(11.24/4.2475)^2/0.0025 \approx 2\,802$. (c) $(1 - 10^{-4})^{10\,000} \approx \mathrm{e}^{-1} = 0.368$. (d) LOLP $= 0.163$; $\mathbb{E}[S] = 4.2475 + 5 \times (0.0135 + 0.0045 + 0.008 + 0.0015 + 0.0005) = 4.3875$ MW; the rounding lowers by 5 MW the six states in which G2 is derated, and the only one that was adequate, G2 derated alone at 85 MW, falls to 80 MW, which is still adequate, while the other five were already deficient.)

Chapter 27

Transmission Reliability and Contingency Analysis

27.1 Transmission contingencies, security and reliability

Chapter 25 measured the adequacy of a generating fleet by the probability distribution of its available capacity, $\Pr(C_t < L_t)$, on a single bus that could carry any generation to any load. A transmission network removes that simplification. Enough generation does not guarantee that the power reaches the customers: a system with 2000 MW of available generation and 1500 MW of demand can still have to curtail load because an important transmission corridor is out of service or would be overloaded. The question of this chapter is

how component outages change the power flows through a transmission network, and how those changes are translated into overload probabilities, expected load curtailment and operating decisions informed by reliability.

The answer combines two models: a probabilistic model of component outages and operating conditions, and a physical network model that determines the power flows and the operational consequences under each of those conditions. The chapter builds both for a three-bus network, first deterministically (which single outages the network survives) and then probabilistically (how likely each outage is and what it costs in curtailed load), and then adds uncertain line ratings, corrective actions, the attribution of risk to individual lines, and the limits of the analysis.

Jargon: contingency, contingency analysis, N-1 and N-k, thermal rating, security, corrective action, redispatch, remedial action scheme, load curtailment, islanding, transfer capability, PTDF, LODF

A *contingency* is an event that changes the configuration or the operating condition of a power network: the loss of a line, a transformer, a generator, a bus section or another facility. *Contingency analysis* computes the state of the network after each contingency of a specified list and checks it against the operating limits. The $N-1$ *criterion* requires a network of N elements to meet its limits after the loss of any one of them; $N-k$ contingencies involve k elements. The *thermal rating* of a line is the largest power (or current) it may carry without its conductor exceeding a temperature limit; a *normal rating* may be carried indefinitely, an *emergency rating* only for a stated time. *Security* is the ability of the system to withstand a contingency without violating its limits. A *corrective action* is taken after a contingency to restore the limits: *redispatch* (a change in the outputs of generators), switching, the action of a *remedial action scheme* (RAS, an automatic scheme that trips generation or load or reconfigures the network when it detects a predefined condition), or *load curtailment*, the controlled interruption of demand. *Islanding* is the separation of the network into parts with no electrical connection between them. The *transfer capability* between two buses is the largest power that can be moved from one to the other without violating a limit. The *power transfer distribution factor* (PTDF) is the change in the flow on a line per unit of power transferred between two buses; the *line outage distribution factor* (LODF) is the change in the flow on a line per unit of pre-outage flow on another line that is taken out of service.

Key idea

A component outage changes the network topology; the topology decides how the power redistributes over the remaining lines; the redistributed flows, compared with the line ratings, decide the corrective action and the curtailment. Transmission risk is the probability-weighted curtailment over the network states, and it ranks outages differently from their severity alone.

This section states the chain from component failure to risk, defines a contingency and the $N-1$ and $N-k$ terminology, and separates the deterministic security question from the probabilistic reliability question.

27.1.1 Deliverability of generation

In the single-bus model of Chapter 25 the shortfall in an interval is $(L_t - C_t)^+$, a function of the total available capacity alone. Section 25.5.2 of that chapter showed the first consequence of adding a network: when the fleet of three generators was split between two areas joined by a tie line rated 60 MW, the loss-of-load probability at a demand of 80 MW rose from 0.235 to 0.388 and the expected shortfall from 5.74 to 8.80 MW, with the same generators and the same forced outage rates. Only the deliverability of the generation had changed. In a meshed network the deliverability depends on more than one tie-line rating: the power divides among parallel paths according to the electrical impedances of the lines, not according to the wishes of the operator, and the loss of one line changes the division over all the others.

Transmission reliability analysis therefore runs through a chain of models, which Figure 27.1 draws. Each box is one stage; the text below each box names the quantity that stage produces, with the symbol this chapter uses for it. Component failures, described by outage indicators and their probabilities, determine the network topology; the topology, through the network's susceptance matrix, determines the power flows; the flows, compared with the line ratings and given the corrective actions available, determine the operational consequence, here the curtailed load; and the consequences, weighted by the probabilities of the network states, give the risk.

Every arrow of the chain is a model, and each must be stated before a risk number can be interpreted.

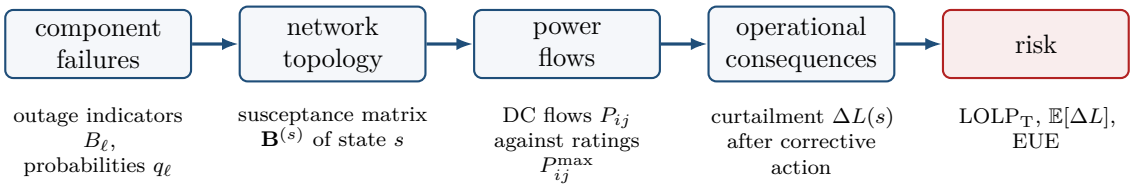


Figure 27.1: The chain of models of transmission reliability analysis. Each stage produces the quantity named below it, which the next stage takes as its input; the symbols are defined in Sections 27.2.2 to 27.4.4.

27.1.2 Transmission contingencies

A contingency is an event that changes the network configuration or its operating condition: the loss of a line, a transformer, a generator, a bus section or another facility. A contingency analysis asks whether the system that results can continue to operate within acceptable limits. The limits typically checked are the thermal ratings of lines and transformers; the bus voltage limits; the reactive-power limits of the generators; the stability requirements of the system; the loadability limits of protective relays, beyond which a relay may interpret a heavy load as a fault and trip the line; and an acceptable level of load interruption, which for the most common contingencies is often none.

In the vocabulary of Chapter 2 of Volume I, a contingency is a component failure, and the limits are the performance criteria that define the required function of the network. That chapter's two-path example (Section 2.4.5) already distinguished three failure definitions for one network: the unavailability of a line, the loss of the load-serving function and the violation of the $N-1$ security requirement, with probabilities 0.05, 0.0025 and 0.0975 from the same component data. The contingency is the component event; whether it is a system failure depends on the limits checked and on the actions allowed after it.

27.1.3 N-1 and N-k contingencies

$N-1$. The loss of one relevant element from a specified initial system condition. A system is $N-1$ secure, for that condition, if after the loss of any single element of the list it still meets its limits, possibly after the corrective actions the criterion allows.

$N-2$ and $N-k$. Combinations or sequences of two unavailable elements, and in general of k elements. The number of $N-k$ contingencies of a network of N elements is the binomial coefficient $\binom{N}{k}$: the three lines of the network of this chapter give 3 single, 3 double and 1 triple contingency, while a network of 1000 elements has 1000 single and 499 500 double contingencies. The growth with k is the reason that exhaustive enumeration stops at small k and that the fast screening methods of Section 27.3 exist.

$N-k$ is a useful shorthand, but formal planning standards are more specific than a count of failed components. North American transmission planning standards, for example, distinguish the normal system with no contingency (category P0), single contingencies (P1), and events such as the fault of a bus section or a breaker, a fault with a stuck breaker or a failed protection system, two overlapping single contingencies, and the loss of a common structure that carries two circuits (P2 to P7), each with its own performance requirements, such as whether load may be interrupted. These categories are not interchangeable with a generic $N-k$ classification: a

single fault that removes two circuits from a common tower is one event with two elements lost, and its requirements differ from those of two independent outages.

27.1.4 Security and reliability

Two different questions are asked of a transmission network. *Deterministic security* asks, for example, whether any remaining element exceeds its emergency rating if line L13 is lost. *Probabilistic reliability* asks what the expected operational consequence is, given the probability that L13 fails, the probability distribution of the line ratings and the corrective actions available. The first is a yes-or-no statement about a specified list of contingencies; the second weights each contingency by its probability and by its consequence. Chapter 25 (Section 25.5.1) drew the same line between adequacy, the existence of enough capacity to serve the demand with outages accounted for, and operating security, the ability to withstand a specified contingency without violating limits.

A line outage can produce an overload without producing any unserved energy. Protection may trip the overloaded line; an operator may redispatch generation; a remedial action scheme may act; or the system may tolerate the loading for a while within an emergency rating, since a conductor heats up over minutes rather than instantly. Therefore

$$\text{post-contingency overload} \neq \text{automatic load curtailment},$$

and every probabilistic contingency study must state its response model: which actions follow an overload, how quickly, and with what effect on the load. In this chapter the response model is stated with each calculation.

27.2 The physical network model and single-outage contingency analysis

This section builds the physical model: a three-bus network, the DC power-flow equations derived from the AC equations under stated assumptions, the base-case flows and margins, and the flows and curtailments after each single-line outage.

27.2.1 The three-bus network

The network of this chapter, which Chapters 29 and 30 reuse, has three buses and three lines forming a triangle. Bus 1 has a generator; bus 2 is a transmission junction with neither generation nor load; bus 3 has the load. The generator injects $P_1^{\mathrm{G}} = 100\,\mathrm{MW}$ and the load withdraws $L_3 = 100\,\mathrm{MW}$, so the net injections are $P_1 = +100\,\mathrm{MW}$, $P_2 = 0$ and $P_3 = -100\,\mathrm{MW}$. Each line is oriented from the lower-numbered bus to the higher: L12 from bus 1 to bus 2, L13 from bus 1 to bus 3 and L23 from bus 2 to bus 3, and a positive flow follows that orientation. Table 27.1 lists the line data: each row is one line, with its series reactance x_{ij} on a 100 MVA base, its susceptance $b_{ij} = 1/x_{ij}$, its thermal rating $P_{ij}^{\max}$, and the base-case flow and margin computed in Section 27.2.3. The values are illustrative.

Table 27.1: Line data of the three-bus network (illustrative), with the base-case DC flows and thermal margins of Section 27.2.3.

Line	Reactance x_{ij}	Susceptance b_{ij}	Rating $P_{ij}^{\max}$	Base flow P_{ij}	Margin g_{ij}
L12	0.1 pu	10 pu	70 MW	50 MW	20 MW
L23	0.1 pu	10 pu	70 MW	50 MW	20 MW
L13	0.2 pu	5 pu	80 MW	50 MW	30 MW

Figure 27.2 draws the network. The circles are the buses, the thick lines the transmission lines, labeled with their reactance and rating; the arrow into bus 1 is the generation and the arrow out of bus 3 the load. The thin arrows inside the triangle are the base-case flows derived in the next two subsections, each in the direction of the line's orientation. The figure shows the physical fact on which the rest of the chapter rests: the 100 MW sent from bus 1 to bus 3 does not travel only on the direct line, but divides between the direct path L13 and the indirect path through bus 2.

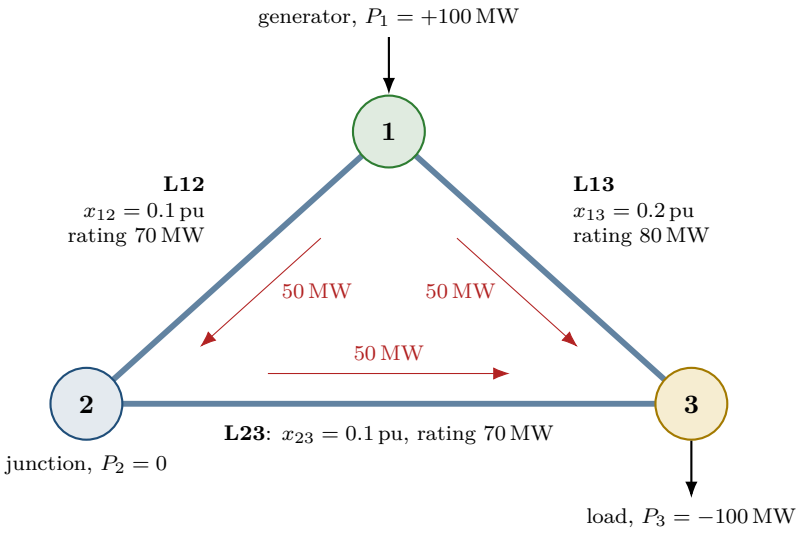


Figure 27.2: The three-bus network (illustrative values): generation at bus 1, a junction at bus 2, the load at bus 3. Reactances are on a 100 MVA base; the red arrows are the base-case DC flows, 50 MW on each line, in the orientations L12 from 1 to 2, L13 from 1 to 3 and L23 from 2 to 3.

27.2.2 DC power flow from the AC equations

Each bus i has a complex voltage of magnitude V_i and phase angle θ_i, in radians. Consider one line between buses i and j with series impedance $r_{ij} + \mathrm{j}x_{ij}$, in per unit, and no shunt elements; its series admittance is $y_{ij} = 1/(r_{ij} + \mathrm{j}x_{ij}) = g_{ij} + \mathrm{j}b'_{ij}$, with $g_{ij} = r_{ij}/(r_{ij}^2 + x_{ij}^2)$ and $b'_{ij} = -x_{ij}/(r_{ij}^2 + x_{ij}^2)$, negative for an inductive line. The current from i to j is $y_{ij}(\tilde{V}_i - \tilde{V}_j)$, where $\tilde{V}_i = V_i \mathrm{e}^{\mathrm{j}\theta_i}$ is the complex voltage, and the complex power sent from i into the line is $\tilde{V}_i$ times the conjugate of that current,

$$P_{ij} + \mathrm{j}Q_{ij} = \tilde{V}_i\, y_{ij}^*\, (\tilde{V}_i^* - \tilde{V}_j^*) = (g_{ij} - \mathrm{j}b'_{ij})(V_i^2 - V_iV_j\cos\theta_{ij} - \mathrm{j}\, V_iV_j\sin\theta_{ij}),$$

with $\theta_{ij} = \theta_i - \theta_j$. Its real part is the active power flow,

$$P_{ij} = g_{ij}\,(V_i^2 - V_iV_j\cos\theta_{ij}) - b'_{ij}\,V_iV_j\sin\theta_{ij}. \tag{27.1}$$

Adding the flow in the opposite direction, P_{ji}, which is the same expression with i and j exchanged, gives the active loss of the line, $P_{ij} + P_{ji} = g_{ij}(V_i^2 + V_j^2 - 2V_iV_j\cos\theta_{ij})$, which for $V_i = V_j = 1$ is $2g_{ij}(1 - \cos\theta_{ij}) \approx g_{ij}\,\theta_{ij}^2$: the losses are of second order in the angle difference.

The DC power-flow model follows from four assumptions, stated as such because each is an approximation with a range of validity: (1) the voltage magnitudes are close to their nominal values, $V_i \approx 1$ pu; (2) the angle differences across lines are small, so that $\sin\theta_{ij} \approx \theta_{ij}$ and $\cos\theta_{ij} \approx 1$; (3) the line resistance is small against the reactance, $r_{ij} \ll x_{ij}$, so that $g_{ij} \approx 0$ and

$-b'_{ij} \approx 1/x_{ij}$; (4) reactive power and the active losses are neglected. Under them equation (27.1) becomes linear,

$$P_{ij} = \frac{\theta_i - \theta_j}{x_{ij}} = b_{ij}\,(\theta_i - \theta_j), \qquad b_{ij} = \frac{1}{x_{ij}}, \tag{27.2}$$

in per unit; in megawatts the right-hand side is multiplied by the system base, $S_{\text{base}} = 100\,\text{MVA}$. The active power flows like a current through a resistive network: the angle differences play the role of voltages and the reactances the role of resistances.

The net injection at each bus is the sum of the flows leaving it on its lines,

$$P_i = \sum_{j \in \mathcal{N}(i)} P_{ij} = \sum_{j \in \mathcal{N}(i)} b_{ij}(\theta_i - \theta_j) = \Big(\sum_{j \in \mathcal{N}(i)} b_{ij}\Big)\theta_i - \sum_{j \in \mathcal{N}(i)} b_{ij}\,\theta_j,$$

where $\mathcal{N}(i)$ is the set of buses connected to i. Collecting the buses gives the matrix form

$$\mathbf{P} = \mathbf{B}\,\boldsymbol{\theta}, \qquad B_{ii} = \sum_{j \in \mathcal{N}(i)} b_{ij}, \quad B_{ij} = -b_{ij}\ (j \neq i), \tag{27.3}$$

with $\mathbf{P}$ the vector of net injections in per unit, $\boldsymbol{\theta}$ the vector of bus angles and $\mathbf{B}$ the DC susceptance matrix. Every row of $\mathbf{B}$ sums to zero, because the injections of a lossless network sum to zero, so $\mathbf{B}$ is singular and the angles are fixed only up to a common constant. One bus is chosen as the angle reference, its angle set to zero, and its row and column removed; the reduced matrix $\mathbf{B}_{\text{r}}$ is invertible when the network is connected.

Remark 27.1. Several symbols of this chapter share letters. The bus angle θ_i is a scalar in radians; the bold $\boldsymbol{\theta}$ is the vector of bus angles here and, in Sections 27.5.1 and 27.6.4, where it is stated, a vector of uncertain model parameters. The bold $\mathbf{B}$ is the DC susceptance matrix; B_ℓ (italic, one subscript) is the outage indicator of line ℓ of Section 27.4.1; B_{ij} in Section 27.3.4 is an element of the susceptance part of the AC admittance matrix. The reactive powers Q_i and Q_{ij} are not the top-event probability Q_{top} of Chapter 10 of Volume I, and the conductor resistance $R(T)$ of Section 27.5.1 is not the reliability function $R(t)$.

For the network of Table 27.1, $b_{12} = 1/0.1 = 10$, $b_{13} = 1/0.2 = 5$ and $b_{23} = 1/0.1 = 10$, all in per unit, and with injections $\mathbf{P} = (1, 0, -1)$ in per unit of 100 MVA, equation (27.3) reads

$$\begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} = \begin{bmatrix} 15 & -10 & -5 \\ -10 & 20 & -10 \\ -5 & -10 & 15 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix}. \tag{27.4}$$

The diagonal entry of bus 1 is $b_{12} + b_{13} = 15$, that of bus 2 is $b_{12} + b_{23} = 20$, and that of bus 3 is $b_{13} + b_{23} = 15$; each row sums to zero.

The DC model gives no information on voltage magnitudes or reactive power, and its error grows with the angle differences and with the ratio r_{ij}/x_{ij}. Both can be checked after the solution: the largest angle difference below is 0.1 rad, for which $\sin 0.1 = 0.09983$, an error of 0.17 % in the linearization; and an illustrative resistance $r_{12} = 0.01$ pu on L12 would give $-b'_{12} = 0.1/0.0101 = 9.90$ instead of 10, and a loss of $g_{12}\,\theta_{12}^2 \approx 0.99 \times 0.05^2 = 0.0025$ pu, or 0.25 MW on a flow of 50 MW. The DC model is therefore adequate for the thermal screening of this chapter, and Section 27.3.4 states what it leaves out.

27.2.3 Base-case flows and margins

Take bus 3 as the reference, $\theta_3 = 0$. Removing its row and column from equation (27.4) leaves

$$\begin{bmatrix} 15 & -10 \\ -10 & 20 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

The second equation gives $\theta_1 = 2\theta_2$, and substituting in the first, $30\theta_2 - 10\theta_2 = 1$, so

$$\theta_2 = 0.05\,\text{rad}, \qquad \theta_1 = 0.10\,\text{rad}.$$

The removed third equation is satisfied, $-5(0.10) - 10(0.05) + 15(0) = -1$, as it must be. By equation (27.2), in megawatts,

$$P_{12} = \frac{0.10 - 0.05}{0.1} \times 100 = 50\,\text{MW},$$
$$P_{23} = \frac{0.05 - 0}{0.1} \times 100 = 50\,\text{MW},$$
$$P_{13} = \frac{0.10 - 0}{0.2} \times 100 = 50\,\text{MW}.$$

The flows balance at every bus: bus 1 sends $50 + 50 = 100\,\text{MW}$, bus 2 receives 50 and passes on 50, bus 3 receives $50 + 50 = 100\,\text{MW}$.

The equal split has a simple explanation. The network offers two parallel paths from bus 1 to bus 3: the direct line L13, of reactance 0.2, and the series path L12–L23, of reactance $0.1 + 0.1 = 0.2$. Parallel paths share a transfer in inverse proportion to their reactances, as parallel resistors share a current, so the share of the direct path is $x_{\text{indirect}}/(x_{\text{direct}} + x_{\text{indirect}}) = 0.2/0.4 = 0.5$, and with equal path reactances the 100 MW divides equally.

The *thermal margin* of a line is the difference between its rating and the magnitude of its flow,

$$g_{ij} = P_{ij}^{\max} - |P_{ij}|, \tag{27.5}$$

in MW; it is the performance-margin function of Chapter 2 of Volume I applied to one line, and $g_{ij} \geq 0$ means the line is within its rating. The last column of Table 27.1 gives $g_{12} = 70 - 50 = 20\,\text{MW}$, $g_{23} = 70 - 50 = 20\,\text{MW}$ and $g_{13} = 80 - 50 = 30\,\text{MW}$. All three lines are within their limits, loaded at 71 %, 71 % and 62.5 % of their ratings.

The base case also fixes the *transfer capability* TC of the intact network, the largest transfer from bus 1 to bus 3 that keeps every line within its rating. Because the flows are linear in the transfer P_{tr} and each line carries half of it, every flow is $P_{\text{tr}}/2$, and the limits require $P_{\text{tr}}/2 \leq 70$ on L12 and L23 and $P_{\text{tr}}/2 \leq 80$ on L13. The series path reaches its 70 MW limits first, when it carries 70 MW and the total transfer is

$$\text{TC(intact)} = 2 \times \min(70, 70, 80) = 140\,\text{MW}.$$

The factor 2 holds only because the two paths have equal reactances; a network with unequal paths has a different factor (an exercise works one). Base-case security, however, does not establish $N-1$ security: a line outage redistributes the power over the surviving lines, and the margins of the base case say nothing about the flows that follow.

27.2.4 Outage of line L13 and its corrective action

Suppose the direct line L13 is disconnected. Removing its susceptance from $\mathbf{B}$ leaves the reduced system

$$\begin{bmatrix} 10 & -10 \\ -10 & 20 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

whose solution is $\theta_2 = 0.1$ and $\theta_1 = 0.2\,\text{rad}$. The flows are $P_{12} = (0.2 - 0.1)/0.1 \times 100 = 100\,\text{MW}$ and $P_{23} = 0.1/0.1 \times 100 = 100\,\text{MW}$. The result is also evident without the matrix: the remaining network is the series path $1 \to 2 \to 3$, and with generation and demand unchanged all 100 MW must cross both L12 and L23. Compared with their 70 MW ratings, both lines are overloaded by $100 - 70 = 30\,\text{MW}$. The contingency creates overloads although the base case was secure; the angle across the path doubles to 0.2 rad, still small enough for the DC model ($\sin 0.2 = 0.1987$).

Corrective action. Assume that the only corrective action is to reduce the transfer from bus 1 to bus 3 by reducing the generation at bus 1 and curtailing the same amount of load at bus 3, with no alternative generation and no imports. The transfer P_{tr} then flows entirely on the series path, and both ratings are respected when $P_{\mathrm{tr}} \le \min(70, 70) = 70\,\mathrm{MW}$. The transfer capability after the outage is 70 MW, and the minimum curtailment under this corrective-action model is

$$\Delta L(\text{L13 out}) = 100 - 70 = 30\,\mathrm{MW}.$$

This is not a model of cascading failure. It is a model of controlled curtailment that restores a feasible operating state; what happens if the overload is not corrected is the subject of the next chapter.

27.2.5 Outages of lines L12 and L23

Suppose instead that L12 is disconnected. The direct path L13 remains, and bus 2 becomes a terminal bus connected only through L23, with no load or generation. The reduced system is $5\theta_1 = 1$ and $10\theta_2 = 0$, so $\theta_1 = 0.2\,\mathrm{rad}$ and $\theta_2 = 0$. No power can flow into a terminal bus that has nowhere to send it, so the flow on L23 is zero in the resulting steady state, and all 100 MW flows on L13, $P_{13} = 0.2/0.2 \times 100 = 100\,\mathrm{MW}$. Against its rating of 80 MW the overload is $100 - 80 = 20\,\mathrm{MW}$. The largest transfer that L13 alone can carry is 80 MW, and the curtailment is

$$\Delta L(\text{L12 out}) = 100 - 80 = 20\,\mathrm{MW}.$$

Outage of L23. The network is symmetric in this respect. Disconnecting L23 leaves bus 2 as a terminal bus on L12, the reduced system gives $\theta_1 = \theta_2 = 0.2\,\mathrm{rad}$, the flow on L12 is zero and L13 again carries the whole transfer, $P_{13} = 100\,\mathrm{MW}$. The required curtailment is the same, $\Delta L(\text{L23 out}) = 100 - 80 = 20\,\mathrm{MW}$.

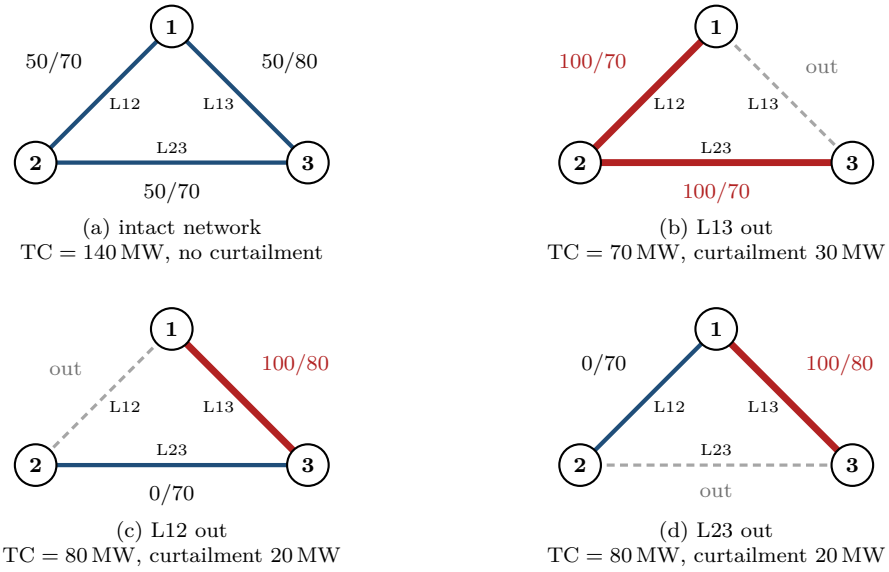
Figure 27.3 draws the four configurations. Each panel is the triangle of Figure 27.2; next to each line is its flow and its rating in MW, written flow/rating; a dashed gray line is out of service; a thick red line is overloaded. Below each panel is the transfer capability of the configuration and the curtailment it requires. The panels show that the outage of the line with the largest rating and the largest base margin, L13, is the one that requires the largest curtailment: what matters is not the margin of the line that is lost but the margins of the lines that receive its flow.

27.2.6 Single-outage results

Table 27.2 collects the results. Each row is one configuration: the contingency, the critical post-outage flow (the flow on the most heavily loaded line, relative to its rating), the rating of that line, the transfer capability of the configuration and the curtailment required to deliver the remaining load within the ratings.

Table 27.2: Single-outage contingency analysis of the three-bus network for a transfer of 100 MW, with curtailment as the only corrective action.

Contingency	Critical post-outage flow	Rating	Transfer capability	Required curtailment
none	50 MW on each line	70–80 MW	140 MW	0 MW
L12 out	L13: 100 MW	80 MW	80 MW	20 MW
L23 out	L13: 100 MW	80 MW	80 MW	20 MW
L13 out	L12 and L23: 100 MW	70 MW	70 MW	30 MW



(a) intact network
TC = 140 MW, no curtailment

(b) L13 out
TC = 70 MW, curtailment 30 MW

(c) L12 out
TC = 80 MW, curtailment 20 MW

(d) L23 out
TC = 80 MW, curtailment 20 MW

Figure 27.3: DC flows of the three-bus network for a transfer of 100 MW from bus 1 to bus 3: (a) intact; (b)–(d) after the outage of one line. Labels give flow/rating in MW; thick red lines are overloaded, dashed gray lines are out of service. TC is the transfer capability of the configuration and the curtailment is 100 MW − TC (illustrative values of Table 27.1).

All three single-line outages violate a rating at the original transfer of 100 MW. The system is therefore not $N-1$ secure for that transfer under these limits and without corrective measures other than curtailment. The table also gives the transfer that would be $N-1$ secure: the smallest post-contingency transfer capability, $\min(80, 80, 70) = 70$ MW. An operator who applied the $N-1$ criterion *preventively*, that is, who limited the pre-contingency operating point so that every single outage left the flows within ratings, would have to keep the transfer at 70 MW and curtail 30 MW even with all three lines in service; one who relied on *corrective* action would run at 100 MW and curtail only after an outage. The choice between the two is an operating policy, and the probabilistic analysis of Section 27.4 is what weighs their costs.

The result of this section is deterministic. It states which outages the network survives and what each costs, but not how often each occurs.

27.3 Distribution factors and AC security constraints

Solving the full power flow for every contingency is easy for three lines and expensive for thousands. This section derives the two linear sensitivity factors that replace those solutions in screening, checks them against the solutions of the previous section, and then states the AC model and the security constraints that the DC model omits.

27.3.1 Sensitivity factors

The DC power flow of equation (27.3) is linear in the injections. Every line flow is therefore a fixed linear combination of the bus injections, and the change in a flow caused by a change in the injections, or by the removal of a line, can be precomputed once from the matrix $\mathbf{B}_\mathrm{r}^{-1}$ of the intact network. For a network of a few thousand buses, a full solution per contingency and per operating point is costly; the precomputed factors reduce each contingency to a few multiplications. Operators and planners use two such factors: the power transfer distribution factor (PTDF), for changes in injections, and the line outage distribution factor (LODF), for the loss of a line. Both are exact within the DC model and approximate for the AC network.

27.3.2 Power transfer distribution factors

With bus 3 as reference, the reduced system gives $\boldsymbol{\theta}_{\mathrm{r}} = \mathbf{B}_{\mathrm{r}}^{-1}\mathbf{P}_{\mathrm{r}}$, so each angle is $\theta_i = \sum_k [\mathbf{B}_{\mathrm{r}}^{-1}]_{ik}\, P_k$, the sum running over the non-reference buses, with $\theta_3 = 0$. Substituting into equation (27.2), the flow on line ℓ from bus i to bus j is

$$P_\ell = \frac{\theta_i - \theta_j}{x_\ell} = \sum_k \frac{[\mathbf{B}_{\mathrm{r}}^{-1}]_{ik} - [\mathbf{B}_{\mathrm{r}}^{-1}]_{jk}}{x_\ell}\, P_k,$$

where the entries of a row for the reference bus are zero. The coefficient of P_k is the *power transfer distribution factor*

$$\mathrm{PTDF}_{\ell,k} = \frac{\partial P_\ell}{\partial P_k} = \frac{[\mathbf{B}_{\mathrm{r}}^{-1}]_{ik} - [\mathbf{B}_{\mathrm{r}}^{-1}]_{jk}}{x_\ell}, \tag{27.6}$$

the change in the flow on line ℓ per unit of power injected at bus k and withdrawn at the reference bus. A transfer that injects power at bus m and withdraws it at bus n is the difference of two such injections, so its factor is

$$\mathrm{PTDF}_{\ell,m\to n} = \frac{\Delta P_\ell}{\Delta P_{m\to n}} = \mathrm{PTDF}_{\ell,m} - \mathrm{PTDF}_{\ell,n}, \tag{27.7}$$

with $\mathrm{PTDF}_{\ell,n} = 0$ when n is the reference bus. A PTDF is dimensionless and its sign follows the orientation of the line.

For the three-bus network the inverse of the reduced matrix is

$$\mathbf{B}_{\mathrm{r}}^{-1} = \begin{bmatrix} 15 & -10 \\ -10 & 20 \end{bmatrix}^{-1} = \frac{1}{200}\begin{bmatrix} 20 & 10 \\ 10 & 15 \end{bmatrix} = \begin{bmatrix} 0.10 & 0.05 \\ 0.05 & 0.075 \end{bmatrix},$$

with determinant $15\times 20 - 10\times 10 = 200$. For an injection at bus 1 (first column), equation (27.6) gives $\mathrm{PTDF}_{12,1} = (0.10 - 0.05)/0.1 = 0.5$, $\mathrm{PTDF}_{13,1} = (0.10 - 0)/0.2 = 0.5$ and $\mathrm{PTDF}_{23,1} = (0.05 - 0)/0.1 = 0.5$. Because bus 3 is the reference, these are the factors of the transfer from bus 1 to bus 3,

$$\mathrm{PTDF}_{12,1\to 3} = \mathrm{PTDF}_{13,1\to 3} = \mathrm{PTDF}_{23,1\to 3} = 0.5,$$

and the base-case flows are $0.5 \times 100 = 50\,\mathrm{MW}$ on each line, as the full solution gave. For an injection at bus 2 (second column), $\mathrm{PTDF}_{12,2} = (0.05 - 0.075)/0.1 = -0.25$, $\mathrm{PTDF}_{13,2} = 0.05/0.2 = 0.25$ and $\mathrm{PTDF}_{23,2} = 0.075/0.1 = 0.75$: power injected at bus 2 and withdrawn at bus 3 goes three quarters directly on L23 and one quarter around through bus 1, against the orientation of L12. By equation (27.7) the transfer from bus 1 to bus 2 has factors $0.5 - (-0.25) = 0.75$ on L12, $0.5 - 0.25 = 0.25$ on L13 and $0.5 - 0.75 = -0.25$ on L23. Table 27.3 collects these factors.

As a use of the factors, a 10 MW increase of the transfer from bus 1 to bus 3 raises every flow by $0.5 \times 10 = 5\,\mathrm{MW}$, which consumes 5 MW of the 20 MW margins of L12 and L23. And if, illustratively, 20 MW of generation were moved from bus 1 to a unit at bus 2, the flow changes would be $20 \times (\mathrm{PTDF}_{\ell,2} - \mathrm{PTDF}_{\ell,1})$, that is -15, -5 and $+5$ MW on L12, L13 and L23, and the flows would become 35, 45 and 55 MW, all without a new power-flow solution.

27.3.3 Line outage distribution factors

The *line outage distribution factor* of a monitored line ℓ for the outage of line m is the change in the flow on ℓ per unit of pre-outage flow on m,

$$\mathrm{LODF}_{\ell,m} = \frac{P_\ell^{(m)} - P_\ell^{(0)}}{P_m^{(0)}}, \tag{27.8}$$

where $P_\ell^{(0)}$ is the base flow on ℓ, $P_\ell^{(m)}$ its flow after line m is removed, and $P_m^{(0)}$ the base flow on the outaged line. It is dimensionless; the post-outage flow is predicted as $P_\ell^{(m)} = P_\ell^{(0)} + \mathrm{LODF}_{\ell,m}\, P_m^{(0)}$.

The LODF follows from the PTDFs without modifying $\mathbf{B}$. Let line m run from bus p to bus q. The outage is simulated in the intact network by injecting an amount t at bus p and withdrawing it at bus q, with t chosen so that line m carries exactly t: then the whole of the line's flow is supplied at one end and removed at the other by the fictitious injections, and the rest of the network sees no exchange with line m, as if it were open. The flow on m after the injection is $P_m^{(0)} + \mathrm{PTDF}_{m,p\to q}\, t$, and setting it equal to t gives

$$t = \frac{P_m^{(0)}}{1 - \mathrm{PTDF}_{m,p\to q}}.$$

The same injection changes the flow on any other line ℓ by $\mathrm{PTDF}_{\ell,p\to q}\, t$, so

$$\mathrm{LODF}_{\ell,m} = \frac{\mathrm{PTDF}_{\ell,p\to q}}{1 - \mathrm{PTDF}_{m,p\to q}}, \qquad \ell \neq m. \tag{27.9}$$

For the outage of L13 ($p = 1$, $q = 3$), $\mathrm{PTDF}_{13,1\to 3} = 0.5$ and

$$\mathrm{LODF}_{12,13} = \frac{0.5}{1 - 0.5} = 1, \qquad \mathrm{LODF}_{23,13} = \frac{0.5}{1 - 0.5} = 1.$$

Equation (27.8) with the full solutions of Section 27.2.4 gives the same values, $\mathrm{LODF}_{12,13} = (100 - 50)/50 = 1$ and likewise for L23: each surviving line absorbs the whole 50 MW that L13 carried. For the outage of L12 ($p = 1$, $q = 2$), with $\mathrm{PTDF}_{12,1\to 2} = 0.75$, the factors are $\mathrm{LODF}_{13,12} = 0.25/0.25 = 1$ and $\mathrm{LODF}_{23,12} = -0.25/0.25 = -1$; for the outage of L23 ($p = 2$, $q = 3$), with $\mathrm{PTDF}_{23,2\to 3} = 0.75$, they are $\mathrm{LODF}_{12,23} = -0.25/0.25 = -1$ and $\mathrm{LODF}_{13,23} = 0.25/0.25 = 1$. Table 27.3 lists all of them beside the flows they predict. Each row is a monitored line; the first two columns are its PTDFs for injections at buses 1 and 2; the next three are its LODFs for the outage of each other line; the last three are the post-outage flows predicted by $P_\ell^{(0)} + \mathrm{LODF}_{\ell,m} P_m^{(0)}$, which agree with the full DC solutions of Section 27.2 to the last digit.

Table 27.3: Distribution factors of the three-bus network (bus 3 the reference) and the post-outage flows they predict, in MW; every predicted flow equals the full DC solution of Figure 27.3.

Line ℓ	$\mathrm{PTDF}_{\ell,1}$	$\mathrm{PTDF}_{\ell,2}$	$\mathrm{LODF}_{\ell,12}$	$\mathrm{LODF}_{\ell,13}$	$\mathrm{LODF}_{\ell,23}$	P_ℓ, L12 out	P_ℓ, L13 out	P_ℓ, L23 out
L12	0.50	−0.25	–	1	−1	out	100	0
L13	0.50	0.25	1	–	1	100	out	100
L23	0.50	0.75	−1	1	–	0	100	out

LODFs are especially useful for screening many line outages quickly, and combinations of them extend the screening to double outages. They have two limits. First, equation (27.9) has the denominator $1 - \mathrm{PTDF}_{m,p\to q}$, which is zero when line m is the only path between its terminals: all of a transfer from p to q then flows on m, and its outage islands the network. In the intact triangle no line is such a bridge, but once L12 is out, L13 carries the whole transfer from bus 1 to bus 3, $\mathrm{PTDF}_{13,1\to 3} = 1$ in that configuration, and the LODF of a subsequent L13 outage is undefined; the islanding must then be detected and treated separately. Second, the factors are DC quantities and do not represent voltage collapse, reactive-power limits, protection dynamics or nonlinear AC behavior.

27.3.4 The AC power-flow model

The DC formulation is convenient, but real transmission security also depends on voltage magnitudes and reactive power. In the full AC model the active and reactive power injected at bus i are

$$P_i = \sum_{j=1}^{n} V_i V_j \left[G_{ij} \cos(\theta_i - \theta_j) + B_{ij} \sin(\theta_i - \theta_j) \right], \tag{27.10}$$

$$Q_i = \sum_{j=1}^{n} V_i V_j \left[G_{ij} \sin(\theta_i - \theta_j) - B_{ij} \cos(\theta_i - \theta_j) \right], \tag{27.11}$$

where n is the number of buses and $G_{ij} + \mathrm{j}B_{ij}$ is element (i, j) of the bus admittance matrix. Equation (27.10) reduces to the DC model under the assumptions of Section 27.2.2: with $G_{ij} = 0$, $V_i = V_j = 1$ and $\sin(\theta_i - \theta_j) \approx \theta_i - \theta_j$, it becomes $P_i = \sum_{j \neq i} B_{ij}(\theta_i - \theta_j)$, and for lossless lines without shunts the off-diagonal B_{ij} equals $1/x_{ij}$, so that $P_i = \sum_j (\theta_i - \theta_j)/x_{ij}$, which is equation (27.3); the DC matrix $\mathbf{B}$ is the negative of the susceptance part of the admittance matrix in that case.

An AC contingency analysis solves equations (27.10) and (27.11) for the post-contingency network and checks constraints such as

$$V_i^{\min} \leq V_i \leq V_i^{\max}, \qquad |S_{ij}| \leq S_{ij}^{\max}, \qquad Q_{\mathrm{G}i}^{\min} \leq Q_{\mathrm{G}i} \leq Q_{\mathrm{G}i}^{\max}, \tag{27.12}$$

where $S_{ij} = P_{ij} + \mathrm{j}Q_{ij}$ is the complex power on line ij with apparent-power rating $S_{ij}^{\max}$ in MVA, and $Q_{\mathrm{G}i}$ is the reactive output of generator i, together with the active-power limits of the generators, the stability requirements and the other operating requirements of the system. The solution must also be physically and operationally feasible: maintaining the voltages after an outage may require more reactive power than a generator can produce, even when every thermal limit appears acceptable, and a network whose DC flows are within ratings can then fail an AC check through a voltage limit, or have no AC solution at all.

27.3.5 Static and dynamic security

A steady-state AC solution, however detailed, describes the operating point that the system settles to after a contingency, if it settles. It does not establish that the system reaches that point. The transition involves the inception and clearing of the fault, the excursions of the generator rotor angles, the recovery of the voltages, the response of inverter controls, the frequency dynamics, the operation of protection, and possibly cascading disconnections. *Static security* assessment checks the post-contingency steady state; *dynamic security* assessment, through time-domain simulation, checks that the system survives the transition (transient stability, voltage stability, frequency stability). Transmission planning therefore uses several types of assessment, steady-state, short-circuit and stability evaluations, each across its defined planning events and operating conditions.

The numerical risk model of the rest of this chapter is explicitly restricted to post-contingency steady-state transfer feasibility under DC power flow: a network state is acceptable if the required transfer can be delivered within the thermal ratings after the stated corrective action. Voltage, reactive-power and stability limits, and the dynamics of reaching the post-contingency state, are outside the model, and the limits of Section 27.6.5 return to what that omission means for the results.

27.4 Probabilistic contingency analysis

The deterministic analysis found that every single outage requires curtailment. This section attaches probabilities to the outages, enumerates every network state, and computes the

transmission loss-of-load probability, the expected curtailment and the contribution of each state, then distinguishes outage frequency from snapshot unavailability and carries the result over a year.

27.4.1 Line outage probabilities

Let the probability of finding each line unavailable at a randomly selected assessment hour be $q_{12} = 0.02$, $q_{23} = 0.03$ and $q_{13} = 0.01$ (illustrative). These are stationary unavailabilities: point-in-time probabilities, the basic-event probabilities q_i of Chapter 10 of Volume I applied to lines. Define the outage indicator of line ℓ,

$$B_\ell = \begin{cases} 1, & \text{line } \ell \text{ is unavailable,} \\ 0, & \text{line } \ell \text{ is available,} \end{cases} \qquad \Pr(B_\ell = 1) = q_\ell, \tag{27.13}$$

as Chapter 25 defined it for a generator. Assume that the outage states of the three lines are independent. A network state s is a vector $\mathbf{b} = (b_{12}, b_{13}, b_{23})$ of indicator values, and its probability is the product of the marginal ones,

$$p_s = \Pr(B_\ell = b_\ell \text{ for every line } \ell) = \prod_\ell q_\ell^{b_\ell} \, (1 - q_\ell)^{1 - b_\ell}. \tag{27.14}$$

These are point-in-time unavailabilities. They are not annual outage frequencies and must not be read as such: $q_{23} = 0.03$ means that L23 is out of service in 3 % of the hours in the long run, not that it fails three times in a hundred years or three times a year. Section 27.4.6 relates the two. The unavailabilities are estimated from suitable outage-and-restoration models, as the steady-state unavailability $\lambda/(\lambda + \mu)$ of Chapter 19 of Volume I, or from observed outage records, as the fraction of the exposure time each line spent out of service.

27.4.2 Enumeration of the network states

With three two-state lines there are $2^3 = 8$ network configurations. Let a transfer of 100 MW be required from bus 1 to bus 3. For each configuration s define the transfer capability $\mathrm{TC}(s)$, the largest feasible transfer under that topology, and the curtailment

$$\Delta L(s) = \max\bigl(0,\ 100\,\mathrm{MW} - \mathrm{TC}(s)\bigr). \tag{27.15}$$

The intact network carries at most 140 MW (Section 27.2.3), the single outages 70 or 80 MW (Table 27.2). The double outages need one more argument. If L12 and L23 are both out, L13 still connects generation and load, bus 2 is isolated with nothing connected to it, and the transfer capability is the 80 MW of L13. If L13 is out together with L12 or L23, no path remains between bus 1 and bus 3: the generator and the load are in different islands, $\mathrm{TC} = 0$ and the whole 100 MW is curtailed. The same holds with all three lines out.

Table 27.4 lists the eight states. Each row is one configuration: the lines unavailable, the state probability from equation (27.14), the transfer capability, the curtailment, and the product of probability and curtailment, which is the state's contribution to the expected curtailment of Section 27.4.4.

The products are written with the factors in the order L12, L23, L13, a factor q_ℓ for each line out and $1 - q_\ell$ for each line in. The eight probabilities sum to 1, as they must for mutually exclusive and exhaustive states, and they are exact: every product of factors with two decimals has six decimals here.

The two islanding states have a structure that Chapter 11 of Volume I would recognize. The connection between bus 1 and bus 3 is a structure function of the three line states, and its minimal cut sets are $\{\mathrm{L13}, \mathrm{L12}\}$ and $\{\mathrm{L13}, \mathrm{L23}\}$: L13 together with either line of the

Table 27.4: The eight network states of the three-bus network for a required transfer of 100 MW, with independent line outages $q_{12} = 0.02$, $q_{23} = 0.03$, $q_{13} = 0.01$ (illustrative).

Lines unavailable	Product	p_s	TC(s)	$\Delta L(s)$	$p_s\,\Delta L(s)$
none	$0.98 \times 0.97 \times 0.99$	0.941094	140 MW	0 MW	0 MW
L13	$0.98 \times 0.97 \times 0.01$	0.009506	70 MW	30 MW	0.28518 MW
L23	$0.98 \times 0.03 \times 0.99$	0.029106	80 MW	20 MW	0.58212 MW
L12	$0.02 \times 0.97 \times 0.99$	0.019206	80 MW	20 MW	0.38412 MW
L23, L13	$0.98 \times 0.03 \times 0.01$	0.000294	0 MW	100 MW	0.02940 MW
L12, L13	$0.02 \times 0.97 \times 0.01$	0.000194	0 MW	100 MW	0.01940 MW
L12, L23	$0.02 \times 0.03 \times 0.99$	0.000594	80 MW	20 MW	0.01188 MW
all three	$0.02 \times 0.03 \times 0.01$	0.000006	0 MW	100 MW	0.00060 MW
Total		1.000000			1.31270 MW

indirect path. Islanding occurs exactly when a minimal cut set is entirely out, with probability $q_{13}\,[1 - (1 - q_{12})(1 - q_{23})] = 0.01 \times 0.0494 = 0.000494$, the sum of the last three islanding rows. The cut sets describe connectivity only; the curtailments of the connected states come from the power-flow model, not from the cut sets.

27.4.3 Transmission loss-of-load probability

In the modeled network curtailment is required whenever $\Delta L(s) > 0$. The probability of needing curtailment, the transmission loss-of-load probability, is the total probability of those states,

$$\mathrm{LOLP_T} = \Pr(\Delta L > 0) = \sum_s p_s\,\mathbf{1}_{\{\Delta L(s) > 0\}}. \tag{27.16}$$

At a transfer of 100 MW every state with at least one line out requires curtailment, so only the intact state avoids it, and

$$\mathrm{LOLP_T} = 1 - 0.941094 = 0.058906,$$

or 5.8906 %. Equivalently, $1 - (1 - q_{12})(1 - q_{23})(1 - q_{13})$, the probability that at least one line is out.

This is a snapshot probability for the network state model: the probability that the network state at a randomly selected instant requires curtailment. It is not automatically the probability of experiencing at least one curtailment during a future hour or year. That probability also depends on how often the states change, which needs a chronological outage and restoration model (Section 27.4.6).

27.4.4 Expected curtailment

The expected curtailment, the expected shortfall of Chapter 25 in this network, is

$$\mathbb{E}[\Delta L] = \sum_s p_s\,\Delta L(s). \tag{27.17}$$

Substituting the rows of Table 27.4,

$$\begin{aligned}\mathbb{E}[\Delta L] &= 0.009506(30) + 0.029106(20) + 0.019206(20) + 0.000294(100) \\ &\quad + 0.000194(100) + 0.000594(20) + 0.000006(100) \\ &= 0.28518 + 0.58212 + 0.38412 + 0.02940 + 0.01940 + 0.01188 + 0.00060 = 1.3127\,\mathrm{MW}.\end{aligned}$$

The same number follows from a grouping of the states by consequence, which gives a closed form valid for any q_ℓ. The curtailment is 30 MW when L13 alone is out, 20 MW when L13 is in and at least one of L12 and L23 is out, and 100 MW when L13 is out with at least one of the others. Writing $u = 1 - (1 - q_{12})(1 - q_{23})$ for the probability that the indirect path is broken,

$$\mathbb{E}[\Delta L] = 30\, q_{13}(1 - u) + 20\,(1 - q_{13})\, u + 100\, q_{13}\, u, \tag{27.18}$$

in MW, and with $u = 1 - 0.98 \times 0.97 = 0.0494$: $30 \times 0.01 \times 0.9506 + 20 \times 0.99 \times 0.0494 + 100 \times 0.01 \times 0.0494 = 0.28518 + 0.97812 + 0.04940 = 1.3127$ MW.

If this snapshot distribution applies to a randomly selected hour and the demand is constant through the hour, the expected unserved energy of that hour is

$$\mathrm{EUE}_{1\,\mathrm{h}} = 1.3127\,\mathrm{MW} \times 1\,\mathrm{h} = 1.3127\,\mathrm{MWh}.$$

The conditional curtailment, by the argument of Chapter 25 (equation (25.11)), is $\mathbb{E}[\Delta L \mid \Delta L > 0] = 1.3127/0.058906 = 22.28$ MW: when curtailment is needed, it averages about a fifth of the transfer. As in the generation case, these numbers represent the expected deficiency under the assumed corrective-action policy. They do not establish an actual customer interruption without the dispatch, the curtailment procedure and the emergency response being specified.

27.4.5 Severity and frequency

The states can be ranked in three ways. By probability, the most likely curtailment state is the outage of L23 alone, $p = 0.029106$, or 2.9106 %. By severity, the worst states are those that disconnect bus 1 from bus 3, with $\Delta L = 100$ MW; the outage of L13 and L23 together is one of them, much less probable (0.000294) but five times as severe. A risk-weighted ranking combines the two through the contribution of each state to the expected curtailment,

$$\mathcal{R}_s = p_s\, \Delta L(s), \tag{27.19}$$

the risk of Chapter 1 of Volume I restricted to one state, with the curtailment as the consequence. For a one-hour snapshot $\mathcal{R}_s$ is the state's contribution to the expected curtailment in MW, or equivalently to the expected unserved energy in MWh over that hour.

Figure 27.4 draws the seven nonzero contributions. The horizontal axis lists the network states by the lines out; the vertical axis is $\mathcal{R}_s$ in MW. Blue bars are single-line outages, red bars multiple outages; the number above each bar is its value. The three single outages dominate: together they contribute 1.25142 MW, or 95.3 % of the expected curtailment, and the outage of L23 alone is the largest contributor because it combines a high probability with a moderate consequence. The three islanding states contribute 0.0494 MW (3.8 %), and the double outage of L12 and L23 0.01188 MW (0.9 %).

For this model the single-line outages therefore dominate the expected curtailment. Severe multiple outages may dominate other risk measures: the probability of complete interruption of the transfer, 0.000494, comes entirely from them, as does every curtailment above 30 MW, the tail of the distribution of ΔL. A ranking by expected value and a ranking by tail consequence answer different questions, and a study must say which it reports.

Principle 27.1. *In probabilistic contingency analysis each network state contributes the product of its probability and its consequence. The most probable contingency, the most severe contingency and the largest contributor to expected curtailment need not be the same, and a ranking by severity alone misallocates attention.*

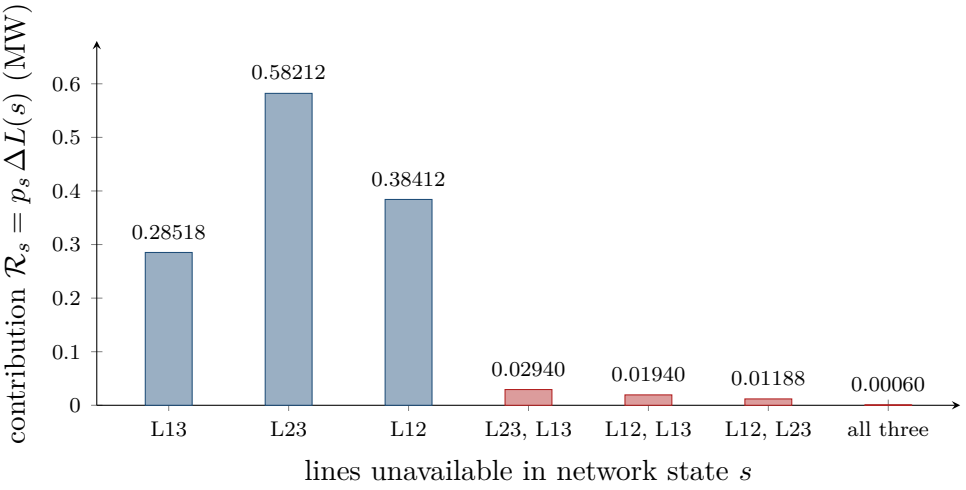


Figure 27.4: Contribution of each network state to the expected curtailment of 1.3127 MW at a transfer of 100 MW (Table 27.4; illustrative unavailabilities $q_{12} = 0.02$, $q_{23} = 0.03$, $q_{13} = 0.01$). Blue: single-line outages; red: multiple outages, of which the first two and the last island the load.

27.4.6 Outage frequency and snapshot unavailability

Suppose outages of line ℓ occur with frequency f_ℓ, in events per year, and each is followed by a restoration. The frequency cannot be substituted for the stationary unavailability q_ℓ: the first counts events per unit time, the second is the fraction of time spent out, and the two differ by the duration of the outages. If the line alternates between in-service periods of mean MTTF_ℓ and outages of mean MTTR_ℓ, it is an alternating renewal process, and its stationary unavailability is the standard result

$$q_\ell = \frac{\mathrm{MTTR}_\ell}{\mathrm{MTTF}_\ell + \mathrm{MTTR}_\ell}, \tag{27.20}$$

which for exponential durations, with outage rate λ_ℓ and restoration rate μ_ℓ, is $q_\ell = \lambda_\ell/(\lambda_\ell + \mu_\ell)$, the steady-state result of Chapter 19. The stationary frequency of outages, the frequency of transitions into the failed state, is the probability of being in service times the outage rate,

$$f_\ell = (1 - q_\ell)\,\lambda_\ell = \frac{1}{\mathrm{MTTF}_\ell + \mathrm{MTTR}_\ell}, \tag{27.21}$$

the failure frequency $f_\mathrm{F} = \pi_\mathrm{up}\lambda$ of Chapter 19 (equation (19.43)). Multiplying the two forms gives the frequency-and-duration relation

$$q_\ell = f_\ell\,\mathrm{MTTR}_\ell, \tag{27.22}$$

which is exact for a stationary process (equation (19.45)), with f_ℓ and MTTR_ℓ in consistent time units. The approximation is the replacement of the frequency by the outage rate: $q_\ell \approx \lambda_\ell\,\mathrm{MTTR}_\ell$, which overstates q_ℓ by the factor $1/(1 - q_\ell)$ and is accurate when outages are rare and short against the in-service periods.

Table 27.5 gives an illustrative set of frequencies and durations consistent with the unavailabilities of Section 27.4.1. Each row is a line, with its outage frequency, its mean outage duration of 87.6 h (about 3.65 days), the unavailability from equation (27.22), and the outage rate $\lambda_\ell = f_\ell/(1-q_\ell)$. For L12, $q_{12} = 2\,\mathrm{yr}^{-1} \times 87.6\,\mathrm{h}/8760\,\mathrm{h} = 0.02$ and $\lambda_{12} = 2/0.98 = 2.041\,\mathrm{yr}^{-1}$, while $\lambda_{12}\,\mathrm{MTTR}_{12} = 0.0204$ overstates q_{12} by 2 %.

Table 27.5: Outage frequency, mean outage duration, unavailability and outage rate of the three lines (illustrative); the restoration rate is $\mu_\ell = 1/87.6\,\text{h} = 100\,\text{yr}^{-1}$ for each line.

Line	Frequency f_ℓ	Duration MTTR_ℓ	$q_\ell = f_\ell\,\text{MTTR}_\ell$	Rate λ_ℓ
L12	$2.0\,\text{yr}^{-1}$	87.6 h	0.02	$2.041\,\text{yr}^{-1}$
L23	$3.0\,\text{yr}^{-1}$	87.6 h	0.03	$3.093\,\text{yr}^{-1}$
L13	$1.0\,\text{yr}^{-1}$	87.6 h	0.01	$1.010\,\text{yr}^{-1}$

Common mistake

Substituting an outage frequency for an unavailability. A line that fails twice a year has $f = 2\,\text{yr}^{-1}$, a number that cannot be a probability; its unavailability is $2\,\text{yr}^{-1}$ times the mean outage duration in years, 0.02 for outages of 87.6 h, and would be ten times smaller for outages of 8.76 h with the same frequency.

Outages lasting several hours or days span many assessment intervals, so counting curtailment *events* and their durations needs a chronological model. Under the two-state Markov model of each line, with independent lines, a curtailment event begins whenever the network leaves the intact state, since at a transfer of 100 MW every other state requires curtailment. The frequency of leaving the intact state is its probability times the sum of the outage rates,

$$f_{\text{curt}} = 0.941094 \times (2.041 + 3.093 + 1.010)\,\text{yr}^{-1} = 0.941094 \times 6.1437 = 5.78 \text{ events per year.}$$

The expected time with curtailment is $0.058906 \times 8760 = 516.0$ h per year, so by the frequency-and-duration relation the mean duration of a curtailment event is $516.0/5.78 = 89.2$ h, slightly longer than one line outage because overlapping outages extend some events.

27.4.7 Expected unserved energy over a year

For each hour t let $\mathbf{b}_t$ be the network state and L_t the demand, and let $\Delta L_t^*(\mathbf{b}_t, L_t)$ be the minimum load curtailment required after enforcing the physical network constraints with the available corrective controls (formulated in Section 27.5.5). The expected unserved energy over a year is then

$$\text{EUE} = \sum_{t=1}^{8760} \mathbb{E}\big[\Delta L_t^*(\mathbf{b}_t, L_t)\big]\,\Delta t, \tag{27.23}$$

with $\Delta t = 1$ h, by linearity of expectation and without any assumption of independence between hours, as in Chapter 25 (equation (25.12)).

If every condition of the illustrative model, the transfer of 100 MW, the ratings and the unavailabilities, stayed the same for 8760 h, the model would give

$$\text{EUE} = 8760\,\text{h} \times 1.3127\,\text{MW} = 11\,499.252\,\text{MWh per year,}$$

with 516.0 h of expected curtailment. Such an extrapolation is unrealistic for most transmission systems. The demand, the generation dispatch, the line ratings, the weather, the planned outages and the restoration dynamics all vary over the year, and a meaningful annual assessment represents those changes hour by hour in equation (27.23). The calculation is also conditional on the response model: with a corrective action that restored service in some states, or a preventive limit on the transfer, the same outages would give a different EUE.

27.5 Uncertain line ratings and risk-informed corrective actions

So far the line ratings were exact constants and the only corrective action was curtailment. This section treats the rating as a random variable derived from the conductor heat balance,

computes the overload probability and the expected curtailment it implies, and then formulates the general corrective-action problem and the expected consequence of a set of contingencies.

27.5.1 Uncertainty of transmission capacity

The thermal capability of an overhead line depends on the ambient air temperature, the wind speed and direction, the solar heating, the conductor temperature and the condition and geometry of the conductor. A line rating is not merely a statistical number: it is a limit derived from thermal, sag, clearance and other operating requirements, which Chapter 1 of Volume I expressed through the conductor heat balance (equation (1.12)),

$$mc_p \frac{\mathrm{d}T}{\mathrm{d}t} = I^2 R(T) + q_{\mathrm{solar}} - q_{\mathrm{conv}} - q_{\mathrm{rad}}, \tag{27.24}$$

with T the conductor temperature, mc_p the heat capacity of the modeled length, I the current, $R(T)$ the temperature-dependent electrical resistance (not the reliability function), q_{solar} the solar gain and q_{conv} and q_{rad} the convective and radiative losses. The steady-state rating is the current that holds the conductor exactly at its temperature limit $T_{\max}$: setting $\mathrm{d}T/\mathrm{d}t = 0$ and $T = T_{\max}$ gives

$$I_{\max} = \sqrt{\frac{q_{\mathrm{conv}}(T_{\max}) + q_{\mathrm{rad}}(T_{\max}) - q_{\mathrm{solar}}}{R(T_{\max})}}, \tag{27.25}$$

and the active-power rating used in the DC model is approximately $P^{\max} \approx \sqrt{3}\, V_{\mathrm{nom}}\, I_{\max}$ at nominal line-to-line voltage V_{nom} and unity power factor. The convective loss grows with the wind speed and falls as the ambient temperature rises, so a cool, windy hour gives a high rating and a hot, still, sunny hour a low one.

Uncertain near-surface weather therefore produces uncertainty in the thermal state of the line and in its feasible current-carrying capacity. In a probabilistic assessment the rating of line ℓ is represented as

$$P_\ell^{\max} = P_\ell^{\max}(Z, \boldsymbol{\theta}), \tag{27.26}$$

where Z is the weather state (ambient temperature, wind, solar radiation), as in Chapter 13 of Volume I, and $\boldsymbol{\theta}$ is a vector of uncertain parameters of the thermal model, such as the absorptivity and emissivity of the conductor surface and the coefficients of its resistance. Here $\boldsymbol{\theta}$ is a parameter vector, not the vector of bus angles of Section 27.2.2.

27.5.2 Overload probability

Let $P_\ell(s)$ be the flow on line ℓ in network state s after the contingency and $P_\ell^{\max}$ its uncertain rating. The overload event is

$$|P_\ell(s)| > P_\ell^{\max}, \tag{27.27}$$

and its probability conditional on the state is

$$\Pr(\text{overload of } \ell \mid s) = \Pr\big(P_\ell^{\max} < |P_\ell(s)| \mid s\big). \tag{27.28}$$

If the rating is independent of the topology and the flow is deterministic given the contingency, the conditioning on s drops from the rating, and the overload probability is the cumulative distribution function of the rating evaluated at the flow, $\Pr(P_\ell^{\max} < x)$ at $x = |P_\ell(s)|$. This is the stress-strength interference of a load against a random capacity, with the flow as the load.

The independence assumption may be wrong. Weather drives both sides: a severe wind event raises the outage probability of exposed lines while it increases the convective cooling of the lines that survive, and a heat wave raises the demand and the flows while it lowers the

ratings. In such cases the overload probability is computed conditionally on the weather and then averaged,

$$\Pr(\text{overload of } \ell \mid s) = \sum_{z} \Pr(P_\ell^{\max} < |P_\ell(s)| \mid Z = z)\ \Pr(Z = z \mid s), \tag{27.29}$$

where the weather distribution is itself conditioned on the network state, since a state with lines out is more likely in storm weather. This is the mechanism of conditional dependence through a shared environment of Chapter 13 (Section 13.4), and weather-conditioned rating and outage models are often the more appropriate ones.

27.5.3 Worked example: an uncertain post-contingency rating

After an outage of L12, line L13 carries $P_{13} = 100\,\text{MW}$ (Section 27.2.5). Suppose its available rating in the hour considered is uncertain and modeled, as an illustrative approximation, by a normal distribution,

$$P_{13}^{\max} \sim \mathcal{N}(90,\ 10^2)\,\text{MW},$$

with mean 90 MW and standard deviation 10 MW. A physical rating model might instead use a truncated or otherwise non-normal distribution, to exclude implausible values such as negative ratings, whose probability here is $\Phi(-9)$, negligible. The overload probability is $\Pr(P_{13}^{\max} < 100)$. Standardizing,

$$z = \frac{100 - 90}{10} = 1, \qquad \Pr(P_{13}^{\max} < 100) = \Phi(1) = 0.8413,$$

where Φ is the standard normal cumulative distribution function.

Figure 27.5 draws this calculation. The horizontal axis is the rating in MW; the curve is the normal density of $P_{13}^{\max}$; the vertical red line is the post-contingency flow of 100 MW, and the shaded area to its left, the probability that the rating is below the flow, is the overload probability 0.8413. The dashed gray line marks the fixed rating of 80 MW used so far. The flow is very likely to exceed the uncertain rating, but in 15.87 % of the hours the cooling is favorable enough for L13 to carry the whole transfer. A fixed 80 MW rating classifies this post-contingency condition as overloaded with certainty.

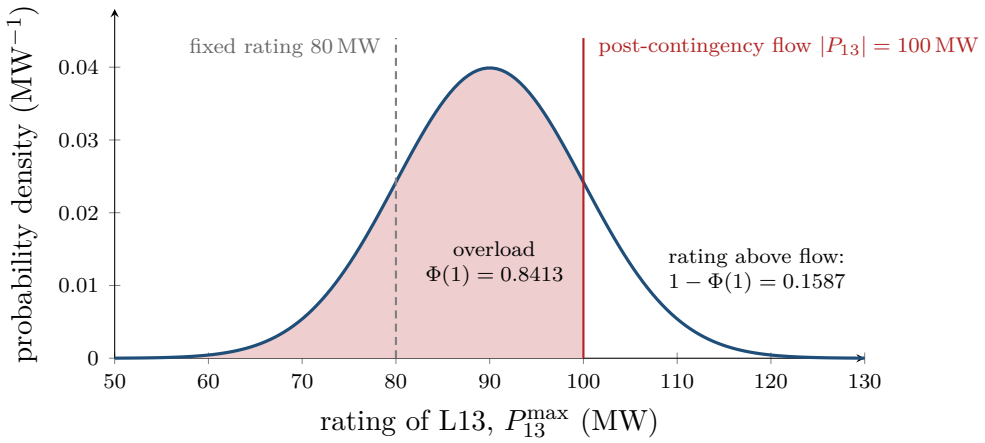


Figure 27.5: Overload probability of L13 after the outage of L12, with the rating modeled as normal, $P_{13}^{\max} \sim \mathcal{N}(90, 10^2)\,\text{MW}$ (illustrative), against the post-contingency flow of 100 MW. The shaded area is $\Pr(P_{13}^{\max} < 100) = \Phi(1) = 0.8413$; the density is $\exp[-(x-90)^2/200]/(10\sqrt{2\pi})$.

27.5.4 Expected curtailment with an uncertain rating

Suppose the only feasible transfer path has the random rating $P^{\max}$ and the required transfer is 100 MW. Under the curtailment model of Section 27.2.4, the required curtailment is

$$\Delta L = (100 - P^{\max})^{+}. \tag{27.30}$$

For a normal rating, $P^{\max} \sim \mathcal{N}(\mu_R, \sigma_R^2)$, its expectation has a closed form. Write $P^{\max} = \mu_R + \sigma_R U$ with U standard normal, and $a = 100$; then $(a - P^{\max})^{+} = \sigma_R\,(z - U)^{+}$ with $z = (a - \mu_R)/\sigma_R$. With φ the standard normal density,

$$\mathbb{E}[(z-U)^{+}] = \int_{-\infty}^{z} (z-u)\,\varphi(u)\,\mathrm{d}u = z\,\Phi(z) - \int_{-\infty}^{z} u\,\varphi(u)\,\mathrm{d}u = z\,\Phi(z) + \varphi(z),$$

because $u\,\varphi(u) = -\mathrm{d}\varphi/\mathrm{d}u$, so that the second integral is $-\varphi(z)$. Multiplying by σ_R,

$$\mathbb{E}[\Delta L] = (a - \mu_R)\,\Phi(z) + \sigma_R\,\varphi(z), \qquad z = \frac{a - \mu_R}{\sigma_R}. \tag{27.31}$$

The symbol φ is used for the standard normal density so that it is not confused with the structure function ϕ of Chapter 9 of Volume I. With $\mu_R = 90$, $\sigma_R = 10$ and $z = 1$, $\Phi(1) = 0.84134$ and $\varphi(1) = 0.24197$, so

$$\mathbb{E}[\Delta L] = 10 \times 0.84134 + 10 \times 0.24197 = 8.4134 + 2.4197 = 10.83\,\text{MW}.$$

This is a probabilistic alternative to treating the thermal rating as an exact constant. A deterministic rating equal to the mean, 90 MW, would give $\Delta L = 10$ MW; the extra 0.83 MW is the effect of the uncertainty itself. Since $x \mapsto (a - x)^{+}$ is convex, Jensen's inequality gives $\mathbb{E}[(a - P^{\max})^{+}] \geq (a - \mathbb{E}[P^{\max}])^{+}$, as for the generation shortfall of Chapter 25 (equation (25.16)): a rating known only in distribution costs more expected curtailment than a rating fixed at its mean.

Table 27.6 applies equation (27.31) to several rating distributions. Each row is a mean and a standard deviation (a standard deviation of zero is a fixed rating), with the overload probability $\Phi(z)$ and the expected curtailment. Figure 27.6 draws the expected curtailment against the mean rating for three standard deviations. The horizontal axis is the mean rating μ_R; the vertical axis is $\mathbb{E}[\Delta L]$ in MW; the black broken line is the fixed rating, $(100 - \mu_R)^{+}$, and the two smooth curves are the normal ratings with $\sigma_R = 5$ and 10 MW. The curves lie above the broken line everywhere and depart from it most near $\mu_R = 100$ MW, where the fixed rating would just carry the transfer and an uncertain one fails to in about half the hours; far below 100 MW they approach the line, since the transfer is then curtailed in almost every outcome.

The single-path result enters the network calculation through the states in which L13 alone carries the transfer: L12 out, L23 out, and both out, of total probability $0.019206 + 0.029106 + 0.000594 = 0.048906$. With the rating of L13 normal, $\mathcal{N}(90, 10^2)$ MW in every hour and independent of the line outages, each of those states contributes 10.833 MW of expected curtailment instead of 20 MW; the state with L13 alone out and the islanding states are unchanged, and the intact state curtails only if $2P_{13}^{\max} < 100$, with probability $\Phi(-4) = 0.00003$. The expected curtailment of the network becomes

$$\mathbb{E}[\Delta L] = 0.28518 + 0.04940 + 0.048906 \times 10.833 + 0.00013 = 0.8645\,\text{MW},$$

against 1.3127 MW with the fixed 80 MW rating, and the transmission loss-of-load probability falls from 0.058906 to 0.0512. The decrease is not an effect of uncertainty: the mean rating of 90 MW is higher than the fixed 80 MW. With a normal rating of mean 80 MW and standard deviation 10 MW, the same calculation gives 1.3241 MW and 0.0591, slightly above the fixed-rating values, which is the effect of uncertainty alone. What a probabilistic rating changes, against a conservative fixed one, is mostly the mean; its spread adds a smaller, always positive, amount.

Table 27.6: Overload probability and expected curtailment for a post-contingency flow of 100 MW on a single path with a normal rating (illustrative); a standard deviation of zero is a fixed rating.

Mean μ_R	Std. dev. σ_R	z	Pr(overload)	$\mathbb{E}[\Delta L]$
80 MW	0	–	1	20.000 MW
80 MW	10 MW	2.0	0.9772	20.085 MW
90 MW	0	–	1	10.000 MW
90 MW	5 MW	2.0	0.9772	10.042 MW
90 MW	10 MW	1.0	0.8413	10.833 MW
90 MW	20 MW	0.5	0.6915	13.956 MW
95 MW	10 MW	0.5	0.6915	6.978 MW
100 MW	10 MW	0.0	0.5000	3.989 MW

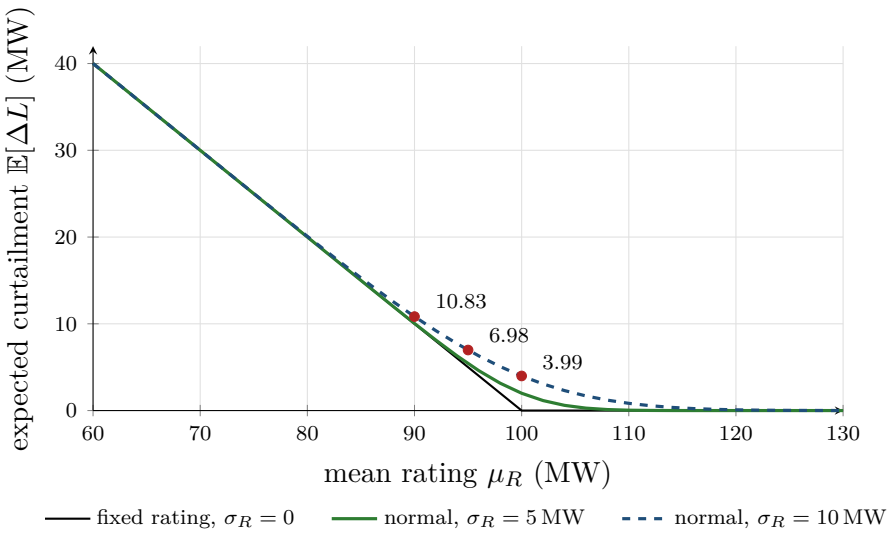


Figure 27.6: Expected curtailment of a 100 MW transfer on a single path, $\mathbb{E}[(100 - P^{\max})^+]$, against the mean rating, for a fixed rating and for normal ratings with standard deviations 5 and 10 MW, from equation (27.31) (illustrative). Dots: the rows of Table 27.6 with $\sigma_R = 10$ MW.

A physical limitation. A thermal rating generally specifies a permissible loading condition, not a random instantaneous trip threshold. Because of the heat capacity mc_p in equation (27.24), the conductor temperature responds to an overload over minutes to tens of minutes, with the thermal time constant τ_{th} that the next chapter uses, and exceeding a rating requires corrective action within a specified time rather than immediate curtailment. Emergency ratings exploit exactly this delay. The equation $\Delta L = (100 - P^{\max})^+$ is therefore a simplified model of dispatch feasibility, not a model of conductor failure or of protection operation.

27.5.5 Redispatch and emergency controls

The network of this chapter has one generator and one load, so the only corrective action available was to curtail. In a real transmission network the corrective actions include generation redispatch, topology switching, demand response, battery discharge, phase-shifting transformers, flexible AC transmission (FACTS) devices, remedial action schemes and controlled load curtailment. The consequence of a contingency depends on which of these are available and on how quickly they can be deployed.

The choice among them is a post-contingency optimization. For network state s with susceptance matrix $\mathbf{B}^{(s)}$, the minimum curtailment is

$$\Delta L^*(s) = \min_{\mathbf{P}^{\mathrm{G}},\, \Delta \boldsymbol{L},\, \boldsymbol{\theta}} \; \sum_n \Delta L_n \tag{27.32}$$

subject to

$$\begin{aligned} \mathbf{B}^{(s)}\boldsymbol{\theta} &= \mathbf{P}^{\mathrm{G}} - \boldsymbol{L} + \Delta \boldsymbol{L}, \\ \left|\frac{\theta_i - \theta_j}{x_{ij}}\right| &\le P_{ij}^{\max}(s) \quad \text{for every line } ij \text{ in service in } s, \\ P_i^{\mathrm{G,min}} \le P_i^{\mathrm{G}} &\le P_i^{\mathrm{G,max}}, \\ 0 \le \Delta L_n &\le L_n, \end{aligned}$$

and any relevant control, ramping or reserve constraints. Here $\mathbf{P}^{\mathrm{G}}$ is the vector of generator outputs, $\boldsymbol{L}$ the vector of bus demands and ΔL_n the demand curtailed at bus n; the first constraint is the DC balance of equation (27.3) with the curtailment acting as a nonnegative injection, and the rating may depend on the state through the weather. With DC power flow and linear constraints the problem is a linear program; with the AC equations it is nonlinear and may be nonconvex. It is the network form of the minimum load shedding of Chapter 25 (equation (25.21)), and for the three-bus network with generation only at bus 1 its solution is the curtailment of equation (27.15).

As an illustration of how a corrective action acts on some contingencies and not others, add to the network a 30 MW unit at bus 2, normally off and available for redispatch (a hypothetical addition). After the outage of L12, bus 1 can send 80 MW over L13 and the new unit 20 MW over L23, within its 70 MW rating, and no load is curtailed. After the outage of L13 it does not help: every megawatt from either bus 1 or bus 2 must cross L23, whose rating still limits the delivery to 70 MW. After the outage of L23 it does not help either: its output can reach bus 3 only through bus 1 and L13, whose 80 MW is already used. When L12 and L13 are both out, the unit supplies 30 MW of the otherwise islanded load, and the curtailment falls from 100 to 70 MW. The redispatch option therefore changes only two rows of Table 27.4, and the expected curtailment falls from 1.3127 to

$$1.3127 - 0.38412 - 0.000194 \times 30 = 0.9228\,\text{MW}.$$

27.5.6 Expected contingency consequence

Let contingency, or network state, s occur with probability p_s during an assessment interval. The expected consequence is

$$\mathbb{E}[\Delta L] = \sum_{s} p_s \, \Delta L^{*}(s), \tag{27.33}$$

and when the consequences are evaluated over a sequence of intervals t, each with its own state probabilities $p_{t,s}$, the expected unserved energy is the double sum

$$\mathrm{EUE} = \sum_{t} \sum_{s} p_{t,s} \, \Delta L^{*}_{t}(s) \, \Delta t, \tag{27.34}$$

which is equation (25.22) of Chapter 25 with the curtailment of problem (27.32). This is the link between contingency enumeration and probabilistic network adequacy.

A model of minimum load shedding assumes that the operator chooses the optimal corrective action in every state, and does so in time. Actual actions are limited by detection delays, control failures, communication outages, operator response times and stability constraints, and those limitations must be included when they affect the intended risk metric. The simplest way is the event-tree structure of Chapter 15 of Volume I (Section 15.4), whose transmission-line tree gave the redispatch branch a success probability of 0.90. Applied to the redispatch of the previous subsection, with the unit at bus 2 starting with probability 0.90 and the curtailment of Table 27.4 applying when it fails, the expected curtailment is

$$0.9 \times 0.9228 + 0.1 \times 1.3127 = 0.9618 \,\mathrm{MW},$$

between the value with no redispatch and the value with perfect redispatch. The corrective action is worth 0.39 MW of expected curtailment if it always works and 0.35 MW at a success probability of 0.90.

27.6 Risk attribution, parameter uncertainty and the limits of contingency analysis

A risk number is most useful when it is broken down by the components and parameters that produce it. This section attributes the expected curtailment to the individual lines, compares improvements of line availability with increases of line rating, tabulates the indices for several parameter sets, and closes with the uncertainty of the parameters and the adequacy of the model itself.

27.6.1 Importance of each line

Suppose the reliability of only one line can be improved. A natural measure of the benefit is the reduction of the expected curtailment obtained by eliminating that line's unavailability while the other line probabilities stay unchanged,

$$\Delta\mathcal{R}^{-}_{\ell} = \mathbb{E}[\Delta L] - \mathbb{E}[\Delta L \mid q_\ell = 0], \tag{27.35}$$

the absolute risk reduction $Q_{\mathrm{sys}} - Q(0_i)$ of Chapter 12 of Volume I (equation (12.17)), with the expected curtailment in the place of the system failure probability. It is an intervention comparison under the stated stochastic model: setting $q_\ell = 0$ removes the line's outage state from the model. It is not the same as conditioning on having observed the line available, which in a model with shared causes, such as weather acting on several lines, would also change the probabilities of the other lines.

The expected curtailment is a linear function of each q_ℓ with the others fixed, since by the law of total probability on the state of line ℓ, $\mathbb{E}[\Delta L] = (1 - q_\ell)\,\mathbb{E}[\Delta L \mid q_\ell = 0] + q_\ell\,\mathbb{E}[\Delta L \mid q_\ell = 1]$.

Its derivative is the Birnbaum-type importance of Chapter 12 (equation (12.12)), in MW per unit of unavailability,

$$I_\ell^{\mathrm{B}} = \frac{\partial\, \mathbb{E}[\Delta L]}{\partial q_\ell} = \mathbb{E}[\Delta L \mid q_\ell = 1] - \mathbb{E}[\Delta L \mid q_\ell = 0], \tag{27.36}$$

and the reduction is $\Delta\mathcal{R}_\ell^- = q_\ell\, I_\ell^{\mathrm{B}}$, as in equation (12.19). Differentiating the closed form of equation (27.18) gives the importances directly: $I_{13}^{\mathrm{B}} = 30(1-u) + 80u = 28.518 + 3.952 = 32.47\,\mathrm{MW}$, $I_{12}^{\mathrm{B}} = (1-q_{23})(20 + 50\,q_{13}) = 0.97 \times 20.5 = 19.885\,\mathrm{MW}$ and $I_{23}^{\mathrm{B}} = (1-q_{12})(20 + 50\,q_{13}) = 0.98 \times 20.5 = 20.09\,\mathrm{MW}$.

For L12, the reduction is checked by the enumeration with $q_{12} = 0$: the remaining curtailment states are L13 alone, $0.97 \times 0.01 = 0.0097$ with 30 MW; L23 alone, $0.03 \times 0.99 = 0.0297$ with 20 MW; and L23 with L13, $0.03 \times 0.01 = 0.0003$ with 100 MW, so $\mathbb{E}[\Delta L \mid q_{12} = 0] = 0.291 + 0.594 + 0.030 = 0.915\,\mathrm{MW}$ and $\Delta\mathcal{R}_{12}^- = 1.3127 - 0.915 = 0.3977\,\mathrm{MW} = 0.02 \times 19.885$. Table 27.7 gives the three lines. Each row is a line, with its unavailability, the expected curtailment after its unavailability is eliminated, the reduction, the Birnbaum-type importance, the Fussell–Vesely importance $\Delta\mathcal{R}_\ell^-/\mathbb{E}[\Delta L]$ and the risk reduction worth $\mathbb{E}[\Delta L]/\mathbb{E}[\Delta L \mid q_\ell = 0]$.

Table 27.7: Importance of each line for the expected curtailment of 1.3127 MW at a transfer of 100 MW (illustrative unavailabilities of Table 27.4).

Line	q_ℓ	$\mathbb{E}[\Delta L \mid q_\ell = 0]$	Reduction $\Delta\mathcal{R}_\ell^-$	Birnbaum I_ℓ^{B}	Fussell–Vesely	RRW
L12	0.02	0.9150 MW	0.3977 MW	19.885 MW	0.303	1.435
L23	0.03	0.7100 MW	0.6027 MW	20.090 MW	0.459	1.849
L13	0.01	0.9880 MW	0.3247 MW	32.470 MW	0.247	1.329

Under these assumptions the ranking by risk reduction is

$$\Delta\mathcal{R}_{23}^- > \Delta\mathcal{R}_{12}^- > \Delta\mathcal{R}_{13}^-,$$

$0.6027 > 0.3977 > 0.3247$ MW. L23 offers the largest modeled benefit because its unavailability is the highest, although the outage of L13 requires the largest single-contingency curtailment and has the largest Birnbaum importance: per unit change of unavailability, L13 matters most, but L13 is rarely out, and eliminating an unavailability of 0.01 buys less than eliminating one of 0.03. A ranking by contingency severity alone, or by sensitivity alone, can therefore differ from a ranking by expected risk reduction, and which ranking is relevant depends on the intervention: a measure that eliminates outages favors L23; one that lowers every unavailability by the same small amount favors L13.

27.6.2 Risk reduction from line capacity

Instead of reducing outage probabilities, one line's rating may be increased, by reconductoring, by dynamic line rating based on measured weather, or by relaxing a clearance limit. Raising the rating of L13 from 80 to 100 MW removes the curtailment whenever L12 or L23 (or both) is out and L13 is in service, since L13 can then carry the whole transfer. It does not help when L13 itself is out, and it does not resolve the islanding caused by L13 together with another line. Raising both L12 and L23 from 70 to 100 MW instead removes the curtailment of an isolated L13 outage, but not of the islanding configurations. The two interventions act on disjoint sets of contingencies.

Table 27.8 quantifies them. Each row is one intervention, with the network states whose consequence improves, the expected curtailment afterward, its reduction and the transmission

loss-of-load probability afterward. Raising L13 to 100 MW removes the contributions $0.38412 + 0.58212 + 0.01188 = 0.97812$ MW, leaving 0.33458 MW, a reduction of three quarters; raising L12 and L23 removes only the 0.28518 MW of the L13 outage. With all three upgrades, only the islanding states remain, 0.0494 MW.

Table 27.8: Expected curtailment and transmission loss-of-load probability after rating increases (illustrative; unavailabilities of Table 27.4, transfer 100 MW).

Intervention	States improved	$\mathbb{E}[\Delta L]$ after	Reduction	$\mathrm{LOLP_T}$ after
none	–	1.3127 MW	–	0.058906
L13 rating $80 \to 100$ MW	L12 out; L23 out; both out	0.3346 MW	0.9781 MW	0.010000
L12, L23 ratings $70 \to 100$ MW	L13 out alone	1.0275 MW	0.2852 MW	0.049400
all three to 100 MW	every connected state	0.0494 MW	1.2633 MW	0.000494

The economic value of a rating increase depends on how often the additional rating is used and on the consequences it avoids, not on the rating gain in MW. A risk-importance study therefore reports both the states whose consequences improve and the total probability-weighted reduction, as the two middle columns of the table do.

27.6.3 Curtailment probability and expected curtailment for several outage probabilities and ratings

The model of this chapter, generation at bus 1, load at bus 3, equal reactances of the direct and the two-line paths, no generation or load at bus 2, independent point-in-time line outages, and curtailment optimized to respect the DC thermal limits, gives the transfer capability of each state in closed form: $2\min(P_{12}^{\max}, P_{23}^{\max}, P_{13}^{\max})$ for the intact network (because the two paths split the transfer equally), $\min(P_{12}^{\max}, P_{23}^{\max})$ with L13 out alone, $P_{13}^{\max}$ with L13 in and at least one other line out, and 0 with L13 out together with another line. With the required transfer L, each state curtails $(L - \mathrm{TC})^+$, and the indices are the transmission loss-of-load probability, the expected curtailment (which for a one-hour interval equals the one-hour EUE in MWh), the probability of complete interruption $\Pr(\Delta L \geq L)$, and the conditional curtailment.

Table 27.9 lists them for the base case and nine variations. The first block varies the required transfer; the second the ratings; the third the unavailabilities; the last row combines a higher transfer with a higher L13 rating. The row with $L = 130$ MW shows the arithmetic: the intact network carries 140 MW and curtails nothing; L13 out alone curtails $130 - 70 = 60$ MW; the three states in which L13 alone carries the transfer curtail $130 - 80 = 50$ MW; the islanding states curtail 130 MW. Hence $\mathbb{E}[\Delta L] = 0.009506 \times 60 + 0.048906 \times 50 + 0.000494 \times 130 = 0.57036 + 2.44530 + 0.06422 = 3.0799$ MW.

Four trends follow. First, raising L13's rating to 100 MW lowers the expected curtailment to a quarter, because L13 can then carry the full transfer whenever it remains connected, even with L12 or L23 out; the remaining curtailment comes from the outage of L13 itself and from islanding. Second, eliminating the unavailability of L23, the line with the highest outage probability, removes 0.6027 MW, almost half of the expected curtailment, and the probability of complete interruption falls from 0.000494 to 0.000200, the islanding probability through the cut set $\{\mathrm{L13}, \mathrm{L12}\}$ alone. Doubling every unavailability roughly doubles both indices and quadruples the probability of complete interruption, which is a product of two unavailabilities.

Third, the indices respond to the operating point in a strongly nonlinear way. At 70 MW the network is $N-1$ secure and only the islanding states curtail; at 80 MW only the outage of L13

Table 27.9: Transmission loss-of-load probability, expected curtailment, probability of complete interruption and conditional curtailment of the three-bus network for several transfers, ratings and unavailabilities (illustrative; base case in the first row).

L (MW)	q_{12}, q_{23}, q_{13}	Ratings (MW)	$\mathrm{LOLP_T}$	$\mathbb{E}[\Delta L]$ (MW)	$\Pr(\Delta L \geq L)$	$\mathbb{E}[\Delta L \mid \Delta L > 0]$ (MW)
100	0.02, 0.03, 0.01	70, 70, 80	0.058906	1.3127	0.000494	22.28
70	0.02, 0.03, 0.01	70, 70, 80	0.000494	0.0346	0.000494	70.00
80	0.02, 0.03, 0.01	70, 70, 80	0.010000	0.1346	0.000494	13.46
130	0.02, 0.03, 0.01	70, 70, 80	0.058906	3.0799	0.000494	52.28
150	0.02, 0.03, 0.01	70, 70, 80	1.000000	13.6689	0.000494	13.67
100	0.02, 0.03, 0.01	70, 70, 100	0.010000	0.3346	0.000494	33.46
100	0.02, 0.03, 0.01	100, 100, 80	0.049400	1.0275	0.000494	20.80
100	0.02, 0, 0.01	70, 70, 80	0.029800	0.7100	0.000200	23.83
100	0.04, 0.06, 0.02	70, 70, 80	0.115648	2.6496	0.001952	22.91
130	0.02, 0.03, 0.01	70, 70, 100	0.058906	2.1018	0.000494	35.68

adds a 10 MW curtailment. At 130 MW the intact network can still carry the transfer, so the loss-of-load probability is unchanged from 100 MW, but the post-contingency curtailments grow and the expected curtailment more than doubles, to 3.08 MW. At 150 MW the intact network itself is short by 10 MW, the loss-of-load probability jumps to 1 and the conditional curtailment falls to 13.67 MW, since most of the curtailment states are now the mild deficit of the intact network. The loss-of-load probability, a step function of the transfer, misses all of the change between 100 and 130 MW that the expected curtailment shows. Fourth, interventions combine with the operating point: at 130 MW the higher L13 rating lowers the expected curtailment from 3.08 to 2.10 MW, a third, against three quarters at 100 MW, because a rating of 100 MW no longer carries the whole transfer. Network reliability changes with both the component condition and the operating point.

27.6.4 Uncertain parameters

The analysis has treated the unavailabilities q_{12}, q_{23}, q_{13}, the ratings and the demand as known. In practice they are uncertain, and they can be collected into a parameter vector

$$\boldsymbol{\theta} = \left(q_{12},\ q_{23},\ q_{13},\ P_{12}^{\max},\ P_{23}^{\max},\ P_{13}^{\max},\ L\right)^{\mathsf{T}},$$

again a vector of parameters, not of bus angles. The risk metric is a function of it, $\mathcal{R}(\boldsymbol{\theta}) = \mathbb{E}[\Delta L \mid \boldsymbol{\theta}]$, the expected curtailment computed with those parameter values; the expectation inside is over the random network states, the aleatory variability of Chapter 1 of Volume I, and the uncertainty of $\boldsymbol{\theta}$ is epistemic. The parameter uncertainty is carried to the risk either by averaging over the parameter distribution, $\mathbb{E}_{\boldsymbol{\theta}}[\mathcal{R}(\boldsymbol{\theta})]$, or by approximating the variance of the risk to first order,

$$\mathrm{Var}[\mathcal{R}] \approx \nabla\mathcal{R}^{\mathsf{T}}\, \boldsymbol{\Sigma}_{\theta}\, \nabla\mathcal{R}, \tag{27.37}$$

where $\nabla\mathcal{R}$ is the gradient of $\mathcal{R}$ with respect to $\boldsymbol{\theta}$ at its mean and $\boldsymbol{\Sigma}_{\theta}$ the covariance matrix of the parameters, when the linearization is suitable.

For the unavailabilities the gradient is already known: its entries are the Birnbaum-type importances of Table 27.7. As an illustration, let the three unavailabilities be independent and uncertain with standard deviations equal to half their values, 0.010, 0.015 and 0.005. Then

$$\begin{aligned}\mathrm{Var}[\mathcal{R}] &\approx 19.885^2(0.010)^2 + 20.09^2(0.015)^2 + 32.47^2(0.005)^2 \\ &= 0.0395 + 0.0908 + 0.0264 = 0.1567\,\mathrm{MW}^2,\end{aligned}$$

a standard deviation of 0.396 MW around 1.3127 MW, about 30 %. Here the linearization is nearly exact: the expected curtailment is multilinear in independent unavailabilities, so its mean over the parameters equals its value at the mean parameters, 1.3127 MW, and the exact variance, 0.15674 MW^2, differs from the first-order value only in the fourth digit. The ratings behave differently. The curtailment is a convex, kinked function of a rating, and Section 27.5.4 showed that averaging over an uncertain rating gives 10.83 MW where the mean rating gives 10 MW. For strongly nonlinear consequences, threshold effects or discrete changes of topology, sampling the joint distribution of the parameters (their posterior, in the sense of Chapter 8 of Volume I) is more reliable than a local Gaussian approximation.

27.6.5 Physical risk and model uncertainty

A transmission reliability model has three kinds of uncertainty that are not interchangeable. *Random physical states*: whether a line is available in an assessment hour, the variability the state enumeration represents. *Uncertain parameters*: incomplete knowledge of the outage probabilities or of the thermal properties of the conductors, the subject of the previous subsection. *Model uncertainty*: whether the DC model, the emergency controls or the failure mechanisms represented adequately describe the physical system.

The distinction matters because the third cannot be reduced by computation. A highly precise Monte Carlo estimate of the EUE does not establish that the underlying physical model is adequate. The model of this chapter assumes, for instance, that every post-contingency topology can be operated up to the listed thermal ratings with no voltage or stability limitation; an AC study might reveal additional constraints, such as a voltage limit at bus 3 when all the power arrives over one long line. Likewise the actual consequences of a brief overload depend on the temperature dynamics, the protection, the remedial controls and the allowable emergency durations, none of which a static model represents. Physical adequacy, whether the power-flow and consequence models are right, and probabilistic adequacy, whether the probabilities are right, must be validated separately.

> **Principle 27.2.** *A risk estimate is conditional on its physical model. Precision in the probabilities cannot compensate for an inadequate power-flow or consequence model, and the two are validated by different evidence.*

27.6.6 Use of probabilistic contingency analysis

Probabilistic contingency analysis is especially useful when several contingencies produce different levels of consequence; when the unavailability of components varies with their condition; when the transmission ratings change with the weather; when the success of corrective actions is uncertain; when operators must prioritize inspections or maintenance; when investment alternatives affect different parts of the network, as the rating increases of Table 27.8 did; and when tail consequences matter more than average performance.

Probabilistic analysis does not replace deterministic compliance studies. A facility may have to satisfy specified $N-1$ or planning-category performance requirements regardless of whether a particular contingency has a low modeled probability, because the requirement is a design rule that protects against the unmodeled as well as the modeled. Probabilistic risk provides additional information for prioritization, planning and operational decisions within, and beyond, those requirements.

The analysis of this chapter took each network state as a static snapshot. Chapter 28 follows a contingency in time: the duration of an overload, the action of protection, correlated weather hazards, and the cascading failures through which an initial outage grows into a larger disturbance.

27.7 What to retain

Definition 27.1 (Probabilistic contingency analysis)**.** Transmission reliability is the probability and size of the load curtailment that the network requires after component outages, computed by enumerating (or sampling) the network states, solving a power flow and a corrective-action model in each, and weighting the curtailments by the state probabilities: $\mathrm{LOLP_T} = \sum_s p_s \, \mathbf{1}_{\{\Delta L(s)>0\}}$, $\mathbb{E}[\Delta L] = \sum_s p_s \, \Delta L^*(s)$ and $\mathrm{EUE} = \sum_t \sum_s p_{t,s} \, \Delta L_t^*(s) \, \Delta t$.

Network topology matters: two systems with equal generation and equal line capacities can have very different reliability, because electrical connectivity and the division of flow among paths decide deliverability. $N-1$ security is a deterministic question of physical feasibility; probabilistic contingency analysis adds the likelihood of each outage and its consequence, and supports expected-curtailment and risk-based comparisons. Unavailability is not outage frequency, and event rates and durations need a chronological model. Uncertain line ratings, which follow from the weather through the conductor heat balance, change the overload and curtailment probabilities materially, and an overload does not by itself imply load shedding, cascading or conductor failure: the corrective-action model must be stated. The contingency with the largest single curtailment need not be the line whose improvement gives the largest expected risk reduction, and a precise risk estimate cannot compensate for an inadequate power-flow or consequence model. Table 27.10 lists the chapter's terms.

Table 27.10: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Contingency	loss of one or more elements	–	–	27.1.2
$N{-}1$, $N{-}k$	loss of one, of k elements	$\binom{N}{k}$ cases	–	27.1.3
DC line flow	linearized active flow	$P_{ij} = (\theta_i - \theta_j)/x_{ij}$	pu, MW	27.2.2
DC susceptance matrix	nodal balance	$\mathbf{P} = \mathbf{B}\boldsymbol{\theta}$	pu	27.2.2
Thermal margin	rating minus flow	$g_{ij} = P_{ij}^{\max} - \lvert P_{ij}\rvert$	MW	27.2.3
Transfer capability	largest feasible transfer	$\mathrm{TC}(s)$	MW	27.2.3
PTDF	flow change per unit transfer	$\partial P_\ell / \partial P_k$	–	27.3.2
LODF	flow change per unit outaged flow	$(P_\ell^{(m)} - P_\ell^{(0)})/P_m^{(0)}$	–	27.3.3
AC injections	full power-flow model	equations (27.10), (27.11)	pu	27.3.4
Static, dynamic security	steady state; transition	–	–	27.3.5
Network-state probability	independent lines	$\prod_\ell q_\ell^{b_\ell}(1-q_\ell)^{1-b_\ell}$	–	27.4.1
Curtailment	load not delivered in a state	$\Delta L(s) = (L - \mathrm{TC}(s))^+$	MW	27.4.2
Transmission LOLP	probability of curtailment	$\sum_s p_s\, \mathbf{1}_{\{\Delta L(s)>0\}}$	–	27.4.3
Expected curtailment	mean curtailment	$\sum_s p_s\, \Delta L(s)$	MW	27.4.4
State contribution	risk of one state	$\mathcal{R}_s = p_s\, \Delta L(s)$	MW	27.4.5
Frequency and duration	unavailability from events	$q_\ell = f_\ell\, \mathrm{MTTR}_\ell$	–	27.4.6
Overload probability	rating below flow	$\Pr(P_\ell^{\max} < \lvert P_\ell(s)\rvert \mid s)$	–	27.5.2
Normal-rating curtailment	expected positive part	$(a-\mu_R)\Phi(z) + \sigma_R\varphi(z)$	MW	27.5.4
Minimum curtailment	optimal corrective action	problem (27.32)	MW	27.5.5
Expected unserved energy	curtailment energy	$\sum_t \sum_s p_{t,s}\, \Delta L_t^*(s)\, \Delta t$	MWh/yr	27.5.6
Risk reduction, Birnbaum	benefit of a line	$\Delta\mathcal{R}_\ell^- = q_\ell\, I_\ell^{\mathrm{B}}$	MW	27.6.1

Further reading. Grainger and Stevenson, *Power System Analysis*, for AC and DC power flow; Wood, Wollenberg and Sheblé, *Power Generation, Operation, and Control*, for contingency analysis and dispatch; Billinton and Allan, *Reliability Evaluation of Power Systems*, for composite generation-transmission reliability; NERC, *TPL-001-5.1 — Transmission System Planning Performance Requirements*, for planning-event categories and performance requirements.

27.8 Exercises

Exercise 27.1. (a) State the defining difference between generation adequacy and transmission reliability. (b) State the flows on L12 and L23 when L13 is disconnected in the three-bus network. (c) State whether exceeding a thermal rating automatically implies immediate customer interruption. (d) State what an LODF estimates. (e) State the difference between outage frequency and unavailability. (f) Give one reason uncertain line ratings matter. (g) Explain why two network states with the same total available transmission capacity can have different consequences. (Answers: (a) transmission reliability requires the network topology and the feasibility of the network flows, not only enough capacity; (b) each carries 100 MW; (c) no, the allowable duration of the overload and the corrective actions decide the consequence; (d) how an outage changes the flows on the surviving lines; (e) the frequency counts outage events per unit time, the unavailability is the fraction of time out of service; (f) weather and thermal conditions change the feasible conductor loading; (g) the topology determines the electrical connectivity and how the power redistributes.)

Exercise 27.2. In the three-bus network, change the reactance of L13 to $x_{13} = 0.1$ pu, keeping every other value of Table 27.1. (a) Solve the DC power flow for the 100 MW transfer (bus 3 the reference) and give the angles and flows. (b) Give the PTDFs of the transfer from bus 1 to bus 3 and the transfer capability of the intact network. (c) Give the curtailment after each single outage, and $\mathrm{LODF}_{12,13}$ from equation (27.9). (d) With the unavailabilities of Section 27.4.1, give the expected curtailment at 100 MW and at 130 MW, and say why the closed form $\mathrm{TC} = 2\min(\cdot)$ of Section 27.6.3 no longer applies. (Answers: (a) $\theta_1 = 1/15 = 0.0667$ and $\theta_2 = 0.0333$ rad; $P_{12} = 33.3$, $P_{13} = 66.7$, $P_{23} = 33.3$ MW. (b) 1/3, 2/3, 1/3 on L12, L13, L23; L13 binds first, $2P_{\mathrm{tr}}/3 = 80$, so TC = 120 MW. (c) 30 MW for L13 out, 20 MW for L12 or L23 out, as before; $\mathrm{LODF}_{12,13} = (1/3)/(1 - 2/3) = 1$. (d) 1.3127 MW, unchanged, because no outage state's curtailment changes; at 130 MW the intact network now curtails 10 MW and $\mathbb{E}[\Delta L] = 0.941094 \times 10 + 3.0799 = 12.49$ MW with $\mathrm{LOLP}_{\mathrm{T}} = 1$; the paths no longer have equal reactances, so they do not split the transfer equally.)

Exercise 27.3. Storm weather, present in 10 % of the hours, raises the line unavailabilities to $q_{12} = 0.10$, $q_{23} = 0.15$, $q_{13} = 0.10$; in the other 90 % they are those of Section 27.4.1. Within each weather state the lines are independent, and the transfer is 100 MW with the ratings of Table 27.1. (a) Compute the transmission LOLP, the expected curtailment and the probability of complete interruption in storm weather, using equation (27.18). (b) Compute the three indices averaged over the weather. (c) Compute the marginal unavailability of each line, and the three indices for independent lines with those marginals. (d) Compare (b) and (c) and explain the difference. (Answers: (a) $1 - 0.90 \times 0.85 \times 0.90 = 0.3115$; with $u = 1 - 0.90 \times 0.85 = 0.235$, $30 \times 0.1 \times 0.765 + 20 \times 0.9 \times 0.235 + 100 \times 0.1 \times 0.235 = 2.295 + 4.230 + 2.350 = 8.875$ MW; $0.1 \times 0.235 = 0.0235$. (b) $0.9 \times 0.058906 + 0.1 \times 0.3115 = 0.08417$; $0.9 \times 1.3127 + 0.1 \times 8.875 = 2.0689$ MW; $0.9 \times 0.000494 + 0.1 \times 0.0235 = 0.002795$. (c) 0.028, 0.042, 0.019; 0.08652, 2.0119 MW, 0.001308. (d) The shared weather makes outages coincide: at the same marginal unavailabilities it lowers the LOLP slightly, raises the expected curtailment by 2.8 %, and more than doubles the probability of complete interruption, the common-cause effect of Chapter 13 of Volume I.)

Exercise 27.4. After an outage of L12, L13 carries 100 MW and its rating is $P_{13}^{\max} \sim \mathcal{N}(95, 10^2)$ MW, independent of the line outages. (a) Compute the overload probability. (b) Compute the expected curtailment from equation (27.31), with $\Phi(0.5) = 0.69146$ and $\varphi(0.5) = 0.35207$. (c) Compare it with the curtailment for a fixed rating of 95 MW and name the inequality. (d) Using the states of Table 27.4, compute the expected curtailment of the network when every state in which L13 alone carries the transfer has this rating, neglecting the intact state's contribution (1.3×10^{-5} MW). (e) Compute the annual EUE if that hour were repeated for 8760 h. (Answers: (a) $\Phi(0.5) = 0.6915$. (b) $5 \times 0.69146 + 10 \times 0.35207 = 3.457 + 3.521 = 6.978$ MW. (c)

5 MW; the uncertain rating adds 1.978 MW, by Jensen's inequality for the convex $(100 - P^{\max})^+$. (d) $0.28518 + 0.04940 + 0.048906 \times 6.978 = 0.6759$ MW. (e) $8760 \times 0.6759 = 5921$ MWh per year.)

Chapter 28

Time-Varying Reliability and Cascading Failures

28.1 Cascading failures and the limits of static contingency analysis

Chapter 27 evaluated transmission contingencies as static network configurations: for each outage it computed the post-contingency power flows, compared them with the line ratings and derived the load curtailment needed to restore a feasible operating state. Real disturbances do not take place in one instant. After a line trips, power redistributes over the surviving network almost immediately; the conductors that now carry more current begin to heat; protection systems monitor the electrical quantities; automatic controls may act within seconds and operators may redispatch generation within minutes; and further elements may disconnect before the system reaches a secure operating condition. The order and the timing of these events can decide the final consequence, so that two physically identical initiating outages can end very differently depending on how fast the system responds. The question of this chapter is

how an initial transmission disturbance evolves in time, and what determines whether the system recovers, sheds load in a controlled way, or suffers a propagating, cascading outage.

To answer it the chapter models failure intensities that vary with time and weather, represents outages that are correlated through a shared weather state, computes the transient temperature of an overloaded conductor and the time left before its limit, compares the timing of protective tripping, remedial controls and restoration, builds and quantifies an event tree of a cascade, states what branching-process models of cascades can and cannot say, and describes how a chronological Monte Carlo simulation estimates the frequency of cascades and the expected unserved energy.

Jargon: contingency, static contingency analysis, cascading failure, thermal rating, protection, remedial action scheme, hybrid dynamical system

A *contingency* is an event that changes the configuration or the operating condition of a network, such as the loss of a line, a transformer, a generator or a bus section. *Static contingency analysis* computes, for each contingency, the steady-state flows of the resulting configuration and checks them against the operating limits; the $N-1$ criterion requires every single contingency to leave the flows within limits. A *cascading failure* (a *cascade*) is, in the definition used by the power-system reliability organizations, the uncontrolled successive loss of system elements, triggered by an incident, which spreads beyond an area predetermined by studies and causes widespread interruption of service. The *thermal rating* of a line is the current or power it may carry without its conductor exceeding an allowed temperature: the *continuous rating* holds indefinitely, an *emergency rating* for a stated limited time. *Protection* is the set of relays and breakers that disconnect an element when measured electrical quantities (current, impedance, voltage, frequency, differential quantities) meet their settings. A *remedial action scheme* (RAS) is an automatic scheme that detects a predefined system condition and executes predetermined corrective actions. A *hybrid dynamical system* is one whose continuous states evolve between discrete events that change its configuration.

Key idea

A cascade is a race between physical processes and responses. The probability that a disturbance stops, is contained by controlled load curtailment, or propagates depends on weather-dependent failure intensities, on the dependence that shared weather creates, on the time the conductors take to heat, and on whether corrective action arrives before protective tripping; event trees, branching models and chronological simulation quantify that race at increasing levels of detail, and the expected unserved energy they produce is often dominated by rare, severe sequences.

This section restates the three-bus network on which the chapter's cascade starts, distinguishes a contingency from a cascade, and states what a static analysis leaves out.

28.1.1 From a contingency to a cascade

The network is the three-bus system of Chapter 27. Bus 1 has generation, bus 2 is a junction with neither generation nor load, and bus 3 has the load; three lines form a triangle. Lines L12 and L23 have reactance $x_{12} = x_{23} = 0.1$ and continuous rating $P_{12}^{\max} = P_{23}^{\max} = 70\,\text{MW}$; line L13 has $x_{13} = 0.2$ and $P_{13}^{\max} = 80\,\text{MW}$ (reactances in per unit on a 100 MVA base). With 100 MW generated at bus 1 and consumed at bus 3, the DC power flow $P_{ij} = (\theta_i - \theta_j)/x_{ij}$, with bus 3 as the angle reference, gives $\theta_1 = 0.10\,\text{rad}$ and $\theta_2 = 0.05\,\text{rad}$, and $P_{12} = P_{23} = P_{13} = 50\,\text{MW}$: the direct path and the two-line path have the same reactance, 0.2, and share the transfer equally. When L13 is lost, the whole transfer must follow the series path, and

$$P_{12} = P_{23} = 100\,\text{MW} > 70\,\text{MW},$$

an overload of 30 MW on each surviving line. Chapter 27 found that restoring both flows within their continuous ratings by reducing the transfer requires a curtailment $\Delta L = 30\,\text{MW}$.

Figure 28.1 draws the network before and after the outage. The circles are the buses, the lines the transmission lines, each labeled with its reactance and continuous rating; the numbers on the lines are their flows. In panel (a) every line carries 50 MW, within its rating. In panel (b) L13 has tripped (dashed, crossed), and L12 and L23 each carry 100 MW against 70 MW ratings.

A static contingency analysis classifies L12 and L23 as overloaded. The classification does not mean that they trip at once. A conductor takes minutes to heat (Section 28.3), and in

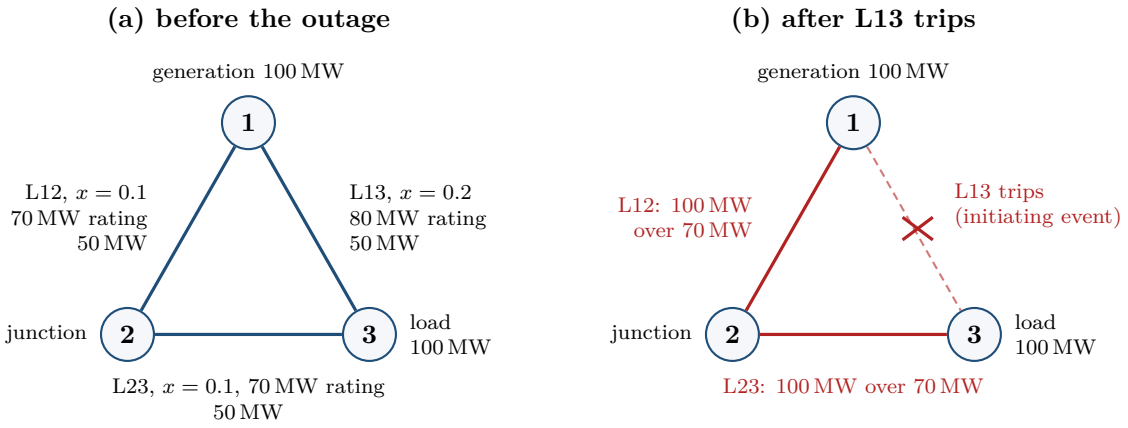


Figure 28.1: The three-bus network of Chapter 27. (a) Base case: the 100 MW transfer from bus 1 to bus 3 splits equally over the direct line L13 and the series path L12–L23, which have the same reactance. (b) After the outage of L13, the initiating event of this chapter's cascade, both surviving lines carry the full 100 MW, 30 MW above their continuous ratings. DC power flow, reactances in per unit on a 100 MVA base.

that time generation can be redispatched or 30 MW of demand curtailed. Three outcomes are possible:

- an *arrested disturbance*: corrective action removes the overload before any further protective tripping;
- a *controlled interruption*: load is deliberately curtailed to keep the system from deteriorating further;
- a *propagating failure*: further elements disconnect, which changes the flows again and may trigger further outages.

Not every pair of successive line trips is a cascade in the formal sense of the jargon box, which requires uncontrolled propagation beyond a predetermined area. In a network of three buses no such area exists, and the two-stage sequence studied in this chapter is a cascade-like propagation used to illustrate the methods; the same methods apply, with more states and more stages, to networks in which a cascade in the formal sense can occur.

28.1.2 Why static contingency analysis is incomplete

Write s for a network topology, the set of lines in service. A static model computes the flow $P_\ell(s)$ of each line ℓ for each topology, and nothing else. A dynamic reliability model must follow the evolution in time of

$$P_\ell(t), \qquad T_{\mathrm{c},\ell}(t), \qquad S_\ell(t), \qquad u(t), \tag{28.1}$$

the flow on line ℓ, its conductor temperature, its status (1 in service, 0 out), and the protective or corrective control actions in force. The flows follow the topology instantly; the temperatures follow the flows with a delay set by the thermal inertia of the conductors; the statuses change at discrete instants, when protection acts or a weather-related failure occurs; and the controls change the injections or the topology. Chapter 3 of Volume I distinguished *adequacy*, enough capacity to meet the demand with outages accounted for, from *security*, the ability to withstand sudden disturbances (Section 3.6.1); a cascade is a failure of security, and its analysis needs the dynamics that an adequacy calculation omits.

Figure 28.2 draws the loop. Each box is one stage of the evolution after the initial outage, and its second line names the state that changes in it. The arrows run from the outage to the power redistribution, the thermal evolution, the protection and control stage, and a new

topology. If another element trips, the return arrow sends the process back to a new power redistribution; if no further event occurs, the exit arrow leaves the loop in a secure state.

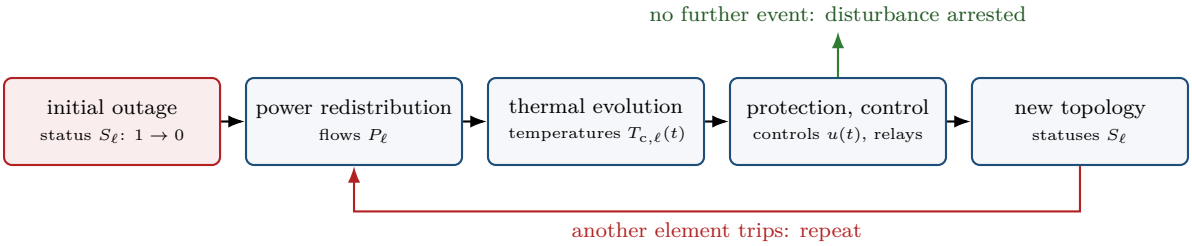


Figure 28.2: A cascade as a hybrid dynamical system. Continuous states (flows, conductor temperatures) evolve between discrete events (trips, control actions) that change the topology; each new topology starts a new power redistribution until no further event occurs.

The process is a hybrid dynamical system in the sense of the jargon box. Its continuous part is the heat balance of each conductor; its discrete part is the sequence of status changes; and the two are coupled, because the temperatures decide when protection might act and each status change resets the flows that drive the temperatures. A static contingency analysis evaluates one pass through the first two boxes and stops. The rest of the chapter adds the missing parts one at a time: the failure intensities that start the sequence (Section 28.2), the thermal evolution (Section 28.3), the race between protection and control and the event tree that aggregates it (Section 28.4), and the models that follow the whole loop (Section 28.5).

28.2 Time-dependent failure intensity and weather-correlated failures

A cascade starts with an outage, and the propensity of a line to fail is not constant: it rises with wind, ice, lightning and heat. This section lets the failure intensity vary with time and weather, computes the outage probability of a line during a storm, and shows the dependence that a shared weather state creates between the outages of two lines.

28.2.1 Time-varying failure intensity

Chapter 5 of Volume I introduced the constant failure rate λ and its exponential reliability function $R(t) = \mathrm{e}^{-\lambda t}$. The outage hazard of an overhead line, however, varies with the wind speed, the ice load, the lightning activity, the condition of the vegetation along the route, the exposure to wildfire, the conductor temperature, the deterioration of its components, and the protection and operating conditions. A general description is a failure intensity conditioned on what is known at time t,

$$\lambda(t \mid \mathcal{F}_t, \mathbf{z}(t)), \tag{28.2}$$

where $\mathcal{F}_t$ is the relevant event history up to t (earlier outages, repairs, protection operations) and $\mathbf{z}(t)$ is the vector of environmental and operating covariates at t (wind speed, ambient temperature, loading). It is the hazard of Chapter 5 with its conditioning made explicit, in yr^{-1} or h^{-1}.

Chapter 5 (Section 5.6.3) showed that for a specified covariate history the reliability is the exponential of minus the integrated hazard, $R(t \mid \mathbf{z}(\cdot)) = \exp[-\int_0^t \lambda(u \mid \mathbf{z}(u))\,\mathrm{d}u]$, equation (5.22). For a line that is in service at t_1, the probability that it remains in service to t_2 is the ratio of the reliabilities at the two times, because $\{T > t_2\} \subseteq \{T > t_1\}$:

$$\Pr(T > t_2 \mid T > t_1, \mathbf{z}(\cdot)) = \frac{R(t_2 \mid \mathbf{z})}{R(t_1 \mid \mathbf{z})} = \exp\left[-\int_{t_1}^{t_2} \lambda(u \mid \mathbf{z}(u))\,\mathrm{d}u\right], \tag{28.3}$$

the integrals from 0 to t_1 cancelling in the ratio. The expression holds for the specified covariate path; when the weather over $[t_1, t_2]$ is itself uncertain, the right-hand side must be averaged over its possible paths, which is not the same as evaluating it at the average path (the average of an exponential is not the exponential of the average, as the temperature example of Section 5.6.4 showed for a hazard that is convex in temperature).

28.2.2 The weather-dependent hazard model

Write Z for the weather state, the variable that describes the weather to which the line is exposed, as in Chapter 13 of Volume I; the covariates $\mathbf{z}_\ell(t)$ of line ℓ are functions of it. A useful statistical formulation is the log-linear (proportional hazards) model of equation (5.21),

$$\lambda_\ell(t \mid Z) = \lambda_{\ell,0} \exp\!\left[\boldsymbol{\gamma}^{\mathsf{T}} \mathbf{z}_\ell(t)\right], \tag{28.4}$$

where $\lambda_{\ell,0}$ is the baseline outage intensity of the line, its intensity when every covariate is zero, and $\boldsymbol{\gamma}$ is the vector of fitted coefficients. (The coefficients are written $\boldsymbol{\gamma}$, as in Chapter 5, because β is the Weibull shape and the common-cause beta factor in this book.) A two-covariate example is

$$\lambda_\ell(t \mid Z) = \lambda_{\ell,0} \exp\!\left[\gamma_v\, v(t) + \gamma_T \left(T_{\mathrm{a}}(t) - T_{\mathrm{ref}}\right)\right], \tag{28.5}$$

with $v(t)$ a wind-related covariate, $T_{\mathrm{a}}(t)$ the ambient temperature, T_{ref} a reference temperature, and γ_v, γ_T the fitted coefficients in the reciprocal units of their covariates. Each unit increase of v multiplies the intensity by e^{γ_v}, whatever the level of the other covariate.

This is a statistical model, not a physical law. Strong winds can raise the outage intensity through contact with vegetation, galloping of the conductors or mechanical damage to towers and fittings; each mechanism has its own dependence on wind speed and direction, and a single wind-speed coefficient fitted to outage records summarizes their combined effect over the conditions in those records. Outside those conditions, or when the mix of mechanisms changes, the coefficient must be revalidated; Chapter 5 made the same point for the temperature-doubling rule.

28.2.3 Worked example: outage hazard during a storm

A line has a baseline outage intensity $\lambda_0 = 0.02\,\mathrm{yr}^{-1}$. During a severe weather period the modeled intensity rises by a factor of twenty, to $\lambda_{\mathrm{storm}} = 0.4\,\mathrm{yr}^{-1}$; in the model of equation (28.4) this corresponds to $\boldsymbol{\gamma}^{\mathsf{T}}\mathbf{z} = \ln 20 = 3.00$ during the storm. The storm lasts $\Delta t = 48\,\mathrm{h}$. With $1\,\mathrm{yr} = 8760\,\mathrm{h}$, the cumulative hazard over the storm is

$$\Lambda_{\mathrm{storm}} = \int_{\mathrm{storm}} \lambda(u)\,\mathrm{d}u = \lambda_{\mathrm{storm}}\,\Delta t = 0.4\,\mathrm{yr}^{-1} \times \frac{48}{8760}\,\mathrm{yr} = 0.002192, \tag{28.6}$$

and by equation (28.3) the probability that the line has an outage during the storm, given that it is in service when the storm begins, is

$$\Pr(\text{outage during the storm}) = 1 - \mathrm{e}^{-\Lambda_{\mathrm{storm}}} = 1 - \mathrm{e}^{-0.002192} = 0.002189,$$

about $0.219\,\%$. Over the same $48\,\mathrm{h}$ at the baseline intensity the probability would be $1 - \mathrm{e}^{-0.02 \times 48/8760} = 0.000110$, so the storm multiplies the probability of an outage in those two days by 19.98, almost exactly the factor of the intensity, because both probabilities are small enough for $1 - \mathrm{e}^{-x} \approx x$.

The probability is nevertheless small, despite a twenty-fold increase in the hazard, because the elevated hazard acts for only $48\,\mathrm{h}$. Over a year that contains one such storm, the cumulative hazard is $0.02 \times 8712/8760 + 0.4 \times 48/8760 = 0.019890 + 0.002192 = 0.022082$, and the annual outage probability rises from $1 - \mathrm{e}^{-0.02} = 1.980\,\%$ to $1 - \mathrm{e}^{-0.022082} = 2.184\,\%$. A forecast of

high hazard therefore matters most for decisions on the timescale of the storm itself, such as the commitment of reserves or the deferral of planned outages over those two days, and much less for annual reliability, unless storms are frequent or long.

Principle 28.1. *The probability of failure over an interval is set by the integrated hazard over that interval, the product of intensity and exposure duration for a constant intensity; neither the peak intensity nor the duration alone determines it.*

28.2.4 Weather-correlated outages

Two transmission lines crossing the same storm region are exposed to the same weather. Their outages can be independent given the realized severity of the storm and yet dependent before that severity is known. Let the shared weather state be

$$Z \in \{\text{normal}, \text{severe}\}, \qquad \Pr(Z = \text{severe}) = 0.10,$$

and let the conditional probability that either line is unavailable during the assessed period be the value of Table 28.1 for each state. Each row of the table is one weather state, with its probability, the outage probability of one line in that state, and the contributions of the state to the marginal and to the joint outage probability computed below. The two outages F_1 and F_2 are assumed independent given Z: $\Pr(F_1 \cap F_2 \mid Z) = \Pr(F_1 \mid Z)\ \Pr(F_2 \mid Z)$. (Here F_1 and F_2 are events; they are unrelated to the cumulative failure probability $F(t)$.)

Table 28.1: Two lines under a shared weather state Z (illustrative values) and the contribution of each state to the marginal and to the joint outage probability.

Weather state z	$\Pr(Z = z)$	$\Pr(F_1 \mid z) = \Pr(F_2 \mid z)$	Contribution to $\Pr(F_i)$	Contribution to $\Pr(F_1 \cap F_2)$
Normal	0.90	0.01	$0.90 \times 0.01 = 0.009$	$0.90 \times 0.01^2 = 0.00009$
Severe	0.10	0.20	$0.10 \times 0.20 = 0.020$	$0.10 \times 0.20^2 = 0.00400$
Total	1		$\Pr(F_i) = 0.029$	$\Pr(F_1 \cap F_2) = 0.00409$

The marginal outage probability of each line follows from the law of total probability,

$$\Pr(F_i) = \sum_z \Pr(F_i \mid Z = z)\ \Pr(Z = z) = 0.90 \times 0.01 + 0.10 \times 0.20 = 0.029,$$

or 2.9 %. An analyst who assumed the two outages to be independent without conditioning on the weather, *marginal independence*, would multiply the marginals,

$$\Pr(F_1 \cap F_2)_{\text{marginal independence}} = 0.029^2 = 0.000841,$$

or 0.0841 %. The correct calculation applies the law of total probability to the intersection and then the conditional independence, which is equation (13.18) of Chapter 13, the weather mixture:

$$\Pr(F_1 \cap F_2) = \sum_z \Pr(Z = z)\ \Pr(F_1 \mid z)\ \Pr(F_2 \mid z) = 0.90 \times 0.01^2 + 0.10 \times 0.20^2 = 0.00409, \quad (28.7)$$

or 0.409 %, 4.86 times the value under marginal independence. The individual outage probabilities of the two lines are the same in both calculations; only the dependence structure differs. Almost all of the joint probability, 0.004 of 0.00409, comes from the severe state, which occurs one time in ten.

Chapter 13 (Section 13.4.3) derived the size of the effect from the law of total covariance: for identical units that are conditionally independent given Z, the covariance of the outage indicators equals the variance of the conditional outage probability $p(Z) = \Pr(F_i \mid Z)$. Here $\mathrm{Var}[p(Z)] = 0.90 \times 0.10 \times (0.20 - 0.01)^2 = 0.003249$, which is exactly $0.00409 - 0.000841$, and the correlation of the two outage indicators is $0.003249/(0.029 \times 0.971) = 0.115$. Observing one outage makes the severe state more plausible, $\Pr(Z = \text{severe} \mid F_1) = 0.020/0.029 = 0.690$ against 0.10 before, so the outage probability of the second line rises to $\Pr(F_2 \mid F_1) = 0.00409/0.029 = 0.141$, almost five times its marginal 0.029. For the cascade, this means that the second outage of a sequence is more likely than its marginal probability suggests precisely when the first one has occurred: the weather that caused the first has not gone away.

Common mistake

Multiplying the marginal outage probabilities of lines exposed to the same weather. Each line's own probability is right, but the product ignores the shared weather state and understates the joint outage probability, here by a factor of almost five. System reliability cannot in general be reconstructed from the outage probabilities of the components alone.

28.3 Conductor thermal dynamics

The time available for corrective action after a line outage is set by how fast the surviving conductors heat. This section states the transient heat balance of a conductor, derives its first-order approximation and the thermal time constant, computes the time an overloaded conductor takes to reach its temperature limit, and checks whether a corrective action arrives in time.

28.3.1 Thermal overload as a transient process

After L13 trips, the flow on L12 rises from 50 MW to 100 MW. If the bus voltages and the power factor stay approximately constant, the current is proportional to the active power flow, and it approximately doubles, $I_{\text{post}} \approx 2I_{\text{pre}}$. The Joule heating I^2R then becomes approximately four times its pre-outage value, for a fixed resistance. The conductor temperature does not jump to a new value: the conductor has thermal mass, and it loses heat by convection and radiation at rates that grow with its temperature. A published standard method gives the current–temperature relationship of bare overhead conductors under steady-state and time-varying conditions; the simplified balance used here has the same structure.

Chapter 1 of Volume I wrote the heat balance of a conductor (Section 1.5.1, equation (1.12)). Per unit length of conductor it reads

$$mc_p \frac{\mathrm{d}T_{\mathrm{c}}}{\mathrm{d}t} = I^2 R(T_{\mathrm{c}}) + q_{\text{solar}} - q_{\text{conv}} - q_{\text{rad}}, \tag{28.8}$$

where m is the conductor mass per unit length (kg/m), c_p its effective specific heat (J/(kg·K)), so that mc_p is the heat capacity per unit length (J/(m·K)); $T_{\mathrm{c}}(t)$ is the conductor temperature (°C, or kelvin for temperature differences); I is the current (A); $R(T_{\mathrm{c}})$ is the resistance per unit length (Ω/m); and q_{solar}, q_{conv} and q_{rad} are the solar heating and the convective and radiative cooling per unit length (W/m). Every term must use the same per-unit-length basis. The same chapter defined the performance-based failure event of the line as the temperature exceeding its limit at some time in the horizon, $F = \{\max_{0 \le t \le H} T_{\mathrm{c}}(t) > T_{\max}\}$, equation (1.13); this section computes when, after the outage, that event would occur.

Remark 28.1. In equation (28.8), $R(T_{\mathrm{c}})$ is the temperature-dependent electrical resistance of the conductor, not the reliability function $R(t)$; and T_{c}, T_{a} and T_{∞} are temperatures, not the time to failure T.

The critical physical state is the conductor temperature, not the ratio of the current to a static rating. A conductor can carry a current above its continuous rating for a limited time without exceeding its temperature limit, which is what an emergency rating expresses. The thermal limit is not the only constraint: relay settings, the clearance of the sagging conductor to the ground and to vegetation, and mechanical limits can be more restrictive, and Section 28.4.1 returns to the first of these.

28.3.2 The first-order approximation and the thermal time constant

For a short interval with approximately constant current and weather, the right-hand side of equation (28.8) is a function of the conductor temperature alone; call it $g(T_c) = I^2R(T_c) + q_{solar} - q_{conv}(T_c) - q_{rad}(T_c)$, in W/m. The *steady-state temperature* T_∞ for the prevailing current and weather is the temperature at which heating and cooling balance, $g(T_\infty) = 0$. A first-order Taylor expansion about T_∞ gives

$$g(T_c) \approx g(T_\infty) + g'(T_\infty)\,(T_c - T_\infty) = -k_{th}\,(T_c - T_\infty),$$

with

$$k_{th} = -g'(T_\infty) = \frac{\mathrm{d}q_{conv}}{\mathrm{d}T_c} + \frac{\mathrm{d}q_{rad}}{\mathrm{d}T_c} - I^2\,\frac{\mathrm{d}R}{\mathrm{d}T_c}\bigg|_{T_c = T_\infty}, \tag{28.9}$$

the effective heat-transfer coefficient per unit length, in W/(m·K): the rate at which net cooling grows with the conductor temperature. It is positive when the losses grow faster with temperature than the Joule heating does, which is the condition for the steady state to be stable. Substituting into equation (28.8) and dividing by mc_p,

$$\frac{\mathrm{d}T_c}{\mathrm{d}t} = \frac{T_\infty - T_c}{\tau_{th}}, \qquad \tau_{th} = \frac{mc_p}{k_{th}}, \tag{28.10}$$

where τ_{th} is the *thermal time constant*, in seconds or minutes: the ratio of the heat the conductor stores per kelvin to the heat it sheds per kelvin per second.

To solve equation (28.10), let $\Delta(t) = T_c(t) - T_\infty$. Since T_∞ is constant, $\mathrm{d}\Delta/\mathrm{d}t = -\Delta/\tau_{th}$, whose solution is $\Delta(t) = \Delta(0)\,\mathrm{e}^{-t/\tau_{th}}$, and therefore

$$T_c(t) = T_\infty + \left[T_c(0) - T_\infty\right]\mathrm{e}^{-t/\tau_{th}}. \tag{28.11}$$

The temperature approaches T_∞ exponentially: after one time constant it has covered $1 - \mathrm{e}^{-1} = 63.2\,\%$ of the step from $T_c(0)$ to T_∞, after three time constants $95.0\,\%$. If T_∞ exceeds the limit T_{max} and $T_c(0) < T_{max}$, setting $T_c(t) = T_{max}$ in equation (28.11) gives $\mathrm{e}^{-t/\tau_{th}} = (T_\infty - T_{max})/(T_\infty - T_c(0))$, and the *time to the thermal limit* is

$$t_{lim} = \tau_{th}\,\ln\frac{T_\infty - T_c(0)}{T_\infty - T_{max}}. \tag{28.12}$$

If $T_\infty \le T_{max}$, the limit is never reached: the post-outage steady state lies within the limit, and the overload in current is not a thermal overload at all.

As an illustration of the magnitudes, a conductor with heat capacity $mc_p = 1350\,\mathrm{J/(m\cdot K)}$ and an effective heat-transfer coefficient $k_{th} = 1.5\,\mathrm{W/(m\cdot K)}$ has $\tau_{th} = 1350/1.5 = 900\,\mathrm{s} = 15\,\mathrm{min}$ (illustrative values). The time constant grows with the conductor's mass and falls with the wind, which raises the convective term of k_{th}. The exponential model is an approximation to the heat balance, valid while the current and the weather stay approximately constant and the radiative loss, which grows with the fourth power of the absolute temperature, can be linearized; it does not replace the detailed standard calculation. Its value is that it makes the interaction between the response time of the controls and the thermal headroom of the conductor explicit.

28.3.3 Worked example: time to the thermal limit

Before the outage, L12 is at $T_c(0) = 60°C$. After L13 trips, its current rises and its new steady-state temperature, under the prevailing weather, is $T_\infty = 110°C$; its time constant is $\tau_{th} = 15\,\text{min}$ and its temperature limit $T_{max} = 90°C$ (illustrative values). The two temperatures are consistent with the doubling of the current under the linear model: the steady-state rise above the ambient is $(q_{solar} + I^2R)/k_{th}$, so a Joule rise of 16.7 K before the outage becomes $4 \times 16.7 = 66.7\,\text{K}$ after it, and both values follow if the ambient temperature plus the solar rise is 43.3°C, for example an ambient of 35°C and a solar rise of 8.3 K. (The rise of the resistance with temperature would make T_∞ somewhat higher.) By equation (28.11),

$$T_c(t) = 110 - 50\,e^{-t/15} \qquad (°C,\ t \text{ in minutes}).$$

The limit is reached when $90 = 110 - 50\,e^{-t/15}$, that is $e^{-t/15} = 20/50 = 0.4$, and by equation (28.12)

$$t_{lim} = 15\,\ln\frac{110-60}{110-90} = -15\,\ln 0.4 = 15 \times 0.9163 = 13.74\,\text{min}.$$

Under these conditions the line stays below 90°C for about 13.7 minutes after the outage, although it carries 43 % more than its continuous rating from the first second.

The figure does not mean that operators may wait 13.7 minutes before acting. Operating obligations, emergency ratings with their own time limits, the behavior of protection and the uncertainty in T_∞ and τ_{th} can require earlier intervention. The uncertainty is large: T_∞ depends on the wind at the conductor, which is rarely measured along the route. Table 28.2 evaluates equation (28.12) and the temperature at eight minutes for several time constants and steady-state temperatures, with $T_c(0) = 60°C$ and $T_{max} = 90°C$; its columns are the time constant, the steady-state temperature, the time to the limit, the temperature eight minutes after the outage, and whether an action completed at eight minutes is in time. The base case is the third row. A time constant of 5 min, as for a light conductor with little heat capacity, leaves 4.6 minutes; a steady state of 130°C, as in calm air and high ambient temperature, leaves 8.4.

Table 28.2: Time to the thermal limit, equation (28.12), and temperature eight minutes after the outage, equation (28.11), for $T_c(0) = 60°C$ and $T_{max} = 90°C$ (illustrative).

τ_{th} (min)	T_∞ (°C)	t_{lim} (min)	$T_c(8\,\text{min})$ (°C)	Action at 8 min in time
5	110	4.58	99.90	no
10	110	9.16	87.53	yes
15	110	13.74	80.67	yes
20	110	18.33	76.48	yes
30	110	27.49	71.70	yes
15	100	20.79	76.53	yes
15	120	10.40	84.80	yes
15	130	8.39	88.93	yes, by 0.4 min

28.3.4 Corrective action before the limit

Suppose redispatch or a controlled reduction of the load finishes $t_A = 8\,\text{min}$ after the initiating outage. At that instant

$$T_c(8) = 110 - 50\,e^{-8/15} = 110 - 50 \times 0.58665 = 80.67°C,$$

below the limit. After the corrective action the flow on L12 is lower, and its new steady-state temperature is $T_{\infty,\mathrm{corr}} = 65°\mathrm{C}$. Equation (28.11) applies again from $t = 8$, with the temperature at that instant as the initial value:

$$T_{\mathrm{c}}(t) = 65 + (80.67 - 65)\,\mathrm{e}^{-(t-8)/15}, \qquad t \geq 8\,\mathrm{min}. \tag{28.13}$$

The conductor begins to cool at once; it is back to 70°C at $t = 8 + 15\ln(15.67/5) = 25.1$ min. Its maximum modeled temperature is 80.67°C, reached at the instant the corrective action takes effect, and under this model an eight-minute corrective action prevents the limit from being exceeded, with $13.74 - 8 = 5.74$ minutes to spare.

Figure 28.3 plots the two trajectories. The horizontal axis is the time after the outage of L13 in minutes and the vertical axis the conductor temperature. The dashed red curve is the trajectory without corrective action, rising toward $T_\infty = 110°\mathrm{C}$; the solid blue curve follows it to the corrective action at eight minutes and then decays toward 65°C. The horizontal dashed line is the limit $T_{\max} = 90°\mathrm{C}$; the dotted lines are the two steady-state temperatures. The uncorrected curve crosses the limit at $t_{\mathrm{lim}} = 13.74$ minutes, marked by the dot on the limit line; the corrected curve peaks at 80.67°C and never reaches it. A temporarily overloaded state is thus managed before it becomes a thermal violation, and the margin is the horizontal distance between the two marked instants.

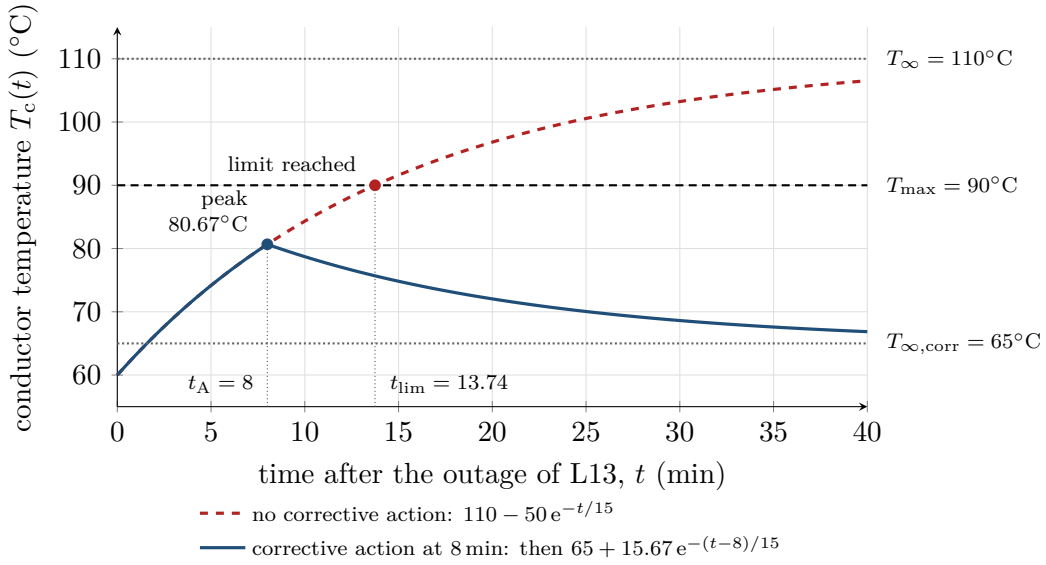


Figure 28.3: Transient temperature of line L12 after the outage of L13, from the first-order model (28.11) with $T_{\mathrm{c}}(0) = 60°\mathrm{C}$, $T_\infty = 110°\mathrm{C}$ and $\tau_{\mathrm{th}} = 15$ min, and after a corrective action at eight minutes that lowers the steady-state temperature to 65°C (illustrative values). Without action the limit is reached at 13.74 min; with it the temperature peaks at 80.67°C.

28.4 Protection, remedial controls and the cascade event tree

Whether the thermal headroom of Section 28.3 is used well depends on which of several timed events happens first. This section separates protection from the thermal limit, describes remedial action schemes, derives the probability that a corrective action wins the race against protective tripping, and then builds, quantifies and ranks the cascade that starts with the outage of L13.

28.4.1 Protection, thermal limits and remedial action schemes

Transmission protection responds to electrical quantities and their settings: impedance, current, differential quantities, voltage and frequency. It does not measure conductor temperature. A distance relay can operate on a heavily loaded line before the conductor reaches its temperature limit, when the load impedance enters its zone; conversely, an emergency overload may be tolerated for a limited time without any protective trip. A reliability model must therefore keep three times apart:

- t_{prot}, the time at which protection would disconnect the element, from the relay settings and the evolving electrical quantities;
- t_{lim}, the time at which the conductor would reach its temperature limit, from equation (28.12);
- t_{A}, the time at which corrective action takes effect.

They represent different mechanisms, and the event sequence depends on their order. Avoiding an unwanted protective disconnection in a given scenario may require $t_{\mathrm{A}} < t_{\mathrm{prot}}$; avoiding a thermal violation requires $t_{\mathrm{A}} < t_{\mathrm{lim}}$. The protection time must come from protection engineering and the applicable operating limits, not from an assumed equation that equates relay operation with conductor temperature.

A remedial action scheme, defined in the jargon box, detects a specified system condition and executes predetermined corrective actions: tripping generators, shedding load in a controlled way, switching the topology, or changing the power transfer. Chapter 15 of Volume I treated it as one heading of a transmission event tree, with a single success probability. A scheme can, however, fail in several distinct ways:

1. the initiating condition is not detected;
2. the control logic does not produce the required action;
3. the communication or output equipment fails;
4. the action is executed too late;
5. the action is executed but is insufficient to restore a secure condition.

The first three are failures on demand of the kind a fault tree under the heading would represent (Section 15.5.1); the fourth is a timing failure, the subject of the next subsection; the fifth is a failure of the success criterion. They should not be collapsed into one probability when their mechanisms, their dependencies or their consequences differ: a scheme that shares sensors, communications or control power with the protection it is meant to pre-empt has a common cause with it, which Chapter 13 of Volume I showed a product of marginal probabilities would miss.

28.4.2 Competing events

Let T_{A} be the random time at which a successful corrective action takes effect and T_{P} the random time of an additional protective trip. A simple condition for the propagation to be arrested is that the corrective action comes first, $T_{\mathrm{A}} < T_{\mathrm{P}}$, and the probability of arrest is $\Pr(T_{\mathrm{A}} < T_{\mathrm{P}})$. The two times are *competing events* (in reliability engineering, competing risks): whichever occurs first decides the branch. Suppose they are independent, T_{A} with density $f_{\mathrm{A}}(t)$ and T_{P} with survival function $R_{\mathrm{P}}(t) = \Pr(T_{\mathrm{P}} > t)$. Conditioning on the time of the corrective action and applying the law of total probability,

$$\Pr(T_{\mathrm{A}} < T_{\mathrm{P}}) = \int_0^{\infty} \Pr(T_{\mathrm{P}} > t \mid T_{\mathrm{A}} = t)\, f_{\mathrm{A}}(t)\, \mathrm{d}t,$$

and by independence the conditional probability is $\Pr(T_{\mathrm{P}} > t) = R_{\mathrm{P}}(t)$, so that

$$\Pr(\text{arrest}) = \Pr(T_{\mathrm{A}} < T_{\mathrm{P}}) = \int_0^{\infty} f_{\mathrm{A}}(t)\, R_{\mathrm{P}}(t)\, \mathrm{d}t. \tag{28.14}$$

This is the dynamic-event-tree branch probability of Chapter 15, $\Pr(T_{\mathrm{G}} < T_{\mathrm{dep}}) = \int \Pr(T_{\mathrm{G}} < \tau \mid T_{\mathrm{dep}} = \tau)\, \mathrm{d}F_{\mathrm{dep}}(\tau)$ (equation (15.21)), in which a restoration time competed with a battery depletion time; and when the protection time is a fixed number t_{P}, it reduces to the recovery-timing model of that chapter, $\Pr(T_{\mathrm{A}} < t_{\mathrm{P}}) = F_{\mathrm{A}}(t_{\mathrm{P}})$, the distribution function of the action time evaluated at the time available (equation (15.18)). The formulation accommodates uncertainty in detection time, communication delay, operator response and protection operation.

Exponential competing times. If T_{A} and T_{P} are exponential with rates λ_{A} and λ_{P} (means $1/\lambda_{\mathrm{A}}$ and $1/\lambda_{\mathrm{P}}$), then $f_{\mathrm{A}}(t) = \lambda_{\mathrm{A}} \mathrm{e}^{-\lambda_{\mathrm{A}} t}$, $R_{\mathrm{P}}(t) = \mathrm{e}^{-\lambda_{\mathrm{P}} t}$, and equation (28.14) gives

$$\Pr(T_{\mathrm{A}} < T_{\mathrm{P}}) = \int_0^{\infty} \lambda_{\mathrm{A}}\, \mathrm{e}^{-(\lambda_{\mathrm{A}} + \lambda_{\mathrm{P}})t}\, \mathrm{d}t = \frac{\lambda_{\mathrm{A}}}{\lambda_{\mathrm{A}} + \lambda_{\mathrm{P}}}. \tag{28.15}$$

The time to the first of the two events is the minimum of two independent exponentials, $\Pr(\min > t) = \mathrm{e}^{-\lambda_{\mathrm{A}} t} \mathrm{e}^{-\lambda_{\mathrm{P}} t}$, exponential with rate $\lambda_{\mathrm{A}} + \lambda_{\mathrm{P}}$; the probability that the corrective action is that first event is its share of the total rate. For a corrective action with mean 8 min and a protective trip with mean 32 min (illustrative),

$$\Pr(T_{\mathrm{A}} < T_{\mathrm{P}}) = \frac{1/8}{1/8 + 1/32} = \frac{4}{5} = 0.80,$$

the success probability of the corrective-action heading used in Section 28.4.3, and the first event occurs after $1/(1/8 + 1/32) = 6.4$ min on average.

Figure 28.4 shows the integral of equation (28.14) for these two exponentials. The horizontal axis is the time after the initiating outage in minutes and the vertical axis a probability density per minute. The dashed blue curve is the density $f_{\mathrm{A}}(t)$ of the corrective-action time, whose area is one; the solid red curve is the integrand $f_{\mathrm{A}}(t)\, R_{\mathrm{P}}(t)$, the density of a corrective action at t with protection not yet operated. The shaded area under it is the probability of arrest, 0.80; the area between the two curves, 0.20, is the probability that protection operates first. The ratio of the solid to the dashed curve, $R_{\mathrm{P}}(t)$, falls with time: late corrective actions are the ones most likely to be pre-empted.

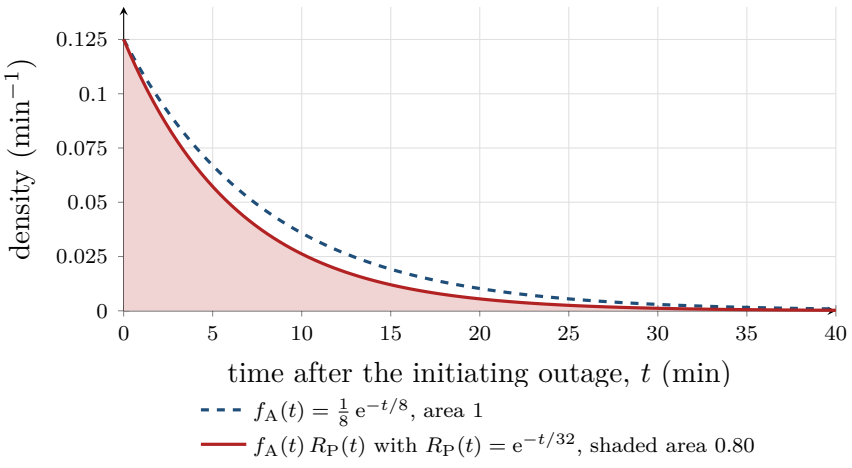


Figure 28.4: The race between corrective action and protective tripping, equation (28.14), for exponential times with means 8 min (corrective action) and 32 min (protective trip), illustrative. The shaded area is $\Pr(T_{\mathrm{A}} < T_{\mathrm{P}}) = \lambda_{\mathrm{A}}/(\lambda_{\mathrm{A}} + \lambda_{\mathrm{P}}) = 0.80$.

Common mistake

Comparing mean times. That the corrective action is faster on average than the protective trip, $\mathbb{E}[T_A] < \mathbb{E}[T_P]$, does not make arrest likely. With exponential means of 8 and 10 min, equation (28.15) gives $\Pr(T_A < T_P) = 0.125/0.225 = 0.556$: the faster mechanism loses the race almost half the time. The branch probability needs the distributions of the two times, not their means.

28.4.3 The cascade scenario: initiating event and three stages

The initiating event I is the loss of L13 in the three-bus network of Figure 28.1, with an illustrative initiating-event frequency

$$f_{\mathrm{IE}} = 0.20\,\mathrm{yr}^{-1},$$

the expected number of L13 outages per year. It is not the 1 % stationary unavailability of L13 used in Chapter 27: a frequency counts events per year and has units, an unavailability is the probability of finding the line out at a random instant and has none (Section 15.1.2). After L13 disconnects, $P_{12} = P_{23} = 100\,\mathrm{MW}$ against continuous ratings of 70 MW, and the disturbance can unfold in three stages. All probabilities and durations of the scenario are hypothetical.

Stage 1: corrective action. Let A be the event that a remedial action reduces the transfer before any further line is lost, with $\Pr(A \mid I) = 0.80$, the value of equation (28.15) for the means of the previous subsection. If it succeeds, the simplified model attributes no unserved energy to the event, $E^{\mathrm{unserved}} = 0$; this assumes the transfer reduction is achieved by redispatch to generation that the three-bus model does not show, for instance a unit near the load. If it fails, with $\Pr(\bar{A} \mid I) = 0.20$, the network stays in its stressed condition. (The event A is unrelated to the availability A.)

Stage 2: additional line trip. Given that the corrective action failed, let B be the event that an additional line trips, with $\Pr(B \mid I, \bar{A}) = 0.60$ and $\Pr(\bar{B} \mid I, \bar{A}) = 0.40$. If no further line trips, a delayed controlled curtailment removes 20 MW for two hours,

$$E^{\mathrm{unserved}}_{\mathrm{controlled}} = 20\,\mathrm{MW} \times 2\,\mathrm{h} = 40\,\mathrm{MWh}.$$

This branch is a simplified emergency response. A curtailment of 20 MW does not by itself restore both surviving lines within their 70 MW continuous ratings in the post-L13 topology, which needs the 30 MW found above; the branch represents a later operating condition, such as an emergency rating of 80 MW allowed for the two hours or an additional corrective measure, and should not be read as the solution of the unmodified DC model. If an additional line does trip, the series path L12–L23 is broken, generation and load are disconnected, and the whole 100 MW is interrupted.

Stage 3: restoration. Given the secondary disconnection, the restoration is fast or prolonged with probability 0.50 each. Let PR denote prolonged restoration. Fast restoration interrupts 100 MW for one hour, prolonged restoration for eight hours:

$$E^{\mathrm{unserved}}_{\mathrm{fast}} = 100\,\mathrm{MW} \times 1\,\mathrm{h} = 100\,\mathrm{MWh}, \qquad E^{\mathrm{unserved}}_{\mathrm{prolonged}} = 100\,\mathrm{MW} \times 8\,\mathrm{h} = 800\,\mathrm{MWh}.$$

28.4.4 The cascade event tree and its sequence probabilities

Figure 28.5 draws the event tree, in the layout of Chapter 15. The initiating event is the box on the left. The three column headings are the corrective action, the absence of a secondary

trip and fast restoration, each phrased so that success (S, upward) is the favorable outcome; each branch carries its conditional probability. To the right, each of the four terminal sequences carries its number k, its end state (green when no load is lost, red otherwise), its conditional probability $\Pr(S_k \mid I)$ and its unserved energy E_k^{unserved}. When the corrective action succeeds, the later headings are not demanded and the line of sequence 1 crosses their columns without a branch; when no secondary trip occurs, the restoration heading is not demanded.

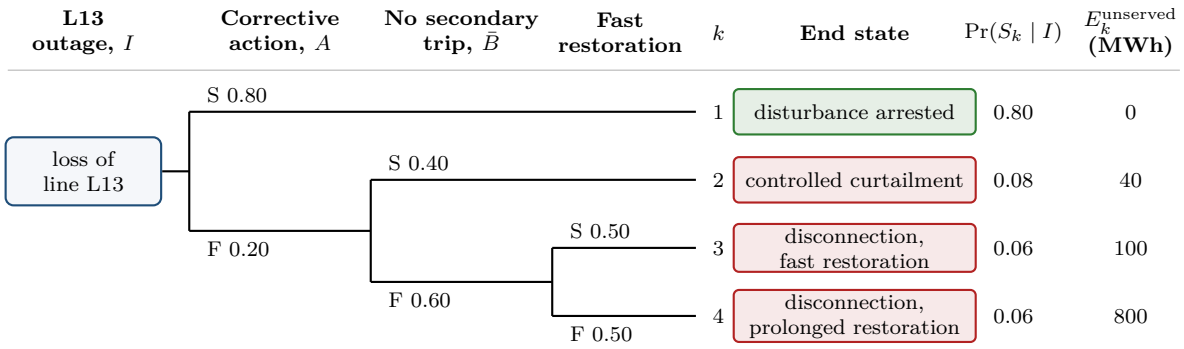


Figure 28.5: Event tree of the cascade that follows the loss of L13, $f_{\text{IE}} = 0.20\,\text{yr}^{-1}$ (illustrative). The secondary-trip heading is demanded only when the corrective action fails, and the restoration heading only after a secondary trip. The four end states are mutually exclusive; their probabilities sum to 1.

By the chain rule of Chapter 15, each sequence probability is the product of the conditional branch probabilities along its path, $\Pr(S_k \mid I) = \prod_j \Pr(B_j \mid I, B_1, \ldots, B_{j-1})$ (equation (15.6)), where B_j is the outcome at the j-th heading on the path:

$$
\begin{aligned}
\Pr(S_1 \mid I) &= \Pr(A \mid I) = 0.80,\\
\Pr(S_2 \mid I) &= \Pr(\bar{A} \mid I)\ \Pr(\bar{B} \mid I, \bar{A}) = 0.20 \times 0.40 = 0.08,\\
\Pr(S_3 \mid I) &= \Pr(\bar{A} \mid I)\ \Pr(B \mid I, \bar{A})\ \Pr(\overline{\text{PR}} \mid I, \bar{A}, B) = 0.20 \times 0.60 \times 0.50 = 0.06,\\
\Pr(S_4 \mid I) &= \Pr(\bar{A} \mid I)\ \Pr(B \mid I, \bar{A})\ \Pr(\text{PR} \mid I, \bar{A}, B) = 0.20 \times 0.60 \times 0.50 = 0.06.
\end{aligned}
$$

They sum to $0.80 + 0.08 + 0.06 + 0.06 = 1$, the completeness check of equation (15.10). Table 28.3 lists, for each sequence, its path, its end state, its conditional probability, its frequency $f_{S_k} = f_{\text{IE}} \Pr(S_k \mid I)$ (equation (15.7)), its unserved energy, and its contribution $\Pr(S_k \mid I)\, E_k^{\text{unserved}}$ to the expected unserved energy per initiating event computed in the next subsection. The frequencies sum to $0.200\,\text{yr}^{-1}$, the initiating-event frequency.

The probability of a secondary disconnection, given the initiating event, is the sum over the sequences that contain it,

$$\Pr(B \mid I) = \Pr(\bar{A} \mid I)\ \Pr(B \mid I, \bar{A}) = 0.20 \times 0.60 = 0.12,$$

so 12 % of L13 outages lead to the disconnection of generation from load in this model. Its expected annual frequency is

$$f_B = f_{\text{IE}}\ \Pr(B \mid I) = 0.20 \times 0.12 = 0.024\,\text{yr}^{-1}, \tag{28.16}$$

one secondary disconnection per $1/0.024 = 41.7$ years on average under the stationary assumptions of the tree. That figure is an average occurrence rate, not a recurrence interval: under the Poisson model of Section 15.1.2 the probability of at least one such disconnection in a given year is $1 - \text{e}^{-0.024} = 0.0237$.

Table 28.3: Sequences of the cascade event tree: conditional probability, annual frequency for $f_{\mathrm{IE}} = 0.20\,\mathrm{yr}^{-1}$, unserved energy per occurrence and contribution to the expected unserved energy per initiating event (illustrative).

k	Path	End state	$\Pr(S_k \mid I)$	f_{S_k} (yr^{-1})	E_k^{unserved} (MWh)	Contribution (MWh)
1	corrective action succeeds	arrested	0.80	0.160	0	0
2	action fails, no secondary trip	controlled curtailment	0.08	0.016	40	3.2
3	action fails, secondary trip, fast restoration	disconnection, 1 h	0.06	0.012	100	6.0
4	action fails, secondary trip, prolonged restoration	disconnection, 8 h	0.06	0.012	800	48.0
	Total		1	0.200		57.2

28.4.5 Expected unserved energy per initiating event and per year

Let E_{ev} be the unserved energy that follows one initiating event; it equals E_k^{unserved} when sequence S_k occurs. By the law of total expectation over the mutually exclusive sequences,

$$\mathbb{E}[E_{\mathrm{ev}} \mid I] = \sum_k \Pr(S_k \mid I)\, E_k^{\mathrm{unserved}}, \tag{28.17}$$

which for the cascade tree is

$$\mathbb{E}[E_{\mathrm{ev}} \mid I] = 0.80 \times 0 + 0.08 \times 40 + 0.06 \times 100 + 0.06 \times 800 = 0 + 3.2 + 6.0 + 48.0 = 57.2\,\mathrm{MWh}.$$

The largest contribution comes from prolonged restoration, $0.06 \times 800 = 48.0\,\mathrm{MWh}$: a sequence that follows only 6 % of initiating events contributes $48.0/57.2 = 83.9\,\%$ of the expected unserved energy per event. Relatively rare severe sequences dominating the expected loss is a characteristic feature of cascade risk.

Over a year, the unserved energy is the sum of the unserved energies of the $N_I(1\,\mathrm{yr})$ initiating events that occur. When their outcomes are independent of one another and of their number, the expectation of the sum is the expected number of events times the expected energy per event (Wald's identity, as in Section 15.3.5), so the contribution of this initiating-event family to the expected unserved energy is

$$\mathrm{EUE} = f_{\mathrm{IE}}\, \mathbb{E}[E_{\mathrm{ev}} \mid I] = \sum_k f_{S_k}\, E_k^{\mathrm{unserved}} = 0.20\,\mathrm{yr}^{-1} \times 57.2\,\mathrm{MWh} = 11.44\,\mathrm{MWh/yr}, \tag{28.18}$$

which is equation (15.14) of Chapter 15 and the EUE of Chapter 25 restricted to one initiating-event family. The formula is exact for the expectation, not an approximation, provided the initiating events follow the stated stationary frequency and the losses of overlapping disturbances are not counted twice; it does not require the annual probability of a cascade to be small. The modeling assumption it rests on is the independence of successive events, which fails if a first disturbance leaves equipment damaged when the next one arrives or if the weather that caused one outage also raises the frequency of the next.

For comparison, Chapter 15 quantified an event tree for a generic line outage with protection, redispatch and RAS headings, $f_{\mathrm{IE}} = 0.60\,\mathrm{yr}^{-1}$, and found an expected unserved energy of

3.588 MWh/yr (Figure 15.4). The cascade tree of this chapter has a third of the initiating frequency and three times the expected unserved energy, because its secondary trip disconnects the whole load and its restoration branch lasts eight hours.

28.4.6 Importance of the protective controls

Suppose the success probability of the corrective action is raised from 0.80 to 0.95, all other conditional branch probabilities unchanged. Its failure probability falls from 0.20 to 0.05. The expected unserved energy per initiating event is proportional to it, because every loss lies below the failure branch:

$$\mathbb{E}[E_{\text{ev}} \mid I] = \Pr(\bar{A} \mid I)\,[0.40 \times 40 + 0.60\,(0.50 \times 100 + 0.50 \times 800)] = \Pr(\bar{A} \mid I) \times 286\,\text{MWh},$$

where 286 MWh is the expected unserved energy given that the corrective action has failed. With $\Pr(\bar{A} \mid I) = 0.05$ it becomes $0.05 \times 286 = 14.3$ MWh per event, and the annual EUE $0.20 \times 14.3 = 2.86$ MWh/yr, a reduction of $11.44 - 2.86 = 8.58$ MWh/yr, or 75 %, equal to the relative reduction of the failure probability. Improving the reliability of the controls is very valuable when failing to arrest a propagation has large downstream consequences.

The importance measures of Chapter 12 of Volume I rank the branches systematically. They were defined there for a system failure probability; with the EUE as the risk measure, which is common practice in probabilistic risk assessment, the definitions carry over with EUE in place of Q_{sys}. For a branch probability p, let EUE(p=1) and EUE(p=0) be the expected unserved energy with that branch made certain or impossible, all else fixed. The Birnbaum importance is $\partial\text{EUE}/\partial p$, which for a risk linear in p equals EUE(p=1) − EUE(p=0); the risk achievement worth is EUE(p=1)/EUE (equation (12.21)); the risk reduction worth EUE/EUE(p=0) (equation (12.22)); and the Fussell–Vesely importance [EUE − EUE(p=0)]/EUE, the fraction of the risk that passes through the branch. Table 28.4 evaluates them for three branches: the failure of the corrective action, $\Pr(\bar{A} \mid I) = 0.20$; the secondary trip, $\Pr(B \mid I, \bar{A}) = 0.60$; and the prolonged restoration, $\Pr(\text{PR} \mid I, \bar{A}, B) = 0.50$. For the secondary trip, for example, EUE(p=1) $= 0.20 \times 0.20 \times 450 = 18.0$ and EUE(p=0) $= 0.20 \times 0.20 \times 40 = 1.6$ MWh/yr, where $450 = 0.5 \times 100 + 0.5 \times 800$.

Table 28.4: Importance measures of three branches of the cascade tree with the annual EUE (11.44 MWh/yr) as the risk measure (illustrative).

Branch	EUE(p=1) (MWh/yr)	EUE(p=0) (MWh/yr)	Birnbaum (MWh/yr)	RAW	RRW	FV
Corrective action fails, $p = 0.20$	57.20	0	57.2	5.00	∞	1.000
Secondary trip, $p = 0.60$	18.00	1.60	16.4	1.57	7.15	0.860
Prolonged restoration, $p = 0.50$	19.84	3.04	16.8	1.73	3.76	0.734

The corrective action dominates every measure. Its Fussell–Vesely importance is one and its risk reduction worth infinite because every unit of unserved energy in the model passes through its failure; its Birnbaum importance, 57.2 MWh/yr per unit of probability, is three and a half times that of the other two. The secondary trip and the prolonged restoration have similar Birnbaum importances but differ in RRW: preventing secondary trips altogether would leave only the controlled-curtailment sequence, 1.6 MWh/yr, while eliminating prolonged restoration would leave 3.04. The ranking says where an improvement of a given size in probability has most effect; it does not say what the improvement costs, which Section 28.6 adds.

Principle 28.2. *In a cascade, the expected unserved energy is driven by the branches that let a disturbance propagate and by the severity of the sequences they lead to. The control that arrests propagation carries all the risk of the sequences below it, and its reliability is often the most valuable parameter of the model.*

28.5 Branching models and chronological Monte Carlo simulation

An event tree follows a cascade of a few stages in detail. This section presents two ways to go further: branching processes, which summarize propagation by one statistical parameter, and chronological Monte Carlo simulation, which follows the whole hybrid system of Figure 28.2 through time.

28.5.1 Branching-process models of cascades

A full model of a transmission cascade tracks individual lines, the topology, the power flows and the protection actions. An aggregated description is sometimes wanted instead, one that answers how many further outages one outage produces on average. A branching process captures the statistical propagation of outages across generations. Let

$$N_0 = \text{number of initiating outages}, \qquad N_n = \text{number of outages in generation } n \geq 1,$$

where the outages of generation $n+1$ are those caused directly by outages of generation n, their *offspring*. Each outage produces a random number of offspring, with mean

$$\lambda_{\mathrm{b}} = \mathbb{E}[\text{number of further outages caused by one outage}], \tag{28.19}$$

the *branching ratio* (mean offspring number, or propagation mean) of the cascade, dimensionless. (It is not a failure rate, despite the letter.) The *cascade size* is the total number of outages in all generations, the initiating ones included,

$$N_{\mathrm{c}} = \sum_{n=0}^{\infty} N_n. \tag{28.20}$$

The simplest model, the Galton–Watson branching process, assumes that every outage produces its offspring independently of every other outage and of the history, with one offspring distribution common to all. Studies of historical sequences of transmission outages have used such models to estimate the number of outages that propagation produces from the counts of outages in successive generations.

Figure 28.6 illustrates the model. Panel (a) is one realization of a Galton–Watson cascade: the circles are outages, each row one generation, and the lines join each outage to the outages it causes; this realization has $N_0 = 1$, $N_1 = 2$, $N_2 = 1$ and $N_3 = 0$, so it stops after the second generation with cascade size $N_{\mathrm{c}} = 4$. Panel (b) is the expected cascade size, derived in the next subsection, as a function of the branching ratio.

28.5.2 Expected cascade size

Given N_n, the number of outages in generation $n+1$ is the sum of the offspring counts of the N_n outages of generation n, each with mean λ_{b} and independent of N_n, so by linearity of expectation

$$\mathbb{E}[N_{n+1} \mid N_n] = \lambda_{\mathrm{b}}\, N_n.$$

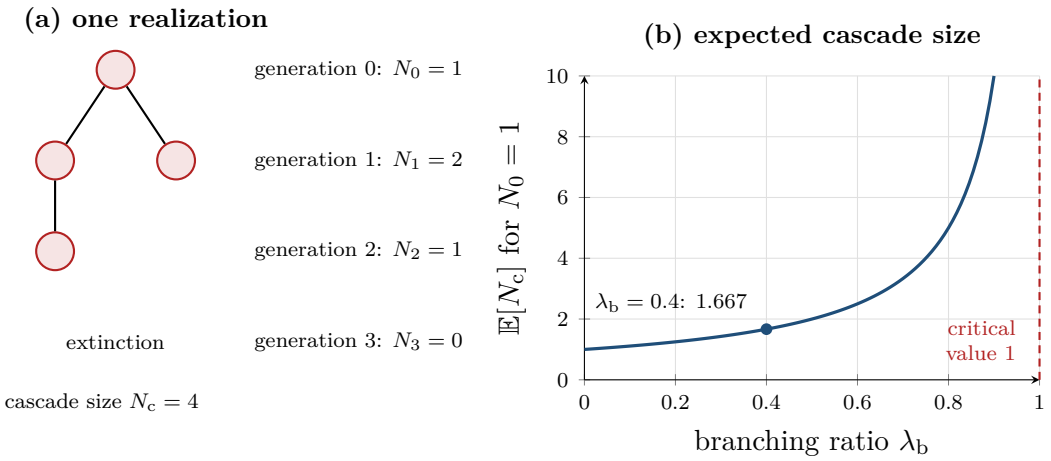


Figure 28.6: Branching-process model of a cascade. (a) One realization of a Galton–Watson process started by one outage. (b) Expected cascade size $\mathbb{E}[N_c] = 1/(1 - \lambda_b)$, equation (28.22), for one initiating outage; it grows without bound as the branching ratio approaches the critical value 1.

Taking the expectation of both sides (the tower property), $\mathbb{E}[N_{n+1}] = \lambda_b \, \mathbb{E}[N_n]$, and by induction from $\mathbb{E}[N_0] = N_0$,

$$\mathbb{E}[N_n] = N_0 \, \lambda_b^n. \tag{28.21}$$

The expected cascade size is the sum of the expected generation sizes; the exchange of expectation and infinite sum in equation (28.20) is allowed because the terms are nonnegative. For $\lambda_b < 1$ the geometric series converges, and

$$\mathbb{E}[N_c] = N_0 \sum_{n=0}^{\infty} \lambda_b^n = \frac{N_0}{1 - \lambda_b}. \tag{28.22}$$

For one initiating outage, $N_0 = 1$, and a branching ratio $\lambda_b = 0.4$,

$$\mathbb{E}[N_c] = \frac{1}{1 - 0.4} = 1.667,$$

which includes the initiating outage; the expected number of subsequent outages is $0.667 = \lambda_b/(1 - \lambda_b)$. At $\lambda_b = 0.6$ the expected size is 2.5, at 0.8 it is 5, at 0.9 it is 10: the size is very sensitive to the branching ratio near one, which is the shape of Figure 28.6(b).

The mean hides a long tail. If the offspring number is Poisson with mean λ_b (an illustrative assumption), the cascade size of a process started by one outage has the Borel distribution, a standard result, $\Pr(N_c = k) = \mathrm{e}^{-\lambda_b k}(\lambda_b k)^{k-1}/k!$ for $k = 1, 2, \ldots$, with variance $\lambda_b/(1 - \lambda_b)^3$. For $\lambda_b = 0.4$ it gives $\Pr(N_c = 1) = \mathrm{e}^{-0.4} = 0.670$, $\Pr(N_c = 2) = 0.180$ and $\Pr(N_c = 3) = 0.072$, so that 7.8 % of cascades reach four outages or more although the mean size is 1.667; the standard deviation of the size is $\sqrt{0.4/0.216} = 1.36$.

28.5.3 The branching ratio and the limits of branching models

Table 28.5 summarizes the behavior of the process by its branching ratio. Each row is one regime, with the mathematical behavior under the standard assumptions (an offspring distribution that is not concentrated on exactly one offspring). Below one, the process is *subcritical*: it dies out with probability one and its expected total size is finite. At one it is *critical*: it still dies out with probability one, but its expected size is infinite. Above one it is *supercritical*: it survives forever with positive probability; for Poisson offspring with $\lambda_b = 1.5$, for example, the extinction probability is the smallest root of $q = \mathrm{e}^{\lambda_b(q-1)}$, $q = 0.417$.

Table 28.5: Behavior of a Galton–Watson process by its branching ratio.

Branching ratio	Regime	Mathematical behavior	Reading for a power network
$\lambda_{\mathrm{b}} < 1$	subcritical	extinction with probability one; $\mathbb{E}[N_{\mathrm{c}}] = N_0/(1-\lambda_{\mathrm{b}})$ finite	propagation dies out; mean size grows sharply as λ_{b} nears one
$\lambda_{\mathrm{b}} = 1$	critical	extinction with probability one, infinite expected size	boundary of strong propagation; heavy-tailed sizes
$\lambda_{\mathrm{b}} > 1$	supercritical	positive probability of unbounded growth	strong modeled propagation, limited in reality by the finite network

For $\lambda_{\mathrm{b}} \geq 1$ the infinite-generation model has no finite expected size, but a real network has finitely many elements, and its cascade cannot produce infinitely many distinct outages. A branching ratio at or above one therefore signals strong modeled propagation, the regime in which large blackouts become likely, not a literal infinite blackout. A branching model may report $\lambda_{\mathrm{b}} = 0.4$, but it generally does not identify:

- which line trips next, and whether the network separates into islands;
- how much demand becomes disconnected;
- whether voltage instability occurs;
- how the line temperatures evolve, and whether the protective controls act;
- which specific intervention best prevents propagation.

Two cascades of three line trips each can have very different consequences depending on where those lines are: in the network of Figure 28.1, a second trip on the series path disconnects the whole load, while in a meshed network the same count of trips might shed nothing. Studies of cascading blackouts describe them as initiating outages followed by a propagation that progressively weakens the system, and stress that initiation and propagation may call for different mitigations; the branching ratio measures the second but cannot locate it. Branching models are useful for aggregate propagation statistics; they are not substitutes for physical network models.

28.5.4 State and events of a chronological simulation

A chronological (sequential) Monte Carlo simulation follows the system through simulated time, sampling the random events and solving the physical model between them. It needs both the continuous and the discrete states. A useful state vector is

$$\mathbf{x}(t) = \big[\, \mathbf{S}(t),\ \mathbf{T}_{\mathrm{c}}(t),\ \boldsymbol{\theta}(t),\ \mathbf{P}_{\mathrm{G}}(t),\ \mathbf{L}(t),\ \mathbf{u}(t) \,\big], \tag{28.23}$$

where $\mathbf{S}(t)$ is the vector of line statuses $S_\ell(t)$, $\mathbf{T}_{\mathrm{c}}(t)$ the conductor temperatures, $\boldsymbol{\theta}(t)$ the bus voltage angles (or the other power-system state variables of an AC model), $\mathbf{P}_{\mathrm{G}}(t)$ the generator outputs, $\mathbf{L}(t)$ the loads $L_i(t)$, and $\mathbf{u}(t)$ the corrective controls in force. The model also needs the event and timing information that drives the discrete part: weather-related failures (from intensities such as equation (28.4)), protection operations, automatic remedial actions, operator responses, and restoration and repair.

A simplified event-driven cascade simulation proceeds as follows.

1. Initialize the network topology, the loads, the weather and the conductor temperatures.
2. For each simulation period, update the weather-dependent failure intensities and sample the time of the next initiating failure.

3. When a failure occurs, remove the failed component and recalculate the network power flows.
4. While the system is disturbed:
   (a) advance the conductor temperatures and the other physical states to the next candidate event time;
   (b) evaluate the protection and corrective-action conditions, and determine which event occurs next and when;
   (c) if a corrective action occurs, update the generation, the load or the topology;
   (d) if protection trips another component, update the topology and recalculate the power flows;
   (e) if no further event is expected, mark the cascade arrested and leave the loop.
5. Record the load interruption and the restoration duration.
6. Continue until the end of the simulation horizon.

The simulation must keep the physical dynamics and the discrete events consistent. If a line trips at time t, the flows must reflect the new topology from t onward, not from the next step of a coarse hourly grid; an hourly snapshot model would miss exactly the minutes in which Figure 28.3 decides the outcome.

28.5.5 Competing event times in the simulation

Step 4(b) of the algorithm decides which event comes first. Suppose an overloaded line is subject to a further outage with a time-varying intensity $\lambda_{\mathrm{trip}}(t)$, measured from the initiating event, while the corrective action occurs at the random time T_{A}. The probability that the line survives to t without tripping is, by equation (28.3),

$$R_{\mathrm{trip}}(t) = \exp\!\left[-\int_0^t \lambda_{\mathrm{trip}}(u)\,\mathrm{d}u\right].$$

If T_{A} is independent of the trip process, given the modeled state and conditions, then conditioning on T_{A} gives $\Pr(\text{correction before trip} \mid T_{\mathrm{A}} = t) = R_{\mathrm{trip}}(t)$, and averaging over T_{A},

$$\Pr(\text{correction before trip}) = \mathbb{E}\big[R_{\mathrm{trip}}(T_{\mathrm{A}})\big], \tag{28.24}$$

the general form of equation (28.14). Both a faster control, which shifts T_{A} toward zero, and a lower propensity for secondary failure, which raises R_{trip}, increase the probability of arresting the propagation.

The intensity can be tied to the thermal model. Suppose, for illustration, that the line cannot trip while its temperature is below the limit and trips at a rate of $0.5\,\mathrm{min}^{-1}$ once the limit is exceeded, so that $\lambda_{\mathrm{trip}}(t) = 0$ for $t < t_{\mathrm{lim}} = 13.74\,\mathrm{min}$ and $0.5\,\mathrm{min}^{-1}$ after, and let T_{A} be exponential with mean $8\,\mathrm{min}$. Splitting the expectation at t_{lim},

$$\mathbb{E}\big[R_{\mathrm{trip}}(T_{\mathrm{A}})\big] = \Pr(T_{\mathrm{A}} < t_{\mathrm{lim}}) + \int_{t_{\mathrm{lim}}}^{\infty} \tfrac{1}{8}\,\mathrm{e}^{-t/8}\,\mathrm{e}^{-0.5(t-t_{\mathrm{lim}})}\,\mathrm{d}t = \big(1 - \mathrm{e}^{-t_{\mathrm{lim}}/8}\big) + \tfrac{1}{5}\,\mathrm{e}^{-t_{\mathrm{lim}}/8},$$

which with $\mathrm{e}^{-13.74/8} = 0.1794$ is $0.8206 + 0.0359 = 0.8565$. In physical systems, however, the trip intensity and the corrective action time need not be independent. Shared communications, the same weather, faulty sensors and the operating conditions can couple them; a storm that raises the trip intensity may also delay the operators. The independence in equation (28.24) is then replaced by a joint model of the two clocks, conditioned on the shared variables.

28.5.6 Estimating cascade risk and the rare-event problem

Suppose N_{yr} independent years are simulated, and let $E^{(j)}$ be the unserved energy in simulated year j. The estimate of the expected unserved energy is the sample mean,

$$\widehat{\mathrm{EUE}} = \frac{1}{N_{\mathrm{yr}}} \sum_{j=1}^{N_{\mathrm{yr}}} E^{(j)}. \tag{28.25}$$

For the annual frequency of cascades, count the events that satisfy a clearly stated cascade definition; with $K^{(j)}$ the number of qualifying cascades in year j,

$$\hat{f}_{\mathrm{casc}} = \frac{1}{N_{\mathrm{yr}}} \sum_{j=1}^{N_{\mathrm{yr}}} K^{(j)}. \tag{28.26}$$

The fraction of simulated years with $E^{(j)} > E_{\mathrm{crit}}$ estimates the annual probability of exceeding a selected consequence threshold, $\Pr(E^{(j)} > E_{\mathrm{crit}})$, the quantity behind the frequency–consequence curves of Chapter 15 of Volume I.

Each estimate is a sample mean, so its standard error is the standard deviation of one year divided by $\sqrt{N_{\mathrm{yr}}}$. For the cascade tree the annual unserved energy is a compound Poisson sum, the sum of a Poisson number of independent event energies, whose variance is a standard result: $\mathrm{Var}[E^{(j)}] = f_{\mathrm{IE}}\,\mathbb{E}[E_{\mathrm{ev}}^2 \mid I]$. With $\mathbb{E}[E_{\mathrm{ev}}^2 \mid I] = 0.08{\times}40^2 + 0.06{\times}100^2 + 0.06{\times}800^2 = 39\,128\,\mathrm{MWh}^2$, the variance is $0.20 \times 39\,128 = 7826\,\mathrm{MWh}^2$ and the standard deviation 88.5 MWh, almost eight times the mean 11.44 MWh. A relative standard error of 10 % requires $88.5/\sqrt{N_{\mathrm{yr}}} = 1.144$, that is $N_{\mathrm{yr}} \approx 5980$ simulated years. For a frequency, the count $K^{(j)}$ is Poisson with mean and variance f, so the relative standard error of $\hat{f}_{\mathrm{casc}}$ is $1/\sqrt{f N_{\mathrm{yr}}}$ and 10 % needs $N_{\mathrm{yr}} = 1/(0.01 f)$ years: 8333 years for the prolonged-restoration sequence of frequency $0.012\,\mathrm{yr}^{-1}$, and 100 000 years for a severe cascade of frequency $10^{-3}\,\mathrm{yr}^{-1}$.

The rare-event challenge. Severe cascades can be very rare, and ordinary Monte Carlo can need an impractical number of simulated years to estimate their probabilities precisely. The usual remedies are

- importance sampling of initiating events and weather;
- stratification by weather state, estimating the risk within each state z and combining, $\mathrm{EUE} = \sum_z \Pr(Z = z)\,\mathrm{EUE}_z$;
- conditional simulation, which simulates only from the initiating event onward and multiplies by its frequency, the structure of equation (28.18);
- splitting methods, which replicate the trajectories that come close to a severe outcome;
- hybrid methods that compute part of the model exactly and sample the rest.

The sampling weights and conditional probabilities must remain valid for the physical model. In importance sampling, a scenario drawn with probability $q(x)$ instead of its true probability $p(x)$ carries the weight $w = p(x)/q(x)$, and the estimate $N_{\mathrm{yr}}^{-1} \sum_j w_j E^{(j)}$ is unbiased. If the severe weather state of Section 28.2.4 is sampled with probability 0.50 instead of 0.10, each severe scenario carries the weight $0.10/0.50 = 0.2$ and each normal one $0.90/0.50 = 1.8$. A simulation that produces thousands of dramatic cascades does not demonstrate a high real risk if its initiating scenarios were deliberately oversampled and those weights were not applied.

Common mistake

Reporting the cascades of a biased simulation as if they were typical. Oversampling makes rare scenarios visible; only the likelihood-ratio weights turn the counts back into frequencies.

28.6 Risk-based decisions, mitigation and model credibility

The models of the previous sections exist to support decisions. This section states what can be optimized with them, computes the value of an additional measurement during an overload, tabulates the annual risk of the cascade for several mitigation and restoration parameter sets, and closes with the pitfalls of cascade models and their engineering context.

28.6.1 What can be optimized

A cascade model can support decisions such as which line should receive additional monitoring; where remedial action schemes are most valuable; which line-rating uncertainty contributes most to the expected risk; how much time a better estimate of the conductor temperature gains; which operator response delays dominate propagation; and which transmission upgrades best reduce severe consequences. Let a denote a proposed intervention. A general risk objective is

$$J(a) = \mathbb{E}[C \mid a] + c_{\mathrm{inv}}(a), \tag{28.27}$$

the expected outage consequence under the intervention plus its investment cost $c_{\mathrm{inv}}(a)$, both annualized; the best intervention minimizes J. For a risk-constrained decision, the cheapest intervention is sought that keeps the probability of a severe event below a tolerance,

$$\min_{a} c_{\mathrm{inv}}(a) \quad \text{subject to} \quad \Pr(E_{\mathrm{ev}} > E_{\max} \mid a) \le \varepsilon, \tag{28.28}$$

where $E_{\max}$ is a consequence threshold and ε the tolerated probability per initiating event.

The two formulations can rank interventions differently. In the cascade tree, raising the corrective-action success to 0.95 lowers the EUE from 11.44 to 2.86 MWh/yr (Section 28.4.6); shortening prolonged restoration from eight to three hours lowers it only to 5.44 MWh/yr (Table 28.6). With $E_{\max} = 500$ MWh and $\varepsilon = 0.01$, however, the better corrective action still leaves $\Pr(E_{\mathrm{ev}} > 500 \mid a) = 0.05 \times 0.60 \times 0.50 = 0.015 > 0.01$, while the faster restoration removes every sequence above 500 MWh, since its longest interruption is now 300 MWh. The best investment is therefore not necessarily the one that most reduces the initiating frequency or even the expected value; an intervention that reduces propagation or restoration time can have the larger effect on severe-event risk.

28.6.2 The value of additional measurements

Suppose L12 is overloaded after the outage of L13, and its actual conductor temperature is uncertain because weather observations along the route are sparse. Two situations are possible: favorable wind cooling that gives substantial thermal headroom, or low wind and high ambient temperature that leave very little time before the limit. A new conductor-temperature or weather measurement can narrow that uncertainty, but it is valuable only if it can change the decision among actions such as continuing the emergency transfer briefly, redispatching immediately, curtailing load, or applying a different emergency response. The expected value of sample information of Chapter 24 (equation (24.18)) states this precisely. With $\mathbf{y}$ the evidence already available, Y the outcome of the new measurement and C the consequence,

$$\mathrm{EVSI} = \min_{a} \mathbb{E}[C \mid \mathbf{y}, a] - \mathbb{E}_{Y}\Big[\min_{a} \mathbb{E}[C \mid \mathbf{y}, Y, a]\Big], \tag{28.29}$$

before the cost of the measurement: the expected loss of the best action now, minus the expected loss of the best action chosen after seeing Y.

An illustrative case shows how the formula works. The weather along L12 is favorable with probability 0.70, giving $T_\infty = 100\,^{\circ}\mathrm{C}$ and $\tau_{\mathrm{th}} = 15$ min and, by equation (28.12), $t_{\mathrm{lim}} = 20.8$ min; or unfavorable with probability 0.30, giving $T_\infty = 130\,^{\circ}\mathrm{C}$ and $\tau_{\mathrm{th}} = 12$ min, $t_{\mathrm{lim}} = 6.7$ min. The

redispatch takes eight minutes. Action a_1 is to continue the transfer while redispatching: in favorable weather the redispatch arrives in time and nothing is lost; in unfavorable weather the limit is reached first, and the system enters the failed-correction branch of the cascade tree, whose expected unserved energy is 286 MWh (Section 28.4.6). Action a_2 is to curtail 30 MW at once until the redispatch is complete, which costs $30\,\text{MW} \times 8/60\,\text{h} = 4\,\text{MWh}$ in either weather. Measured in unserved energy,

$$\mathbb{E}[C \mid a_1] = 0.30 \times 286 = 85.8\,\text{MWh}, \qquad \mathbb{E}[C \mid a_2] = 4\,\text{MWh},$$

so without a measurement the best action is to curtail, with an expected loss of 4 MWh. With perfect knowledge of the weather the operator would continue in favorable weather and curtail in unfavorable weather, with expected loss $0.30 \times 4 = 1.2\,\text{MWh}$; the expected value of perfect information is $4 - 1.2 = 2.8\,\text{MWh}$ per event. A real measurement is imperfect. Continuing is the better action only if the probability of unfavorable weather after the measurement is below $4/286 = 0.0140$. A measurement that flags unfavorable weather with sensitivity 0.90 and specificity 0.90 leaves, after a negative reading, $\Pr(\text{unfavorable} \mid Y^-) = 0.03/0.66 = 0.045$, above the threshold; the decision is to curtail whatever the reading, and the EVSI is zero. With sensitivity and specificity 0.99, a negative reading (probability 0.696) leaves $0.003/0.696 = 0.0043$, below the threshold, so the operator continues after it and curtails after a positive one; the expected loss is $0.304 \times 4 + 0.003 \times 286 = 1.216 + 0.858 = 2.07\,\text{MWh}$, and the EVSI is $4 - 2.07 = 1.93\,\text{MWh}$ per overload event. A measurement is worth its cost only through the decisions it can change.

28.6.3 Mitigation and restoration: a worked table

The cascade tree can be recomputed for any set of its parameters. With f_{IE} the initiating frequency, $s = \Pr(A \mid I)$ the corrective-action success, $b = \Pr(B \mid I, \bar{A})$ the secondary-trip probability, $\ell = \Pr(\text{PR} \mid I, \bar{A}, B)$ the prolonged-restoration probability and h the prolonged restoration duration in hours, the sequence probabilities are

$$\Pr(S_2 \mid I) = (1-s)(1-b), \qquad \Pr(S_3 \mid I) = (1-s)\,b\,(1-\ell), \qquad \Pr(S_4 \mid I) = (1-s)\,b\,\ell,$$

the expected unserved energy per initiating event is $\mathbb{E}[E_{\text{ev}} \mid I] = 40\ \Pr(S_2 \mid I) + 100\ \Pr(S_3 \mid I) + 100\,h\ \Pr(S_4 \mid I)$ in MWh, the annual EUE is f_{IE} times that, and the frequency of secondary disconnections is $f_B = f_{\text{IE}}(1-s)\,b$. Table 28.6 evaluates the model for the base case and six variants, named by the column headings: the corrective-action success raised to 0.95 (better action); the prolonged restoration shortened from eight to three hours (shorter outage); the secondary-trip probability halved to 0.30 (fewer trips); the prolonged-restoration probability lowered to 0.20 (fewer long); the initiating frequency doubled to $0.40\,\text{yr}^{-1}$; and the first two improvements together (both). The rows above the first rule are the inputs, the rows below it the outputs; the last row is the probability per initiating event of a sequence above 500 MWh.

The shorter-outage column shows the arithmetic. The sequence probabilities are unchanged, 0.08, 0.06 and 0.06, and only the energy of the prolonged sequence changes, from 800 to $100\,\text{MW} \times 3\,\text{h} = 300\,\text{MWh}$: $\mathbb{E}[E_{\text{ev}} \mid I] = 0.08 \times 40 + 0.06 \times 100 + 0.06 \times 300 = 3.2 + 6.0 + 18.0 = 27.2\,\text{MWh}$, and the annual EUE is $0.20 \times 27.2 = 5.44\,\text{MWh/yr}$.

Both improvements suggested by the base case reduce the EUE, through different mechanisms. The better corrective action reduces the probability of reaching the severe sequences, and it reduces the frequency of secondary disconnections fourfold. The faster restoration leaves every probability and every frequency unchanged and reduces the severity of the sequences once they occur. The distinction matters when investments are compared: one acts on propagation, the other on consequence, and a reliability target stated as a frequency of disconnections would credit only the first. The three single changes to the propagation and restoration branches (fewer secondary trips, fewer prolonged restorations, faster restoration) give similar annual EUE,

Table 28.6: Annual risk of the cascade tree for several mitigation and restoration parameter sets (illustrative). Changed inputs are in bold.

Quantity	Base	Better action	Shorter outage	Fewer trips	Fewer long	Double f_{IE}	Both
f_{IE} (yr^{-1})	0.20	0.20	0.20	0.20	0.20	**0.40**	0.20
Corrective-action success s	0.80	**0.95**	0.80	0.80	0.80	0.80	**0.95**
Secondary trip b	0.60	0.60	0.60	**0.30**	0.60	0.60	0.60
Prolonged restoration ℓ	0.50	0.50	0.50	0.50	**0.20**	0.50	0.50
Prolonged duration h (h)	8	8	**3**	8	8	8	**3**
$\Pr(S_2 \mid I)$, curtailment	0.080	0.020	0.080	0.140	0.080	0.080	0.020
$\Pr(S_3 \mid I)$, fast	0.060	0.015	0.060	0.030	0.096	0.060	0.015
$\Pr(S_4 \mid I)$, prolonged	0.060	0.015	0.060	0.030	0.024	0.060	0.015
$\mathbb{E}[E_{\mathrm{ev}} \mid I]$ (MWh)	57.2	14.3	27.2	32.6	32.0	57.2	6.8
Annual EUE (MWh/yr)	11.44	2.86	5.44	6.52	6.40	22.88	1.36
f_B, secondary disconnections (yr^{-1})	0.024	0.006	0.024	0.012	0.024	0.048	0.006
$\Pr(E_{\mathrm{ev}} > 500\,\mathrm{MWh} \mid I)$	0.060	0.015	0	0.030	0.024	0.060	0

between 5.4 and 6.5 MWh/yr, but different severe-event probabilities. Doubling the initiating frequency doubles every annual output and leaves every conditional one unchanged. The last column shows that improvements to propagation and to restoration multiply: the EUE falls by 88 %, more than either alone.

28.6.4 Modeling pitfalls

Table 28.7 lists the pitfalls that most often undermine a time-dependent or cascade reliability model, each with the reason it matters and the section of this chapter that treats it. Most of them are the substitution of a simpler quantity for the one the question requires: an overload for a trip, a marginal probability for a joint one, an unavailability for a frequency, a mean time for a distribution, a count of outages for their location.

28.6.5 Standards and engineering context

Three published references frame the material of this chapter. A standard method for the current–temperature relationship of bare overhead conductors computes, under steady-state and time-varying conditions, what equation (28.8) approximates; detailed rating and time-to-limit calculations follow it rather than the first-order model. The definition of cascading used by the power-system reliability organizations, restated in the jargon box, distinguishes uncontrolled widespread propagation from ordinary sequential or controlled outages, so that the two-stage sequence of this chapter's network is cascade-like but not a cascade in that sense. Recent transmission planning requirements for extreme temperature events, in force with phased implementation, require planners to study the system under extreme heat and cold. A probabilistic, weather-conditioned study of the kind built in this chapter can inform such planning, but it does not by itself fulfill the specific obligations of a planning standard, which prescribe their own events, conditions and performance criteria. The further reading at the end of the chapter names the three documents.

This chapter closes the part of the book on power-system reliability. The models of adequacy, contingencies and cascades all assumed that the state of the network, its flows, temperatures and statuses, is known; Chapter 29 asks which states and failure conditions can actually be inferred from the measurements available.

Table 28.7: Pitfalls of time-dependent and cascade reliability models.

Pitfall	Why it matters	Section
Equating an overload with an immediate trip	Thermal loading, relay behavior and emergency operating requirements are distinct	28.4.1
Treating weather-related failures as independent	Shared weather can strongly increase the joint outage probability	28.2.4
Treating a snapshot unavailability as an event frequency	They have different units and meanings	28.4.3
Ignoring response-time distributions	The mean control time does not determine the probability of successful intervention	28.4.2
Using constant line ratings in extreme conditions	Weather and conductor temperature change the transfer capability	28.3.3
Ignoring protection and RAS dependencies	Shared sensors, communications and control power create common failure modes	28.4.1
Applying branching statistics without power-flow models	Aggregate propagation cannot locate consequences	28.5.3
Ignoring rare-event simulation error	Severe-event probabilities may be poorly estimated	28.5.6
Treating a two-line sequence as a validated system-wide cascade	Formal cascading criteria involve uncontrolled propagation beyond the studied area	28.1.1

28.7 What to retain

Definition 28.1 (Time-varying reliability and cascading failure)**.** A cascading failure is a time-dependent process in which an initiating outage redistributes flows, conductors heat with their thermal time constant, and protection, corrective controls and restoration compete in time; its risk is the initiating-event frequency times the expected consequence over the sequences that follow, computed with weather-dependent intensities, the dependence a shared weather state creates, and the distributions, not the means, of the competing times.

Cascading failures are time-dependent processes whose consequences depend on the order and timing of failures, protective operations and corrective actions. Failure hazards change with weather and operating conditions, and the integrated exposure, not the instantaneous hazard, determines the failure probability over an interval. Shared weather creates dependence between outages. Thermal inertia gives time to respond, but emergency operation and protection must be modeled separately. Event trees connect initiation, propagation, control reliability and restoration consequences; severe outcomes can dominate the expected unserved energy even when they are rare. Branching models quantify aggregate propagation but need physical network models for location-specific consequences, and chronological simulation is necessary when timing, restoration, thermal dynamics and controls materially affect the risk. The best mitigation may target initiation, propagation or restoration, depending on which mechanism carries the risk being managed. Table 28.8 lists the chapter's terms, their meaning and units, and the subsection that defines each.

Table 28.8: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Cascading failure	uncontrolled successive loss of elements	–	–	28.1.1
Hybrid dynamical system	continuous states between discrete events	$P_\ell(t)$, $T_{\mathrm{c},\ell}(t)$, $S_\ell(t)$, $u(t)$	varies	28.1.2
Time-varying intensity	hazard given history and covariates	$\lambda(t \mid \mathcal{F}_t, \mathbf{z}(t))$	yr^{-1}	28.2.1
Weather-dependent hazard	log-linear in weather covariates	$\lambda_{\ell,0}\exp[\boldsymbol{\gamma}^{\mathsf{T}}\mathbf{z}_\ell(t)]$	yr^{-1}	28.2.2
Outage probability over an exposure	from the integrated hazard	$1-\exp[-\int \lambda(u)\,\mathrm{d}u]$	–	28.2.3
Weather mixture	joint outage under shared weather	$\sum_z \Pr(z)\prod_i \Pr(F_i \mid z)$	–	28.2.4
Thermal time constant	stored over shed heat per kelvin	$\tau_{\mathrm{th}} = mc_p/k_{\mathrm{th}}$	min	28.3.2
First-order thermal response	exponential approach to T_∞	$T_\infty + [T_{\mathrm{c}}(0)-T_\infty]\mathrm{e}^{-t/\tau_{\mathrm{th}}}$	°C	28.3.2
Time to the thermal limit	time before $T_{\max}$ is reached	$\tau_{\mathrm{th}}\ln\frac{T_\infty - T_{\mathrm{c}}(0)}{T_\infty - T_{\max}}$	min	28.3.2
Remedial action scheme	automatic predetermined corrective action	–	–	28.4.1
Competing events	first of two random times decides	$\int f_{\mathrm{A}} R_{\mathrm{P}}\,\mathrm{d}t = \frac{\lambda_{\mathrm{A}}}{\lambda_{\mathrm{A}}+\lambda_{\mathrm{P}}}$	–	28.4.2
Expected unserved energy per event	sequence-weighted energy	$\sum_k \Pr(S_k \mid I)\,E_k^{\mathrm{unserved}}$	MWh	28.4.5
Annual EUE of a family	frequency times expected energy	$f_{\mathrm{IE}}\,\mathbb{E}[E_{\mathrm{ev}} \mid I]$	MWh/yr	28.4.5
Branching ratio	mean offspring per outage	λ_{b}	–	28.5.1
Cascade size	total outages, all generations	$\mathbb{E}[N_{\mathrm{c}}] = N_0/(1-\lambda_{\mathrm{b}})$	–	28.5.2
Chronological Monte Carlo	event-driven simulation in time	$\widehat{\mathrm{EUE}} = N_{\mathrm{yr}}^{-1}\sum_j E^{(j)}$	MWh/yr	28.5.6
Importance sampling weight	corrects an oversampled scenario	$w = p(x)/q(x)$	–	28.5.6
Expected value of sample information	loss reduction from a measurement	equation (28.29)	MWh	28.6.2

Further reading. IEEE 738-2023, *Current-Temperature Relationship of Bare Overhead Conductors*, for the conductor heat balance under steady and time-varying conditions; NERC, *Glossary of Terms*, including the definition of Cascading; Dobson and Newman, *Cascading Blackout Overall Structure and Implications for Sampling and Mitigation*, for initiation and propagation of cascades and their sampling; NERC TPL-008-1, *Extreme Temperature Transmission Planning*, for planning requirements under extreme heat and cold.

28.8 Exercises

Exercise 28.1. (a) State what makes a cascading-reliability analysis different from a static $N-1$ analysis, and what determines the outage probability of a line whose failure intensity is raised for only a few days. (b) A line with baseline intensity $0.02\,\mathrm{yr}^{-1}$ is exposed to a storm that raises its intensity to $1.0\,\mathrm{yr}^{-1}$ for 72 h. Compute the cumulative hazard and the outage

probability over the storm, and the probability over the same 72 h at the baseline intensity. (c) Two lines share a weather state that is severe with probability 0.05; each line's outage probability over the assessed period is 0.01 in normal and 0.30 in severe weather, independently given the weather. Compute the marginal outage probability of each line, the joint probability under marginal independence and the correct joint probability. (Answers: (a) it models the timing and propagation of subsequent events, not only the post-contingency steady state; the time integral of the failure intensity over the period. (b) $1.0 \times 72/8760 = 0.008219$, probability $1 - \mathrm{e}^{-0.008219} = 0.008185$, against 0.000164 at baseline. (c) $0.95 \times 0.01 + 0.05 \times 0.30 = 0.0245$; $0.0245^2 = 0.000600$; $0.95 \times 0.0001 + 0.05 \times 0.09 = 0.004595$, 7.66 times larger.)

Exercise 28.2. (a) State why a conductor does not reach its new steady-state temperature immediately after an overload begins. (b) After an outage, a line at $T_\mathrm{c}(0) = 60\,^\circ\mathrm{C}$ has $T_\infty = 120\,^\circ\mathrm{C}$, $\tau_\mathrm{th} = 10\,\mathrm{min}$ and $T_\mathrm{max} = 90\,^\circ\mathrm{C}$. Compute the time to the limit and the temperature five and eight minutes after the outage. (c) State whether a corrective action completed at eight minutes arrives in time, and find the smallest time constant for which it would, with the same temperatures. (Answers: (a) the conductor's thermal capacitance mc_p makes the response transient, with time constant mc_p/k_th. (b) $t_\mathrm{lim} = 10 \ln(60/30) = 6.93\,\mathrm{min}$; $T_\mathrm{c}(5) = 120 - 60\,\mathrm{e}^{-0.5} = 83.61\,^\circ\mathrm{C}$ and $T_\mathrm{c}(8) = 120 - 60\,\mathrm{e}^{-0.8} = 93.04\,^\circ\mathrm{C}$. (c) No, the limit is passed 1.07 minutes earlier; it requires $\tau_\mathrm{th} \ln 2 \geq 8$, that is $\tau_\mathrm{th} \geq 11.54\,\mathrm{min}$.)

Exercise 28.3. (a) State what determines whether a corrective action arrests the propagation of a disturbance. (b) The corrective-action time and the protective-trip time are independent exponentials with means 5 min and 20 min. Compute the probability of arrest and the mean time to the first of the two events. (c) In the cascade tree of Section 28.4.4, the prolonged restoration is lengthened to 24 h, all else at the base values. Compute $\mathbb{E}[E_\mathrm{ev} \mid I]$, the annual EUE, and the share of the EUE due to prolonged restoration. (Answers: (a) whether it acts before the consequential protective trip or failure event, which requires the distributions of the competing times, not only their means. (b) $0.2/(0.2 + 0.05) = 0.80$; mean $1/0.25 = 4\,\mathrm{min}$. (c) $3.2 + 6.0 + 0.06 \times 2400 = 153.2\,\mathrm{MWh}$; $0.20 \times 153.2 = 30.64\,\mathrm{MWh/yr}$; $144/153.2 = 94.0\,\%$.)

Exercise 28.4. (a) State what a branching ratio below one implies for a Galton–Watson cascade, and why ordinary Monte Carlo is often inefficient for estimating very severe cascades. (b) Compute the expected cascade size and the expected number of subsequent outages for one initiating outage and $\lambda_\mathrm{b} = 0.6$. (c) A cascade class has a frequency of $0.024\,\mathrm{yr}^{-1}$. Compute the number of simulated years that gives a relative standard error of 10 % on its estimated frequency, and the number for a class of frequency $2 \times 10^{-4}\,\mathrm{yr}^{-1}$. (Answers: (a) extinction with probability one and a finite expected total size, $N_0/(1 - \lambda_\mathrm{b})$; the severe outcomes are too rare to be observed often, so their estimates have large relative error. (b) $1/(1 - 0.6) = 2.5$, of which 1.5 subsequent outages. (c) $N_\mathrm{yr} = 1/(0.01 f) = 4167$ years; 500 000 years.)

Part VIII

State Estimation, Observability, and Model Credibility

Chapter 29

Observability and Measurement Sufficiency

29.1 States, measurements and observability

Chapters 25 to 28 computed the reliability of a power system from models of generation availability, transmission outages, network constraints and cascading sequences. Those calculations took the condition of the system as known: power flows were computed from specified injections and a specified topology, and line ratings and outage probabilities were assigned. An operator does not see the condition of the system directly. What is available are measurements of power flows, voltages, currents, temperatures and equipment status, some of them missing, noisy, delayed or wrong. Before a reliability index is computed or an abnormal condition is diagnosed, a more basic question must be answered,

whether a given set of measurements determines the physical state of the system, and how confidently it distinguishes the operating and failure conditions that matter.

That is the problem of observability. It is closely related to state estimation, sensor placement, fault diagnosis and parameter identifiability, and the chapter keeps it distinct from each of them. It defines the state and the measurement model and derives the rank condition for observability; works the condition on the three-bus transmission network of Chapter 27 with three sets of line-flow measurements; separates observability from the precision of the estimate, from measurement redundancy and from bad-data detection; treats sensor placement and the difference between observing states and diagnosing sensor faults; extends the condition to nonlinear and dynamic models and to unknown parameters; and closes by asking what a measurement set must provide for a particular reliability decision.

Jargon: state, measurement model, observability, unobservable direction, state estimation, measurement redundancy, critical measurement, bad-data detection, identifiability

The *state* of a system is the smallest set of variables that, with the model and its known parameters, determines every quantity of interest at one instant; collected in a column it is the *state vector* $\mathbf{x}$, of dimension n. A *measurement model* $\mathbf{y} = h(\mathbf{x}) + \boldsymbol{\varepsilon}$ states how each of the m measured quantities depends on the state, with an additive error. A system is *observable* from a measurement set when distinct states produce distinct noise-free measurements, so that the state can in principle be recovered from them. An *unobservable direction* is a change of state that leaves every modeled measurement unchanged. *State estimation* computes the most plausible state from noisy measurements, usually by weighted least squares. *Measurement redundancy* is the excess of measurements over what observability requires, counted by the degrees of freedom $m - n$; a *critical measurement* is one whose removal makes an observable system unobservable. *Bad-data detection* is the test of whether the redundant measurements are consistent with the model and with one another, the power-system name for the residual tests of fault detection. *Identifiability* is the corresponding property of the model's parameters: whether distinct parameter values produce distinct measurements.

Key idea

A system is observable when no change of state is invisible to the measurements; for a linear model that is full column rank of the measurement Jacobian. Observability is necessary for state estimation but not sufficient for an engineering conclusion: the precision of the estimate, the redundancy that allows bad data to be detected, the separation of sensor faults and parameters, and the sufficiency of the evidence for the decision at hand are further properties of the same measurement model.

This section defines the state of a power network, writes the measurement equation, shows with two bus angles why a measurement can determine a combination of states but not the states themselves, and derives the rank condition for linear observability.

29.1.1 The physical state of a power network

In power-system analysis the state vector conventionally contains the bus voltage magnitudes V_i, in per unit, and the bus voltage angles θ_i, in radians. Only differences of angles have physical effect: every line flow depends on $\theta_i - \theta_j$, so adding the same constant to all the angles changes no flow, no injection and no measurement. One bus is therefore chosen as the *reference bus*, its angle fixed at zero, and the remaining angles are measured from it. For an N-bus AC network with bus 1 as reference, the state vector is

$$\mathbf{x} = \begin{bmatrix} \theta_2 & \cdots & \theta_N & V_1 & \cdots & V_N \end{bmatrix}^{\mathsf{T}}, \qquad n = 2N - 1, \tag{29.1}$$

when all the voltage magnitudes and all the nonreference angles are estimated. In the DC power-flow approximation of Chapter 27, the magnitudes are fixed at 1 per unit and the losses neglected, so the unknown state reduces to the $N - 1$ nonreference angles; with bus N as reference,

$$\mathbf{x} = \begin{bmatrix} \theta_1 & \cdots & \theta_{N-1} \end{bmatrix}^{\mathsf{T}}, \qquad n = N - 1. \tag{29.2}$$

The choice of reference bus is a convention; it changes the numerical values of the angles but not any flow.

The state is a property of the physical system at one instant. The network equations that predict flows and injections from it are the model, and the meter readings, with their errors, are

the evidence. Observability is a statement about the relation between the second and the third: whether the model, applied to the measurements available, leaves only one state consistent with them.

29.1.2 The measurement equation

Let the m measured quantities be collected in the measurement vector $\mathbf{y} = [y_1, \dots, y_m]^{\mathsf{T}}$. A general measurement model is

$$\mathbf{y} = h(\mathbf{x}, \boldsymbol{\theta}) + \boldsymbol{\varepsilon}, \tag{29.3}$$

where $\mathbf{x}$ is the unknown physical state, $\boldsymbol{\theta}$ the vector of model parameters, such as line reactances or sensor gains, h the physical relation between state and measured quantities, and $\boldsymbol{\varepsilon}$ the measurement error, taken to have mean zero and covariance matrix $\mathbf{R} = \mathbb{E}[\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^{\mathsf{T}}]$. The power-system literature writes the measurement vector $\mathbf{z}$; this book keeps $\mathbf{y}$, its symbol for observations in every chapter. Until Section 29.5.5 the parameters are treated as known and the model is written $\mathbf{y} = h(\mathbf{x}) + \boldsymbol{\varepsilon}$.

Remark 29.1. Three symbols of this chapter must not be confused with symbols of earlier chapters. The bold $\mathbf{R}$ is the covariance matrix of the measurement errors, not the reliability function $R(t)$. The bold $\mathbf{H}$ introduced below is the measurement Jacobian, not the operating horizon H. And the bold $\boldsymbol{\theta}$ is the vector of model parameters, whereas θ_i, with a bus index, is the voltage angle of bus i.

The observability problem asks whether two distinct states can produce measurements that the model cannot tell apart. It is posed for the noise-free model: if two states give the same $h(\mathbf{x})$, no precision of the meters will separate them, and if they give different values, sufficiently precise meters will.

29.1.3 A combination of states against the states themselves

Let the state be two bus angles, $\mathbf{x} = [\theta_1, \theta_2]^{\mathsf{T}}$, and let the only measurement be their difference,

$$y = \theta_1 - \theta_2.$$

If $y = 0.05\,\text{rad}$, the states $(\theta_1, \theta_2) = (0.10, 0.05)$, $(0.15, 0.10)$ and $(0.20, 0.15)$ are all consistent with it, and so is every state on the line $\theta_1 - \theta_2 = 0.05$ in the plane of the two angles. The measurement determines one combination of the states, the difference, and leaves the other, the common level of the two angles, completely free.

Additional information is needed to fix that level: a measurement of one angle, a reference that fixes one of them by convention, or a second measurement that depends on the angles in a different combination. The example contains the whole distinction of this chapter in miniature. Observing a combination of states is not observing the individual states, and whether a measurement set suffices depends on which of them is required.

29.1.4 Linear observability and the rank condition

Suppose the measurement model is linear,

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}, \qquad \mathbf{H} \in \mathbb{R}^{m \times n}, \tag{29.4}$$

where $\mathbf{H}$ is the measurement matrix: its row i, written $\mathbf{h}_i$, states how measurement i responds to each state variable. For a nonlinear model, $\mathbf{H}$ is the Jacobian $\partial h / \partial \mathbf{x}$ of Section 29.5.1; in either case it is called the measurement Jacobian.

In the noise-free case $\mathbf{y} = \mathbf{H}\mathbf{x}$. Two states $\mathbf{x}_1$ and $\mathbf{x}_2$ give the same measurements exactly when $\mathbf{H}(\mathbf{x}_1 - \mathbf{x}_2) = \mathbf{0}$, that is, when their difference lies in the *null space* of $\mathbf{H}$, the set of vectors

$\mathbf{v}$ with $\mathbf{H}\mathbf{v} = \mathbf{0}$. The state is therefore uniquely recoverable if and only if the null space contains only the zero vector. The rank–nullity theorem of linear algebra states that the rank of $\mathbf{H}$ (the number of its linearly independent columns, equal to the number of its linearly independent rows) plus the dimension of its null space equals the number of columns n. The null space is trivial when the rank is n, so the condition for linear observability is

$$\operatorname{rank}(\mathbf{H}) = n, \tag{29.5}$$

full column rank. It requires at least n measurements, $m \geq n$, but the count is not enough: the rows must also be independent.

If $\operatorname{rank}(\mathbf{H}) < n$, the null space contains a nonzero vector $\mathbf{v}$ with

$$\mathbf{H}\mathbf{v} = \mathbf{0}, \tag{29.6}$$

and then, for every scalar α,

$$\mathbf{H}(\mathbf{x} + \alpha\mathbf{v}) = \mathbf{H}\mathbf{x} + \alpha\,\mathbf{H}\mathbf{v} = \mathbf{H}\mathbf{x}.$$

The measurements cannot distinguish states that differ along $\mathbf{v}$. Such a $\mathbf{v}$ is an unobservable direction: a change of the physical state that produces no change in the modeled measurements. In the example of Section 29.1.3, $\mathbf{H} = [1, -1]$, of rank one, and $\mathbf{v} = [1, 1]^{\mathsf{T}}$ is the unobservable direction, the common level of the two angles.

When the rank is n, the state can be computed from the measurements. The matrix $\mathbf{H}^{\mathsf{T}}\mathbf{H}$ is then invertible, because for any $\mathbf{v} \neq \mathbf{0}$, $\mathbf{v}^{\mathsf{T}}\mathbf{H}^{\mathsf{T}}\mathbf{H}\mathbf{v} = \|\mathbf{H}\mathbf{v}\|^2 > 0$; multiplying $\mathbf{y} = \mathbf{H}\mathbf{x}$ by $\mathbf{H}^{\mathsf{T}}$ and solving gives $\mathbf{x} = (\mathbf{H}^{\mathsf{T}}\mathbf{H})^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{y}$, which reduces to $\mathbf{x} = \mathbf{H}^{-1}\mathbf{y}$ when $\mathbf{H}$ is square. The same invertibility, with weights, is what the estimator of Section 29.3.1 requires.

Principle 29.1. *A linear measurement model observes the state when its Jacobian has full column rank. A measurement set that leaves a nonzero null-space direction cannot detect any change of state along that direction, however precise its meters.*

29.2 Observability of the three-bus network

The rank condition is now applied to the transmission network that Chapter 27 used for contingency analysis. This section restates the network and its DC model, and computes the measurement Jacobian, its rank and its null space for three sets of line-flow measurements.

29.2.1 The network and its DC model

The network has three buses joined by three lines in a triangle. Figure 29.1 draws it. Bus 1, at the top, holds the generation; bus 2, at the lower left, is a junction with no generation or load; bus 3, at the lower right, holds the load and is the angle reference. Each line is labeled with its name, its reactance x_{ij} in per unit on a 100 MVA base and its thermal rating $P_{ij}^{\max}$; the arrows inside the triangle are the base-case flows, and the small squares are the three candidate flow meters, one per line, whose readings are y_{12}, y_{23} and y_{13}.

The DC power-flow model of Chapter 27 is restated here with the equation it relies on. The active power flowing from bus i to bus j is

$$P_{ij} = S_{\text{base}}\,\frac{\theta_i - \theta_j}{x_{ij}}, \tag{29.7}$$

in MW, with angles in radians, x_{ij} in per unit and the system base $S_{\text{base}} = 100\,\text{MVA}$; in per unit the factor S_{base} drops out and $P_{ij} = (\theta_i - \theta_j)/x_{ij} = b_{ij}(\theta_i - \theta_j)$, with the line susceptance

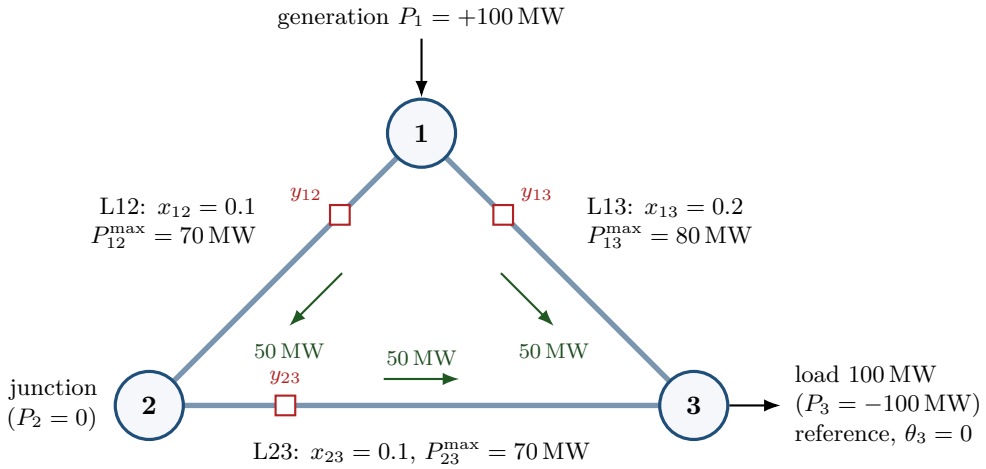


Figure 29.1: The three-bus transmission network of Chapter 27: reactances in per unit on a 100 MVA base, thermal ratings, base-case flows (green arrows, from the DC solution $\theta_1 = 0.10$, $\theta_2 = 0.05$ rad) and the three candidate line-flow meters (red squares). Bus 3 is the angle reference.

$b_{ij} = 1/x_{ij}$. The approximation assumes voltage magnitudes near 1 per unit, small angle differences and resistance small against reactance, and neglects losses. Collecting the injections, $\mathbf{P} = \mathbf{B}\boldsymbol{\theta}$ with the DC susceptance matrix

$$\mathbf{B} = \begin{bmatrix} 15 & -10 & -5 \\ -10 & 20 & -10 \\ -5 & -10 & 15 \end{bmatrix},$$

whose off-diagonal entries are $-b_{ij}$ and whose diagonal entries are the sums of the susceptances at each bus ($b_{12} = b_{23} = 10$, $b_{13} = 5$). With $P_1 = +1$, $P_2 = 0$ and $P_3 = -1$ per unit and $\theta_3 = 0$, the first two rows give $15\theta_1 - 10\theta_2 = 1$ and $-10\theta_1 + 20\theta_2 = 0$, so $\theta_2 = \theta_1/2$, $10\theta_1 = 1$, and

$$\theta_1 = 0.10\,\text{rad}, \qquad \theta_2 = 0.05\,\text{rad}.$$

With $\theta_3 = 0$ the state of the DC model is $\mathbf{x} = [\theta_1, \theta_2]^{\mathsf{T}}$, $n = 2$, and equation (29.7) gives the three line flows as linear functions of it:

$$P_{12} = 1000\,(\theta_1 - \theta_2), \qquad P_{23} = 1000\,\theta_2, \qquad P_{13} = 500\,\theta_1, \tag{29.8}$$

in MW with angles in radians, the coefficients being $100/0.1 = 1000$ MW/rad and $100/0.2 = 500$ MW/rad. At the base state all three flows are 50 MW: the direct path L13 and the indirect path L12–L23 have the same reactance, 0.2 per unit, so the 100 MW transfer divides equally. Table 29.1 collects the line data and the measurement row $\mathbf{h}_i$ that each flow meter contributes to $\mathbf{H}$.

29.2.2 Measurement set A: one line-flow meter

Suppose only the flow on L12 is measured, $y_{12} = P_{12} + \varepsilon$. By equation (29.8),

$$y_{12} = \begin{bmatrix} 1000 & -1000 \end{bmatrix} \mathbf{x} + \varepsilon, \qquad \mathbf{H}_A = \begin{bmatrix} 1000 & -1000 \end{bmatrix}.$$

A single nonzero row has rank one, $\operatorname{rank}(\mathbf{H}_A) = 1 < n = 2$, and the state is not observable from this measurement set. The null space is spanned by

$$\mathbf{v} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \qquad \mathbf{H}_A\mathbf{v} = 1000 - 1000 = 0.$$

Table 29.1: Lines of the three-bus network, their base-case flows and the measurement row of a flow meter on each, $\mathbf{h}_i = \partial P_{ij}/\partial[\theta_1, \theta_2]$, in MW/rad.

Line	x_{ij} (pu)	b_{ij} (pu)	$P_{ij}^{\max}$ (MW)	Base flow (MW)	Row $\mathbf{h}_i$
L12	0.1	10	70	50	$[1000,\ -1000]$
L23	0.1	10	70	50	$[0,\ 1000]$
L13	0.2	5	80	50	$[500,\ 0]$

The L12 meter cannot see a change that raises both angles by the same amount. The states $(0.10, 0.05)$ and $(0.15, 0.10)$ both give $P_{12} = 50\,\text{MW}$, although the second gives $P_{23} = 100\,\text{MW}$ and $P_{13} = 75\,\text{MW}$ instead of 50 MW each. The ambiguity has a physical meaning: the meter on L12 measures the angle difference between buses 1 and 2, and the common level of θ_1 and θ_2 relative to the reference bus, which sets the flows on the other two lines, does not enter it.

Figure 29.2(a) draws the situation in the plane of the two angles, in milliradians. The horizontal axis is θ_1 and the vertical axis θ_2; the solid line is the set of states consistent with $y_{12} = 50\,\text{MW}$, that is $\theta_1 - \theta_2 = 50\,\text{mrad}$; the dots are the three states of Section 29.1.3; the arrow is the null-space direction $\mathbf{v}$. Every point on the solid line explains the measurement equally well, and the arrow points along it. The dotted horizontal line, $\theta_2 = 50\,\text{mrad}$, is what a meter on L23 adds, the subject of the next subsection.

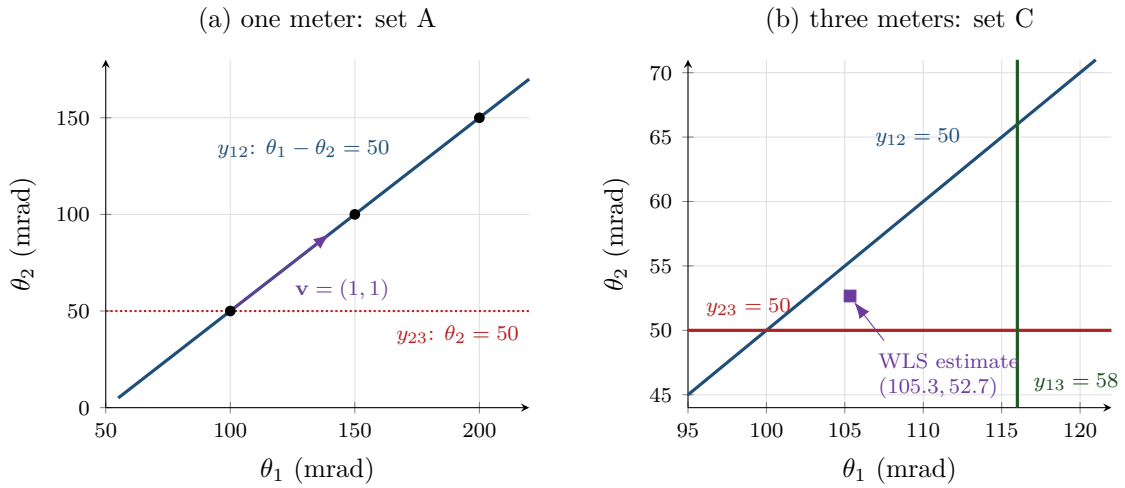


Figure 29.2: Measurements as lines in the plane of the two state angles (mrad), noise-free DC model of equation (29.8). (a) The L12 meter reading 50 MW fixes only $\theta_1 - \theta_2$; the states $(100, 50)$, $(150, 100)$ and $(200, 150)$ explain it equally, and $\mathbf{v} = (1, 1)$ is the unobservable direction. A meter on L23 (dotted) intersects it at one point. (b) Three meters reading 50, 50 and 58 MW define three lines that do not meet in one point; the weighted least-squares estimate of Section 29.3.5 lies inside the triangle they form.

Remark 29.2. The rank count above uses the flow meters only. The bus without generation or load, bus 2, supplies information of its own: its injection is known to be zero, $P_2 = P_{23} - P_{12} = 0$, a *zero-injection* (virtual) measurement in the vocabulary of power-system state estimation. Its row is $[-1000,\ 2000]$, and together with the L12 row the determinant is $1000 \times 2000 - (-1000)(-1000) = 10^6 \neq 0$, so the L12 meter plus the knowledge that bus 2 is a junction would make the state observable. The examples and the sufficiency map of this chapter use the line-flow meters alone, without injection information, so that the effect of each meter is seen in isolation.

29.2.3 Measurement set B: two line-flow meters

Now let the flows on L12 and L23 both be measured. The measurement matrix stacks the two rows of Table 29.1,

$$\mathbf{H}_B = \begin{bmatrix} 1000 & -1000 \\ 0 & 1000 \end{bmatrix}, \qquad \det(\mathbf{H}_B) = 1000 \times 1000 - (-1000) \times 0 = 10^6,$$

in $(\mathrm{MW/rad})^2$. A square matrix with a nonzero determinant is invertible, so $\operatorname{rank}(\mathbf{H}_B) = 2 = n$ and the state is observable. In the noise-free case the two equations are solved in order. The L23 meter gives θ_2 directly, and the L12 meter then gives θ_1:

$$\theta_2 = \frac{y_{23}}{1000}, \qquad \theta_1 = \frac{y_{12}}{1000} + \theta_2 = \frac{y_{12} + y_{23}}{1000}.$$

With $y_{12} = y_{23} = 50\,\mathrm{MW}$, $\theta_2 = 0.05\,\mathrm{rad}$ and $\theta_1 = 0.10\,\mathrm{rad}$, the base state, and with it the flow on the unmetered line, $P_{13} = 500 \times 0.10 = 50\,\mathrm{MW}$. In Figure 29.2(a) the L23 meter adds the dotted line $\theta_2 = 50\,\mathrm{mrad}$, which crosses the L12 line at the single point $(100, 50)$. Its row $[0, 1000]$ is not a multiple of the L12 row, so it constrains the direction $\mathbf{v}$ that the first meter left free: $\mathbf{h}_{23}\mathbf{v} = 1000 \neq 0$.

29.2.4 Measurement set C: three line-flow meters

With all three flows measured, $\mathbf{y} = [y_{12}, y_{23}, y_{13}]^{\mathsf{T}}$ and

$$\mathbf{H}_C = \begin{bmatrix} 1000 & -1000 \\ 0 & 1000 \\ 500 & 0 \end{bmatrix}.$$

Its rank cannot exceed its number of columns, and two of its rows are already independent, so $\operatorname{rank}(\mathbf{H}_C) = 2$. Adding the L13 meter does not increase the number of observable state dimensions, because the state was already observable; it adds one measurement more than observability needs, $m - n = 3 - 2 = 1$ degree of freedom of *measurement redundancy*. The third row is a combination of the other two: $\mathbf{h}_{13} = \frac{1}{2}(\mathbf{h}_{12} + \mathbf{h}_{23}) = \frac{1}{2}[1000, 0]$, which is the network fact $P_{13} = (P_{12} + P_{23})/2$ that Section 29.3.5 uses.

Redundancy is a different property from observability. A third meter can reduce the uncertainty of the estimated state, allow inconsistent measurements to be detected through residuals, and keep the state observable when one meter is lost, all without adding a state dimension. Two measurements can be sufficient to estimate two states; the third makes the estimate more precise and checkable. The next section quantifies each of these effects.

Common mistake

Equating observability with measurement redundancy, or either with the number of meters. Observability asks whether the measurements determine the state at all and is decided by the rank of $\mathbf{H}$; redundancy counts the measurements beyond that, $m - n$, and is what makes errors detectable. Set B is observable without redundancy; two meters on the same line give $m = n = 2$ but rank one, so the state is not observable although the two readings can be checked against each other.

29.3 Estimation uncertainty, measurement redundancy and bad data

Observability says that the state can be recovered from perfect measurements. This section derives what happens with noisy ones: the weighted least-squares estimate and its covariance,

the information matrix and the geometry of the confidence region, the difference between structural and numerical observability, and the use of redundant measurements to detect bad data and to find the measurements whose loss would make the state unobservable.

29.3.1 Full rank is necessary but not sufficient

With noise, $\mathbf{y} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}$ with $\mathbb{E}[\boldsymbol{\varepsilon}] = \mathbf{0}$ and $\mathrm{Cov}(\boldsymbol{\varepsilon}) = \mathbf{R}$, the exact state cannot be recovered even when $\mathrm{rank}(\mathbf{H}) = n$; it is estimated, with an uncertainty that depends on the meters. Let $\mathbf{W} = \mathbf{R}^{-1}$ be the *weight matrix*, which gives each measurement a weight inversely proportional to its error variance. The *weighted least-squares* (WLS) estimate minimizes the objective

$$J(\mathbf{x}) = (\mathbf{y} - \mathbf{H}\mathbf{x})^{\mathsf{T}}\mathbf{W}(\mathbf{y} - \mathbf{H}\mathbf{x}), \tag{29.9}$$

the sum of squared measurement discrepancies, each in units of its own standard deviation when $\mathbf{R}$ is diagonal. Its gradient is $\partial J/\partial \mathbf{x} = -2\mathbf{H}^{\mathsf{T}}\mathbf{W}(\mathbf{y} - \mathbf{H}\mathbf{x})$. Setting it to zero gives the normal equations $\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\,\hat{\mathbf{x}} = \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}$. Define the *gain matrix*

$$\mathbf{G} = \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}, \tag{29.10}$$

of dimension $n \times n$. It is invertible exactly when $\mathbf{H}$ has full column rank, since $\mathbf{v}^{\mathsf{T}}\mathbf{G}\mathbf{v} = \|\mathbf{W}^{1/2}\mathbf{H}\mathbf{v}\|^2$ vanishes only for $\mathbf{v}$ in the null space of $\mathbf{H}$. Then

$$\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} = (\mathbf{H}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{H})^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{y}. \tag{29.11}$$

Observability is thus exactly the condition for the WLS estimate to exist.

The covariance of the estimate follows by substitution. Since $\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}(\mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}) = \mathbf{x} + \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\boldsymbol{\varepsilon}$, the estimate is unbiased and its error is $\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\boldsymbol{\varepsilon}$, so

$$\mathrm{Cov}(\hat{\mathbf{x}}) = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\,\mathbf{R}\,\mathbf{W}\mathbf{H}\mathbf{G}^{-1} = \mathbf{G}^{-1}(\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H})\mathbf{G}^{-1},$$

using $\mathbf{W}\mathbf{R}\mathbf{W} = \mathbf{W}$, and therefore

$$\mathrm{Cov}(\hat{\mathbf{x}}) = \mathbf{G}^{-1} = (\mathbf{H}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{H})^{-1}. \tag{29.12}$$

This expression connects which meters exist (the rows of $\mathbf{H}$), what they measure (the entries of each row), how noisy they are ($\mathbf{R}$) and how precisely the physical state can be inferred. For Gaussian errors the WLS estimate is also the maximum likelihood estimate: the log-likelihood is $\ell(\mathbf{x}) = -\tfrac{1}{2}J(\mathbf{x})$ plus a constant, whose matrix of negative second derivatives is $\mathbf{G}$. Section 7.5.2 of Chapter 7 of Volume I defined the Fisher information of a parameter as the expected negative second derivative of the log-likelihood, $\mathcal{I}(\lambda) = \mathbb{E}[-\mathrm{d}^2\ell/\mathrm{d}\lambda^2]$, and the standard error of its estimate as approximately $1/\sqrt{\mathcal{I}}$. The gain matrix is the matrix form of that quantity for the state, $\mathcal{I}(\mathbf{x}) = \mathbf{G}$, and for a linear model with Gaussian errors the relation $\mathrm{Cov}(\hat{\mathbf{x}}) = \mathcal{I}(\mathbf{x})^{-1}$ is exact rather than a large-sample approximation.

29.3.2 Two against three measurements

Let every flow meter have an independent Gaussian error with standard deviation $\sigma_y = 2\,\mathrm{MW}$, so that $\mathbf{R} = 4\,\mathbf{I}$ in MW^2 and $\mathbf{W} = 0.25\,\mathbf{I}$ in MW^{-2}. With the meters of set B the estimate is the noise-free solution applied to the readings, $\hat{\theta}_2 = y_{23}/1000$ and $\hat{\theta}_1 = (y_{12} + y_{23})/1000$. Its variances follow from the variance of a sum of independent errors,

$$\mathrm{Var}(\hat{\theta}_2) = \frac{4}{10^6} = 4 \times 10^{-6}\,\mathrm{rad}^2, \qquad \mathrm{Var}(\hat{\theta}_1) = \frac{4+4}{10^6} = 8 \times 10^{-6}\,\mathrm{rad}^2,$$

and $\mathrm{Cov}(\hat{\theta}_1, \hat{\theta}_2) = 4 \times 10^{-6}\,\mathrm{rad}^2$, the variance of the shared reading y_{23}. The standard deviations are $\sigma_{\theta_1} = 2.828 \times 10^{-3}\,\mathrm{rad}$ and $\sigma_{\theta_2} = 2.000 \times 10^{-3}\,\mathrm{rad}$. The same result follows from equation (29.12): $\mathbf{G}_B = \mathbf{H}_B^{\mathsf{T}}\mathbf{H}_B/4$ has entries 250 000, $-250\,000$ and $500\,000\,\mathrm{rad}^{-2}$, its determinant is $6.25 \times 10^{10}\,\mathrm{rad}^{-4}$, and its inverse is $[8, 4; 4, 4] \times 10^{-6}\,\mathrm{rad}^2$.

Adding the L13 meter and using the WLS estimate, the gain matrix becomes

$$\mathbf{G}_C = \mathbf{H}_C^{\mathsf{T}}\mathbf{W}\mathbf{H}_C = \frac{1}{4}\begin{bmatrix} 1000^2 + 500^2 & -1000^2 \\ -1000^2 & 1000^2 + 1000^2 \end{bmatrix} = \begin{bmatrix} 312\,500 & -250\,000 \\ -250\,000 & 500\,000 \end{bmatrix}\mathrm{rad}^{-2},$$

which is $\mathbf{G}_B$ plus the contribution $\mathbf{h}_{13}^{\mathsf{T}}\mathbf{h}_{13}/4$ of the new meter, $62\,500\,\mathrm{rad}^{-2}$ in the $\theta_1\theta_1$ entry only, since L13 responds to θ_1 alone. Its determinant is $312\,500 \times 500\,000 - 250\,000^2 = 9.375 \times 10^{10}\,\mathrm{rad}^{-4}$, and the inverse of a 2×2 matrix, the swapped diagonal and negated off-diagonal divided by the determinant, gives

$$\mathrm{Cov}(\hat{\mathbf{x}}) = \mathbf{G}_C^{-1} = \frac{1}{9.375 \times 10^{10}}\begin{bmatrix} 500\,000 & 250\,000 \\ 250\,000 & 312\,500 \end{bmatrix} = \begin{bmatrix} 5.333 & 2.667 \\ 2.667 & 3.333 \end{bmatrix} \times 10^{-6}\,\mathrm{rad}^2,$$

so that $\sigma_{\theta_1} = 2.309 \times 10^{-3}\,\mathrm{rad}$ and $\sigma_{\theta_2} = 1.826 \times 10^{-3}\,\mathrm{rad}$. Table 29.2 sets the two cases side by side.

Table 29.2: Standard deviations and variances of the estimated angles with two meters (L12, L23) and three meters (L12, L23, L13), each with $\sigma_y = 2\,\mathrm{MW}$.

State	Two meters (mrad)	Three meters (mrad)	Change of std. deviation	Change of variance
σ_{θ_1}	2.828	2.309	$-18.4\,\%$	$-33.3\,\%$
σ_{θ_2}	2.000	1.826	$-8.7\,\%$	$-16.7\,\%$

The third meter improves the precision of both angles without changing the rank: the variance of $\hat{\theta}_1$ falls by a third, from 8 to $5.333 \times 10^{-6}\,\mathrm{rad}^2$, and that of $\hat{\theta}_2$ by a sixth. The gain is larger for θ_1, the angle the new meter measures; θ_2 gains only through its correlation with θ_1. Figure 29.3 draws the two columns of standard deviations as bars, in milliradians: for each angle, the left bar is the two-meter value and the right bar the three-meter value.

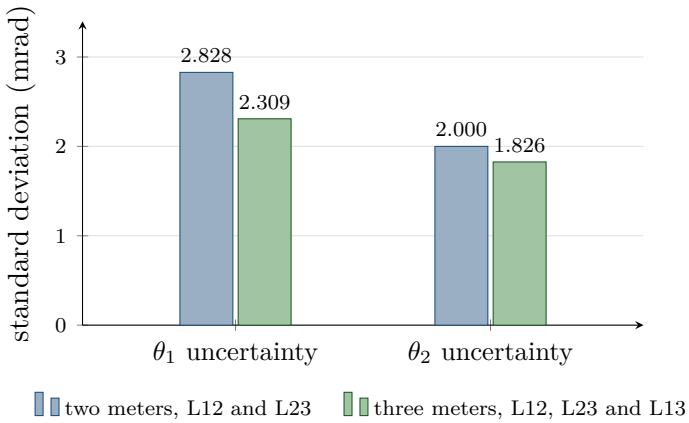


Figure 29.3: Standard deviation of the estimated bus angles with two and three line-flow meters, each with $\sigma_y = 2\,\mathrm{MW}$, from $\mathrm{Cov}(\hat{\mathbf{x}}) = (\mathbf{H}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{H})^{-1}$. The redundant third meter improves both estimates without changing the rank of $\mathbf{H}$.

29.3.3 The information matrix and the geometry of estimation

The gain matrix $\mathbf{G}$ is also called the *information matrix* of the measurement set: it is the Fisher information of the state, and its inverse is the covariance of the estimate when the assumptions of the estimator hold. Its structure is read from its eigenvalues. Write $\tilde{\mathbf{H}} = \mathbf{W}^{1/2}\mathbf{H}$ for the Jacobian with each row divided by the standard deviation of its meter, so that $\mathbf{G} = \tilde{\mathbf{H}}^{\mathsf{T}}\tilde{\mathbf{H}}$; the eigenvalues of $\mathbf{G}$ are then the squares $\sigma_1^2 \geq \sigma_2^2 \geq \cdots \geq \sigma_n^2$ of the singular values of $\tilde{\mathbf{H}}$, and its eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ are orthogonal directions in the state space. (The singular values carry a numerical subscript; σ_y and σ_{θ_1} remain standard deviations.) A state change $\delta\mathbf{x}$ along $\mathbf{v}_i$ changes the weighted measurements by $\|\tilde{\mathbf{H}}\delta\mathbf{x}\| = \sigma_i\|\delta\mathbf{x}\|$, and the variance of the estimate along that direction is $1/\sigma_i^2$. A large eigenvalue means the measurements constrain that direction strongly; a small one means weak information; a zero eigenvalue means the direction is unobservable. The ratio $\kappa = \sigma_1/\sigma_n$ is the *condition number* of the weighted Jacobian: the ratio of the best-constrained to the worst-constrained direction, 1 when every direction is equally well measured and infinite when one is not measured at all.

For set C, the eigenvalues of $\mathbf{G}_C$ solve $\sigma^4 - 812\,500\,\sigma^2 + 9.375 \times 10^{10} = 0$, the characteristic equation of a 2×2 matrix with that trace and determinant, which gives $\sigma_1^2 = 673\,250$ and $\sigma_2^2 = 139\,250\,\mathrm{rad}^{-2}$. The weak direction is $\mathbf{v}_2 = (0.822, 0.570)$, in which both angles rise together; along it the standard deviation of the estimate is $1/\sqrt{139\,250} = 2.680\,\mathrm{mrad}$. The strong direction, $\mathbf{v}_1 = (0.570, -0.822)$, is close to the angle difference that the L12 meter measures, with a standard deviation of 1.219 mrad. For set B the eigenvalues are 654 508 and $95\,492\,\mathrm{rad}^{-2}$, the weak direction $(0.851, 0.526)$ and its standard deviation 3.236 mrad. The weak direction of both sets lies close to $\mathbf{v} = (1, 1)$, the direction that set A could not see at all; the L13 meter raises the smallest eigenvalue by 46 %, from 95 492 to $139\,250\,\mathrm{rad}^{-2}$. The condition numbers are $\kappa = \sqrt{654\,508/95\,492} = 2.618$ for set B and $\sqrt{673\,250/139\,250} = 2.199$ for set C.

For Gaussian errors the estimate's confidence region is an ellipse centered on $\hat{\mathbf{x}}$, the set $(\mathbf{x} - \hat{\mathbf{x}})^{\mathsf{T}}\mathbf{G}(\mathbf{x} - \hat{\mathbf{x}}) \leq \chi^2_{2,0.95} = 5.991$ for 95 %, since the quadratic form is chi-square with two degrees of freedom by the whitening argument of Section 22.2.4 of Chapter 22. Its axes lie along the eigenvectors and its semi-axes are $\sqrt{5.991}/\sigma_i$. Figure 29.4 draws three such ellipses. The horizontal axis is the error of $\hat{\theta}_1$ and the vertical axis the error of $\hat{\theta}_2$, both in milliradians; the outer dashed ellipse belongs to the two meters on L12 and L13, the middle one to set B, and the inner one to set C. The ellipse of L12 and L13 is the longest and narrowest, with semi-axes 14.30 and 3.35 mrad and $\kappa = 4.27$: that pair constrains the angle difference well and the common level poorly. Set B has semi-axes 7.92 and 3.03 mrad; the redundant meter of set C shrinks the long axis to 6.56 mrad and leaves the short one almost unchanged, at 2.98 mrad, because it adds information mainly along the weak direction.

29.3.4 Structural and numerical observability

The ellipses motivate a distinction. *Structural observability* asks whether the model and the arrangement of the meters permit the state to be determined in principle; it is the rank condition, which depends only on which entries of $\mathbf{H}$ are nonzero and how the rows combine. *Numerical observability* asks whether the measurements provide enough independent information, with acceptable conditioning and uncertainty, for the estimate to be useful; it depends on the values in $\mathbf{H}$ and $\mathbf{R}$ and is judged by the smallest eigenvalue of $\mathbf{G}$, the condition number and the resulting covariance. A system can have full rank and still be almost unobservable in practice.

Let

$$\mathbf{H} = \begin{bmatrix} 1 & 1 \\ 1 & 1.001 \end{bmatrix},$$

with two states and two measurements in arbitrary consistent units. Its determinant is $1.001 - 1 = 0.001$, so it has rank two and the state is structurally observable. But the rows are

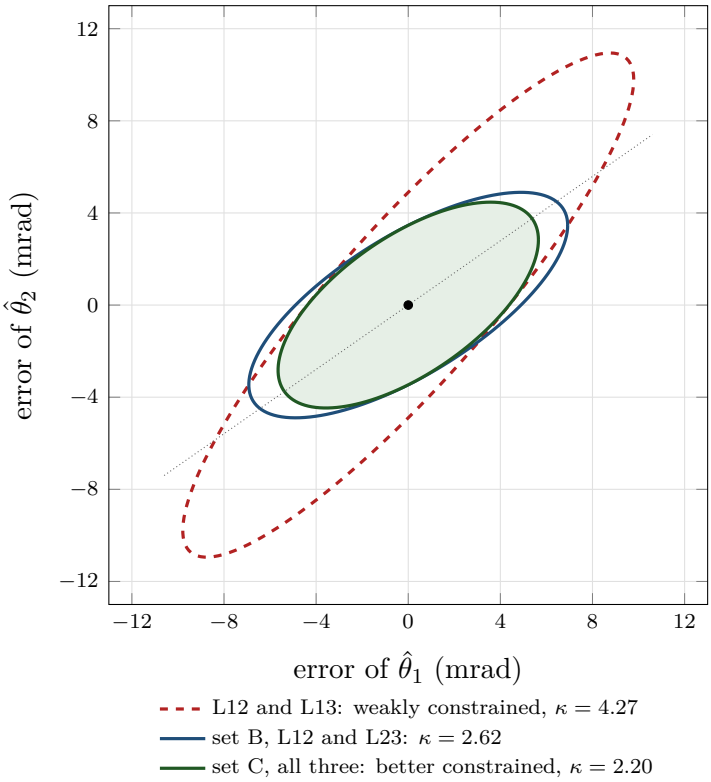


Figure 29.4: Ninety-five per cent confidence ellipses of the estimated bus angles, $(\mathbf{x}-\hat{\mathbf{x}})^{\mathsf{T}}\mathbf{G}(\mathbf{x}-\hat{\mathbf{x}}) \le 5.991$, for three meter sets with $\sigma_y = 2\,\mathrm{MW}$ each. The ellipses are drawn from the Cholesky factor $\mathbf{L}$ of $\mathrm{Cov}(\hat{\mathbf{x}})$ as $\sqrt{5.991}\,\mathbf{L}\,[\cos t, \sin t]^{\mathsf{T}}$; $\mathrm{Cov}(\hat{\mathbf{x}})$ is $[16, 16; 16, 20]$, $[8, 4; 4, 4]$ and $[5.33, 2.67; 2.67, 3.33]\,\mathrm{mrad}^2$. The dotted line is the weak eigenvector $(0.822, 0.570)$ of set C.

nearly identical, and the inverse,

$$\mathbf{H}^{-1} = \frac{1}{0.001}\begin{bmatrix} 1.001 & -1 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} 1001 & -1000 \\ -1000 & 1000 \end{bmatrix},$$

multiplies small measurement errors into large state errors. The estimated sum $\hat{x}_1 + \hat{x}_2 = y_1$ inherits the error of one measurement, while the estimated difference $\hat{x}_1 - \hat{x}_2 = 2001\,y_1 - 2000\,y_2$, with independent errors of standard deviation σ on each measurement, has standard deviation $\sigma\sqrt{2001^2 + 2000^2} = 2829\,\sigma$. With $\sigma = 0.01$ the sum is known to ± 0.01 and the difference to ± 28.3. The singular values of $\mathbf{H}$ are 2.0005 and 0.000500, and its condition number is 4002, against 2.2 for the three line-flow meters of set C.

A binary classification into observable and unobservable is therefore insufficient for engineering decisions. The question is not only whether $\mathbf{G}$ is invertible but how large its smallest eigenvalue is relative to the precision the decision requires; a nearly singular gain matrix is the numerical sign of a nearly unobservable direction. Chapter 22 met the same distinction for fault signatures: in Section 22.5.3 two temperature sensors sharing a calibration gave a signature matrix of full rank and a condition number of 38, with a bias estimate ten times less precise than with independent calibration.

29.3.5 Measurement redundancy and a consistency residual

With two states and three flow measurements there is one redundant degree of freedom, and the network equations make it explicit. By equation (29.8), $P_{12} + P_{23} = 1000\,\theta_1 = 2 \times 500\,\theta_1$, so

$$P_{13} = \frac{P_{12} + P_{23}}{2} \tag{29.13}$$

for these reactances and this reference. A consistency residual can therefore be formed without estimating the angles first,

$$e_{\mathrm{c}} = y_{12} + y_{23} - 2\,y_{13}, \tag{29.14}$$

in MW. It is the measurement vector multiplied by $\mathbf{a}^{\mathsf{T}} = [1, 1, -2]$, a vector with $\mathbf{a}^{\mathsf{T}}\mathbf{H}_C = [1000 - 1000,\ -1000 + 1000] = \mathbf{0}$, so the state cancels: $e_{\mathrm{c}} = \mathbf{a}^{\mathsf{T}}\boldsymbol{\varepsilon}$. Under an accurate model and unbiased measurements its mean is zero. It is the analytical redundancy of Section 22.5.4 of Chapter 22, here supplied by the network's Kirchhoff relations instead of an energy balance.

Suppose the readings are $y_{12} = 50\,\mathrm{MW}$, $y_{23} = 50\,\mathrm{MW}$ and $y_{13} = 58\,\mathrm{MW}$. Then

$$e_{\mathrm{c}} = 50 + 50 - 2 \times 58 = -16\,\mathrm{MW},$$

and the measurements are inconsistent with the DC model by that amount. Figure 29.2(b) shows the same fact geometrically: the three readings define three lines in the plane of the angles that no longer pass through one point but form a small triangle.

The WLS estimate of equation (29.11) with $\mathbf{R} = 4\,\mathbf{I}$ is $\hat{\theta}_1 = 0.10533$ and $\hat{\theta}_2 = 0.05267\,\mathrm{rad}$, the square in the figure, at which all three flows equal 52.67 MW. The measurement residuals $\mathbf{e} = \mathbf{y} - \mathbf{H}\hat{\mathbf{x}}$ are $(-2.67, -2.67, +5.33)\,\mathrm{MW}$, and the minimized objective is $J(\hat{\mathbf{x}}) = (2.667^2 + 2.667^2 + 5.333^2)/4 = 10.67$. The estimate spreads the inconsistency over the three meters in proportion to how their rows enter $\mathbf{a}$.

29.3.6 Significance of an inconsistency

Whether $-16\,\mathrm{MW}$ is significant depends on its uncertainty. With independent errors of standard deviation 2 MW on each meter,

$$\mathrm{Var}(e_{\mathrm{c}}) = \mathrm{Var}(y_{12}) + \mathrm{Var}(y_{23}) + 2^2\,\mathrm{Var}(y_{13}) = 4 + 4 + 16 = 24\,\mathrm{MW}^2,$$

so $\sigma_c = \sqrt{24} = 4.899\,\text{MW}$ and the standardized residual is

$$z_c = \frac{-16}{4.899} = -3.27.$$

Its square is $J_c = 16^2/24 = 10.67$, which under the healthy model is chi-square with one degree of freedom (equation (22.8) of Chapter 22 with $m = 1$). The 1 % threshold is $\chi^2_{1,0.99} = 6.635$, equivalently $|z_c| > 2.576$, and the probability of a value at least as large under the healthy model is 0.0011. The healthy consistency model is rejected at the 1 % level.

That J_c equals the minimized WLS objective $J(\hat{\mathbf{x}})$ is no coincidence. Every residual vector $\mathbf{e} = \mathbf{y} - \mathbf{H}\mathbf{x}$ satisfies $\mathbf{a}^\mathsf{T}\mathbf{e} = \mathbf{a}^\mathsf{T}\mathbf{y} = e_c$, and when the redundancy is one, every vector with that property is a residual of some $\mathbf{x}$. Minimizing $\mathbf{e}^\mathsf{T}\mathbf{W}\mathbf{e}$ subject to $\mathbf{a}^\mathsf{T}\mathbf{e} = e_c$ with a Lagrange multiplier gives $\mathbf{e} = \mathbf{R}\mathbf{a}\,e_c/(\mathbf{a}^\mathsf{T}\mathbf{R}\mathbf{a})$ and the minimum $e_c^2/(\mathbf{a}^\mathsf{T}\mathbf{R}\mathbf{a})$. With $\mathbf{a}^\mathsf{T}\mathbf{R}\mathbf{a} = 4\times(1+1+4) = 24$ this is $256/24 = 10.67$ and the residuals are $4\times(1, 1, -2)\times(-16)/24 = (-2.67, -2.67, +5.33)\,\text{MW}$, as found. With $m - n$ degrees of freedom the minimized objective is chi-square with $m - n$ degrees of freedom; the chi-square test of bad-data detection rests on it, and Chapter 30 develops it in general.

What the test establishes is that at least one of the following is wrong: a measurement, the assumed topology, a line parameter, the timing of the measurements, or the physical model itself. It does not identify the L13 meter as faulty. The residuals of the L12 and L23 meters are half as large as that of L13 but have proportionally smaller variances: dividing each residual by its own standard deviation, which Chapter 30 derives as the diagonal of the residual covariance $\mathbf{\Omega}$ ($\Omega_{ii} = 0.667$, 0.667 and $2.667\,\text{MW}^2$ here), gives normalized residuals of -3.27, -3.27 and $+3.27$, equal in magnitude. With one degree of redundancy, the evidence points to an inconsistency and not to its location. The sizes of error that would explain it differ by meter: a bias of $+8\,\text{MW}$ on L13, or of $-16\,\text{MW}$ on L12 or on L23. A 1 % test with its threshold $|e_c| > 2.576 \times 4.899 = 12.6\,\text{MW}$ responds to half as large a bias on the L13 meter as on the other two, because L13 enters e_c with weight two.

29.3.7 Critical measurements

A measurement is *critical*, relative to an observable measurement set, when its removal makes the state unobservable. In set B both measurements are critical: removing either leaves one row for two states, rank one. In set C no single measurement is critical: removing any one of the three leaves two independent rows, since any two of $\mathbf{h}_{12}$, $\mathbf{h}_{23}$ and $\mathbf{h}_{13}$ are linearly independent. Set C remains observable after the loss of any one meter, which makes the estimation robust to a single sensor failure.

A critical measurement has a second property that matters for bad data. With $m = n$, as in set B, the estimate fits every measurement exactly and every residual is identically zero; more generally, the residual of a critical measurement is always zero, whatever its error, because no other measurement carries information about the state direction it alone determines. A gross error on a critical measurement is therefore undetectable: it is absorbed into the estimated state. In set C each pair of measurements is a *critical pair*, whose joint removal makes the state unobservable, and the standard result is that bad data on a member of a critical pair can be detected but not identified, which is what the equal normalized residuals of Section 29.3.6 showed. The ability to estimate the state after removing one measurement does not imply the ability to tell which measurement was corrupted when an inconsistency is detected; that requires further diagnostic information and assumptions, the subject of Section 29.4.3.

Principle 29.2. *Observability is decided by the rank of* $\mathbf{H}$*, precision by the eigenvalues of* $\mathbf{G}$*, and the ability to detect bad data by the redundancy beyond observability. A critical measurement contributes to observability but cannot be checked; a measurement with only one degree of redundancy around it can be checked but its error cannot be located.*

29.4 Sensor placement and measurement sufficiency

Which meters to install is a design decision with several objectives that do not always agree. This section writes sensor placement as an optimization problem, states the three standard criteria of estimation precision, separates state observability from the diagnosability of sensor faults with the example of a bias on the L13 meter, and collects the results in a map of what each measurement set supports.

29.4.1 Sensor placement as an optimization problem

Let the candidate sensors be indexed by $i = 1, \dots, M$. Sensor i contributes a measurement row $\mathbf{h}_i$ and has error variance $\sigma_{y,i}^2$, hence weight $w_i = 1/\sigma_{y,i}^2$. Let $s_i \in \{0, 1\}$ indicate whether it is installed. With independent errors, the information matrix of the installed set is the sum of the contributions of its sensors,

$$\mathbf{G}(\mathbf{s}) = \sum_{i=1}^{M} s_i \, w_i \, \mathbf{h}_i^{\mathsf{T}} \mathbf{h}_i, \tag{29.15}$$

which is equation (29.10) written row by row: $\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H} = \sum_i w_i \mathbf{h}_i^{\mathsf{T}} \mathbf{h}_i$ for diagonal $\mathbf{W}$. Each sensor adds a matrix of rank one, which is why one meter can raise the information along one direction only, as the L13 meter did in Section 29.3.2. When prior information about the state is available, with covariance $\mathbf{\Sigma}_0$ (for example from a forecast of the injections), the information of the prior adds to that of the measurements,

$$\mathbf{G}_{\text{post}}(\mathbf{s}) = \mathbf{\Sigma}_0^{-1} + \mathbf{G}(\mathbf{s}), \tag{29.16}$$

the inverse of the posterior covariance for a linear model with Gaussian prior and errors. Both expressions assume a linear model and independent Gaussian errors.

A placement problem might minimize the cost of the installed sensors, c_i for sensor i, subject to observability,

$$\min_{\mathbf{s}} \sum_{i=1}^{M} c_i s_i \quad \text{subject to} \quad \operatorname{rank}\big(\mathbf{H}(\mathbf{s})\big) = n, \tag{29.17}$$

where $\mathbf{H}(\mathbf{s})$ stacks the rows of the installed sensors. For the three-bus network with equal costs, any two of the three line-flow meters solve it. The rank constraint can be strengthened to require observability after the loss of any one sensor, which for the three-bus network forces all three meters, or replaced by a requirement on precision, the subject of the next subsection.

29.4.2 Three design objectives

The informativeness of a sensor set is quantified by a scalar function of its information matrix. Three are standard in the design of experiments, and Table 29.3 lists them with the geometric meaning of each in terms of the confidence ellipsoid of Section 29.3.3.

Table 29.3: The standard optimality criteria for sensor placement.

Criterion	Objective	Interpretation
A-optimality	minimize $\operatorname{tr}(\mathbf{G}^{-1})$	reduce the total estimation variance, the sum of the state variances
D-optimality	maximize $\ln\det\mathbf{G}$	reduce the volume of the confidence ellipsoid, which is proportional to $(\det\mathbf{G})^{-1/2}$
E-optimality	maximize σ_n^2, the smallest eigenvalue of $\mathbf{G}$	improve the weakest information direction, the longest axis of the ellipsoid

The criteria may select different sensors. One candidate may strongly improve a state that is already well estimated, another the least observable direction. Table 29.4 works an illustrative case. A meter on L12 is installed; one more instrument may be added, chosen among a meter on L23, a meter on L13 (both with $\sigma_y = 2\,\mathrm{MW}$), a meter of the load injection at bus 3 ($P_3 = -500\,\theta_1 - 1000\,\theta_2$, also 2 MW), and a phasor measurement unit that measures the angle of bus 1 relative to the reference bus with a standard deviation of 1 mrad (row $[1, 0]$ in rad). The last two are illustrative additions to the network of Figure 29.1. The columns give the standard deviations of the two angles and the three criteria. One row of the arithmetic, the load meter: $\mathbf{G} = \frac{1}{4}[1000, -1000]^\mathsf{T}[1000, -1000] + \frac{1}{4}[-500, -1000]^\mathsf{T}[-500, -1000]$ has entries 312 500, $-125\,000$ and $500\,000\,\mathrm{rad}^{-2}$, so $\det \mathbf{G} = 1.406 \times 10^{11}\,\mathrm{rad}^{-4}$, the inverse has trace $(312\,500 + 500\,000)/1.406 \times 10^{11} = 5.78\,\mathrm{mrad}^2$, and the eigenvalues are 562 500 and $250\,000\,\mathrm{rad}^{-2}$.

Table 29.4: Choice of a second instrument to add to a meter on L12 (illustrative noise levels). Units: mrad for standard deviations, mrad^2 for the trace, rad^{-4} for the determinant, rad^{-2} for the smallest eigenvalue. Bold: the best value of each criterion.

Added instrument	σ_{θ_1}	σ_{θ_2}	A: $\mathrm{tr}\,\mathbf{G}^{-1}$	D: $\det \mathbf{G}$	E: σ_2^2
Flow meter on L23, 2 MW	2.83	2.00	12.00	6.25×10^{10}	95 500
Flow meter on L13, 2 MW	4.00	4.47	36.00	1.56×10^{10}	29 300
Load meter at bus 3, 2 MW	1.89	1.49	**5.78**	1.41×10^{11}	**250 000**
Angle of bus 1, 1 mrad	1.00	2.24	6.00	$\mathbf{2.50 \times 10^{11}}$	191 000

The A- and E-criteria choose the load meter; the D-criterion chooses the angle measurement. The angle measurement fixes θ_1 to 1 mrad and so gives the smallest ellipse area, but it leaves θ_2 at 2.24 mrad and its weakest direction at $1/\sqrt{191\,000} = 2.29\,\mathrm{mrad}$; the load meter spreads its information more evenly, with a weakest direction of 2.00 mrad and a condition number of 1.5. Of the two line meters, the one on L23 is four times better than the one on L13 by the D-criterion: both rows meet the L12 row at 45°, but the L13 row is half as long, because a line of twice the reactance carries half the flow per radian of angle, and the determinant scales with the square of that length. If the objective is the detection of faulty measurements rather than the precision of the state, yet another placement may be preferable: redundancy around each meter matters there, and the cheapest observable set has none.

29.4.3 Observability against fault diagnosability

Suppose the model is augmented with unknown sensor biases. Following Section 22.5.2 of Chapter 22, which wrote the effect of faults on the residuals as $\mathbf{S}_f\mathbf{f}$ with a fault-signature matrix $\mathbf{S}_f$, the measurement equation becomes

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{S}_f\mathbf{b} + \boldsymbol{\varepsilon}, \tag{29.18}$$

where $\mathbf{b}$ is the vector of the biases admitted as unknowns and column j of $\mathbf{S}_f$ marks the meter that bias j affects. The augmented unknown vector is $\mathbf{x}_\mathrm{a} = [\mathbf{x}^\mathsf{T}, \mathbf{b}^\mathsf{T}]^\mathsf{T}$ and the augmented Jacobian is

$$\mathbf{H}_\mathrm{aug} = \begin{bmatrix} \mathbf{H} & \mathbf{S}_f \end{bmatrix}. \tag{29.19}$$

Estimating the physical state and every admitted bias without further assumptions requires $\mathbf{H}_\mathrm{aug}$ to have full column rank, a stronger requirement than $\mathrm{rank}(\mathbf{H}) = n$.

Suppose the L12 and L23 meters are trusted and the L13 meter is suspected of a bias b_{13}. Then $\mathbf{x}_\mathrm{a} = [\theta_1, \theta_2, b_{13}]^\mathsf{T}$ and

$$\mathbf{H}_\mathrm{aug} = \begin{bmatrix} 1000 & -1000 & 0 \\ 0 & 1000 & 0 \\ 500 & 0 & 1 \end{bmatrix},$$

whose determinant, expanded along the last column, is $1 \times \det(\mathbf{H}_B) = 10^6 \neq 0$: full column rank. The two angles are estimated from L12 and L23 and the remaining discrepancy on L13 is attributed to b_{13}, provided the stated assumptions hold. For the readings of Section 29.3.5, $\hat{\theta}_1 = 0.10$, $\hat{\theta}_2 = 0.05\,\text{rad}$, the predicted L13 flow is $50\,\text{MW}$ and $\hat{b}_{13} = 58 - 50 = 8\,\text{MW}$, with variance $\text{Var}(y_{13}) + 500^2\,\text{Var}(\hat{\theta}_1) = 4 + 250\,000 \times 8 \times 10^{-6} = 6\,\text{MW}^2$; the ratio $8/\sqrt{6} = 3.27$ reproduces the consistency statistic. The augmented model spends the one degree of redundancy on the bias, and nothing is left to check it.

If unknown biases are admitted on all three meters simultaneously, the unknowns are θ_1, θ_2, b_{12}, b_{23} and b_{13}, five against three observations, and $\mathbf{H}_{\text{aug}}$, of size 3×5, cannot have rank five: the augmented system is underdetermined. With exactly one bias admitted but its location unknown, the data cannot choose either. Table 29.5 fits each single-bias hypothesis to the readings $(50, 50, 58)\,\text{MW}$: each uses the two other meters to fix the angles and assigns the remainder to the suspected one. Every hypothesis reproduces the three readings exactly, with zero residual, so their likelihoods are equal; they are observationally equivalent, and only prior probabilities or other evidence, such as a calibration record, can separate them. In the language of Section 22.5.2, the three bias signatures project onto the one-dimensional space of residuals as the multiples 1, 1 and -2 of the same vector $\mathbf{a}$.

Table 29.5: The three single-bias explanations of the readings $y_{12} = 50$, $y_{23} = 50$ and $y_{13} = 58\,\text{MW}$. Each fixes the angles from the two meters it trusts.

Biased meter	$\hat{\theta}_1$ (rad)	$\hat{\theta}_2$ (rad)	Predicted flow (MW)	Bias (MW)
L13	0.100	0.050	$P_{13} = 50$	$+8$
L12	0.116	0.050	$P_{12} = 66$	-16
L23	0.116	0.066	$P_{23} = 66$	-16

The diagnostic conclusion therefore depends on which faults are allowed and which measurements are trusted. Trusting L12 and L23 makes the L13 bias estimable; admitting a bias anywhere makes the location of the fault unobservable from these three meters. That is the distinction between state observability and *fault diagnosability*, the observability of the fault parameters of the augmented model.

29.4.4 A measurement sufficiency map

An engineering report can state, question by question, what each measurement set supports. Table 29.6 does it for the three sets of this chapter, using the flow meters only. Each row is an engineering question and each column a measurement set; an entry states whether the set supports the answer, with the subsection that shows it.

The last row deserves emphasis. More measurements add information, but they do not automatically resolve every competing fault hypothesis: the three-meter set detects a single bias but cannot locate it. A measurement system can be adequate for state estimation and inadequate for a definitive diagnosis of sensor faults, and the map makes that visible before a conclusion is drawn from it.

29.5 Nonlinear observability, dynamic observability and identifiability

The rank condition was derived for a linear, static model with known parameters. This section extends it in three directions: to nonlinear measurement models, where it holds only locally; to

Table 29.6: Measurement sufficiency map of the three-bus network for three sets of line-flow meters (DC model, no injection measurements).

Engineering question	L12 only	L12 and L23	All three flows	See
Infer $\theta_1 - \theta_2$	yes	yes	yes	29.2.2
Infer both nonreference bus angles	no	yes	yes	29.2.3
Estimate all three line flows from the known topology	no	yes	yes	29.2.3
Test the redundant flows for consistency	no	no	yes	29.3.6
Continue state estimation after any single meter is lost	no	no	yes	29.3.7
Uniquely diagnose an arbitrary unknown sensor bias	no	no	not in general	29.4.3

dynamic models, where a sequence of measurements can reveal states that one instant hides; and to unknown parameters, where it becomes the question of identifiability.

29.5.1 Measurement Jacobians

In AC state estimation the measurement model $\mathbf{y} = h(\mathbf{x}) + \boldsymbol{\varepsilon}$ is nonlinear: line flows and injections depend on products of voltage magnitudes and on sines and cosines of angle differences. Around an operating point $\mathbf{x}_0$, a first-order Taylor expansion gives

$$h(\mathbf{x}) \approx h(\mathbf{x}_0) + \mathbf{H}(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0), \qquad \mathbf{H}(\mathbf{x}_0) = \left.\frac{\partial h}{\partial \mathbf{x}}\right|_{\mathbf{x}_0}, \tag{29.20}$$

where $\mathbf{H}(\mathbf{x}_0)$, the $m \times n$ matrix of first partial derivatives with entries $\partial h_i / \partial x_j$, is the *measurement Jacobian*. The local condition for observability is

$$\operatorname{rank}\big[\mathbf{H}(\mathbf{x}_0)\big] = n. \tag{29.21}$$

When it holds and h is continuously differentiable, a suitable subset of n measurements has an invertible Jacobian at $\mathbf{x}_0$, and the inverse function theorem guarantees that those measurements determine the state uniquely within some neighborhood of $\mathbf{x}_0$. That is *local* observability. It does not establish *global* uniqueness: a nonlinear model can map distant states to the same measurements. The DC model of Section 29.2.1 is the linearization of the AC flow equations at flat voltage magnitudes and small angles, which is why its Jacobian rows are constants.

29.5.2 A nonlinear ambiguity

Let a sensor report $y = x^2$ for a scalar state x. The reading $y = 4$ is produced by $x = 2$ and by $x = -2$. The Jacobian is $H(x) = 2x$; at $x = 2$ it equals 4, of rank one, so the state is locally observable there, and a small change of x near 2 changes y. Globally the measurement does not distinguish positive from negative states. At $x = 0$ the Jacobian vanishes and even local observability is lost: a small change of x produces a change of y of second order only. The example separates local from global identifiability of the physical state.

The power-flow equations contain the same ambiguity. With voltage magnitudes of 1 per unit, the AC active-power flow on L12 is $P_{12} = S_{\text{base}} \sin(\theta_1 - \theta_2)/x_{12} = 1000 \sin\delta$ MW, with $\delta = \theta_1 - \theta_2$. A reading of 50 MW gives $\sin\delta = 0.05$, satisfied by $\delta = 0.0500$ rad and by $\delta = \pi - 0.0500 = 3.0916$ rad. The Jacobian $1000\cos\delta$ is $+998.7$ MW/rad at the first solution, close to the DC coefficient of 1000, and -998.7 MW/rad at the second. Both solutions are locally observable; the second, beyond the steady-state stability limit of $\pi/2$, is excluded by operating knowledge rather than by the measurement. Figure 29.5 draws both cases. In panel

(a) the horizontal axis is the state x and the vertical axis the reading; the parabola is $y = x^2$, the dashed line the reading $y = 4$, and the dots its two solutions. In panel (b) the horizontal axis is the angle difference in radians and the vertical axis the L12 flow; the curve is $1000 \sin \delta$, the dashed line 50 MW, and the dots its two solutions near the ends of the interval.

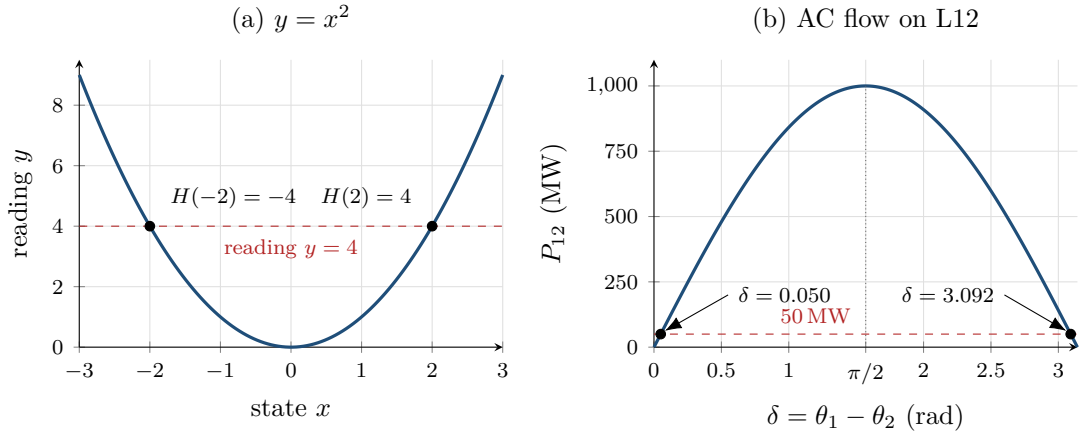


Figure 29.5: Local observability without global uniqueness. (a) The reading $y = x^2 = 4$ is produced by $x = \pm 2$; the Jacobian $2x$ is nonzero at both and zero at $x = 0$. (b) The AC flow $P_{12} = 1000 \sin \delta$ MW ($V_1 = V_2 = 1$ per unit, $x_{12} = 0.1$ per unit, 100 MVA base) equals 50 MW at $\delta = 0.050$ and $\delta = \pi - 0.050$ rad; the second lies beyond the steady-state limit $\pi/2$ (dotted).

In practice the state estimation of a power system also depends on the types of measurement (flows, injections, voltage magnitudes), on the operating conditions, on the topology, and on the possibility of several nonlinear solutions; a flat start near the expected operating point, and the exclusion of physically implausible solutions, are part of the estimator's design.

29.5.3 Numerical convergence of state estimation

A nonlinear WLS estimator minimizes

$$J(\mathbf{x}) = [\mathbf{y} - h(\mathbf{x})]^{\mathsf{T}} \mathbf{W} [\mathbf{y} - h(\mathbf{x})], \tag{29.22}$$

which has no closed-form minimizer. The *Gauss–Newton* iteration replaces h at the current iterate $\mathbf{x}^{(k)}$ by its linearization, equation (29.20), with $\mathbf{H} = \mathbf{H}(\mathbf{x}^{(k)})$. The objective then becomes the linear WLS problem of Section 29.3.1 in the step $\Delta\mathbf{x} = \mathbf{x} - \mathbf{x}^{(k)}$, with $\mathbf{y} - h(\mathbf{x}^{(k)})$ in place of $\mathbf{y}$, and its solution is

$$\Delta\mathbf{x} = (\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H})^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}[\mathbf{y} - h(\mathbf{x}^{(k)})], \qquad \mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \Delta\mathbf{x}. \tag{29.23}$$

Each step requires the gain matrix at the current iterate to be invertible, that is, local observability at every iterate, not only at the solution.

Table 29.7 runs the iteration on the example $y = x^2$ with the reading $y = 4$. With one measurement and one state, equation (29.23) reduces to $\Delta x = (4 - x^2)/(2x)$. The columns are the starting point and the first five iterates. One step of the arithmetic, from $x^{(0)} = 1$: $\Delta x = (4 - 1)/2 = 1.5$, so $x^{(1)} = 2.5$.

From a positive start the iteration converges to $+2$, from a negative start to -2, and from $x^{(0)} = 0$ it cannot start, because the Jacobian is zero there. A start close to the singular point, 0.5, overshoots to 4.25 before converging. The underlying state is locally observable at both solutions, yet the estimate returned depends on the initial guess, and the algorithm can fail at a point of lost local observability. In a power-system estimator the same roles are played by poor conditioning, a poor initial guess, unmodeled topology errors, and a model that does

Table 29.7: Gauss–Newton iterates for the reading $y = x^2 = 4$ from five starting points.

$x^{(0)}$	$x^{(1)}$	$x^{(2)}$	$x^{(3)}$	$x^{(4)}$	$x^{(5)}$	Limit
1.0	2.5000	2.0500	2.0006	2.0000	2.0000	$+2$
3.0	2.1667	2.0064	2.0000	2.0000	2.0000	$+2$
0.5	4.2500	2.5956	2.0683	2.0011	2.0000	$+2$
-1.0	-2.5000	-2.0500	-2.0006	-2.0000	-2.0000	-2
0	$H = 0$: gain matrix singular, no step					none

not fit the data. Observability, numerical convergence and model adequacy are therefore three different properties, and a converged estimate establishes only the second.

29.5.4 Dynamic observability

A state may be unobservable from the measurements of one instant and observable from a sequence of them. Consider the discrete-time linear system

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k, \qquad y_k = \mathbf{C}\mathbf{x}_k, \tag{29.24}$$

with state-transition matrix $\mathbf{A}$ ($n \times n$) and output matrix $\mathbf{C}$ (here $1 \times n$, one measurement per instant). The measurements from $k = 0$ to $n-1$ are $y_0 = \mathbf{C}\mathbf{x}_0$, $y_1 = \mathbf{C}\mathbf{A}\mathbf{x}_0$, …, $y_{n-1} = \mathbf{C}\mathbf{A}^{n-1}\mathbf{x}_0$: a linear measurement model for the initial state whose Jacobian is the *observability matrix*

$$\mathcal{O} = \begin{bmatrix} \mathbf{C} \\ \mathbf{C}\mathbf{A} \\ \vdots \\ \mathbf{C}\mathbf{A}^{n-1} \end{bmatrix}. \tag{29.25}$$

By the rank condition of Section 29.1.4, the initial state, and with it every later state, is observable from n successive measurements if and only if

$$\operatorname{rank}(\mathcal{O}) = n. \tag{29.26}$$

Further measurements cannot raise the rank: by the Cayley–Hamilton theorem of linear algebra, $\mathbf{A}^n$ and every higher power are linear combinations of $\mathbf{I}, \mathbf{A}, \ldots, \mathbf{A}^{n-1}$, so the test with n block rows is complete. This is the standard observability condition of linear system theory.

Take two states, $\mathbf{x}_k = [x_{1,k}, x_{2,k}]^{\mathsf{T}}$, with

$$\mathbf{A} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \qquad \mathbf{C} = \begin{bmatrix} 1 & 0 \end{bmatrix},$$

so that $x_{1,k+1} = x_{1,k} + x_{2,k}$, $x_{2,k+1} = x_{2,k}$, and only the first state is measured, $y_k = x_{1,k}$. The model describes a level x_1 that changes each step by a constant rate x_2, such as a degradation state and its drift (Chapter 21). At a single instant the rate is hidden. But $\mathbf{C}\mathbf{A} = [1, 1]$, so

$$\mathcal{O} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \qquad \det \mathcal{O} = 1, \qquad \operatorname{rank}(\mathcal{O}) = 2,$$

and the system is dynamically observable. The first measurement gives $y_0 = x_{1,0}$ and the second $y_1 = x_{1,0} + x_{2,0}$, hence

$$x_{2,0} = y_1 - y_0. \tag{29.27}$$

The hidden state affects how the measured state evolves, and the evolution reveals it. Observers and dynamic state estimators, among them the Kalman filter, rest on this mechanism. Had the rate been measured instead, $\mathbf{C} = [0, 1]$, then $\mathbf{C}\mathbf{A} = [0, 1]$ and $\mathcal{O}$ would have two equal rows, rank one: the level never influences the rate, and no sequence of rate measurements reveals it.

The practical limitation. A state estimated from differences of noisy measurements can be very uncertain. With independent measurement errors of standard deviation σ, equation (29.27) gives $\mathrm{Var}(\hat{x}_{2,0}) = 2\sigma^2$, twice the variance of a single reading. More samples help: with K readings $y_k = x_{1,0} + k\,x_{2,0} + \varepsilon_k$, $k = 0, \ldots, K-1$, the rows of the measurement matrix are $[1, k]$, which are the rows of the observability matrix extended to K instants, and the WLS slope has variance $\sigma^2 / \sum_k (k - \bar{k})^2 = 12\,\sigma^2/[K(K^2-1)]$, which is $2\sigma^2$ for $K = 2$, $0.1\,\sigma^2$ for $K = 5$ and $0.012\,\sigma^2$ for $K = 10$. Dynamic observability also assumes that the state-transition model $\mathbf{A}$ is accurate. If the rate in fact changes between samples, or an unmeasured input drives the level, the estimate of x_2 absorbs the error, and an apparently observable system yields unreliable estimates.

29.5.5 States and parameters

Observability concerns the recovery of the evolving state $\mathbf{x}(t)$ from measurements. *Identifiability* concerns whether the model parameters $\boldsymbol{\theta}$ of equation (29.3) can be uniquely determined from the available experiments and observations: whether $h(\mathbf{x}, \boldsymbol{\theta}_1) = h(\mathbf{x}, \boldsymbol{\theta}_2)$ for every operating condition implies $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_2$. The two questions are related but different. The state may be observable under known parameters while some parameters remain unidentifiable, and Chapter 31 treats the second question in detail.

In the DC network, the flow on L13 is

$$P_{13} = \frac{100}{x_{13}}\,\theta_1 .$$

If both the reactance x_{13} and the angle θ_1 are unknown and only P_{13} is measured, the observation constrains their ratio and not their values: $P_{13} = 50\,\mathrm{MW}$ is produced by $(\theta_1, x_{13}) = (0.10, 0.2)$, by $(0.12, 0.24)$, and by every pair with $\theta_1 = 0.5\,x_{13}$. Many parameter–state pairs give the same flow.

The rank condition extends to the joint unknown $[\mathbf{x}^\mathsf{T}, \boldsymbol{\theta}^\mathsf{T}]^\mathsf{T}$ through the joint Jacobian

$$\begin{bmatrix} \mathbf{H} & \mathbf{S}_\theta \end{bmatrix} = \begin{bmatrix} \dfrac{\partial h}{\partial \mathbf{x}} & \dfrac{\partial h}{\partial \boldsymbol{\theta}} \end{bmatrix}, \tag{29.28}$$

where $\mathbf{S}_\theta = \partial h / \partial \boldsymbol{\theta}$ is the sensitivity matrix of the measurements to the parameters. Its rank and singular values state which combinations of states and parameters the measurements constrain locally. For the single L13 meter at $\theta_1 = 0.10$ and $x_{13} = 0.2$, $\partial P_{13}/\partial\theta_1 = 100/x_{13} = 500\,\mathrm{MW/rad}$ and $\partial P_{13}/\partial x_{13} = -100\,\theta_1/x_{13}^2 = -250\,\mathrm{MW}$ per unit of reactance; the joint row $[500, -250]$ has rank one, and its null direction $(\delta\theta_1, \delta x_{13}) \propto (0.5, 1)$ is the proportional change of both. With all three flows measured and x_{13} unknown, the joint Jacobian is

$$\begin{bmatrix} 1000 & -1000 & 0 \\ 0 & 1000 & 0 \\ 500 & 0 & -250 \end{bmatrix},$$

with determinant $-250 \times 10^6 \neq 0$: the angles and the reactance are jointly identifiable. The L12 and L23 meters fix the angles, and the L13 meter, no longer redundant, fixes x_{13}: from $\hat{x}_{13} = 100\,\hat{\theta}_1 / y_{13}$ and $\sigma_y = 2\,\mathrm{MW}$, its standard deviation is 0.0098 per unit, 4.9 % of 0.2. The degree of freedom that detected bad data in Section 29.3.6 has been spent on the parameter, and the L13 meter has become critical.

29.5.6 Excitation

Having more observations than parameters is not sufficient for identifiability. The observations must respond differently to the unknown quantities, and two parameters that always enter as

a product, a ratio or another inseparable combination remain unidentifiable even with many repeated measurements at the same operating condition. Let

$$y = a\,u + b + \varepsilon,$$

with a known input u, unknown parameters $\boldsymbol{\theta} = [a, b]^{\mathsf{T}}$ (slope and intercept, symbols local to this example), and independent errors of standard deviation σ. If every experiment uses the same input $u = u_0$, the observations constrain $a\,u_0 + b$ and cannot separate a from b. The Fisher information of K observations is

$$\mathcal{I}(\boldsymbol{\theta}) = \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta = \frac{1}{\sigma^2}\begin{bmatrix} \sum_k u_k^2 & \sum_k u_k \\ \sum_k u_k & K \end{bmatrix}, \qquad \det \mathcal{I}(\boldsymbol{\theta}) = \frac{K\sum_k (u_k - \bar{u})^2}{\sigma^4},$$

which is zero when all u_k are equal: the information matrix is singular, and the parameters are not identifiable however many observations are taken. When the inputs differ, $\mathrm{Var}(\hat{a}) = \sigma^2 / \sum_k (u_k - \bar{u})^2$. With $K = 4$ and $\sigma = 1$, inputs $(1, 1, 1, 1)$ give a singular matrix; inputs $(0.5, 0.75, 1.25, 1.5)$ give $\mathrm{Var}(\hat{a}) = 1/0.625 = 1.6$; inputs $(0, 0.5, 1.5, 2)$ give $1/2.5 = 0.4$. Spreading the inputs is *excitation*, a question of experimental design.

Figure 29.6 draws both cases. In panel (a) the horizontal axis is the reactance x_{13} and the vertical axis the angle θ_1; the solid line is the set of pairs that give $P_{13} = 50\,\mathrm{MW}$, and the shaded band is the $\pm 2\sigma$ range of θ_1 estimated from the L12 and L23 meters, which cuts the line in a short segment around $x_{13} = 0.2$. In panel (b) the horizontal axis is the input u and the vertical axis the output y; the two lines $y = 2u + 1$ and $y = u + 2$ both pass through the point $(1, 3)$, so observations at $u = 1$ alone (squares) cannot choose between them, while observations at other inputs (dots) fall on one line only.

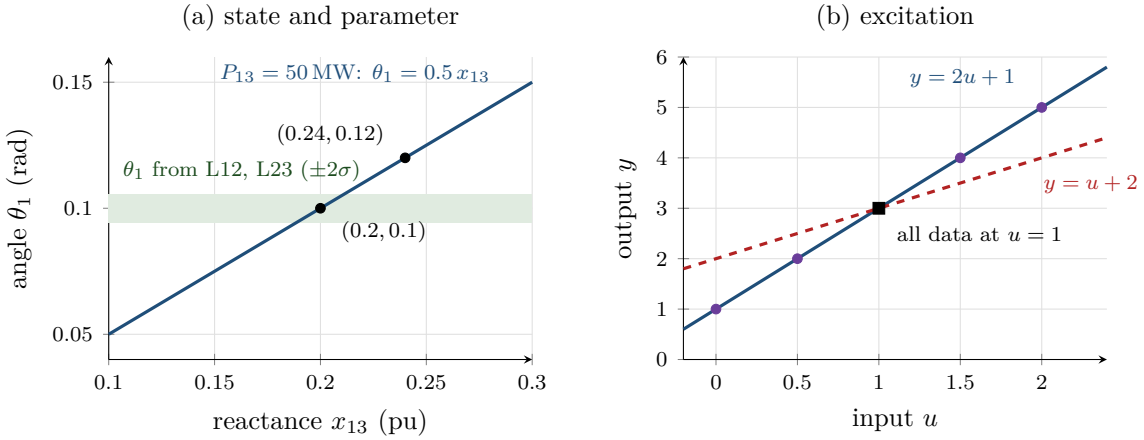


Figure 29.6: Identifiability. (a) A single L13 meter reading 50 MW fixes only the ratio $\theta_1/x_{13} = 0.5$ (solid line); the angle estimated from the L12 and L23 meters, $0.100 \pm 2 \times 0.00283\,\mathrm{rad}$ (band), selects a short segment of it and so identifies x_{13}. (b) Two parameter sets $(a, b) = (2, 1)$ and $(1, 2)$ predict the same output at $u = 1$ (square); noise-free observations at varied inputs (dots) lie on the first line only. Illustrative values.

In a power system, changes of dispatch, or the natural variation of the operating conditions over hours and seasons, provide the excitation that identifies parameters such as line reactances or the coefficients of a conductor thermal model. Deliberate operating changes made for the purpose must satisfy the same safety and security constraints as any other operating decision.

29.6 Observability in reliability inference and engineering decisions

The final question is what the observability of a measurement set means for the reliability conclusions drawn from it. This section writes a risk probability as an integral over the posterior of the state and the rating, shows that a decision quantity can be observable when the state is not and the converse, classifies the evidence behind a conclusion, and closes with a worked test of rank, conditioning and critical measurements for every set of the three-bus meters.

29.6.1 Observability and risk calculations

Suppose a transmission operator needs the probability that the flow on line ℓ exceeds its rating, given the measurements, $\Pr(|P_\ell| > P_\ell^{\max} \mid \mathbf{y})$, where P_ℓ is the actual flow, $P_\ell^{\max}$ the actual rating, which depends on weather, and $\mathbf{y}$ the available measurements. A point estimate of flow and rating is not always enough. When both are uncertain, the probability is the integral of the exceedance indicator over their joint posterior,

$$\Pr(|P_\ell| > P_\ell^{\max} \mid \mathbf{y}) = \int \mathbf{1}\{|P_\ell(\mathbf{x})| > P_\ell^{\max}\}\, p(\mathbf{x}, P_\ell^{\max} \mid \mathbf{y})\, \mathrm{d}\mathbf{x}\, \mathrm{d}P_\ell^{\max}, \tag{29.29}$$

which requires a joint posterior over the physical state and the rating. The posterior is the prior updated by the likelihood of the measurements, as in Chapter 8 of Volume I. Along the directions the measurements observe, the likelihood dominates; along unobservable or weakly observed directions, the posterior is essentially the prior.

An illustrative calculation shows the effect. With the meters of set B, the estimated L13 flow is $\hat{P}_{13} = 500\,\hat{\theta}_1 = 50\,\mathrm{MW}$ with variance $500^2 \times 8 \times 10^{-6} = 2\,\mathrm{MW}^2$. Let the rating of L13, which no flow meter measures, have the prior $P_{13}^{\max} \sim \mathcal{N}(80, 8^2)\,\mathrm{MW}$ to represent weather uncertainty. With independent Gaussian errors, the margin $P_{13}^{\max} - P_{13}$ has mean 30 MW and variance $2 + 64 = 66\,\mathrm{MW}^2$, and

$$\Pr(P_{13} > P_{13}^{\max} \mid \mathbf{y}) = \Phi\!\left(\frac{-30}{\sqrt{66}}\right) = \Phi(-3.69) = 1.1 \times 10^{-4}.$$

Of the $66\,\mathrm{MW}^2$, 64 come from the prior of the rating. The probability is computable and the state is observable, but the answer is dominated by an assumption. Such a risk probability should not be presented as strongly supported by data when the observations barely constrain the variables that drive it; when measurements are insufficient, the result may depend heavily on prior assumptions.

29.6.2 Decision quantities under partial observability

An unobservable state can still permit an observable decision quantity. Suppose only L12 is metered, set A, so that the individual angles are unobservable. The flow on L12, however, is measured directly. If the question is whether L12 exceeds its known rating of 70 MW, the reading $y_{12} = 50\,\mathrm{MW}$ with $\sigma_y = 2\,\mathrm{MW}$ answers it: the flow is ten standard deviations below the rating, and with a flat prior the posterior exceedance probability is $1 - \Phi(10) = 7.6 \times 10^{-24}$. Full-state observability was not needed.

What matters is whether the decision-relevant quantity is sufficiently constrained, and that is again a rank question. A quantity $q = \mathbf{c}\,\mathbf{x}$, linear in the state with row $\mathbf{c}$, is estimable from a measurement set when $\mathbf{c}$ is a combination of the rows of $\mathbf{H}$, equivalently when $\mathbf{c}\,\mathbf{v} = 0$ for every unobservable direction $\mathbf{v}$. For set A, $\mathbf{v} = (1, 1)$: the L12 flow $[1000, -1000]$ passes the test, but the flow on L12 after an outage of L13, which is $P_{12} + P_{13}$ because all of the transfer then takes the path through bus 2 (Chapter 27, line outage distribution factor 1), has row

$[1500, -1000]$ and $\mathbf{c}\,\mathbf{v} = 500 \neq 0$. The N−1 question for that outage cannot be answered from the L12 meter. With set B it can: the post-contingency flow is estimated at 100 MW with variance $[1500, -1000]\ \mathrm{Cov}(\hat{\mathbf{x}})\,[1500, -1000]^{\mathsf{T}} = 18 - 12 + 4 = 10\,\mathrm{MW}^2$, a standard deviation of 3.2 MW, and the overload of L12 beyond 70 MW is certain for practical purposes.

Conversely, a fully observable state does not make an uncertain conductor rating identifiable, as Section 29.6.1 showed. The adequacy of a measurement system must be assessed relative to the actual question.

Principle 29.3. *Measurement sufficiency is a property of a measurement set and a question together. A decision quantity is observable when its sensitivity to the state vanishes along every unobservable direction; full-state observability is neither necessary for that nor sufficient when the decision also depends on parameters the measurements do not reach.*

29.6.3 A classification of evidence

A reliability inference report can classify each of its conclusions by the evidence behind it. Table 29.8 gives six classes, each with its meaning and an example from this chapter or from Chapter 22.

Table 29.8: Classification of the evidence behind an inferred quantity.

Class	Meaning	Example
Directly measured	the target quantity has direct sensor evidence	the L12 flow from its own meter
Inferred and well constrained	several measurements and a credible model determine the target	both bus angles from independent flow meters (set C)
Partially constrained	some combinations are known, others remain uncertain	the angle difference without an absolute reference (set A)
Prior dominated	the data provide little information about the target	an unmeasured coefficient of a conductor thermal model; the L13 rating of Section 29.6.1
Not identifiable under the model	distinct explanations are observationally equivalent	a real temperature deviation against an unknown bias with one sensor; the location of a single bias in Table 29.5
Model inadequate	no candidate state or parameter value explains the observations adequately	a persistently inconsistent energy balance

The classes separate computing an answer from having enough evidence to trust it. Every row produces a number when the calculation is run; only the first two produce a number that the measurements support. The third and fourth need a statement of what has been assumed, the fifth needs further evidence of a different kind, and the sixth needs a revised model before any number is reported.

29.6.4 Sufficiency evaluated for each question

A system can be observable for one purpose and inadequate for another. With the L12 and L23 meters, the two nonreference angles of the three-bus network are estimated under the DC model. Suppose the goal changes to finding which line's thermal-rating uncertainty contributes most to the probability of a post-contingency overload. Additional information is then needed: conductor temperatures or the weather observations that drive the ratings; the parameters of

the conductor thermal model; the condition of each line and the uncertainty of its rating; the outage-state probabilities; and the correlations among weather, failures and thermal cooling, which Chapter 28 showed can dominate the risk.

Adding the third flow meter improves the precision of the angles without materially resolving the rating uncertainty. In the calculation of Section 29.6.1 it reduces the variance of $\hat{P}_{13}$ from 2 to 1.33 MW^2 and the exceedance probability from 1.1×10^{-4} to 1.0×10^{-4}; an instrument that reduced the standard deviation of the rating from 8 to 2 MW would reduce it to about 10^{-34}. The value of a measurement depends on the question it is intended to answer. That is why sensor placement for state estimation, for fault diagnosis and for reliability-risk assessment may produce different optimal instrumentation plans, and why a later chapter on the value of information treats the choice of measurements as a decision.

29.6.5 An observability test for the three-bus meters

Every subset of the three line-flow meters can be tested at once. Table 29.9 lists the seven nonempty subsets. For each it gives the number of measurements m, the rank of $\mathbf{H}$, the unobservable direction when the rank is one, the redundancy $m - n$ and condition number κ of $\mathbf{H}$ when the rank is two (equal to that of the weighted Jacobian, since every meter has $\sigma_y = 2$ MW), the standard deviations of the two estimated angles, the critical measurements, and whether the state stays observable after the loss of any one meter. One row of the arithmetic, L12 and L13: $\mathbf{H} = [1000, -1000; 500, 0]$ has determinant $5 \times 10^5 \neq 0$, so rank two; $\mathbf{H}^{\mathsf{T}}\mathbf{H}$ has trace 2.25×10^6 and determinant 2.5×10^{11}, eigenvalues 2.133×10^6 and 1.172×10^5, singular values 1460 and 342, so $\kappa = 4.27$; $\hat{\theta}_1 = y_{13}/500$ has standard deviation $2/500 = 4.00$ mrad, and $\hat{\theta}_2 = \hat{\theta}_1 - y_{12}/1000$ has $\sqrt{16 + 4} = 4.47$ mrad; both meters are critical.

Table 29.9: Observability test for every subset of the three line-flow meters, $\sigma_y = 2$ MW each, DC model with $\mathbf{x} = [\theta_1, \theta_2]^{\mathsf{T}}$. Standard deviations in mrad.

Meters	m	Rank	Unobservable direction, or κ	$m-n$	σ_{θ_1}	σ_{θ_2}	Critical measurements	Survives one loss
L12	1	1	$(1, 1)$	–	–	–	–	no
L23	1	1	$(1, 0)$	–	–	–	–	no
L13	1	1	$(0, 1)$	–	–	–	–	no
L12, L23	2	2	$\kappa = 2.62$	0	2.83	2.00	L12, L23	no
L12, L13	2	2	$\kappa = 4.27$	0	4.00	4.47	L12, L13	no
L23, L13	2	2	$\kappa = 2.00$	0	4.00	2.00	L23, L13	no
L12, L23, L13	3	2	$\kappa = 2.20$	1	2.31	1.83	none	yes

The table reproduces the three steps of the sets A, B and C and adds the others. With L12 alone the rank is one; removing L23 from set B leaves exactly that. Installing L13 with L12 restores rank two, because L12 and L13 provide independent state information, but the pair is the worst conditioned of the three pairs and gives the least precise estimate: its rows $[1000, -1000]$ and $[500, 0]$ meet at 45°, as the L12 and L23 rows do, but the L13 row is half as long, whereas the L23 and L13 rows are orthogonal and measure the two angles separately. With all three meters the rank stays at two, the precision improves, no measurement is critical, and the state survives the loss of any single meter. The distinction between acquiring a new dimension of information and acquiring precision or redundancy is the central one of measurement-system design. The estimation itself, the residual tests and the location of bad data in general networks are the subject of the next chapter.

29.7 What to retain

Definition 29.1 (Observability and measurement sufficiency)**.** A state is observable from a measurement set when no change of state leaves the modeled measurements unchanged, which for a linear or linearized model is $\operatorname{rank}(\mathbf{H}) = n$. A measurement set is sufficient for a question when the quantity that question needs is constrained by the measurements, with precision given by $\operatorname{Cov}(\hat{\mathbf{x}}) = (\mathbf{H}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{H})^{-1}$, and not by the prior.

The main conclusions are eight. Observability determines whether the available measurements distinguish the physical states. Full rank establishes the structural sufficiency of the information for the linear state problem but does not guarantee precise estimates. Measurement noise and numerical conditioning determine the estimation uncertainty. Measurement redundancy improves bad-data detection and resilience to the loss of a sensor, even when it does not increase the observable state dimension, and a critical measurement can never be checked. State observability is different from fault diagnosability and from parameter identifiability. Nonlinear systems can be locally observable without being globally identifiable. Dynamic measurements can reveal states that are hidden in an instantaneous snapshot. And the engineering question determines what information is sufficient: full-state observability is not always necessary, while an observable state may still be inadequate for a risk calculation involving uncertain parameters or thresholds. Table 29.10 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 29.10: The terms of observability and measurement sufficiency and where each is defined.

Term	Expression	Units	Meaning	Defined in
State vector	$\mathbf{x} = [\theta_1, \dots, \theta_{N-1}]^{\mathsf{T}}$ (DC)	rad	nonreference bus angles	29.1.1
Measurement model	$\mathbf{y} = h(\mathbf{x}, \boldsymbol{\theta}) + \boldsymbol{\varepsilon}$, $\mathrm{Cov}(\boldsymbol{\varepsilon}) = \mathbf{R}$	as $\mathbf{y}$	measurements as functions of the state	29.1.2
Observability	$\mathrm{rank}(\mathbf{H}) = n$	–	state recoverable from measurements	29.1.4
Unobservable direction	$\mathbf{H}\mathbf{v} = \mathbf{0}$, $\mathbf{v} \neq \mathbf{0}$	–	state change no meter sees	29.1.4
WLS estimate	$\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}$	as $\mathbf{x}$	minimizer of $J(\mathbf{x})$	29.3.1
Gain (information) matrix	$\mathbf{G} = \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}$, $\mathrm{Cov}(\hat{\mathbf{x}}) = \mathbf{G}^{-1}$	rad^{-2}	Fisher information of the state	29.3.1
Condition number	$\kappa = \sigma_1/\sigma_n$ of $\mathbf{W}^{1/2}\mathbf{H}$	–	numerical observability	29.3.3
Measurement redundancy	$m - n$	–	measurements beyond observability	29.2.4
Consistency residual	$e_{\mathrm{c}} = y_{12} + y_{23} - 2y_{13}$	MW	bad-data detection statistic	29.3.5
Critical measurement	removal gives rank $< n$	–	observability rests on it; residual zero	29.3.7
A-, D-, E-optimality	$\mathrm{tr}\,\mathbf{G}^{-1}$, $\ln\det\mathbf{G}$, σ_n^2	varies	sensor-placement criteria	29.4.2
Augmented Jacobian	$[\mathbf{H} \;\; \mathbf{S}_f]$	–	state and sensor biases	29.4.3
Measurement Jacobian	$\mathbf{H}(\mathbf{x}_0) = \partial h/\partial \mathbf{x}\vert_{\mathbf{x}_0}$	varies	local linear model	29.5.1
Gauss–Newton step	$\Delta\mathbf{x} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}[\mathbf{y} - h(\mathbf{x})]$	as $\mathbf{x}$	nonlinear WLS iteration	29.5.3
Observability matrix	$\mathcal{O} = [\mathbf{C}; \mathbf{C}\mathbf{A}; \dots; \mathbf{C}\mathbf{A}^{n-1}]$	–	dynamic observability	29.5.4
Identifiability	$\mathrm{rank}[\mathbf{H} \;\; \mathbf{S}_\theta]$	–	parameters recoverable	29.5.5
Excitation	$\det\mathcal{I} \propto \sum_k (u_k - \bar{u})^2$	–	input variation for identification	29.5.6
Exceedance probability	$\Pr(\lvert P_\ell\rvert > P_\ell^{\max} \mid \mathbf{y})$	–	risk from a joint posterior	29.6.1

Further reading. Abur and Expósito, *Power System State Estimation: Theory and Implementation*: state estimation, measurement Jacobians, observability and bad-data analysis. Monticelli, *State Estimation in Electric Power Systems: A Generalized Approach*: observability, topology processing and estimation. Chen, *Linear System Theory and Design*: observability and controllability of dynamic systems.

29.8 Exercises

Exercise 29.1. (a) State the linear observability condition for n unknown states. (The measurement matrix $\mathbf{H}$ has full column rank n.) (b) State what a nonzero vector in the null space of $\mathbf{H}$ represents. (A state change that produces no change in the modeled measurements.) (c) State what happens when a redundant flow measurement is added to an already observable system. (It can reduce the estimation uncertainty and allow consistency checking without increasing the rank.) (d) Explain why a full-rank measurement matrix can still yield unreliable estimates. (It may be poorly conditioned, or the measurements may be very noisy, so that $\mathbf{G}^{-1}$ is large along some direction.)

Exercise 29.2. In the three-bus network, the flows on L12 and L13 are measured, each with $\sigma_y = 2\,\mathrm{MW}$. (a) Write $\mathbf{H}$, find its rank and state whether the state is observable.

($\mathbf{H} = [1000, -1000; 500, 0]$, determinant 5×10^5, rank two, observable.) (b) Compute the gain matrix and the standard deviations of the two estimated angles, and compare them with set B. ($\mathbf{G} = [312\,500, -250\,000; -250\,000, 250\,000]\,\mathrm{rad}^{-2}$; $\mathrm{Cov}(\hat{\mathbf{x}}) = [16, 16; 16, 20]\,\mathrm{mrad}^2$; $\sigma_{\theta_1} = 4.00$ and $\sigma_{\theta_2} = 4.47\,\mathrm{mrad}$, against 2.83 and 2.00 for set B.) (c) Name the critical measurements and compute the condition number. (Both meters; $\kappa = 4.27$.) (d) For the noise-free readings $y_{12} = y_{13} = 50\,\mathrm{MW}$, compute the angles and the flow on L23. ($\theta_1 = 0.10$, $\theta_2 = 0.05\,\mathrm{rad}$, $P_{23} = 50\,\mathrm{MW}$.)

Exercise 29.3. All three flows are measured with $\sigma_y = 2\,\mathrm{MW}$, and the readings are $y_{12} = 48$, $y_{23} = 53$ and $y_{13} = 56\,\mathrm{MW}$. (a) Compute the consistency residual, its standardized value and the chi-square statistic, and state the conclusion at the 5 % and 1 % levels. ($e_\mathrm{c} = -11\,\mathrm{MW}$; $z_\mathrm{c} = -2.25$; $J_\mathrm{c} = 5.04$, with tail probability 0.025: rejected at 5 %, threshold 3.84, not at 1 %, threshold 6.63.) (b) Compute the bias that each single-bias hypothesis requires. (L13: $\hat{\theta}_1 = 0.101$, $\hat{\theta}_2 = 0.053\,\mathrm{rad}$, predicted flow 50.5 MW, bias +5.5 MW; L12 or L23: $-11\,\mathrm{MW}$.) (c) State whether the data identify the faulty meter. (No: each hypothesis fits the three readings exactly, so the data detect an inconsistency but cannot locate it.)

Exercise 29.4. (a) For the two-state system of Section 29.5.4 with $\mathbf{C} = [1, 0]$, the first two readings are $y_0 = 3.0$ and $y_1 = 3.4$, each with error standard deviation 0.1. Estimate the hidden rate and its standard deviation, and the standard deviation with ten readings. ($\hat{x}_{2,0} = 0.4$; $0.1\sqrt{2} = 0.141$; $0.1\sqrt{12/990} = 0.011$.) State the benefit of dynamic observability. (Measurements over time can reveal states hidden at a single instant.) (b) Find the rank of $\mathcal{O}$ when $\mathbf{C} = [0, 1]$. (One: the measured rate never depends on the level.) (c) State what distinguishes state observability from parameter identifiability. (Observability concerns recovering the system's states; identifiability concerns determining its model parameters.) (d) For $y = au + b + \varepsilon$ with $\sigma = 1$ and inputs $(0, 1, 2, 3)$, compute $\mathrm{Var}(\hat{a})$. ($\sum (u - \bar{u})^2 = 5$, so 0.2.) (e) State whether a decision-relevant quantity can be inferred without full-state observability. (Yes, if that quantity is sufficiently constrained, as the L12 overload question is by the L12 meter alone.)

Chapter 30

State Estimation and Bad-Data Detection

30.1 State estimation and its two outputs

Chapter 29 asked whether a set of measurements contains enough independent information to determine the physical state of a network. This chapter takes up a different problem. Even when the state is observable, the measurements generally do not agree with one another exactly. A transmission operator may see 50 MW entering a corridor, 50 MW on an intermediate line, and 58 MW on a third line that the network model says should carry 50 MW. Averaging the readings is not physically meaningful, because the three flows are tied together by the network equations; discarding the largest discrepancy may discard the one correct reading. An estimation procedure is needed that respects the network physics and the uncertainty of each measurement, and that separates two tasks the operator easily conflates: *finding* an inconsistency and *identifying its cause*. The question of this chapter is

when measurements disagree with a physical model of the system, how the actual operating state is estimated, and how it is determined whether the discrepancy comes from measurement errors, from an incorrect network topology, or from an inadequately modeled physical condition.

The chapter answers it in five steps. It formulates the state estimation problem for the alternating-current (AC) and direct-current (DC) network models and derives the weighted least-squares estimator and its covariance, worked on the three-bus network of the earlier chapters. It derives the residuals of the estimate and their covariance, and from them the global chi-square test and the normalized residuals that detect and identify bad data. It shows where those tests fail: critical measurements, interacting errors that conform to the network model, and errors of the assumed topology. It treats robust estimators that limit the influence of gross errors, and Bayesian state estimation, in which competing explanations of an inconsistency receive posterior probabilities. Finally, it carries those probabilities into the overload probability of a line, the quantity a reliability decision needs.

Jargon: state estimation, state vector, SCADA, PMU, measurement model, WLS, gain matrix, residual sensitivity matrix, bad data, bad-data detection and identification, chi-square test, normalized residual, LNR test, measurement redundancy, critical measurement, interacting and conforming bad data, topology error, topology processing, robust estimator

State estimation is the computation of the most probable operating state of a network from a redundant set of noisy measurements and a model of the network. The *state vector* $\mathbf{x}$ holds the quantities that, with the network model, determine every other electrical quantity: in a transmission network, the bus voltage angles and magnitudes. The measurements come from the *supervisory control and data acquisition* (SCADA) system, which reports power flows, injections, voltage magnitudes and breaker statuses every few seconds, and from *phasor measurement units* (PMUs), which measure voltage and current phasors, angles included, with a common time reference. The *measurement model* $h(\mathbf{x})$ predicts each measurement from the state. *Weighted least squares* (WLS) is the estimator that minimizes the sum of squared measurement residuals, each weighted by the inverse of its error variance; its *gain matrix* $\mathbf{G}$ is the matrix of its normal equations. *Bad data* are measurements with gross errors, much larger than their stated uncertainty. *Bad-data detection* decides whether the measurement set contains bad data; *bad-data identification* decides which measurements are bad. The *chi-square test* detects through the weighted sum of squared residuals; the *normalized residual* divides each residual by its own standard deviation, and the *largest normalized residual* (LNR) *test* identifies the measurement with the largest one; the *residual sensitivity matrix* $\mathbf{S}$ maps measurement errors into residuals. *Measurement redundancy* is the excess of the number of measurements m over the number of states n. A *critical measurement* is one whose removal makes the state unobservable. Bad data in several measurements are *interacting* when their residuals are strongly correlated, and *conforming* when the errors are mutually consistent with the network model, so that they partly or wholly cancel in the residuals. A *topology error* is an error in the assumed connectivity of the network, such as a wrong breaker status; *topology processing* builds the network model from the breaker and switch statuses. A *robust estimator* limits the influence of measurements whose residuals are large compared with their stated uncertainty.

Key idea

A state estimator returns two things: the state most consistent with the measurements and the network model, and evidence on whether that model explains the measurements at all. Residual tests detect an inconsistency; they do not, by themselves, say whether a sensor, a breaker status or a line parameter is responsible.

This section separates the measurements from the state they depend on, and describes what a state estimator takes in and what it returns.

30.1.1 Measurements are not the state

The electrical state of an AC transmission network of N buses is described by the voltage phase angle θ_i (rad) and the voltage magnitude V_i (per unit) of every bus. Angles are defined only relative to one another, so the angle of one reference bus is fixed at zero and removed. With bus 1 as the reference, the state vector of a conventional AC state estimator is

$$\mathbf{x} = \begin{bmatrix} \theta_2 & \cdots & \theta_N & V_1 & \cdots & V_N \end{bmatrix}^{\mathsf{T}}, \tag{30.1}$$

of dimension $n = 2N - 1$. The state is not what the instruments report. SCADA delivers active and reactive power injections at buses, active and reactive power flows on branches,

voltage magnitudes, and the statuses of circuit breakers and switches; only PMUs measure angles directly. The measurements are related to the state through the network equations. For a lossless line of series reactance x_{ij} (per unit), the active power flowing from bus i to bus j is

$$P_{ij} = \frac{V_i V_j}{x_{ij}} \sin(\theta_i - \theta_j), \tag{30.2}$$

in per unit of the system power base S_{base}; a complete AC model adds the line resistance and shunt admittance and the analogous expressions for reactive power.

Each measurement is the model's prediction plus an error. Stacking the m measurements in a vector,

$$\mathbf{y} = h(\mathbf{x}; \boldsymbol{\theta}) + \boldsymbol{\varepsilon}, \tag{30.3}$$

where $\mathbf{y}$ holds the measured quantities, $\mathbf{x}$ is the unknown electrical state, $\boldsymbol{\theta}$ collects the network parameters and topology (line reactances, transformer taps, breaker statuses), h is the physical measurement model, one function per measurement, and $\boldsymbol{\varepsilon}$ is the vector of measurement errors. The measurement vector is written $\mathbf{y}$, the symbol this book uses for observations; the power-system literature writes it $\mathbf{z}$. The parameter vector $\boldsymbol{\theta}$ is not the set of bus angles θ_i, which belong to the state. For the derivations of the next two sections the parameters and the topology are assumed known and are dropped from the notation, $\mathbf{y} = h(\mathbf{x}) + \boldsymbol{\varepsilon}$; Section 30.4 returns to the case in which they are wrong.

30.1.2 The role of the state estimator

Figure 30.1 draws the state estimator with its inputs and outputs. The three boxes on the left are its inputs: the telemetered measurements, the topology and parameters of the network, and the models of measurement uncertainty. The box in the center is the estimator, which infers the network state most consistent with the measurements and the model. The two boxes on the right are its outputs.

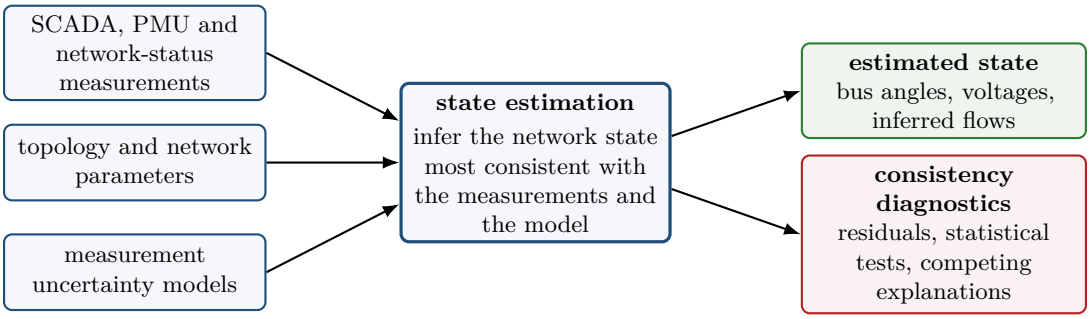


Figure 30.1: The state estimator, its three inputs and its two outputs. The estimated state feeds monitoring and security analysis; the consistency diagnostics are where bad-data detection begins.

The first output is an estimate of the physical operating state, from which every flow, injection and voltage of the network follows, measured or not. The second is evidence on whether the assumed model adequately explains the observations: the residuals between measurements and estimate, the statistical tests built on them, and the competing explanations of any inconsistency. The second output is where bad-data detection begins, and it is the one through which state estimation enters reliability engineering: an estimated flow is a usable input to a risk calculation only when the model that produced it has been checked against the evidence. The distinction of this book between the physical system, the model of it and the evidence is explicit here. The network and its true flows are the system; the measurement model $h(\mathbf{x})$, with its topology and parameters, is the model; the telemetered values, with their errors, are the evidence.

30.2 Weighted least-squares estimation

Weighted least squares is the standard state estimator of power systems. This section states the linear DC measurement model, derives the WLS objective from the Gaussian likelihood and its solution from the normal equations, derives the covariance of the estimate, and works the estimator on the three-bus network.

30.2.1 The linear DC measurement model

Under the DC power-flow approximation used for the contingency analysis of Chapter 27, voltage magnitudes are set to one per unit, line resistances are neglected and angle differences are small, so that $\sin(\theta_i - \theta_j) \approx \theta_i - \theta_j$ in equation (30.2). The flow on line ij becomes, in MW,

$$P_{ij} = S_{\text{base}} \, \frac{\theta_i - \theta_j}{x_{ij}}, \tag{30.4}$$

linear in the angles, and the state reduces to the $n = N - 1$ angles of the non-reference buses. Every measurement of a flow or an injection is then a linear combination of the states, and the measurement model is

$$\mathbf{y} = \mathbf{H}\,\mathbf{x} + \boldsymbol{\varepsilon}, \tag{30.5}$$

where $\mathbf{H}$, of dimension $m \times n$, is the measurement matrix. In general $\mathbf{H} = \partial h / \partial \mathbf{x}$ is the Jacobian of the measurement model, which for the DC model is constant. The errors are assumed to have zero mean and a known covariance,

$$\mathbb{E}[\boldsymbol{\varepsilon}] = \mathbf{0}, \qquad \operatorname{Cov}(\boldsymbol{\varepsilon}) = \mathbf{R}, \tag{30.6}$$

and when the measurement errors are independent of one another $\mathbf{R}$ is diagonal, $\mathbf{R} = \operatorname{diag}(\sigma_1^2, \sigma_2^2, \ldots, \sigma_m^2)$, with σ_i the standard deviation of the error of measurement i in its own units. The variances need not be equal: a more precise flow or voltage measurement has a smaller σ_i and, as the next subsection shows, receives more weight.

Remark 30.1. The bold $\mathbf{R}$ is the covariance matrix of the measurement errors, not the reliability function $R(t)$; the bold $\mathbf{H}$ is the measurement Jacobian, not the operating horizon H; and the bold $\mathbf{S}$ of Section 30.3 is the residual sensitivity matrix, not the shortfall S_t of Chapter 25 nor the FMEA severity S. The hypotheses $H_0, H_1, \ldots$ of Section 30.6 are written in italic and are distinct from the matrix $\mathbf{H}$.

30.2.2 The WLS objective from the Gaussian likelihood

Suppose the errors are jointly Gaussian, $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$. Then, for a given state, the measurements are Gaussian with mean $\mathbf{Hx}$ and covariance $\mathbf{R}$, and their likelihood is the multivariate Gaussian density

$$p(\mathbf{y} \mid \mathbf{x}) = \frac{1}{(2\pi)^{m/2}\,|\mathbf{R}|^{1/2}} \exp\!\Big[-\tfrac{1}{2}(\mathbf{y} - \mathbf{H}\mathbf{x})^{\mathsf{T}} \mathbf{R}^{-1} (\mathbf{y} - \mathbf{H}\mathbf{x})\Big]. \tag{30.7}$$

The normalizing factor does not depend on $\mathbf{x}$, and the logarithm is increasing, so maximizing the likelihood over $\mathbf{x}$ is equivalent to minimizing minus twice its logarithm without the constant,

$$J(\mathbf{x}) = (\mathbf{y} - \mathbf{H}\mathbf{x})^{\mathsf{T}} \mathbf{W} (\mathbf{y} - \mathbf{H}\mathbf{x}), \qquad \mathbf{W} = \mathbf{R}^{-1}. \tag{30.8}$$

This is the *weighted least-squares objective*, dimensionless, and $\mathbf{W}$ is the *weight matrix*. The WLS estimate is therefore the maximum likelihood estimate of the state under Gaussian errors, and the weights are fixed by the measurement uncertainty, not chosen freely.

Different measurements, different weights. For independent measurements the weight matrix is diagonal, with entries $1/\sigma_i^2$, and the objective is a sum of squared standardized discrepancies,

$$J(\mathbf{x}) = \sum_{i=1}^{m} \frac{[y_i - h_i(\mathbf{x})]^2}{\sigma_i^2}, \tag{30.9}$$

in which $h_i(\mathbf{x})$ is the prediction of measurement i, here the i-th row of $\mathbf{H}\mathbf{x}$. A discrepancy of 5 MW contributes 25 to J when the expected error of that measurement is 1 MW, and 0.25 when it is 10 MW: the same discrepancy is a hundred times more significant on the precise instrument. A covariance matrix that misstates the precision of the measurements therefore misplaces the estimate and, as Section 30.3 shows, every test built on it.

30.2.3 The normal equations and the solution

Expanding the objective,

$$J(\mathbf{x}) = \mathbf{y}^{\mathsf{T}}\mathbf{W}\mathbf{y} - 2\,\mathbf{x}^{\mathsf{T}}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} + \mathbf{x}^{\mathsf{T}}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\,\mathbf{x},$$

using the symmetry of $\mathbf{W}$ to combine the two cross terms. The gradient with respect to $\mathbf{x}$ is

$$\frac{\partial J}{\partial \mathbf{x}} = -2\,\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} + 2\,\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\,\mathbf{x} = -2\,\mathbf{H}^{\mathsf{T}}\mathbf{W}(\mathbf{y} - \mathbf{H}\mathbf{x}), \tag{30.10}$$

and the Hessian is $2\,\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}$. Setting the gradient to zero at the estimate $\hat{\mathbf{x}}$ gives

$$\mathbf{H}^{\mathsf{T}}\mathbf{W}(\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}) = \mathbf{0} \quad \Longleftrightarrow \quad \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\,\hat{\mathbf{x}} = \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y},$$

the *normal equations*. Their matrix is the *gain matrix*

$$\mathbf{G} = \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}, \tag{30.11}$$

of dimension $n \times n$, symmetric, and also called the information matrix. For any vector $\mathbf{v}$, $\mathbf{v}^{\mathsf{T}}\mathbf{G}\mathbf{v} = (\mathbf{H}\mathbf{v})^{\mathsf{T}}\mathbf{W}(\mathbf{H}\mathbf{v}) \geq 0$, with equality only when $\mathbf{H}\mathbf{v} = \mathbf{0}$. The gain matrix is therefore positive definite, and the objective strictly convex with a unique minimum, exactly when $\mathbf{H}$ has full column rank n, which is the observability condition of Chapter 29. Then

$$\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}. \tag{30.12}$$

In numerical practice $\mathbf{G}^{-1}$ is not formed: the normal equations are solved by a sparse Cholesky factorization of $\mathbf{G}$, or by an orthogonal factorization of $\mathbf{W}^{1/2}\mathbf{H}$ when $\mathbf{G}$ is ill-conditioned.

For the nonlinear AC model the same reasoning is applied repeatedly. Around a current iterate $\mathbf{x}^{(k)}$, the model is linearized, $h(\mathbf{x}^{(k)} + \Delta\mathbf{x}) \approx h(\mathbf{x}^{(k)}) + \mathbf{H}(\mathbf{x}^{(k)})\,\Delta\mathbf{x}$, with $\mathbf{H}(\mathbf{x}^{(k)})$ the Jacobian evaluated there; the objective becomes quadratic in $\Delta\mathbf{x}$, and its minimum solves

$$\mathbf{G}(\mathbf{x}^{(k)})\,\Delta\mathbf{x}^{(k)} = \mathbf{H}^{\mathsf{T}}(\mathbf{x}^{(k)})\,\mathbf{W}\,[\mathbf{y} - h(\mathbf{x}^{(k)})], \qquad \mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \Delta\mathbf{x}^{(k)}, \tag{30.13}$$

the Gauss–Newton iteration, repeated until the update is negligible. For the linear DC model it converges in one step to equation (30.12).

30.2.4 The estimation covariance

Substituting the model $\mathbf{y} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}$ into equation (30.12),

$$\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\,\mathbf{x} + \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\boldsymbol{\varepsilon} = \mathbf{x} + \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\boldsymbol{\varepsilon}.$$

The estimation error $\hat{\mathbf{x}} - \mathbf{x}$ is a linear function of the measurement errors with mean zero, so the estimate is unbiased, and its covariance is

$$\mathrm{Cov}(\hat{\mathbf{x}} - \mathbf{x}) = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\,\mathbf{R}\,\mathbf{W}\mathbf{H}\,\mathbf{G}^{-1} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\,\mathbf{G}^{-1} = \mathbf{G}^{-1}, \tag{30.14}$$

using $\mathbf{WRW} = \mathbf{W}$ and the symmetry of $\mathbf{G}$. The inverse of the gain matrix is the formal covariance of the state estimate. It depends on which sensors exist, on what they measure through $\mathbf{H}$, and on how noisy they are through $\mathbf{W}$: it is the quantity a sensor-placement study of Chapter 29 evaluates. It is also conditional on the model. It assumes the topology, the line parameters and the zero mean of the errors are right, and it contains no allowance for a wrong breaker status, a wrong reactance or a biased sensor; those require additional modeling, which the rest of the chapter supplies.

30.2.5 The three-bus network and its measurements

The example is the three-bus network of Chapters 27 and 29, drawn in Figure 30.2. The circles are the buses; the lines between them are L12, L13 and L23, each labeled with its series reactance in per unit on $S_{\text{base}} = 100\,\text{MVA}$; the small squares mark the three flow meters, y_{12}, y_{23} and y_{13}, which measure the active power flowing from the lower-numbered to the higher-numbered bus. Bus 1 holds the generation, bus 3 the load and the angle reference.

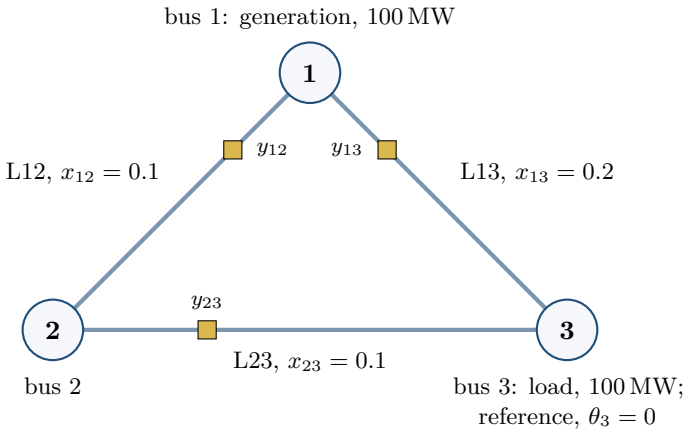


Figure 30.2: The three-bus transmission network, with series reactances in per unit on a 100 MVA base and the three active-power flow meters (squares). Flows are positive from bus 1 to bus 2, from bus 2 to bus 3 and from bus 1 to bus 3.

Bus 3 is the reference, $\theta_3 = 0$, and the state is the pair of remaining angles,

$$\mathbf{x} = \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \quad \text{(rad)}.$$

By equation (30.4), with $S_{\text{base}}/x_{12} = 100/0.1 = 1000\,\text{MW/rad}$, $S_{\text{base}}/x_{23} = 1000\,\text{MW/rad}$ and $S_{\text{base}}/x_{13} = 500\,\text{MW/rad}$, the three flows are

$$P_{12} = 1000\,(\theta_1 - \theta_2), \qquad P_{23} = 1000\,\theta_2, \qquad P_{13} = 500\,\theta_1,$$

and the measurement matrix of the three meters, rows in the order y_{12}, y_{23}, y_{13}, is

$$\mathbf{H} = \begin{bmatrix} 1000 & -1000 \\ 0 & 1000 \\ 500 & 0 \end{bmatrix} \text{ MW/rad}. \tag{30.15}$$

This is measurement set C of Chapter 29: its rank is two, so the state is observable, and with $m = 3$ measurements and $n = 2$ states it has one degree of redundancy, $m - n = 1$. The base case of Chapter 27 transfers 100 MW from bus 1 to bus 3 at $\theta_1 = 0.10$ rad and $\theta_2 = 0.05$ rad, which gives $P_{12} = P_{23} = P_{13} = 50$ MW; the line ratings are $P_{12}^{\max} = P_{23}^{\max} = 70$ MW and $P_{13}^{\max} = 80$ MW.

The SCADA system now reports the values of Table 30.1. The first two agree with the base case; the third reports 58 MW on L13. Each meter has a nominal uncertainty of 2 MW (one standard deviation), and the errors are assumed independent, zero-mean and Gaussian, so that

$$\mathbf{y} = \begin{bmatrix} 50 \\ 50 \\ 58 \end{bmatrix} \text{ MW}, \qquad \mathbf{R} = \begin{bmatrix} 4 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 4 \end{bmatrix} \text{ MW}^2, \qquad \mathbf{W} = \mathbf{R}^{-1} = \tfrac{1}{4}\,\mathbf{I}\ \text{MW}^{-2}.$$

The task is to estimate the two angles and to decide whether the three reported flows are mutually consistent.

Table 30.1: The three telemetered flows of the example and their nominal uncertainty.

Measurement	Symbol	Observed value	Standard deviation σ_i
L12 active-power flow	y_{12}	50 MW	2 MW
L23 active-power flow	y_{23}	50 MW	2 MW
L13 active-power flow	y_{13}	58 MW	2 MW

30.2.6 The gain matrix and the estimate

With $\mathbf{W} = \frac{1}{4}\mathbf{I}$, the gain matrix is one quarter of $\mathbf{H}^{\mathsf{T}}\mathbf{H}$, whose entries are the sums of products of the columns of $\mathbf{H}$: the $(1,1)$ entry is $1000^2 + 0^2 + 500^2 = 1\,250\,000$, the $(1,2)$ and $(2,1)$ entries are $1000(-1000) + 0(1000) + 500(0) = -1\,000\,000$, and the $(2,2)$ entry is $(-1000)^2 + 1000^2 + 0^2 = 2\,000\,000$, all in MW^2/rad^2. Multiplying by $\frac{1}{4}\,\text{MW}^{-2}$,

$$\mathbf{G} = \begin{bmatrix} 312\,500 & -250\,000 \\ -250\,000 & 500\,000 \end{bmatrix} \text{rad}^{-2}.$$

The right-hand side of the normal equations is

$$\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} = \frac{1}{4}\begin{bmatrix} 1000(50) + 0(50) + 500(58) \\ -1000(50) + 1000(50) + 0(58) \end{bmatrix} = \begin{bmatrix} 19\,750 \\ 0 \end{bmatrix} \text{rad}^{-1},$$

so the normal equations read

$$\begin{bmatrix} 312\,500 & -250\,000 \\ -250\,000 & 500\,000 \end{bmatrix} \begin{bmatrix} \hat{\theta}_1 \\ \hat{\theta}_2 \end{bmatrix} = \begin{bmatrix} 19\,750 \\ 0 \end{bmatrix}.$$

The second equation, $-250\,000\,\hat{\theta}_1 + 500\,000\,\hat{\theta}_2 = 0$, gives $\hat{\theta}_2 = \hat{\theta}_1/2$. Substituted into the first, $(312\,500 - 125\,000)\,\hat{\theta}_1 = 187\,500\,\hat{\theta}_1 = 19\,750$, so

$$\hat{\theta}_1 = 0.105333\,\text{rad}, \qquad \hat{\theta}_2 = 0.052667\,\text{rad}.$$

Both angles lie above the physically consistent base case, 0.10 and 0.05 rad: the high L13 reading has pulled the whole estimate upward.

The covariance of the estimate, equation (30.14), is the inverse of the gain matrix. Its determinant is $312\,500 \times 500\,000 - 250\,000^2 = 9.375 \times 10^{10}\,\text{rad}^{-4}$, and

$$\mathbf{G}^{-1} = \frac{1}{9.375 \times 10^{10}} \begin{bmatrix} 500\,000 & 250\,000 \\ 250\,000 & 312\,500 \end{bmatrix} = \begin{bmatrix} 5.3333 & 2.6667 \\ 2.6667 & 3.3333 \end{bmatrix} \times 10^{-6}\,\text{rad}^2. \tag{30.16}$$

The standard deviations of the estimated angles are $\sqrt{5.3333 \times 10^{-6}} = 0.002309\,\text{rad}$ and $\sqrt{3.3333 \times 10^{-6}} = 0.001826\,\text{rad}$, against 0.002828 and 0.002000 rad with the two meters L12 and L23 alone (Chapter 29). The third meter adds no observable direction but improves the precision.

30.2.7 The estimated flows

The estimated measurements are $\hat{\mathbf{y}} = \mathbf{H}\hat{\mathbf{x}}$:

$$\hat{P}_{12} = 1000\,(0.105333 - 0.052667) = 52.667\,\text{MW},$$
$$\hat{P}_{23} = 1000 \times 0.052667 = 52.667\,\text{MW},$$
$$\hat{P}_{13} = 500 \times 0.105333 = 52.667\,\text{MW}.$$

Table 30.2 sets them against the observations. Each row is a meter; the columns give the observed flow, the estimated flow, the difference between them, and the standard deviation of the estimated flow, the square root of the diagonal of $\text{Cov}(\mathbf{H}\hat{\mathbf{x}}) = \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}$. One entry of the arithmetic: $\mathbf{h}_3\mathbf{G}^{-1}\mathbf{h}_3^{\mathsf{T}} = 500^2 \times 5.3333 \times 10^{-6} = 1.3333\,\text{MW}^2$ for the third row $\mathbf{h}_3 = (500, 0)$, a standard deviation of 1.155 MW.

Table 30.2: Observed and estimated flows of the three-bus example, their difference, and the standard deviation of each estimated flow.

Line	Observed (MW)	Estimated (MW)	Observed − estimated (MW)	Std. dev. of estimate (MW)
L12	50.000	52.667	−2.667	1.826
L23	50.000	52.667	−2.667	1.826
L13	58.000	52.667	+5.333	1.155

The estimator has produced a network-consistent compromise. The three estimated flows satisfy the network equations exactly, which no subset of the observations does: bus 1 sends $\hat{P}_{12} + \hat{P}_{13} = 105.33\,\text{MW}$, bus 2 passes on what it receives, and bus 3 receives 105.33 MW. Every estimated flow is more precise than any single meter, whose standard deviation is 2 MW, because each draws on all three readings. The observations, however, sit at nonzero distances from the estimate. Those distances are the starting point of bad-data detection.

30.3 Residual analysis and bad-data detection

The differences between the measurements and the estimate carry the evidence of inconsistency. This section derives the residual sensitivity matrix and the residual covariance, the global chi-square test, and the normalized residuals, and shows on the three-bus example why the normalized residuals detect the inconsistency without identifying the measurement responsible.

30.3.1 The residual vector and the residual sensitivity matrix

The *measurement residual* is the observation minus its estimate,

$$\mathbf{e} = \mathbf{y} - h(\hat{\mathbf{x}}) = \mathbf{y} - \mathbf{H}\hat{\mathbf{x}}, \tag{30.17}$$

the residual of Chapter 22 with the healthy prediction replaced by the estimate. For the example it is the fourth column of Table 30.2,

$$\mathbf{e} = \begin{bmatrix} -2.6667 \\ -2.6667 \\ 5.3333 \end{bmatrix}\ \text{MW}.$$

The largest residual belongs to L13, and it is tempting to conclude that the L13 meter is wrong. A residual is not a diagnosis. The estimate was built from all three measurements, and WLS distributes any inconsistency over the whole measurement set; how it does so is fixed by the measurement matrix and the weights, and must be understood before a residual is read.

Substituting equation (30.12) into equation (30.17),

$$\mathbf{e} = \mathbf{y} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} = \left(\mathbf{I} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\right)\mathbf{y}.$$

The matrix in parentheses is the *residual sensitivity matrix*,

$$\mathbf{S} = \mathbf{I} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}, \qquad \mathbf{e} = \mathbf{S}\,\mathbf{y}, \tag{30.18}$$

of dimension $m \times m$, dimensionless. It has three properties. First, it annihilates the measurement matrix,

$$\mathbf{S}\mathbf{H} = \mathbf{H} - \mathbf{H}\mathbf{G}^{-1}(\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}) = \mathbf{H} - \mathbf{H}\mathbf{G}^{-1}\mathbf{G} = \mathbf{0},$$

so that, if the model $\mathbf{y} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}$ is correct,

$$\mathbf{e} = \mathbf{S}\mathbf{H}\mathbf{x} + \mathbf{S}\boldsymbol{\varepsilon} = \mathbf{S}\,\boldsymbol{\varepsilon}\,: \tag{30.19}$$

the residual depends only on the measurement errors and not on the state. Element S_{ij} is the change of residual i per unit error of measurement j. The conclusion holds only under the assumed model; if the topology or the parameters are wrong, the residual also carries the model error, the subject of Section 30.4. Second, $\mathbf{S}$ is idempotent, $\mathbf{S}^2 = \mathbf{S}$: writing $\mathbf{K} = \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}$, $\mathbf{K}^2 = \mathbf{H}\mathbf{G}^{-1}(\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H})\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W} = \mathbf{K}$, hence $(\mathbf{I}-\mathbf{K})^2 = \mathbf{I} - 2\mathbf{K} + \mathbf{K} = \mathbf{I} - \mathbf{K}$. Third, its trace is the redundancy: by the cyclic property of the trace, $\operatorname{tr}\mathbf{K} = \operatorname{tr}(\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}) = \operatorname{tr}\mathbf{I}_n = n$, so $\operatorname{tr}\mathbf{S} = m - n$.

For the example, $\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}$ is computed row by row from equation (30.16); for instance its $(1,2)$ entry is $(1000, -1000)\,\mathbf{G}^{-1}(0, 1000)^{\mathsf{T}} = 10^6 \times (2.6667 - 3.3333) \times 10^{-6} = -0.6667\,\mathrm{MW}^2$. In full,

$$\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}} = \begin{bmatrix} 10/3 & -2/3 & 4/3 \\ -2/3 & 10/3 & 4/3 \\ 4/3 & 4/3 & 4/3 \end{bmatrix} \mathrm{MW}^2, \qquad \mathbf{S} = \mathbf{I} - \tfrac{1}{4}\,\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}} = \frac{1}{6}\begin{bmatrix} 1 & 1 & -2 \\ 1 & 1 & -2 \\ -2 & -2 & 4 \end{bmatrix}.$$

The trace of $\mathbf{S}$ is $(1+1+4)/6 = 1 = m - n$, and $\mathbf{S}\mathbf{y}$ reproduces the residual: its first entry is $(50 + 50 - 2 \times 58)/6 = -2.667\,\mathrm{MW}$ and its third $(-100 + 4 \times 58)/6 = (-2)(-16)/6 = 5.333\,\mathrm{MW}$.

30.3.2 Residual covariance

By equation (30.19) the residual has mean zero under the correct model and covariance

$$\boldsymbol{\Omega} = \operatorname{Cov}(\mathbf{e}) = \mathbf{S}\,\mathbf{R}\,\mathbf{S}^{\mathsf{T}}.$$

The product $\mathbf{S}\mathbf{R}$ is already symmetric,

$$\mathbf{S}\mathbf{R} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{R} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}},$$

and multiplying it on the right by $\mathbf{S}^{\mathsf{T}} = \mathbf{I} - \mathbf{W}\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}$,

$$\mathbf{S}\mathbf{R}\mathbf{S}^{\mathsf{T}} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}} + \mathbf{H}\mathbf{G}^{-1}(\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H})\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}.$$

Hence the residual covariance is

$$\boldsymbol{\Omega} = \mathbf{S}\,\mathbf{R} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}, \tag{30.20}$$

in the squared units of the measurements. The second form has a direct reading: $\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}$ is the covariance of the estimated measurements $\hat{\mathbf{y}}$, and since $\mathbf{y} = \hat{\mathbf{y}} + \mathbf{e}$ with $\hat{\mathbf{y}}$ and $\mathbf{e}$ uncorrelated ($\mathrm{Cov}(\hat{\mathbf{y}}, \mathbf{e}) = \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{R}\mathbf{S}^{\mathsf{T}} = \mathbf{H}\mathbf{G}^{-1}(\mathbf{S}\mathbf{H})^{\mathsf{T}} = \mathbf{0}$), the covariance of the measurements splits into the part explained by the estimate and the part left in the residual. For the example,

$$\boldsymbol{\Omega} = 4\,\mathbf{S} = \begin{bmatrix} 0.6667 & 0.6667 & -1.3333 \\ 0.6667 & 0.6667 & -1.3333 \\ -1.3333 & -1.3333 & 2.6667 \end{bmatrix} \mathrm{MW}^2. \tag{30.21}$$

The residual variances are $0.6667\,\mathrm{MW}^2$ for L12 and L23 and $2.6667\,\mathrm{MW}^2$ for L13. The correlation coefficient of the L12 and L13 residuals is $-1.3333/\sqrt{0.6667 \times 2.6667} = -1$, and that of L12 and L23 is $+1$: the residuals are perfectly correlated, and in this arrangement $e_{12} = e_{23}$ and $e_{13} = -2e_{12}$ whatever the measurements. $\boldsymbol{\Omega}$ has rank one, and the residual vector moves in a single direction. That is what three measurements and two estimated states leave: one independent residual dimension, the redundancy $m - n = 1$.

Residual variance is not measurement variance. Each meter has error variance $\sigma_i^2 = 4\,\mathrm{MW}^2$, yet the residual variances are 0.67 and $2.67\,\mathrm{MW}^2$. The estimator adjusts the state to fit the measurements, and part of every measurement error is absorbed into the estimated state: by equation (30.20) the absorbed part of measurement i is $(\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}})_{ii}$, $3.33\,\mathrm{MW}^2$ for L12 and L23 and $1.33\,\mathrm{MW}^2$ for L13. Only the component of the error that no state adjustment can explain remains in the residual. A measurement on which the estimate leans heavily, one with high *leverage*, keeps little of its own error in its residual, and the L12 and L23 meters are of that kind here. Dividing every residual by the original measurement standard deviation is therefore not the appropriate normalization, as Section 30.3.5 shows.

30.3.3 The global chi-square test

At the optimum the WLS objective is the weighted sum of squared residuals,

$$J(\hat{\mathbf{x}}) = \mathbf{e}^{\mathsf{T}}\mathbf{W}\,\mathbf{e} = \sum_{i=1}^{m} \frac{e_i^2}{\sigma_i^2}, \tag{30.22}$$

the second form for independent measurements. For the example,

$$J(\hat{\mathbf{x}}) = \frac{(-2.6667)^2}{4} + \frac{(-2.6667)^2}{4} + \frac{5.3333^2}{4} = 1.7778 + 1.7778 + 7.1111 = 10.667.$$

Its distribution under the correct model follows from equation (30.19). Chapter 22 showed, in Section 22.2.4, that for a residual vector with a known positive-definite covariance $\boldsymbol{\Sigma}_e$ of dimension m, the statistic $\mathbf{e}^{\mathsf{T}}\boldsymbol{\Sigma}_e^{-1}\mathbf{e}$ is χ^2_m. That result cannot be applied with $\boldsymbol{\Omega}$, which is singular, and J weights the residuals by $\mathbf{W}$, not by an inverse of $\boldsymbol{\Omega}$. The distribution is obtained by whitening the measurement errors instead. Let $\mathbf{R} = \mathbf{A}\mathbf{A}^{\mathsf{T}}$ be a Cholesky factorization ($\mathbf{A} = 2\,\mathbf{I}\,\mathrm{MW}$ in the example), $\mathbf{w} = \mathbf{A}^{-1}\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$ the whitened errors and $\tilde{\mathbf{H}} = \mathbf{A}^{-1}\mathbf{H}$ the whitened measurement matrix. Then $\mathbf{W} = \mathbf{A}^{-\mathsf{T}}\mathbf{A}^{-1}$, $\mathbf{G} = \tilde{\mathbf{H}}^{\mathsf{T}}\tilde{\mathbf{H}}$, and the whitened residual is

$$\mathbf{A}^{-1}\mathbf{e} = \mathbf{A}^{-1}\mathbf{S}\mathbf{A}\,\mathbf{w} = \left(\mathbf{I} - \tilde{\mathbf{H}}(\tilde{\mathbf{H}}^{\mathsf{T}}\tilde{\mathbf{H}})^{-1}\tilde{\mathbf{H}}^{\mathsf{T}}\right)\mathbf{w} = (\mathbf{I} - \boldsymbol{\Pi})\,\mathbf{w},$$

where $\boldsymbol{\Pi}$ is the orthogonal projection onto the n-dimensional column space of $\tilde{\mathbf{H}}$. The matrix $\mathbf{I} - \boldsymbol{\Pi}$ is symmetric and idempotent, the orthogonal projection onto the $(m-n)$-dimensional complement, and

$$J(\hat{\mathbf{x}}) = (\mathbf{A}^{-1}\mathbf{e})^{\mathsf{T}}(\mathbf{A}^{-1}\mathbf{e}) = \mathbf{w}^{\mathsf{T}}(\mathbf{I} - \boldsymbol{\Pi})\,\mathbf{w}.$$

An orthogonal projection of rank $m-n$ can be written $\mathbf{Q}\,\mathrm{diag}(1,\dots,1,0,\dots,0)\mathbf{Q}^{\mathsf{T}}$ with $\mathbf{Q}$ orthogonal and $m-n$ ones. The rotated vector $\mathbf{v}=\mathbf{Q}^{\mathsf{T}}\mathbf{w}$ is again $\mathcal{N}(\mathbf{0},\mathbf{I}_m)$, so $J(\hat{\mathbf{x}})=\sum_{j=1}^{m-n} v_j^2$, a sum of $m-n$ squared independent standard normal variables:

$$J(\hat{\mathbf{x}}) \sim \chi^2_{m-n} \quad \text{under the correct linear Gaussian model,} \tag{30.23}$$

provided the m measurements give a full-rank estimator of the n states and $\mathbf{R}$ is correctly specified. Its mean is $m-n=\mathrm{tr}\,\mathbf{S}$: each estimated state uses up one degree of freedom of the measurements.

In the example $m=3$ and $n=2$, so $J\sim\chi^2_1$. The 99th percentile of χ^2_1 is $\chi^2_{1,0.99}=6.635$, and the rule

$$J(\hat{\mathbf{x}}) > \chi^2_{m-n,\,1-\alpha} \tag{30.24}$$

declares the measurement set inconsistent at significance level α. Here $10.667>6.635$: the measurements fail the global consistency test at the 1 % level. The upper-tail probability is $\Pr(\chi^2_1>10.667)=2[1-\Phi(\sqrt{10.667})]=2[1-\Phi(3.266)]=2\times 0.000545=0.00109$, using that a χ^2_1 variable is the square of a standard normal one. Figure 30.3 draws that tail probability. The horizontal axis is the statistic J, the vertical axis the probability $\Pr(\chi^2_1>J)$ on a logarithmic scale; the dashed lines mark the 1 % level and its threshold 6.635, and the dot is the example.

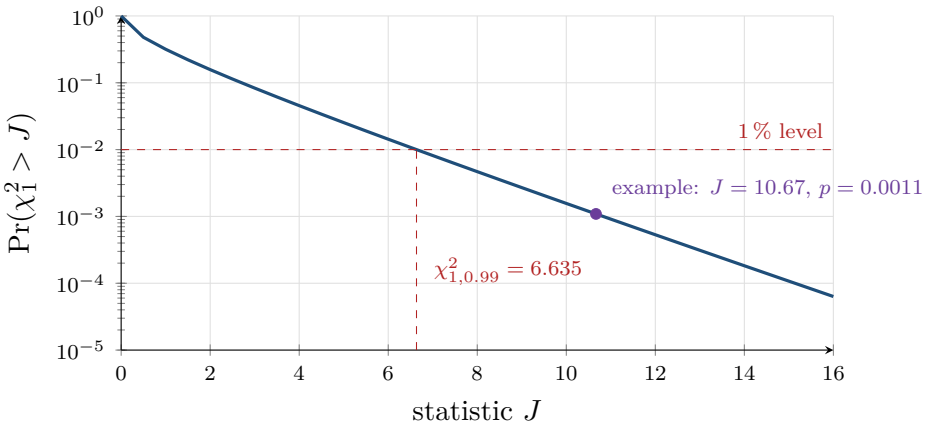


Figure 30.3: Upper-tail probability of the chi-square distribution with one degree of freedom, $\Pr(\chi^2_1>J)=2[1-\Phi(\sqrt{J})]$, on a logarithmic scale. The global test rejects the measurement set at the 1 % level when $J>6.635$; the three-bus example gives $J=10.667$ and a tail probability of 0.00109.

The tail probability means that a statistic this large would occur in about 0.11 % of cases if the model were correct and the errors unbiased and Gaussian with the stated variances. It does not mean that a particular meter is faulty with probability 0.99891: a frequentist tail probability is a statement about the data under one model, not a posterior probability of a fault hypothesis. The posterior requires the alternatives and their priors, which Section 30.6 supplies.

30.3.4 Interpreting the test

The global chi-square test asks whether the aggregate disagreement among the measurements is unusually large relative to the uncertainty that the measurement model predicts. It does not identify the source. A large J is consistent with a gross error of a meter, with a calibration bias of a meter, with an incorrect network topology, with incorrect line parameters, with measurements taken at different instants and therefore not describing one operating state, with an inadequate physical model (for example a DC approximation applied where losses or voltage deviations matter), and with an incorrectly specified noise covariance, since J scales with $\mathbf{W}$. The correct

output of the test is therefore a statement about the measurement set: *the measurements are inconsistent with the assumed model at the selected significance level.* Which of the explanations holds is a further question.

Principle 30.1. *A residual test detects an inconsistency between the measurements and the model; it does not identify its cause. Gross sensor errors, biases, topology errors, parameter errors, timing errors and misstated uncertainties all inflate the same statistic.*

30.3.5 Normalized residuals and the largest normalized residual test

Identification needs a statistic per measurement. Under the correct model the residual e_i is Gaussian with mean zero and variance Ω_{ii}, so the *normalized residual*

$$e_i^{\mathrm{N}} = \frac{e_i}{\sqrt{\Omega_{ii}}} \tag{30.25}$$

is standard normal, dimensionless. It is the standardized residual z of Chapter 22 with the right variance: the variance of the residual, not of the measurement. For the example, $\sqrt{\Omega_{11}} = \sqrt{\Omega_{22}} = \sqrt{0.6667} = 0.8165\,\mathrm{MW}$ and $\sqrt{\Omega_{33}} = \sqrt{2.6667} = 1.6330\,\mathrm{MW}$, and

$$e_{12}^{\mathrm{N}} = \frac{-2.6667}{0.8165} = -3.266, \qquad e_{23}^{\mathrm{N}} = -3.266, \qquad e_{13}^{\mathrm{N}} = \frac{5.3333}{1.6330} = +3.266.$$

The *largest normalized residual test* compares the largest $|e_i^{\mathrm{N}}|$ with a threshold, commonly 3; if it is exceeded, the corresponding measurement is declared bad, removed (or corrected), and the estimation and the test are repeated until no normalized residual exceeds the threshold. It is a standard result that, for a single gross error in a measurement that is neither critical nor a member of a critical pair (Section 30.3.6), and with a correct model, the largest normalized residual falls in expectation on the erroneous measurement; the size of the error is then estimated by

$$\hat{b}_i = \frac{R_{ii}}{\Omega_{ii}}\, e_i = \frac{e_i}{S_{ii}}, \tag{30.26}$$

since a gross error b on measurement i adds $S_{ii}b$ to e_i, and $\Omega_{ii} = S_{ii}R_{ii}$ for diagonal $\mathbf{R}$.

Table 30.3 collects the quantities per meter. The columns are the residual, the standard deviation of the measurement error, the standard deviation of the residual, the residual divided by the measurement standard deviation, the normalized residual, and the gross error that equation (30.26) would attribute to that meter alone. Figure 30.4 draws the second, third, fourth and fifth columns. In panel (a) the bars are standard deviations in MW for each meter, the measurement error in gray and the residual in blue; in panel (b) the bars are the residuals scaled two ways, by the measurement standard deviation in gray and by the residual standard deviation in blue, and the dashed lines are the LNR threshold ± 3.

Table 30.3: Residuals, their standard deviations, the two scalings and the gross error attributable to each meter, for the observation $\mathbf{y} = (50, 50, 58)\,\mathrm{MW}$.

Meter	e_i (MW)	σ_i (MW)	$\sqrt{\Omega_{ii}}$ (MW)	e_i/σ_i	e_i^{N}	$\hat{b}_i$ (MW)
L12	-2.667	2	0.816	-1.333	-3.266	-16
L23	-2.667	2	0.816	-1.333	-3.266	-16
L13	$+5.333$	2	1.633	$+2.667$	$+3.266$	$+8$

Two things appear. Scaled by the measurement standard deviation, the residuals single out L13 at 2.67 and keep every value below 3; that scaling is wrong, because it ignores the

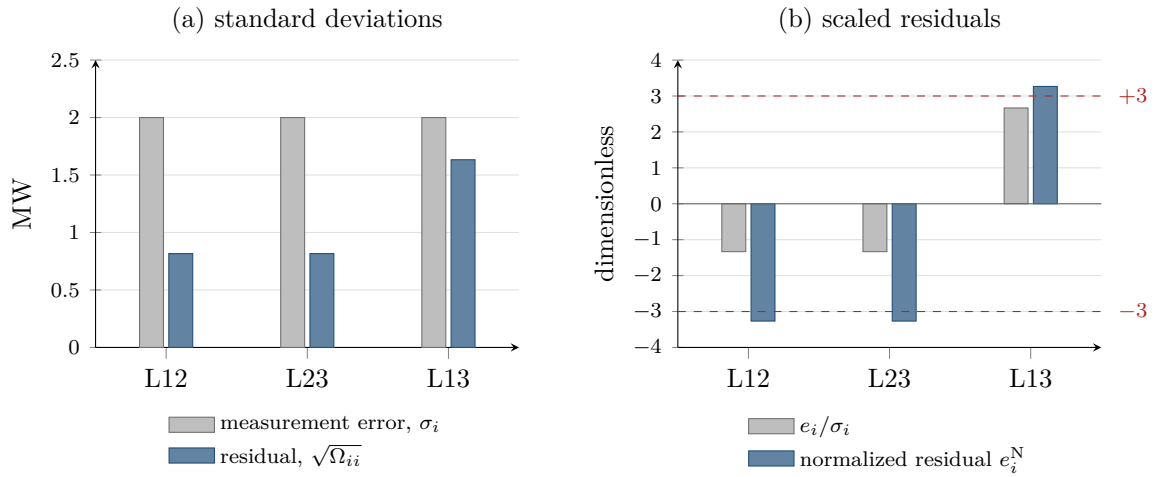


Figure 30.4: Residual analysis of the three-bus example. (a) The residual standard deviation $\sqrt{\Omega_{ii}}$ is smaller than the measurement standard deviation $\sigma_i = 2\,\mathrm{MW}$, and differs between meters. (b) Dividing by σ_i makes L13 look like the outlier and keeps every value below 3; dividing by $\sqrt{\Omega_{ii}}$ gives three normalized residuals of equal magnitude, 3.266, all beyond the LNR threshold.

variance absorbed by the estimate, and it understates every residual. Correctly normalized, the three magnitudes are identical, 3.266, and all three exceed the threshold. Although L13 has the largest raw residual, the largest normalized residual test does not identify it. The test has detected the inconsistency and stops there.

Common mistake

Normalizing each residual by the standard deviation of its measurement, e_i/σ_i, or reading the largest raw residual as the bad measurement. The residual of a measurement with high leverage is small even when that measurement is wrong; the normalized residual uses the residual's own variance Ω_{ii}, and the raw residual depends on the measurement configuration as much as on the error.

30.3.6 Why the normalized residuals tie

The model has one degree of redundancy, and every residual is generated by the same underlying inconsistency. The network physics imposes one relation among the three flows. Around the loop $1 \to 2 \to 3 \to 1$ the angle differences sum to zero, $(\theta_1 - \theta_2) + (\theta_2 - \theta_3) - (\theta_1 - \theta_3) = 0$; replacing each difference by $x_{ij}P_{ij}/S_{\mathrm{base}}$ from equation (30.4) gives $0.1\,P_{12} + 0.1\,P_{23} - 0.2\,P_{13} = 0$, that is,

$$P_{12} + P_{23} - 2\,P_{13} = 0, \qquad \text{or} \quad \mathbf{a}^{\mathsf{T}}\mathbf{H} = \mathbf{0} \text{ with } \mathbf{a} = \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix}. \tag{30.27}$$

Applied to the measurements it defines the *consistency (parity) residual* of Chapter 29,

$$e_{\mathrm{c}} = \mathbf{a}^{\mathsf{T}}\mathbf{y} = y_{12} + y_{23} - 2\,y_{13}, \tag{30.28}$$

which for the observations is $50 + 50 - 2 \times 58 = -16\,\mathrm{MW}$. It contains no state, because $\mathbf{a}^{\mathsf{T}}\mathbf{H} = \mathbf{0}$, and under the correct model its variance is $\mathbf{a}^{\mathsf{T}}\mathbf{R}\mathbf{a} = 4(1 + 1 + 4) = 24\,\mathrm{MW}^2$. The residual sensitivity matrix of the example is $\mathbf{S} = \mathbf{a}\mathbf{a}^{\mathsf{T}}/6$, so

$$\mathbf{e} = \mathbf{S}\mathbf{y} = \mathbf{a}\,\frac{e_{\mathrm{c}}}{6}, \qquad \boldsymbol{\Omega} = \tfrac{2}{3}\,\mathbf{a}\mathbf{a}^{\mathsf{T}}, \qquad e_i^{\mathrm{N}} = \frac{a_i e_{\mathrm{c}}/6}{\sqrt{2/3}\,|a_i|} = \operatorname{sign}(a_i)\,\frac{e_{\mathrm{c}}}{\sqrt{24}}\,:$$

the single inconsistency e_c is distributed among the three residuals in proportion to $\mathbf{a}$, every normalized residual has magnitude $|e_c|/\sqrt{24} = 16/4.899 = 3.266$, and $J = e_c^2/24 = 256/24 = 10.667$ is its square. With one degree of redundancy the normalized residuals always tie, whatever the measurements and the weights.

The tie expresses an ambiguity in the data. Every column of $\mathbf{S}$ is a multiple of $\mathbf{a}$, so a gross error on any single meter produces a residual vector in the same direction, and each of the three single-error explanations of Table 30.4 fits the observations exactly. Each row assumes one meter wrong by the amount of equation (30.26), corrects it, and gives the state and the L13 flow that the corrected measurements imply. The data cannot tell the rows apart, yet two of them put the L13 flow at 58 MW and one at 50 MW. Pairs of measurements like these, whose simultaneous removal makes the state unobservable, are called *critical pairs*: removing L12 and L23 leaves only L13, of rank one, and likewise for every other pair. Bad data in a measurement of a critical pair can be detected but not identified.

Table 30.4: The three single-error explanations of the observation $\mathbf{y} = (50, 50, 58)$ MW. Each corrects one meter by its estimated gross error and fits the network model exactly.

Explanation	$\hat{b}_i$ (MW)	Corrected $\mathbf{y}$ (MW)	θ_1 (rad)	θ_2 (rad)	P_{13} (MW)
L12 meter reads low	-16	$(66, 50, 58)$	0.116	0.050	58
L23 meter reads low	-16	$(50, 66, 58)$	0.116	0.066	58
L13 meter reads high	$+8$	$(50, 50, 50)$	0.100	0.050	50

30.3.7 Removing the suspicious measurement

Suppose L13 is judged faulty and its measurement removed. The remaining measurements $y_{12} = 50$ and $y_{23} = 50$ MW form measurement set B of Chapter 29, and they determine the state exactly: $\hat{\theta}_2 = y_{23}/1000 = 0.05$ rad and $\hat{\theta}_1 = (y_{12} + y_{23})/1000 = 0.10$ rad, so that $\hat{P}_{13} = 500 \times 0.10 = 50$ MW, the original base case. The fit is perfect, and that perfection is not evidence. With $m = n = 2$ there is no redundancy left, $\mathbf{S} = \mathbf{0}$, and any two readings are fitted exactly, a wrong one included. Removing L12 instead and keeping L23 and L13 fits just as perfectly, at the first row of Table 30.4: $\theta_1 = 58/500 = 0.116$ rad, $\theta_2 = 0.05$ rad and an inferred L12 flow of 66 MW. A perfect fit after a measurement is removed is not independent proof that the removed measurement was the faulty one; the result rests on the choice to trust the two that were kept.

30.4 Critical measurements, interacting bad data and topology errors

Residual tests can only see errors that leave a trace in the residuals. This section identifies three situations in which they do not, or in which the trace is misleading: errors in critical measurements, several errors that conform to the network model, and errors of the assumed topology rather than of any sensor.

30.4.1 Critical measurements

A measurement is *critical* if removing it makes the state unobservable. Chapter 29 showed that the L12 and L23 measurements together determine the two angles, and that removing either leaves one measurement for two states. In the two-measurement configuration both measurements are therefore critical. A critical measurement has no residual redundancy that

could reveal its error under the measurement model: its row of $\mathbf{S}$ is zero, its residual is identically zero and $\Omega_{ii} = 0$, so its normalized residual is undefined. A bias in it shifts the estimated state without producing any inconsistency. If the L12 meter of set B reads 55 MW instead of 50, the estimate becomes $\hat{\theta}_1 = (55 + 50)/1000 = 0.105$ rad and $\hat{\theta}_2 = 0.05$ rad, the inferred L13 flow 52.5 MW, and both residuals remain zero. The three-meter set C has no critical measurement, since removing any one meter leaves two independent rows, but every pair of its meters is a critical pair, which is why the previous section could detect but not identify. Observability alone is not enough for reliable bad-data processing: detection needs redundancy, and identification needs more.

30.4.2 Multiple interacting errors

Suppose that several meters are biased at once, possibly through a common cause such as a shared calibration reference or a shared transducer supply, the common-cause coupling of Chapter 13 of Volume I applied to instruments. The measurement model becomes

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{b} + \boldsymbol{\varepsilon}, \qquad \mathbf{b} = \begin{bmatrix} b_{12} \\ b_{23} \\ b_{13} \end{bmatrix}, \tag{30.29}$$

with $\mathbf{b}$ the vector of gross errors (biases) in MW. The consistency residual of equation (30.28) receives the bias contribution

$$\mathbf{a}^{\mathsf{T}}\mathbf{b} = b_{12} + b_{23} - 2\,b_{13}.$$

For $b_{12} = 4$, $b_{23} = 0$ and $b_{13} = 2$ MW it is $4 + 0 - 2 \times 2 = 0$. The biases are not zero, but the consistency residual, and with it every residual, does not detect them. Starting from the base case, the meters report $(54, 50, 52)$ MW; WLS returns $\hat{\theta}_1 = 0.104$ and $\hat{\theta}_2 = 0.050$ rad, the estimated flows equal the readings, all residuals are zero, and $J = 0$. The estimate of the L13 flow is 2 MW high, with nothing to show it. The corrupted measurements lie exactly on the set of measurement vectors the model can produce. Bad data of this kind are *conforming*: the errors are mutually consistent with the network model, and here fully so, which makes them *undetectable* by any residual test on these measurements. Errors in several measurements are *interacting* when their residuals are strongly correlated, $|\Omega_{ij}|/\sqrt{\Omega_{ii}\Omega_{jj}}$ near one, as every pair is in this example. Interacting errors that do not conform are detectable, but may be too small to cross the threshold: $\mathbf{b} = (4, 0, -2)$ MW gives $\mathbf{a}^{\mathsf{T}}\mathbf{b} = 8$ MW and $J = 8^2/24 = 2.67$, below 6.635.

30.4.3 The general condition for undetectable bad data

With the bias of equation (30.29), the residual is, by the argument of equation (30.19),

$$\mathbf{e} = \mathbf{S}(\mathbf{b} + \boldsymbol{\varepsilon}) = \mathbf{S}\mathbf{b} + \mathbf{S}\boldsymbol{\varepsilon},$$

and the bias leaves a trace in the residual only through $\mathbf{S}\mathbf{b}$. The bias is undetectable when $\mathbf{S}\mathbf{b} = \mathbf{0}$, that is, by equation (30.18), when

$$\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\,\mathbf{b} = \mathbf{b}. \tag{30.30}$$

This condition holds exactly when $\mathbf{b}$ lies in the column space of $\mathbf{H}$. If $\mathbf{b} = \mathbf{H}\mathbf{c}$ for some vector $\mathbf{c}$ of angle shifts, then $\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}\mathbf{c} = \mathbf{H}\mathbf{G}^{-1}\mathbf{G}\mathbf{c} = \mathbf{H}\mathbf{c} = \mathbf{b}$. Conversely, if equation (30.30) holds, then $\mathbf{b} = \mathbf{H}\mathbf{c}$ with $\mathbf{c} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{b}$. In that case the corrupted measurements are

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{H}\mathbf{c} + \boldsymbol{\varepsilon} = \mathbf{H}(\mathbf{x} + \mathbf{c}) + \boldsymbol{\varepsilon},$$

perfectly consistent with a different state, and the estimator returns $\hat{\mathbf{x}} = \mathbf{x} + \mathbf{c} + \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\boldsymbol{\varepsilon}$, shifted by $\mathbf{c}$, with residuals no larger than usual. The biases of the previous subsection are of this

form, $(4, 0, 2) = \mathbf{H}\,(0.004, 0)^{\mathsf{T}}$, a shift of θ_1 by 0.004 rad. The algebra is the same for accidental coherent errors and for measurement data manipulated deliberately, which the literature calls false data injection; and the condition is a rank property of the kind Chapter 22 stated for fault signatures (Section 22.5.2): a bias whose signature on the measurements coincides with that of a state change cannot be separated from it.

Principle 30.2. *Residual consistency alone cannot distinguish every state error from every measurement error. A combination of measurement errors that lies in the column space of* $\mathbf{H}$ *is absorbed into the estimated state; only independent information (authenticated or diverse sensors, further model constraints, temporal consistency) can reveal it.*

30.4.4 A sensor fault is not the only explanation

Everything so far has assumed that the topology is correct. Suppose instead that a breaker status is wrong: line L13 is physically disconnected, but the network model still treats it as connected. The estimator then forms its measurement equations from the wrong network, with a wrong susceptance matrix and a wrong $\mathbf{H}$. This is a *topology error*, a fault of the model rather than of any sensor, and the wrong model can produce large residuals even when every meter is working correctly. A large residual is evidence that the measurements and the model disagree; it does not say which of the two is wrong.

30.4.5 The L13-disconnected example

Let L13 be open while the system continues to transfer 100 MW from bus 1 to bus 3 through L12 and L23. The physical flows are $P_{12} = 100$ MW, $P_{23} = 100$ MW and $P_{13} = 0$, and let the meters report them exactly, $\mathbf{y} = (100, 100, 0)$ MW. The estimator still uses the closed-triangle matrix $\mathbf{H}$ of equation (30.15), which imposes the loop relation (30.27). For these measurements the consistency residual is $100 + 100 - 2 \times 0 = 200$ MW, against a standard deviation of $\sqrt{24} = 4.9$ MW. The WLS solution follows as before: $\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} = (25\,000, 0)$, so $\hat{\theta}_2 = \hat{\theta}_1/2$ and $187\,500\,\hat{\theta}_1 = 25\,000$, giving $\hat{\theta}_1 = 0.13333$ and $\hat{\theta}_2 = 0.06667$ rad and estimated flows of 66.67 MW on every line. The residuals are $(33.33, 33.33, -66.67)$ MW, $J = 200^2/24 = 1666.7$, and the three normalized residuals have magnitude $\sqrt{1666.7} = 40.8$. The inconsistency is enormous relative to 2 MW meter errors, and no meter is faulty.

Figure 30.5 compares the two networks. Panel (a) is the physical network, with L13 dashed and open and the three meter readings beside the lines; panel (b) is the network the estimator assumes, with the estimated flows and, in parentheses, the residuals. Table 30.5 gives the results of three ways of processing the same readings: the closed model with all meters, the closed model with the L13 meter discarded as bad data, and the correct open model. The columns are the estimated angles, the estimated flows and the chi-square statistic.

Table 30.5: Three ways of processing the readings $\mathbf{y} = (100, 100, 0)$ MW taken with L13 open. Angles in rad, flows in MW.

Processing	$\hat{\theta}_1$	$\hat{\theta}_2$	$\hat{P}_{12}$	$\hat{P}_{23}$	$\hat{P}_{13}$	J
Closed model, all meters	0.1333	0.0667	66.67	66.67	66.67	1666.7
Closed model, L13 meter discarded	0.2000	0.1000	100.00	100.00	100.00	0 (no redundancy)
Open model, all meters	0.2000	0.1000	100.00	100.00	0	0

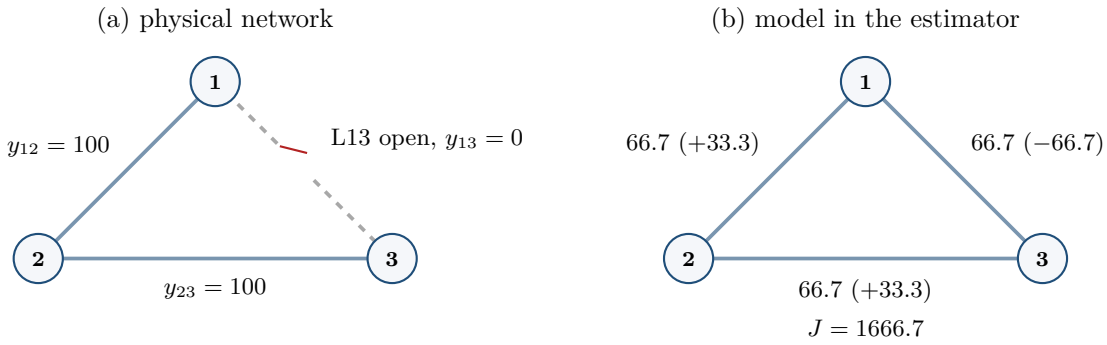


Figure 30.5: A topology error. (a) Line L13 is open and the meters report the true flows, in MW. (b) The estimator assumes L13 closed; its estimated flows and, in parentheses, the residuals are in MW. Every meter is correct, yet the chi-square statistic is 1666.7.

The second row shows how a topology error can be compounded by bad-data processing. The L13 meter has the largest raw residual and, like the others, a normalized residual of 40.8; discarding it as bad data leaves L12 and L23, which are fitted exactly. The estimated angles are then the true ones, but the model still contains L13 and assigns it $500 \times 0.2 = 100\,\text{MW}$, above its 80 MW rating: the estimator reports a spurious overload on a line that is not in service, and the only meter that told the truth has been discarded.

30.4.6 Correct topology restores physical consistency

With L13 disconnected, the flow on L13 is zero whatever the angles, and the L13 meter measures that zero. The measurement matrix of the open network has the third row $(0, 0)$,

$$\mathbf{H}_{\text{open}} = \begin{bmatrix} 1000 & -1000 \\ 0 & 1000 \\ 0 & 0 \end{bmatrix},$$

and the two remaining lines obey $P_{12} = 1000\,(\theta_1 - \theta_2)$ and $P_{23} = 1000\,\theta_2$. The solution is $\theta_2 = 100/1000 = 0.10$ and $\theta_1 = 0.10 + 100/1000 = 0.20\,\text{rad}$, so that $P_{12} = P_{23} = 100\,\text{MW}$ and $P_{13} = 0$, the third row of Table 30.5. All residuals vanish and $J = 0$; the redundancy is still one, now carried by the L13 meter, whose reading is a check on the open status of the line. The measurements are physically consistent with the open-line topology. A state estimation system must therefore not classify every large residual as a sensor fault; the fault may lie in the assumed topology.

30.4.7 Topology processing

Because a wrong breaker status changes $\mathbf{H}$, a practical transmission state-estimation workflow treats the topology as an input to be verified, in seven steps:

1. collect the statuses of breakers and switches;
2. construct the bus-branch or node-breaker network model from them;
3. check the network connectivity and the assignment of equipment to buses;
4. form the measurement equations using the inferred topology;
5. solve the state estimation problem;
6. analyze the residuals and other consistency checks;

7. investigate candidate topology errors when the fit is inadequate.

Incorrect breaker statuses, node connectivity or circuit assignments can produce widespread inconsistencies, often concentrated on the measurements around one substation. Topology identification is therefore closely coupled to state estimation; generalized state estimation includes the statuses of selected breakers among the unknowns, so that the same residual analysis applies to them.

30.5 Robust state estimation

The WLS estimate is the maximum likelihood estimate under Gaussian errors, and gross errors are not Gaussian. This section shows why WLS is sensitive to them, defines the least absolute value and Huber estimators, derives the iteratively reweighted least-squares algorithm that computes them, and asks whether limiting the influence of a measurement amounts to diagnosing it.

30.5.1 Sensitivity of WLS to gross errors

The WLS objective (30.9) penalizes each standardized residual quadratically. A large residual receives a penalty that grows with its square, and the estimator moves the state to reduce it. The estimate (30.12) is a linear function of the measurements, so the influence of one measurement on the estimate is unbounded. In the three-bus example with $y_{12} = y_{23} = 50\,\mathrm{MW}$, the normal equations give $\hat{\theta}_2 = \hat{\theta}_1/2$ and $187\,500\,\hat{\theta}_1 = (50\,000 + 500\,y_{13})/4$, so that

$$\hat{P}_{12} = \hat{P}_{23} = \hat{P}_{13} = \frac{100 + y_{13}}{3}\,\mathrm{MW}: \tag{30.31}$$

every MW of error on the L13 meter moves every estimated flow by a third of a MW, without limit. A single measurement with a very large error can substantially move the estimated state. *Robust estimators* replace the quadratic penalty by one that grows more slowly for large residuals, to reduce the influence of observations inconsistent with the assumed noise model.

30.5.2 Least absolute value estimation

The *least absolute value* (LAV) estimator minimizes the sum of absolute standardized residuals,

$$\min_{\mathbf{x}} \sum_{i=1}^{m} \frac{|y_i - h_i(\mathbf{x})|}{\sigma_i}, \tag{30.32}$$

so that large residuals receive a linear rather than a quadratic penalty. For linear measurement equations it is a linear program: write each residual as the difference of two nonnegative variables, $y_i - \mathbf{h}_i\mathbf{x} = s_i^+ - s_i^-$, and minimize $\sum_i (s_i^+ + s_i^-)/\sigma_i$ subject to these equalities and $s_i^{\pm} \geq 0$. An optimal basic solution of that program fits n of the measurements exactly and leaves the others to carry the residuals.

In the three-bus example every achievable residual vector satisfies $\mathbf{a}^{\mathsf{T}}\mathbf{e} = \mathbf{a}^{\mathsf{T}}\mathbf{y} = -16\,\mathrm{MW}$, because $\mathbf{a}^{\mathsf{T}}\mathbf{H} = \mathbf{0}$. With equal σ_i, LAV minimizes $|e_{12}| + |e_{23}| + |e_{13}|$ subject to $e_{12} + e_{23} - 2e_{13} = -16$. Placing the whole inconsistency on L13, $e_{13} = 8$, costs 8; placing it on L12 or on L23 costs 16. LAV fits L12 and L23 exactly, attributes 8 MW to L13, and estimates the base case, 50 MW on every line. That looks like an identification, but it is not one. If the readings had been $\mathbf{y} = (34, 50, 50)\,\mathrm{MW}$, with the L12 meter 16 MW low and the L13 meter correct, the consistency residual and the residual vector would be exactly the same, and LAV would again fit L12 and L23, attribute 8 MW to L13, and estimate $(34, 50, 42)\,\mathrm{MW}$ against true flows of 50 MW, wrong on L12 and L13. LAV chooses L13 because of its coefficient in the loop relation: placing the

inconsistency on meter i costs $|e_c|/|a_i|$, least for the meter with the largest $|a_i|$, which is also the meter of lowest leverage (1/3 on the diagonal of $\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}$, against 5/6 for L12 and L23). The choice reflects the geometry of the measurement set, not evidence against L13. LAV methods have further limitations: the solution can be nonunique when two choices cost the same, and it can be unstable when the redundancy is small or some measurements have high leverage.

30.5.3 Huber estimation

The Huber estimator combines the two penalties. For a standardized residual u, the Huber loss with threshold $\delta > 0$ is

$$\rho_\delta(u) = \begin{cases} \tfrac{1}{2}u^2, & |u| \le \delta, \\ \delta|u| - \tfrac{1}{2}\delta^2, & |u| > \delta, \end{cases} \tag{30.33}$$

quadratic near zero and linear in the tails, continuous with a continuous derivative at $|u| = \delta$; the constant $-\tfrac{1}{2}\delta^2$ joins the two pieces. The robust objective is

$$\min_{\mathbf{x}} \sum_{i=1}^{m} \rho_\delta\left(\frac{y_i - h_i(\mathbf{x})}{\sigma_i}\right). \tag{30.34}$$

Small residuals are treated as WLS treats them; large residuals receive a reduced influence. The threshold δ, dimensionless and in standard deviations, sets the transition between the quadratic and the linear regions, typically between 1 and 2. (The symbol δ is the Huber threshold in this section only.) Figure 30.6 compares the losses and the weights of the next subsection. In panel (a) the horizontal axis is the standardized residual u and the vertical axis the loss, for WLS ($u^2/2$), LAV ($|u|$) and Huber with $\delta = 1.5$; in panel (b) the vertical axis is the weight $w(u)$ that each estimator gives a residual, and the dots mark the converged weights of the example worked below.

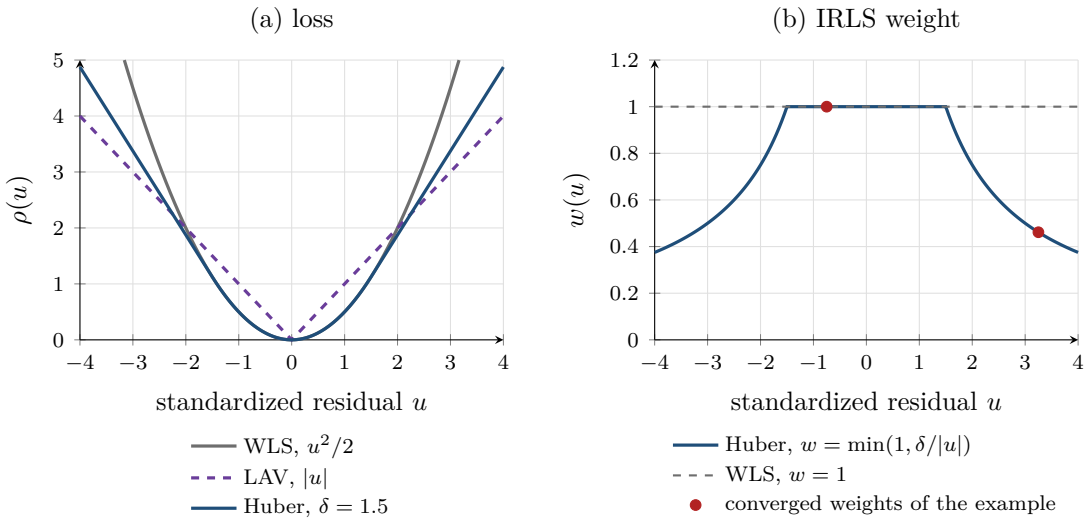


Figure 30.6: Losses and weights of three estimators as functions of the standardized residual u. (a) WLS penalizes quadratically, LAV linearly, and Huber quadratically for $|u| \le \delta = 1.5$ and linearly beyond, equation (30.33). (b) The IRLS weight, 1 for WLS and $\min(1, \delta/|u|)$ for Huber; the dots are the converged weights of L12 and L23 ($u = -0.75$, $w = 1$) and of L13 ($u = 3.25$, $w = 0.4615$) in Table 30.6.

30.5.4 Iteratively reweighted least squares

A common method for solving the robust problem is *iteratively reweighted least squares* (IRLS), which follows from the stationarity condition of equation (30.34). Write $u_i(\mathbf{x}) = [y_i - h_i(\mathbf{x})]/\sigma_i$ and $\psi = \rho_\delta'$, so that $\psi(u) = u$ for $|u| \le \delta$ and $\psi(u) = \delta\,\mathrm{sign}(u)$ otherwise. The gradient of the objective vanishes when

$$\sum_{i=1}^{m} \psi(u_i)\,\frac{1}{\sigma_i}\,\frac{\partial h_i}{\partial \mathbf{x}} = \mathbf{0}.$$

Writing $\psi(u_i) = w_i u_i$ with the weight $w_i = \psi(u_i)/u_i$, this is $\sum_i (w_i/\sigma_i^2)\,[y_i - h_i(\mathbf{x})]\,\mathbf{h}_i^{\mathsf{T}} = \mathbf{0}$, which for a linear model is

$$\mathbf{H}^{\mathsf{T}}\mathbf{W}_w(\mathbf{y} - \mathbf{H}\mathbf{x}) = \mathbf{0}, \qquad \mathbf{W}_w = \mathrm{diag}\left(\frac{w_i}{\sigma_i^2}\right):$$

the normal equations of a WLS problem whose weights are the original ones multiplied by w_i. For the Huber loss

$$w_i = \begin{cases} 1, & |u_i| \le \delta, \\ \delta/|u_i|, & |u_i| > \delta. \end{cases} \tag{30.35}$$

The weights depend on the solution, which gives the algorithm:

1. set $w_i = 1$ for every measurement, so that the first iterate is the WLS estimate;
2. solve the WLS problem with weights w_i/σ_i^2, giving $\mathbf{x}^{(k)}$;
3. compute the standardized residuals $u_i = [y_i - h_i(\mathbf{x}^{(k)})]/\sigma_i$;
4. update the weights by equation (30.35);
5. repeat steps 2 to 4 until the change of $\mathbf{x}^{(k)}$ is below a tolerance.

For the Huber loss each iteration does not increase the objective, and the iteration converges to its minimum.

Table 30.6 runs the algorithm on the three-bus example, $\mathbf{y} = (50, 50, 58)$ MW, $\sigma_i = 2$ MW and $\delta = 1.5$. Each row is an iteration; the columns are the weight of the L13 meter used in that iteration (the weights of L12 and L23 stay at one, since their $|u_i|$ never exceeds 1.5), the estimated angle $\hat\theta_1$, the estimated flows, which remain equal on the three lines, the L13 residual and its standardized value, and the Huber objective. One step of the arithmetic: the WLS iterate leaves $u_{13} = 5.333/2 = 2.667 > 1.5$, so $w_{13} = 1.5/2.667 = 0.5625$ in iteration 1.

Table 30.6: IRLS iterations of the Huber estimator ($\delta = 1.5$) for the three-bus example. Iteration 0 is the WLS estimate.

Iteration k	w_{13}	$\hat\theta_1$ (rad)	$\hat P_{ij}$ (MW)	e_{13} (MW)	u_{13}	Objective
0	1	0.105333	52.667	5.333	2.667	4.6528
1	0.5625	0.103512	51.756	6.244	3.122	4.3289
2	0.4805	0.103099	51.550	6.450	3.225	4.3131
3	0.4651	0.103019	51.509	6.491	3.245	4.3125
4	0.4622	0.103004	51.502	6.498	3.249	4.3125
converged	0.4615	0.103000	51.500	6.500	3.250	4.3125

The converged solution can be checked directly. At the minimum, $\psi(u_i)$ must be proportional to a_i, the entries of the loop vector, because the achievable residuals satisfy $\mathbf{a}^{\mathsf{T}}\mathbf{e} = -16$. With L13 in the linear region, $\psi(u_{13}) = 1.5$, so $\psi(u_{12}) = \psi(u_{23}) = -0.75$, that is $u_{12} = u_{23} = -0.75$ and $e_{12} = e_{23} = -1.5$ MW; the constraint then gives $e_{13} = (16-3)/2 = 6.5$ MW and $u_{13} = 3.25 > 1.5$,

consistent. The estimated flows are 51.5 MW, between the WLS estimate of 52.67 and the base case of 50 MW, and the objective is $2 \times \frac{1}{2}(0.75)^2 + 1.5 \times 3.25 - 1.125 = 4.3125$. The influence of the L13 reading is now bounded: by the same argument, every reading $y_{13} \geq 54.5$ MW (the point where the WLS residual of L13 reaches $u_{13} = 1.5$) leaves the Huber estimate at 51.5 MW, whereas the WLS estimate of equation (30.31) keeps rising.

30.5.5 Robust estimation and diagnosis

Robust estimation limits the influence of some large residuals under its chosen loss; it does not necessarily solve the diagnosis problem. Three cases show the limits. If a real topology error produces many large residuals, as in Section 30.4.5, downweighting them conceals the evidence that the model is inadequate, and the estimator returns a state for a network that does not exist. If the bad measurements are conforming, lying in the column space of the model as in Section 30.4.3, a robust estimator sees no large residual and does not react at all. And, as the LAV and Huber solutions of the example show, a robust estimator decides which measurement to downweight by its loss function and the leverage of the measurements, which can pick the wrong measurement when the residuals cannot separate the alternatives. Robust estimation is therefore one tool within a larger inference and model-validation workflow.

Common mistake

Reading the measurement that a robust estimator downweights as the identified bad measurement. In the three-bus example both LAV and Huber downweight L13 whichever meter is actually wrong, because L13 has the largest coefficient in the only consistency relation; the choice reflects the geometry of the measurement set, not evidence against that meter.

30.6 Bayesian state estimation and reliability decisions

A point estimate and a test statistic leave the competing explanations of an inconsistency unweighed. This section turns the state estimate into a posterior distribution, assigns posterior probabilities to the explanations as Chapter 22 did for the pump, assembles the diagnostic workflow, computes the estimator's response to the suspicious reading over a range of values, and carries the result into the overload probability of a line.

30.6.1 From point estimates to posterior distributions

Classical WLS produces a state estimate $\hat{\mathbf{x}}$ and, under its assumptions, the covariance $\mathbf{G}^{-1}$. Bayesian state estimation starts instead from Bayes' theorem as Chapter 8 of Volume I stated it, posterior proportional to likelihood times prior, applied to the state:

$$p(\mathbf{x} \mid \mathbf{y}) \propto p(\mathbf{y} \mid \mathbf{x})\, p(\mathbf{x}), \tag{30.36}$$

where $p(\mathbf{x})$ is the prior distribution of the state (for example from a forecast or from the previous estimate), $p(\mathbf{y} \mid \mathbf{x})$ the measurement likelihood of equation (30.7), and $p(\mathbf{x} \mid \mathbf{y})$ the posterior distribution of the state. Let the prior be Gaussian, $\mathbf{x} \sim \mathcal{N}(\mathbf{x}_0, \boldsymbol{\Sigma}_0)$, and the measurements linear and Gaussian, $\mathbf{y} \mid \mathbf{x} \sim \mathcal{N}(\mathbf{H}\mathbf{x}, \mathbf{R})$. Minus twice the logarithm of the product is, up to a constant,

$$(\mathbf{y}-\mathbf{H}\mathbf{x})^{\mathsf{T}}\mathbf{W}(\mathbf{y}-\mathbf{H}\mathbf{x})+(\mathbf{x}-\mathbf{x}_0)^{\mathsf{T}}\boldsymbol{\Sigma}_0^{-1}(\mathbf{x}-\mathbf{x}_0) = \mathbf{x}^{\mathsf{T}}(\mathbf{G}+\boldsymbol{\Sigma}_0^{-1})\mathbf{x} - 2\,\mathbf{x}^{\mathsf{T}}(\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}+\boldsymbol{\Sigma}_0^{-1}\mathbf{x}_0)+\text{const},$$

a quadratic in $\mathbf{x}$. Completing the square identifies a Gaussian posterior with covariance and mean

$$\boldsymbol{\Sigma}_{\text{post}} = (\boldsymbol{\Sigma}_0^{-1} + \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H})^{-1}, \qquad \bar{\mathbf{x}} = \boldsymbol{\Sigma}_{\text{post}}(\boldsymbol{\Sigma}_0^{-1}\mathbf{x}_0 + \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}). \tag{30.37}$$

The posterior precision is the prior precision plus the gain matrix: the prior and the measurements add information. A diffuse prior, $\boldsymbol{\Sigma}_0^{-1} \to \mathbf{0}$, returns the WLS estimate and its covariance $\mathbf{G}^{-1}$ when the measurements make the state observable.

A proper prior also regularizes directions the measurements leave unobservable, but it does not make them observable. With only the L12 meter (measurement set A of Chapter 29), $\mathbf{G} = 250\,000 \left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right] \text{rad}^{-2}$ is singular, and WLS has no solution. With a prior standard deviation of 0.01 rad on each angle, $\boldsymbol{\Sigma}_0^{-1} = 10^4\,\mathbf{I}$, the posterior precision has eigenvalue 510 000 along the direction $(1,-1)/\sqrt{2}$ and 10 000 along $(1,1)/\sqrt{2}$: the posterior standard deviation is 0.0014 rad along the angle difference that L12 measures, and 0.01 rad, exactly the prior, along the common shift of both angles that it cannot see. The posterior is proper in both directions, but in the second all the information came from the prior. A Bayesian estimator must keep the information supplied by the data distinct from the information imposed by prior assumptions.

30.6.2 Competing diagnostic hypotheses

When the measurements are inconsistent, the explanations of Section 30.3.4 can be written as hypotheses, each with its own measurement model. For the three-bus example:

$$
\begin{aligned}
&H_0 : \text{healthy measurements and correct topology},\\
&H_{\mathrm{S}} : \text{bias of one meter, } H_{\mathrm{S},12},\ H_{\mathrm{S},23} \text{ or } H_{\mathrm{S},13},\\
&H_{\mathrm{T}} : \text{L13 status incorrect (line open, modeled closed)},\\
&H_{\mathrm{P}} : \text{line-parameter error}.
\end{aligned}
$$

Each specifies a likelihood $p(\mathbf{y} \mid \mathbf{x}, H_i)$, and, as in equation (22.23) of Chapter 22, the posterior probability of a hypothesis is

$$
\Pr(H_i \mid \mathbf{y}) = \frac{\Pr(H_i)\, p(\mathbf{y} \mid H_i)}{\sum_j \Pr(H_j)\, p(\mathbf{y} \mid H_j)}, \qquad p(\mathbf{y} \mid H_i) = \int p(\mathbf{y} \mid \mathbf{x}, H_i)\, p(\mathbf{x} \mid H_i)\, \mathrm{d}\mathbf{x}, \tag{30.38}
$$

where $p(\mathbf{y} \mid H_i)$ is the model evidence (marginal likelihood) of Chapter 8, the likelihood averaged over the unknown state and over any parameter the hypothesis introduces. This is more informative than selecting the hypothesis with the smallest fitted residual: a hypothesis with additional free parameters, such as an unknown bias, can fit almost any measurement set, whereas the evidence averages its fit over the prior of those parameters and so penalizes the excess flexibility, provided the priors are defensible.

The example can be computed. Take a prior on the state that is flat over a wide range and the same under every hypothesis, and give each bias hypothesis a Gaussian prior on the size of the bias, $b_i \sim \mathcal{N}(0, \tau^2)$ with $\tau = 10$ MW (illustrative), which raises the error variance of meter i from 4 to 104 MW^2. For the hypotheses that keep the closed topology, the evidence reduces to the density of the consistency residual $e_{\mathrm{c}} = \mathbf{a}^{\mathsf{T}}\mathbf{y}$. Changing variables from $\mathbf{y}$ to e_{c} and two coordinates that determine $\mathbf{x}$, integration over the flat prior removes the latter, and the Jacobian of the change is common to every hypothesis sharing $\mathbf{H}$. Since $e_{\mathrm{c}} = \mathbf{a}^{\mathsf{T}}\mathbf{b} + \mathbf{a}^{\mathsf{T}}\boldsymbol{\varepsilon}$ does not depend on the state, it is $\mathcal{N}(0, 24)$ under H_0 and $\mathcal{N}(0, 24 + a_i^2\tau^2)$ under $H_{\mathrm{S},i}$: variance 124 MW^2 for a bias on L12 or L23 and $24 + 4 \times 100 = 424$ MW^2 for a bias on L13. Under H_{T} the L13 meter measures the flow of an open line, which is zero, and a reading of 58 MW lies 29 standard deviations away: its evidence is smaller than that of H_0 by a factor of order 10^{-180}. The parameter hypothesis H_{P} is discussed below. Table 30.7 collects the computation with illustrative priors 0.90 for H_0, 0.02 for each meter bias and 0.04 for the topology error. The columns are the prior, the variance of e_{c}, its density at $e_{\mathrm{c}} = -16$ MW, the ratio of that density to the one under H_0, and the posterior. One row of the arithmetic, $H_{\mathrm{S},13}$: $\sqrt{24/424} \times \exp(256/48 - 256/848) = 0.2379 \times \mathrm{e}^{5.0314} = 36.44$, and the weights 0.90×1, 0.02×32.46 (twice) and 0.02×36.44 sum to 2.927.

Table 30.7: Posterior probabilities of the explanations of the observation $\mathbf{y} = (50, 50, 58)$ MW, from the density of the consistency residual $e_c = -16$ MW. Priors and the bias scale $\tau = 10$ MW illustrative.

Hypothesis	Prior	$\mathrm{Var}(e_c)$ (MW^2)	$p(e_c \mid H_i)$ (MW^{-1})	Ratio to H_0	Posterior
H_0: healthy	0.90	24	3.93×10^{-4}	1	0.307
$H_{S,12}$: L12 meter bias	0.02	124	1.28×10^{-2}	32.46	0.222
$H_{S,23}$: L23 meter bias	0.02	124	1.28×10^{-2}	32.46	0.222
$H_{S,13}$: L13 meter bias	0.02	424	1.43×10^{-2}	36.44	0.249
H_T: L13 open	0.04	–	–	$\sim 10^{-180}$	0.000

The posterior has moved most of the probability away from the healthy hypothesis, as the chi-square test said it should, but it divides that probability almost evenly among the three meters. The small preference for L13 does not come from the data, which fit the three bias hypotheses equally well; it comes from the bias prior, since an L13 bias needs only 8 MW to explain the inconsistency while an L12 or L23 bias needs 16 MW. This is the tie of the normalized residuals restated as probabilities. For the readings of the topology example, $\mathbf{y} = (100, 100, 0)$ MW, the same calculation reverses: H_T fits exactly, every closed-topology hypothesis needs an error of 100 MW or more, ten bias standard deviations, and the posterior of H_T is indistinguishable from one.

The parameter hypothesis H_P, a wrong reactance of L13 for example, changes the third row of $\mathbf{H}$ from $(500, 0)$ to $(S_{base}/x_{13}, 0)$. At a single operating point it produces the same residual pattern as an L13 meter bias, since $x_{13} = 0.172$ would make $P_{13} = 58$ MW at $\theta_1 = 0.10$ rad, and the two cannot be distinguished from one scan. Across scans at different loadings they can: a parameter error scales with the flow, a bias does not. Whether such parameters can be inferred is the subject of the next chapter. The posterior probabilities remain conditional on the completeness and adequacy of the hypothesis set, as Section 22.6.3 of Chapter 22 showed for the pump.

30.6.3 The diagnostic workflow

When the residuals are unusually large, the methods of the chapter assemble into a disciplined workflow. Figure 30.7 draws it as seven stages read from top to bottom; the box names each stage and the text on the right states what it does. The arrow on the left returns from the last stage to the first: a further measurement or check, once made, re-enters the workflow as new evidence.

The workflow connects state estimation to the questions of the earlier chapters on condition monitoring and on the value of measurements. Stages 1 and 2 remove the explanations that need no statistics; stages 3 and 4 detect; stages 5 and 6 diagnose, and stage 6 is the one most often omitted, since it asks whether the competing explanations can be separated at all; stage 7 decides. The purpose is not merely to obtain an angle estimate. It is to establish which physical conclusions remain credible under imperfect measurements.

30.6.4 One suspicious measurement over a range of values

The response of the estimator to the suspicious reading can be followed in closed form. Hold the L12 and L23 readings at 50 MW and vary the reported L13 flow y_{13}, with $\sigma_i = 2$ MW throughout. By equation (30.31) every estimated flow is $\hat{P} = (100 + y_{13})/3$, so that $\hat{\theta}_1 = \hat{P}/500$ and $\hat{\theta}_2 = \hat{P}/1000$; the residuals are $e_{12} = e_{23} = 50 - \hat{P}$ and $e_{13} = y_{13} - \hat{P}$; the statistic is $J = (100 - 2y_{13})^2/24$ with tail probability $\Pr(\chi^2_1 > J)$; and the normalized residuals are

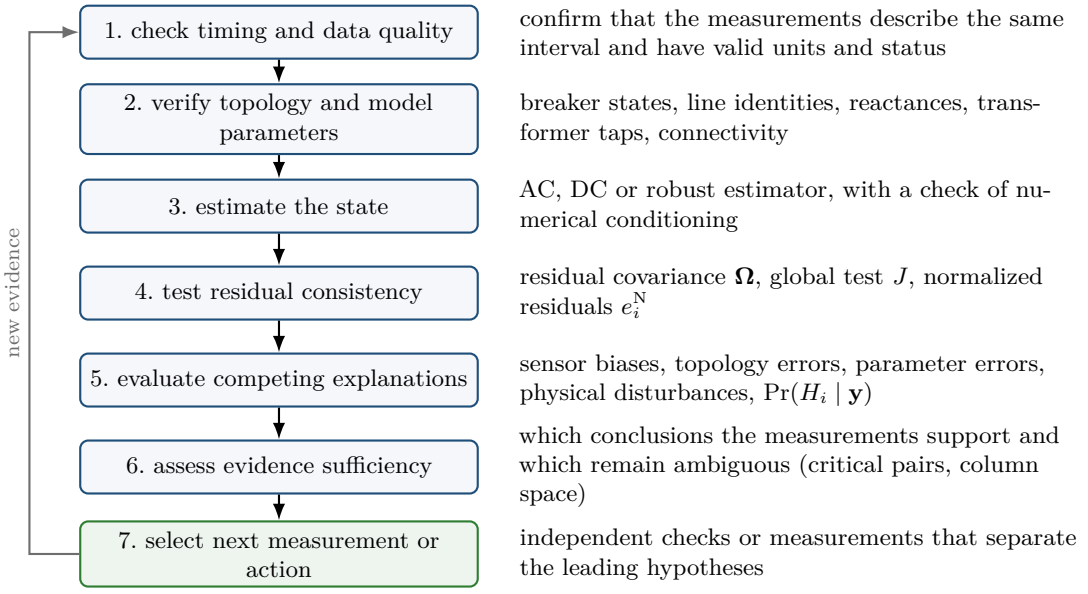


Figure 30.7: The diagnostic workflow for inconsistent measurements, from data quality to the choice of the next measurement or action. A further check selected at stage 7 returns its outcome to stage 1.

$e_{12}/\sqrt{2/3}$ and $e_{13}/\sqrt{8/3}$. Table 30.8 evaluates the model for the reading of the example, 58 MW, for the consistent reading, 50 MW, and for seven others between 40 and 65 MW. One row of the arithmetic, $y_{13} = 54$ MW: $\hat{P} = 154/3 = 51.333$ MW, $\hat{\theta}_1 = 0.102667$ rad, $e_{12} = -1.333$ MW, $e_{13} = 2.667$ MW, $J = (100 - 108)^2/24 = 2.667$, and $e_{12}^{\mathrm{N}} = -1.333/0.8165 = -1.633$.

Table 30.8: WLS estimate, residuals, chi-square statistic and normalized residuals of the three-bus example as functions of the reported L13 flow, with $y_{12} = y_{23} = 50$ MW and $\sigma_i = 2$ MW ($\hat{\theta}_2 = \hat{\theta}_1/2$; $e_{23} = e_{12}$). Bold: rejected at the 1 % level.

y_{13} (MW)	$\hat{\theta}_1$ (rad)	$\hat{P}$ (MW)	e_{12} (MW)	e_{13} (MW)	J	$\Pr(\chi_1^2 > J)$	e_{12}^{N}	e_{13}^{N}
40	0.093333	46.667	3.333	−6.667	**16.667**	4.5×10^{-5}	4.082	−4.082
44	0.096000	48.000	2.000	−4.000	6.000	0.0143	2.449	−2.449
50	0.100000	50.000	0.000	0.000	0.000	1	0.000	0.000
54	0.102667	51.333	−1.333	2.667	2.667	0.102	−1.633	1.633
56	0.104000	52.000	−2.000	4.000	6.000	0.0143	−2.449	2.449
57	0.104667	52.333	−2.333	4.667	**8.167**	0.0043	−2.858	2.858
58	0.105333	52.667	−2.667	5.333	**10.667**	0.0011	−3.266	3.266
60	0.106667	53.333	−3.333	6.667	**16.667**	4.5×10^{-5}	−4.082	4.082
65	0.110000	55.000	−5.000	10.000	**37.500**	9.1×10^{-10}	−6.124	6.124

Figure 30.8 draws the same model continuously. In panel (a) the horizontal axis is the reported L13 flow and the vertical axis the statistic J; the dashed line is the 1 % threshold 6.635 and the dotted lines the readings at which J reaches it. In panel (b) the vertical axis is the estimated L13 flow; the dotted line is the reported flow itself, and the dashed line an illustrative weather-dependent rating of L13 of 55 MW, used in the next subsection.

Three trends appear. At $y_{13} = 50$ MW all residuals vanish and the estimate returns to the base case. As the reading rises, WLS adjusts the state to accommodate it, moving every estimated flow by one third of the change and distributing the inconsistency over all three residuals in the fixed proportions $1 : 1 : -2$; the test rejects the set once $|y_{13} - 50| > \sqrt{24 \times 6.635}/2 = 6.31$ MW, that is below 43.69 or above 56.31 MW. And at every reading the normalized residual magnitudes

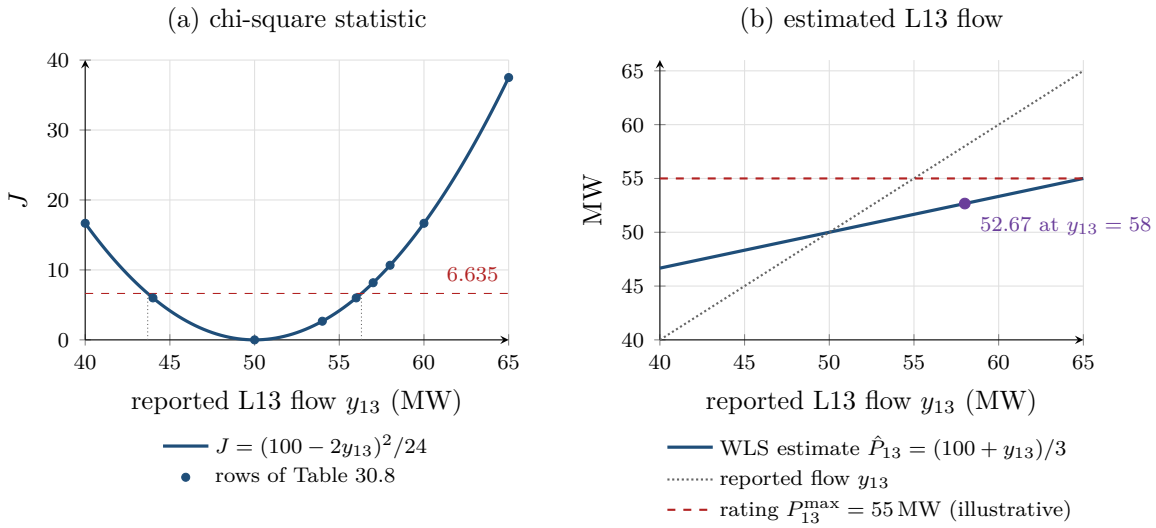


Figure 30.8: Response of the WLS estimator to the reported L13 flow, with $y_{12} = y_{23} = 50\,\text{MW}$ and $\sigma_i = 2\,\text{MW}$. (a) The statistic $J = (100 - 2y_{13})^2/24$ exceeds the 1 % threshold for readings more than 6.31 MW from 50 MW. (b) The estimate moves one third of the way toward the reported value, $\hat{P}_{13} = (100 + y_{13})/3$, and reaches the illustrative 55 MW rating only at $y_{13} = 65\,\text{MW}$.

remain tied: the statistical test detects the incompatibility without identifying which meter is responsible.

30.6.5 Connection to reliability decisions

Suppose the state estimator reports $\hat{P}_{13} = 52.67\,\text{MW}$ and the weather-dependent rating of L13 for the present conditions is estimated at $P_{13}^{\max} = 55\,\text{MW}$, an illustrative value well below its 80 MW rating under the conditions assumed in Chapter 27. The apparent margin, in the sense of the performance-margin function g of Chapter 1 of Volume I, is

$$g_{13} = P_{13}^{\max} - \hat{P}_{13} = 55 - 52.67 = 2.33\,\text{MW}.$$

The measurement set is statistically inconsistent, and the margin should not be trusted without further assessment. Under the healthy model the estimated L13 flow has standard deviation 1.155 MW (Table 30.2), and the probability of an overload is $1 - \Phi(2.33/1.155) = 1 - \Phi(2.02) = 0.022$. That figure describes the uncertainty of the healthy Gaussian model and nothing else; it does not describe the uncertainty that a possible meter bias, topology error or parameter error adds. An operational risk calculation needs instead the overload probability given the evidence, averaged over the competing explanations,

$$\Pr(P_{13} > P_{13}^{\max} \mid \mathbf{y}) = \sum_i \Pr(H_i \mid \mathbf{y})\, \Pr(P_{13} > P_{13}^{\max} \mid \mathbf{y}, H_i), \tag{30.39}$$

which assumes that the hypotheses form a sufficiently complete, mutually exclusive set. A complete calculation would also carry the uncertainty of the rating itself.

Table 30.9 evaluates equation (30.39) for the hypotheses of Table 30.7. Under each hypothesis the posterior of the state is the WLS solution with that hypothesis's covariance (the variance of the biased meter raised to $104\,\text{MW}^2$), and the posterior of P_{13} is Gaussian with the mean and standard deviation given. The columns are the posterior probability of the hypothesis, the posterior mean and standard deviation of the L13 flow under it, the overload probability under it, and its contribution to the total. One row of the arithmetic, $H_{\text{S},13}$: with weights

$(1/4, 1/4, 1/104)$ the normal equations give $127\,404\,\theta_1 = 12\,779$, so $\theta_1 = 0.100302$ rad and $P_{13} = 50.15$ MW with standard deviation 1.401 MW, and $1 - \Phi((55 - 50.15)/1.401) = 0.00027$.

Table 30.9: Overload probability of L13 ($P_{13}^{\max} = 55$ MW, illustrative) for the observation $\mathbf{y} = (50, 50, 58)$ MW, by hypothesis and in total, equation (30.39).

Hypothesis	$\Pr(H_i \mid \mathbf{y})$	Mean P_{13} (MW)	Std. dev. (MW)	Pr(overload $\mid \mathbf{y}, H_i$)	Contribution
H_0: healthy	0.307	52.67	1.155	0.0217	0.0067
$H_{\mathrm{S},12}$: L12 meter bias	0.222	56.97	1.867	0.854	0.1894
$H_{\mathrm{S},23}$: L23 meter bias	0.222	56.97	1.867	0.854	0.1894
$H_{\mathrm{S},13}$: L13 meter bias	0.249	50.15	1.401	0.0003	0.0001
H_{T}: L13 open	0.000	0	–	0	0
Total	1				0.386

The overload probability given the evidence is 0.386, eighteen times the 0.022 of the healthy model. The two hypotheses in which the L13 meter is right, and one of the others wrong, put the L13 flow near 57 MW, above the rating, and together they carry almost half of the posterior probability. The result changes by orders of magnitude depending on which hypothesis is true, which is why the diagnosis is part of the risk calculation and not a preliminary to it; and the most valuable next measurement is one that separates the meter hypotheses, such as an independent reading of the L13 flow or of a bus injection, in the sense of the expected information gain of Section 22.5.5.

> **Principle 30.3.** *A good fit is not necessarily a correct model. A state estimator that converges numerically can still give an unreliable basis for operational decisions if the measurement model is wrong, and an overload probability computed under the healthy model alone omits the largest part of the risk when the measurements are inconsistent.*

Bad-data detection and model credibility are therefore integral to reliability engineering. Whether the equations and measurements suffice to determine the physical parameters as well as the state, and whether the model represents the system at all, is the subject of the next chapter, Chapter 31.

30.7 What to retain

Definition 30.1 (State estimation and bad-data processing)**.** State estimation infers the operating state most consistent with redundant measurements and a network model; for Gaussian errors it is the weighted least-squares estimate $\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}$ with covariance $\mathbf{G}^{-1}$. Its residuals $\mathbf{e} = \mathbf{S}\mathbf{y}$, with covariance $\mathbf{\Omega} = \mathbf{S}\mathbf{R}$, detect inconsistency through $J \sim \chi^2_{m-n}$ and the normalized residuals, but identify its cause only when the competing explanations leave distinguishable traces in the measurements.

The main conclusions are nine. State estimation uses physical equations and measurement uncertainty to reconcile noisy observations into one network-consistent state. Residuals provide evidence of inconsistency, not necessarily a unique diagnosis. The residual covariance differs from the measurement-error covariance, because the estimate absorbs part of every error. The global chi-square test detects aggregate inconsistency under stated statistical assumptions. Normalized residuals can fail to isolate a corrupted measurement, especially when redundancy is limited, and always tie when it is one. Critical measurements and conforming errors in the column space of $\mathbf{H}$ produce undetectable errors. Topology and parameter errors can mimic sensor faults. Robust estimation limits the influence of some outliers but does not establish that

the physical model is correct. And probabilistic inference can carry competing explanations into operational risk calculations, where they can change the answer by an order of magnitude. Table 30.10 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 30.10: The terms of state estimation and bad-data processing and where each is defined.

Term	Expression	Units	Meaning	Defined in
State vector	$\mathbf{x} = (\theta_2, \dots, \theta_N, V_1, \dots, V_N)$	rad, pu	fixes the electrical state	30.1.1
Measurement model	$\mathbf{y} = h(\mathbf{x}) + \boldsymbol{\varepsilon}$, $\mathrm{Cov}(\boldsymbol{\varepsilon}) = \mathbf{R}$	as $\mathbf{y}$	predicts each measurement	30.1.1
Measurement Jacobian	$\mathbf{H} = \partial h / \partial \mathbf{x}$	varies	sensitivity to the states	30.2.1
WLS objective	$J = (\mathbf{y} - \mathbf{H}\mathbf{x})^{\mathsf{T}}\mathbf{W}(\mathbf{y} - \mathbf{H}\mathbf{x})$	–	$-2\times$ log-likelihood	30.2.2
Gain matrix	$\mathbf{G} = \mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{H}$	rad^{-2}	normal-equation matrix	30.2.3
WLS estimate	$\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y}$	rad	maximum likelihood state	30.2.3
Estimation covariance	$\mathrm{Cov}(\hat{\mathbf{x}} - \mathbf{x}) = \mathbf{G}^{-1}$	rad^2	precision under the model	30.2.4
Residual	$\mathbf{e} = \mathbf{y} - \mathbf{H}\hat{\mathbf{x}}$	as $\mathbf{y}$	observed minus estimated	30.3.1
Residual sensitivity matrix	$\mathbf{S} = \mathbf{I} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}$, $\mathbf{e} = \mathbf{S}\boldsymbol{\varepsilon}$	–	maps errors to residuals	30.3.1
Residual covariance	$\boldsymbol{\Omega} = \mathbf{S}\mathbf{R} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}$	MW^2	covariance of residuals	30.3.2
Chi-square test	$J(\hat{\mathbf{x}}) \sim \chi^2_{m-n}$	–	global detection	30.3.3
Measurement redundancy	$m - n$	–	degrees of freedom	30.3.3
Normalized residual	$e_i^{\mathrm{N}} = e_i / \sqrt{\Omega_{ii}}$	–	in its own std. deviations	30.3.5
LNR test	largest $\lvert e_i^{\mathrm{N}} \rvert > 3$	–	bad-data identification	30.3.5
Estimated gross error	$\hat{b}_i = R_{ii} e_i / \Omega_{ii}$	as y_i	size of one bad datum	30.3.5
Critical measurement, pair	removal makes $\mathbf{x}$ unobservable	–	undetectable; unidentifiable	30.4.1
Interacting, conforming	$\lvert\Omega_{ij}\rvert / \sqrt{\Omega_{ii}\Omega_{jj}} \approx 1$; $\mathbf{S}\mathbf{b} \approx \mathbf{0}$	–	correlated; consistent errors	30.4.2
Undetectable bad data	$\mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{b} = \mathbf{b}$, i.e. $\mathbf{b} = \mathbf{H}\mathbf{c}$	as $\mathbf{y}$	absorbed into the state	30.4.3
Topology error	wrong $\mathbf{H}$ from a wrong status	–	mimics sensor faults	30.4.4
LAV, Huber estimators	$\sum_i \lvert u_i \rvert$; $\sum_i \rho_\delta(u_i)$	–	robust estimators	30.5.2
IRLS	WLS with weights w_i/σ_i^2, iterated	–	solves robust estimators	30.5.4
Posterior of the state	$\boldsymbol{\Sigma}_{\mathrm{post}} = (\boldsymbol{\Sigma}_0^{-1} + \mathbf{G})^{-1}$	rad^2	Bayesian state estimate	30.6.1
Hypothesis posterior	$\Pr(H_i \mid \mathbf{y}) \propto \Pr(H_i)\, p(\mathbf{y} \mid H_i)$	–	weight of an explanation	30.6.2
Overload probability	$\sum_i \Pr(H_i \mid \mathbf{y})\ \Pr(\text{overload} \mid \mathbf{y}, H_i)$	–	risk over explanations	30.6.5

Further reading. Abur and Expósito, *Power System State Estimation: Theory and Implementation*: weighted least squares, observability, normalized residuals and bad-data processing.

Monticelli, *State Estimation in Electric Power Systems: A Generalized Approach*: topology processing, generalized state estimation and bad-data analysis. Lista, *Statistical Methods for Data Analysis in Particle Physics*: general likelihood and goodness-of-fit principles applicable to model validation.

30.8 Exercises

Exercise 30.1. (a) State what determines the weights in conventional Gaussian WLS state estimation. (The measurement-error covariance: $\mathbf{W} = \mathbf{R}^{-1}$, so a measurement with error standard deviation σ_i has weight $1/\sigma_i^2$.) (b) Compute the contribution to J of a discrepancy of 5 MW on a meter with $\sigma_i = 1$ MW and on one with $\sigma_i = 10$ MW. (25 and 0.25.) (c) Explain why the optimized WLS residual can be smaller than the original sensor error. (The estimated state absorbs part of the measurement error; only the part no state adjustment can explain remains, $\mathbf{\Omega} = \mathbf{R} - \mathbf{H}\mathbf{G}^{-1}\mathbf{H}^{\mathsf{T}}$.) (d) State the purpose of a global chi-square bad-data test. (To test the aggregate consistency of the measurements with the assumed measurement model; it neither identifies the faulty sensor nor proves the topology correct.) (e) State which variance normalizes an individual WLS residual. (The diagonal entry Ω_{ii} of the residual covariance, not σ_i^2.)

Exercise 30.2. (a) Explain why all three normalized residual magnitudes tie in the three-bus example. (There is only one independent residual direction, $\mathbf{a} = (1, 1, -2)^{\mathsf{T}}$; every normalized residual equals $\pm e_{\mathrm{c}}/\sqrt{24}$.) (b) State whether an incorrect breaker status can cause large residuals when all flow sensors are healthy. (Yes, because the measurement model is built on the wrong topology; with L13 open and modeled closed, $J = 1666.7$ with exact meters.) (c) State what robust estimation does. (It reduces the influence of some large residuals under its chosen loss function; it does not guarantee correct topology identification or make every state identifiable.) (d) Explain why a Bayesian posterior over topology and sensor faults is useful. (It represents the competing explanations and propagates them into operational risk, such as the overload probability of a line; it does not guarantee that the hypothesis set is complete or remove the need for model validation.)

Exercise 30.3. In the three-bus example the L13 meter is replaced by a more precise one, $\sigma_{13} = 1$ MW, while L12 and L23 keep $\sigma = 2$ MW; the readings are again $\mathbf{y} = (50, 50, 58)$ MW. (a) Compute the gain matrix, the estimate and the estimated flows. ($\mathbf{G} = \left[\begin{smallmatrix} 500\,000 & -250\,000 \\ -250\,000 & 500\,000 \end{smallmatrix}\right]$ rad^{-2}, $\mathbf{H}^{\mathsf{T}}\mathbf{W}\mathbf{y} = (41\,500, 0)$; $\hat{\theta}_1 = 0.110667$, $\hat{\theta}_2 = 0.055333$ rad; all three flows 55.333 MW.) (b) Compute the residuals, J and its tail probability. ($(-5.333, -5.333, 2.667)$ MW; $J = 21.333$; 3.9×10^{-6}.) (c) Compute the residual variances and the normalized residuals, and state whether the LNR test now identifies a meter. ($\Omega_{ii} = 1.333$, 1.333 and 0.333 MW^2; $e^{\mathrm{N}} = (-4.619, -4.619, +4.619)$; no: the magnitudes still tie, since the redundancy is still one, and the estimated gross errors are again -16, -16 and $+8$ MW.) (d) Explain why the estimate moved toward the L13 reading. (Its weight rose fourfold, so the estimate trusts it more; the precision of a meter changes the estimate and the test, not the identifiability of the error.)

Exercise 30.4. (a) With only the L12 and L23 meters (measurement set B), the L12 meter reads 55 MW instead of the true 50 MW, and L23 reads 50 MW. Compute the estimate, the inferred L13 flow and the residuals. ($\hat{\theta}_1 = 0.105$, $\hat{\theta}_2 = 0.050$ rad; $\hat{P}_{13} = 52.5$ MW; both residuals zero: both measurements are critical.) (b) With all three meters, decide whether the bias vector $\mathbf{b} = (2, 2, 2)$ MW is detectable, and if not find the state shift that absorbs it. ($\mathbf{a}^{\mathsf{T}}\mathbf{b} = 2 + 2 - 4 = 0$ and $\mathbf{b} = \mathbf{H}\mathbf{c}$ with $\mathbf{c} = (0.004, 0.002)$ rad: undetectable.) (c) For $\mathbf{y} = (50, 50, 62)$ MW and $\sigma_i = 2$ MW, compute the WLS estimate, the first IRLS weights of the Huber estimator with $\delta = 1.5$, and the converged Huber estimate. (WLS: all flows 54 MW, $\mathbf{u} = (-2, -2, 4)$; first weights $(0.75, 0.75, 0.375)$; converged: all flows 51.5 MW, residuals

$(-1.5, -1.5, 10.5)$ MW, $w_{13} = 1.5/5.25 = 0.286$, the same estimate as for $y_{13} = 58$ MW.) (d) With the priors and bias scale of Table 30.7, compute the posterior probabilities of H_0, $H_{\mathrm{S},12}$, $H_{\mathrm{S},23}$ and $H_{\mathrm{S},13}$ for $\mathbf{y} = (50, 50, 65)$ MW. ($e_{\mathrm{c}} = -30$ MW; ratios to H_0 of 1.62×10^6, 1.62×10^6 and 1.14×10^7; posterior 3×10^{-6}, 0.110, 0.110 and 0.779: the larger inconsistency favors the L13 bias, which needs 15 MW against 30 MW for the others.)

Chapter 31

Identifiability, Uncertainty, and Model Adequacy

31.1 Parameter estimation and identifiability

Chapter 29 asked whether the states of a system can be recovered from its measurements, and Chapter 30 estimated those states by weighted least squares and tested the residuals for bad data. Both chapters took the parameters of the measurement model as known: the line reactances, the network topology, the variances of the meters. The physical models used in reliability work also contain parameters that are not known and must be estimated from operating data: the heat losses of a data center room, the efficiency of a cooling plant, the convective cooling of an overhead conductor, the internal resistance of a battery. An optimization algorithm fitted to such data may converge, return its parameters to six decimal places and leave small residuals. None of this establishes that the parameters have been determined correctly. The question of this chapter is

whether the available measurements determine the physical parameters of interest uniquely and with useful accuracy, and whether the model, with any values of those parameters, represents the system.

The chapter answers it in five steps. It separates structural from practical identifiability and works both on small thermal examples. It defines the sensitivity matrix of a model with respect to its parameters, reads identifiability from its rank, and applies it to a data center thermal balance measured under two experimental designs. It derives the parameter covariance from the Fisher information, reads the weakly constrained directions from the singular values of the scaled sensitivity matrix, and adds prior information through the Gaussian linear posterior. It tests whether an identifiable model fits at all, introduces the model discrepancy term, and compares competing physical explanations by their Bayesian evidence. Finally, it applies the ideas to line ratings, battery storage and reliability predictions, tabulates the effect of the experimental design, and assembles the checks into a modeling workflow.

Jargon: identifiability (structural, practical), sensitivity matrix, Fisher information, Cramér–Rao bound, condition number, weakly constrained directions, model adequacy, model discrepancy, evidence, Bayes factor, Laplace approximation, design of experiments

A parameter of a model is *identifiable* when the observations determine it uniquely. *Structural identifiability* asks this with ideal, noise-free data and the specified model and experiment: no two distinct admissible parameter vectors may produce identical observations. *Practical identifiability* asks whether the actual, noisy data determine the parameters with useful precision. The *sensitivity matrix* $\mathbf{S}_\theta = \partial h/\partial \boldsymbol{\theta}$ holds the derivatives of the predicted measurements with respect to the parameters. The *Fisher information* $\mathcal{I}(\boldsymbol{\theta})$ measures how sharply the likelihood is curved about the true parameters; for Gaussian errors it is $\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta$, and the *Cramér–Rao bound* states that no unbiased estimator has a covariance smaller than its inverse. The *condition number* κ is the ratio of the largest to the smallest singular value of the scaled sensitivity matrix; the parameter combinations along the small singular values are the *weakly constrained directions*, also called sloppy directions. *Model adequacy* is the property that the model, at some admissible parameter values, explains the observations within their stated uncertainty; the *model discrepancy* $\delta(\cdot)$ is the difference between the physical system and the model at the best parameter values. The *evidence* (marginal likelihood) of a model is the probability of the data averaged over the prior of its parameters; the *Bayes factor* is the ratio of the evidences of two models; the *Laplace approximation* evaluates the evidence from the curvature of the posterior at its mode. *Design of experiments* chooses the operating conditions and measurements so that the parameters of interest become identifiable and precise; varying the inputs for this purpose is called excitation.

Key idea

A credible parameter estimate requires three separate findings: the parameters are identifiable under the model and the experiment, the actual data constrain them precisely enough for the purpose, and the model explains the data within their stated uncertainty. A converged fit with small standard errors establishes none of the three by itself.

This section states the three questions, writes parameter estimation as an inverse problem, and defines structural and practical identifiability with one small example of each.

31.1.1 Three questions about an estimated parameter

A parameter estimate is credible only if three questions have been answered, and they are distinct.

1. *Identifiability.* Could the unknown quantities be distinguished from the available observations, even with perfect measurements?
2. *Estimation uncertainty.* If they can be distinguished, how accurately can they be estimated from the actual, noisy measurements?
3. *Model adequacy.* Does the proposed physical model explain the observations to an acceptable degree?

The first is a property of the model and the experimental design; the second adds the measurement errors; the third compares the model with the physical system through the evidence. The book's recurring distinction between the physical system, the model of it and the evidence is therefore the organizing idea of the chapter: identifiability is a property of the model and the experiment, uncertainty is a property of the evidence under the model, and adequacy is the relation between the model and the system.

Identifiability is related to the observability of Chapter 29 but is not the same question. Observability concerns the states $\mathbf{x}$ of a system, such as the bus voltage angles of a network,

given a model whose parameters are known. Identifiability concerns the parameters $\boldsymbol{\theta}$ of the model itself. A state may be observable under known parameters while some parameters remain unidentifiable. In the DC network of the earlier chapters, the flow on line L13 is $P_{13} = (\theta_1 - \theta_3)/x_{13}$; with bus 3 as the angle reference, $\theta_3 = 0$. If both the reactance x_{13} and the angle θ_1 are unknown and only P_{13} is measured, the measurement constrains the ratio θ_1/x_{13} and not the two individually: every pair with the same ratio gives the same flow.

Remark 31.1. The bold $\boldsymbol{\theta}$ of this chapter is the vector of unknown model parameters; the scalar θ_i is the voltage angle of bus i in the network chapters, a state. The bold $\mathbf{R}$ is the covariance matrix of the measurement errors and the italic $R(T)$ and R_{int} are electrical resistances; none of them is the reliability function R. The sensitivity matrix $\mathbf{S}_\theta$ is not the residual sensitivity matrix $\mathbf{S}$ of Chapter 30, and the thermal conductance G is not the gain matrix $\mathbf{G}$.

31.1.2 Parameter estimation as an inverse problem

A forward physical model predicts the measurements from the parameters and the operating conditions:

$$\mathbf{y} = h(\boldsymbol{\theta}, \mathbf{u}) + \boldsymbol{\varepsilon}, \tag{31.1}$$

where $\mathbf{y}$ is the vector of m measurements (the book's symbol for observations, written $\mathbf{z}$ in much of the power-system literature), $\boldsymbol{\theta}$ is the vector of n_θ unknown physical parameters, $\mathbf{u}$ collects the known operating inputs, h is the physical model, and $\boldsymbol{\varepsilon}$ is the measurement error, with zero mean and covariance $\mathbf{R}$ and weight matrix $\mathbf{W} = \mathbf{R}^{-1}$ as in Chapter 30. The *forward problem* computes $\mathbf{y}$ from $\boldsymbol{\theta}$ and $\mathbf{u}$; it has one answer. The *inverse problem* recovers $\boldsymbol{\theta}$ from $\mathbf{y}$ and $\mathbf{u}$; its answer may not exist, may not be unique, and may depend sensitively on the errors in $\mathbf{y}$.

Figure 31.1 draws the two directions. The boxes on the left are the unknown parameters and the known inputs; the middle box is the model; the error enters at the summing node; the box on the right holds the measurements. Solid arrows are the forward problem; the dashed arrow returning below is the inverse problem. The three notes beneath attach the chapter's three questions to the parts of the diagram they concern: identifiability to the model with the error removed, estimation uncertainty to the error, and adequacy to the model itself.

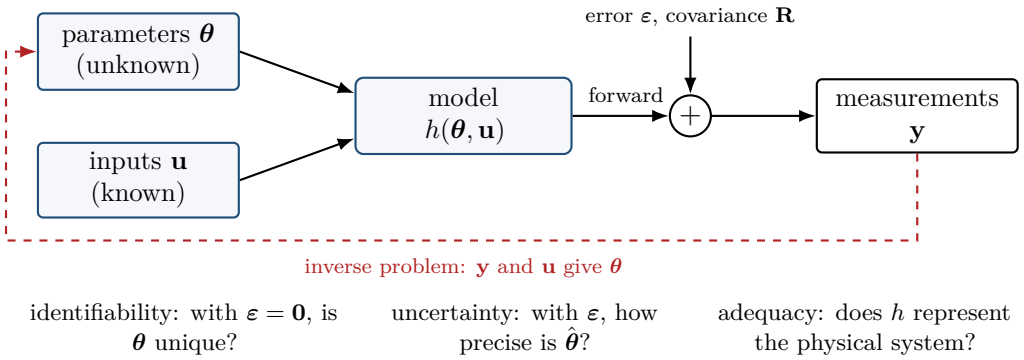


Figure 31.1: Parameter estimation as an inverse problem. The forward model of equation (31.1) maps parameters and inputs to measurements; estimation runs the other way. The three questions of the chapter concern the model without error, the error, and the model itself.

31.1.3 Structural identifiability

A model is *globally structurally identifiable* under a specified experiment if distinct admissible parameter vectors cannot generate identical ideal observations: for parameter vectors $\boldsymbol{\theta}_1$ and

$\boldsymbol{\theta}_2$ in the admissible domain,

$$h(\boldsymbol{\theta}_1, \mathbf{u}) = h(\boldsymbol{\theta}_2, \mathbf{u}) \text{ for the inputs } \mathbf{u} \text{ of the experiment} \quad \Longrightarrow \quad \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2. \tag{31.2}$$

The model is *locally* structurally identifiable at $\boldsymbol{\theta}$ if the implication holds for all $\boldsymbol{\theta}_2$ in some neighborhood of $\boldsymbol{\theta}_1 = \boldsymbol{\theta}$. Structural identifiability assumes ideal data and the specified mathematical model. It asks whether the experimental design contains, in principle, enough independent information; it says nothing about noise, and it is a property of the model and the experiment together, not of the model alone.

Example: two thermal conductances. A room exchanges heat with the outdoor air through two parallel paths, its walls and its roof, with thermal conductances G_1 and G_2 in kW/K. Writing the temperature difference as $\Delta T = T_{\text{out}} - T_{\text{in}}$ in kelvin, the heat flows into the room through the two paths are

$$Q_1 = G_1 \, \Delta T, \qquad Q_2 = G_2 \, \Delta T,$$

and the total heat flow through the envelope is

$$Q = Q_1 + Q_2 = (G_1 + G_2) \, \Delta T. \tag{31.3}$$

Figure 31.2 draws the room as a balance node. The box on the left is the outdoor air at T_{out}, the box in the middle the room at T_{in}; the two arrows between them are the wall and roof paths. The arrow from below and the arrow to the right belong to the data center balance of Section 31.2.4 and are ignored here. If only Q and ΔT are measured, the data determine the total conductance

$$G_{\text{tot}} = G_1 + G_2$$

and nothing else. The pairs $(G_1, G_2) = (2, 3)\,\text{kW/K}$ and $(1, 4)\,\text{kW/K}$ produce the same heat flow at every temperature difference, and so does every pair $(G_1 + c,\ G_2 - c)$. Even infinitely many perfect measurements at different temperatures cannot separate the conductances under this model: equation (31.3) is invariant under the shift c, and condition (31.2) fails for every experiment that measures only Q and ΔT. That is a structural identifiability problem. The remedy is a different measurement, for example a heat-flux meter on the roof that observes Q_2 alone.

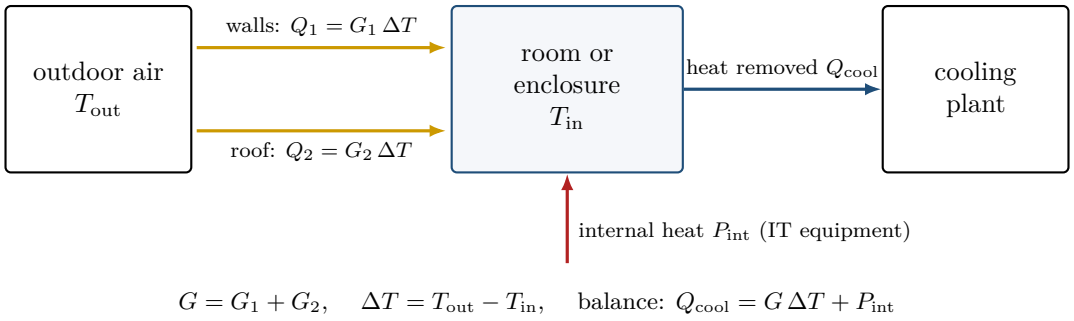


Figure 31.2: A room or equipment enclosure as a quasi-steady heat-balance node. Heat enters from the outdoor air through two parallel envelope paths and from the internal equipment, and is removed by the cooling plant. Measuring only the envelope total Q and ΔT identifies $G_1 + G_2$, not G_1 and G_2 (equation (31.3)); measuring Q_{cool} and ΔT is the experiment of equation (31.9).

31.1.4 Practical identifiability

Suppose now that the model is structurally identifiable but the measurements contain noise. A parameter may influence the measurements only weakly, or two parameters may have nearly indistinguishable effects. The parameters are then identifiable in principle and difficult to estimate accurately: the question of *practical identifiability*. Consider two measurements of two parameters,

$$y_1 = \theta_1 + \theta_2, \qquad y_2 = \theta_1 + 1.001\,\theta_2.$$

The equations are independent, and in exact arithmetic with perfect measurements both parameters can be recovered. Subtracting the first equation from the second gives $y_2 - y_1 = 0.001\,\theta_2$, so that

$$\theta_2 = 1000\,(y_2 - y_1), \qquad \theta_1 = y_1 - \theta_2 = 1001\,y_1 - 1000\,y_2. \tag{31.4}$$

An error of only 0.01 in the difference $y_2 - y_1$ produces an error of 10 in θ_2, and an equal and opposite error in θ_1. If the two measurements carry independent errors of standard deviation σ, the standard deviation of $\hat{\theta}_2$ is $1000\sqrt{2}\,\sigma = 1414\,\sigma$ and that of $\hat{\theta}_1$ is $\sqrt{1001^2 + 1000^2}\,\sigma = 1415\,\sigma$. The problem is mathematically identifiable and numerically fragile; the singular values of Section 31.3.4 measure the fragility, and for this matrix their ratio is about 4000.

31.1.5 Structural and practical identifiability compared

Table 31.1 sets the two notions side by side. Each row is an aspect of the analysis; the middle column gives it for structural identifiability and the right column for practical identifiability. The table's last row is the operational consequence: a structural problem is solved only by a different measurement or a different experimental input, while a practical problem yields to more precise, more varied or longer data.

Table 31.1: Structural and practical identifiability.

Aspect	Structural identifiability	Practical identifiability
Main question	Can the parameters be recovered uniquely in principle?	Can the parameters be estimated with useful precision?
Ideal measurements assumed	yes	no; the actual uncertainty matters
Main tools	symmetries, invariances, rank, algebraic and differential methods	covariance, singular values, likelihood profiles, posterior geometry
Possible conclusion	parameter combinations fundamentally inseparable	parameters distinguishable but uncertain
Usual remedy	add a different measurement or experimental input	improve precision, excitation, sensor quality or duration

Principle 31.1. *Identifiability is a property of the model and the experiment together. It is decided before any data are collected, and no quantity of data from an experiment that cannot separate two parameters will separate them.*

31.2 Sensitivity matrices and the thermal balance experiments

Identifiability is decided in practice by linearizing the model about a parameter vector and examining the resulting matrix. This section defines the sensitivity matrix, reads local identifiability from its rank, explains why it must be scaled before its numbers are compared, and applies it to a data center thermal balance measured under two experimental designs.

31.2.1 Local linearization and the sensitivity matrix

Let the model predict m observations from n_θ unknown parameters, $\hat{\mathbf{y}} = h(\boldsymbol{\theta})$, with the known inputs $\mathbf{u}$ held fixed and omitted from the notation. Around a reference parameter vector $\boldsymbol{\theta}_0$, a first-order Taylor expansion gives

$$h(\boldsymbol{\theta}_0 + \Delta\boldsymbol{\theta}) \approx h(\boldsymbol{\theta}_0) + \mathbf{S}_\theta \, \Delta\boldsymbol{\theta}, \qquad [\mathbf{S}_\theta]_{ij} = \left.\frac{\partial h_i}{\partial \theta_j}\right|_{\boldsymbol{\theta}_0}, \tag{31.5}$$

where $\mathbf{S}_\theta = \partial h / \partial \boldsymbol{\theta}$ is the $m \times n_\theta$ *parameter sensitivity matrix*, the Jacobian of the model with respect to its parameters. Its column j describes how all the predicted observations change when parameter j alone changes; its row i describes how observation i responds to every parameter. For a model linear in its parameters the expansion is exact and $\mathbf{S}_\theta$ does not depend on $\boldsymbol{\theta}_0$. The matrix plays for parameters the role that the measurement Jacobian $\mathbf{H} = \partial h / \partial \mathbf{x}$ of Chapter 30 plays for states; when states and parameters are both unknown, the augmented matrix $[\,\mathbf{H} \;\; \mathbf{S}_\theta\,]$ describes which combinations of the two the measurements constrain.

31.2.2 Rank and parameter distinguishability

If the sensitivity matrix is rank deficient, $\operatorname{rank}(\mathbf{S}_\theta) < n_\theta$, its columns are linearly dependent and some perturbation $\Delta\boldsymbol{\theta} \neq \mathbf{0}$ satisfies

$$\mathbf{S}_\theta \, \Delta\boldsymbol{\theta} = \mathbf{0}. \tag{31.6}$$

To first order, the parameters can then move along $\Delta\boldsymbol{\theta}$ without changing any predicted observation: $\Delta\boldsymbol{\theta}$ is a *locally unidentifiable direction*, a vector of the null space of $\mathbf{S}_\theta$. In the two-conductance example, with measurements at temperature differences $\Delta T_1, \ldots, \Delta T_m$ and $\boldsymbol{\theta} = (G_1, G_2)^\mathsf{T}$, row i of $\mathbf{S}_\theta$ is $(\Delta T_i, \; \Delta T_i)$; the two columns are equal, the rank is one for any choice of temperatures, and the null direction is $\Delta\boldsymbol{\theta} \propto (1, -1)^\mathsf{T}$, the shift c of Section 31.1.3.

If $\mathbf{S}_\theta$ has full column rank, n_θ, at $\boldsymbol{\theta}_0$ and the model is continuously differentiable, the model is locally identifiable there: no other parameter vector in a small neighborhood produces the same ideal observations. Full local rank does not prove global structural identifiability, because distant parameter values may still produce identical outputs. The single observation $h(\theta) = \theta^2$ has the nonzero sensitivity 2θ at every $\theta \neq 0$, yet θ and $-\theta$ are indistinguishable; only the admissible domain, here $\theta > 0$, restores uniqueness. Conversely, full rank does not imply useful precision: the matrix of equation (31.4), with rows $(1, 1)$ and $(1, 1.001)$, has rank two and nearly parallel columns. A large number of observations does not help if they do not respond differently to the unknowns: two parameters that always appear as a product, a ratio or a sum remain unidentifiable however many measurements are taken at the same operating condition.

31.2.3 Scaling of parameters and measurements

Parameters usually carry different units. In the thermal models of this chapter a conductance G is in kW/K, a heat input P in kW and a thermal capacitance C in kWh/K. The column of $\mathbf{S}_\theta$ for G is then in kelvin and the column for P is dimensionless, and a derivative with respect to G cannot be compared naively with one with respect to P: changing the unit of P from

kilowatts to watts multiplies one column by 1/1000 and changes every conclusion drawn from the raw numbers. The remedy is to express the parameters relative to meaningful scales,

$$\boldsymbol{\xi} = \mathbf{D}_\theta^{-1}(\boldsymbol{\theta} - \boldsymbol{\theta}_0), \qquad \mathbf{S}_\xi = \frac{\partial h}{\partial \boldsymbol{\xi}} = \mathbf{S}_\theta \mathbf{D}_\theta, \tag{31.7}$$

where $\mathbf{D}_\theta$ is a diagonal matrix of parameter scales, such as nominal values (so that ξ_j is a relative change) or prior standard deviations (so that ξ_j is a change in standard deviations). The measurements are scaled by their uncertainty in the same way, by *whitening* with the error covariance:

$$\mathbf{S}_\mathrm{w} = \mathbf{R}^{-1/2}\,\mathbf{S}_\theta\,\mathbf{D}_\theta, \tag{31.8}$$

where $\mathbf{R}^{-1/2}$ is the inverse of a matrix square root of $\mathbf{R}$, $\mathrm{diag}(1/\sigma_i)$ for independent errors with standard deviations σ_i. Each row of $\mathbf{S}_\mathrm{w}$ then measures a response in units of that measurement's standard deviation, and each column a parameter change in units of its scale; the matrix is dimensionless. The singular values of $\mathbf{S}_\mathrm{w}$, Section 31.3.4, describe the parameter information meaningfully; those of the raw $\mathbf{S}_\theta$ do not. Scaling is not a cosmetic step: singular values and condition numbers depend on it.

31.2.4 A data center thermal balance

Consider a room or equipment enclosure of a data center exchanging heat with the outdoor environment, the node of Figure 31.2 with its envelope paths lumped into one effective conductance $G = G_1 + G_2$. Under quasi-steady conditions, with the stored heat of the room neither growing nor shrinking, the heat removed by the cooling plant equals the heat entering from the equipment and through the envelope:

$$Q_\mathrm{cool} = P_\mathrm{int} + G\,(T_\mathrm{out} - T_\mathrm{in}),$$

where Q_cool is the cooling heat removal in kW, P_int the internal heat generation of the IT and auxiliary equipment in kW, G the effective envelope conductance in kW/K, and T_out and T_in the outdoor and indoor temperatures. With $\Delta T = T_\mathrm{out} - T_\mathrm{in}$ as before,

$$Q_\mathrm{cool} = G\,\Delta T + P_\mathrm{int}, \qquad \boldsymbol{\theta} = \begin{bmatrix} G \\ P_\mathrm{int} \end{bmatrix}. \tag{31.9}$$

The enclosure is one room served by a cooling plant such as the three-chiller plant of the running case; its loads of tens of kilowatts are illustrative. The model is linear in its parameters. For measurements at temperature differences $\Delta T_1, \ldots, \Delta T_m$, row i of the sensitivity matrix is $(\partial Q_{\mathrm{cool},i}/\partial G,\ \partial Q_{\mathrm{cool},i}/\partial P_\mathrm{int}) = (\Delta T_i,\ 1)$. The simple linear inverse problem already shows every concept of the chapter.

31.2.5 Experiment A: one temperature difference

Suppose the room is held at $T_\mathrm{in} = 22^\circ\mathrm{C}$ while $T_\mathrm{out} = 32^\circ\mathrm{C}$, so that $\Delta T = 10\,\mathrm{K}$, and the measured cooling load is $Q_\mathrm{cool} = 60\,\mathrm{kW}$. The model gives one equation in two unknowns,

$$60 = 10\,G + P_\mathrm{int}.$$

Possible solutions include $G = 2\,\mathrm{kW/K}$ with $P_\mathrm{int} = 40\,\mathrm{kW}$, $G = 3\,\mathrm{kW/K}$ with $P_\mathrm{int} = 30\,\mathrm{kW}$, and infinitely many others along the line $P_\mathrm{int} = 60 - 10\,G$. Now suppose the measurement is repeated 100 times under identical conditions. The sensitivity matrix has 100 identical rows $(10,\ 1)$; its rank is one, its singular values are $\sqrt{100 \times 101} = 100.5$ and 0, and its null direction is $\Delta\boldsymbol{\theta} \propto (1, -10)^\mathsf{T}$: raising G by $1\,\mathrm{kW/K}$ and lowering P_int by $10\,\mathrm{kW}$ leaves every prediction unchanged. With a measurement error of standard deviation $\sigma_Q = 1\,\mathrm{kW}$, the hundred repetitions

estimate the combination $10\,G + P_{\text{int}}$ with a standard deviation of $\sigma_Q/\sqrt{100} = 0.1\,\text{kW}$, and they say nothing about G and P_{int} separately.

Figure 31.3(a) shows why. The horizontal axis is the temperature difference ΔT in kelvin and the vertical axis the cooling load in kilowatts; each straight line is the model of equation (31.9) for one parameter vector, with slope G and intercept P_{int}. The single measured point at $\Delta T = 10\,\text{K}$ lies on all three lines drawn, and on every other line through it. Repetition makes the point more precise; it does not choose among the lines.

Common mistake

Treating repeated measurements as new information about every parameter. A hundred readings at the same operating condition reduce the noise in the one combination that condition measures, here $10\,G + P_{\text{int}}$, and leave the other combination exactly as unknown as before. Repeated measurements of the same relationship create no new independent parameter information.

31.2.6 Experiment B: a changed operating condition

Now collect measurements at three different temperature differences,

measurement	1	2	3
ΔT (K)	0	10	20
Q_{cool} (kW)	40	60	80

so that

$$\mathbf{y} = \begin{bmatrix} 40 \\ 60 \\ 80 \end{bmatrix}, \qquad \mathbf{S}_\theta = \begin{bmatrix} 0 & 1 \\ 10 & 1 \\ 20 & 1 \end{bmatrix}. \tag{31.10}$$

The columns are not proportional, so $\operatorname{rank}(\mathbf{S}_\theta) = 2$ and both parameters are identifiable under this experimental design. The first observation, at $\Delta T = 0$, gives $P_{\text{int}} = 40\,\text{kW}$ directly; the difference between the second and the first gives

$$G = \frac{60 - 40}{10\,\text{K}}\,\text{kW} = 2\,\text{kW/K};$$

and the third observation confirms the same relationship, $2 \times 20 + 40 = 80\,\text{kW}$. Therefore

$$\hat{\boldsymbol{\theta}} = \begin{bmatrix} \hat{G} \\ \hat{P}_{\text{int}} \end{bmatrix} = \begin{bmatrix} 2\,\text{kW/K} \\ 40\,\text{kW} \end{bmatrix}.$$

Varying the input has made the parameters structurally identifiable.

Figure 31.3(b) draws the result on the same axes. The three measured points lie on one line, the line with $G = 2$ and $P_{\text{int}} = 40$; the other two lines of panel (a), drawn thin, pass through the middle point and miss the outer two.

Why the changed condition works. The single operating condition of Experiment A constrained only one combined heat-flow term. Varying the temperature difference made the conductance contribution $G\,\Delta T$ change from one measurement to the next while the internal heat generation stayed constant, and the observations then separate the two. This is the principle of *experimental excitation*: the inputs are varied so that the parameters act differently on the measurements. In a power system, changing the dispatch or using naturally varied operating conditions can supply information about some network and thermal parameters, but any deliberate operating change must respect the security constraints of the system; a test that risks a limit violation is not an acceptable way to learn a parameter.

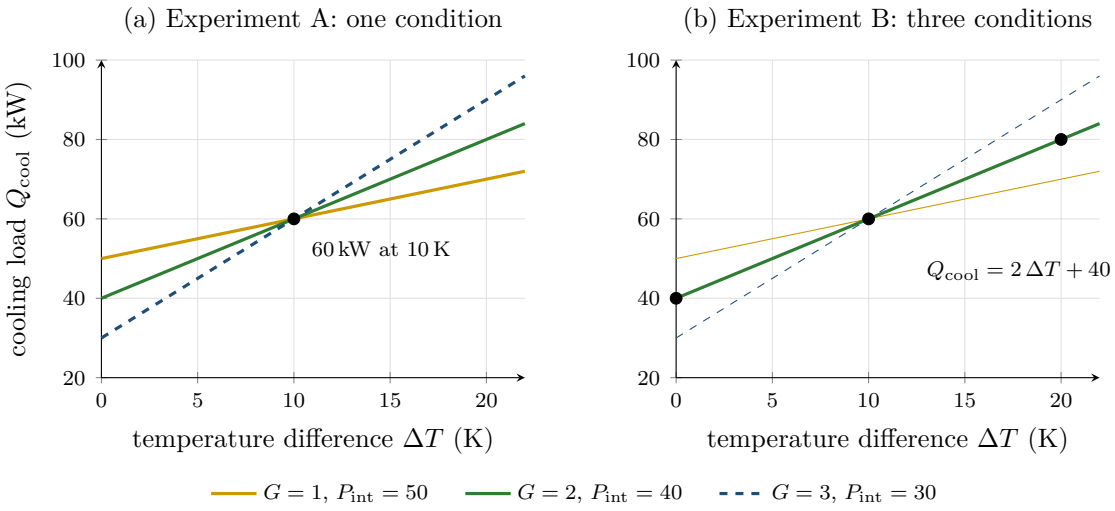


Figure 31.3: The thermal balance $Q_{\mathrm{cool}} = G\,\Delta T + P_{\mathrm{int}}$ (equation (31.9)) for three parameter vectors, G in kW/K and P_{int} in kW (illustrative). (a) One measurement condition lies on every line through the point; the parameters are not identifiable. (b) Three conditions, $\Delta T = 0$, 10 and 20 K, single out the line $G = 2$, $P_{\mathrm{int}} = 40$.

Principle 31.2. *Excitation is worth more than repetition. Parameters are separated by measurements taken at operating conditions under which they act differently, and repeated measurements of the same relationship refine only the combination that relationship already constrains.*

31.3 Fisher information, singular values and Bayesian parameter estimation

Experiment B determines the parameters uniquely with perfect measurements, but real cooling-load measurements contain errors. This section derives how precisely the parameters are known from the Fisher information, interprets the covariance and its correlation as an uncertainty ellipse, reads the weakly constrained directions from the singular values of the scaled sensitivity matrix, compares two designs that are both identifiable, and adds prior information through the Gaussian linear posterior.

31.3.1 From identifiability to confidence intervals

Assume independent Gaussian measurement errors, $\varepsilon_i \sim \mathcal{N}(0, \sigma_Q^2)$ with $\sigma_Q = 1\,\mathrm{kW}$, so that $\mathbf{R} = \sigma_Q^2 \mathbf{I}$ and $\mathbf{W} = \mathbf{R}^{-1} = \mathbf{I}/\sigma_Q^2$. For the linear model $\mathbf{y} = \mathbf{S}_\theta \boldsymbol{\theta} + \boldsymbol{\varepsilon}$ the likelihood is

$$p(\mathbf{y} \mid \boldsymbol{\theta}) \propto \exp\!\left[-\tfrac{1}{2}(\mathbf{y} - \mathbf{S}_\theta \boldsymbol{\theta})^{\mathsf{T}} \mathbf{W} (\mathbf{y} - \mathbf{S}_\theta \boldsymbol{\theta})\right], \tag{31.11}$$

and the log-likelihood is $\ell(\boldsymbol{\theta}) = -\tfrac{1}{2} J(\boldsymbol{\theta}) + \text{const}$, where $J(\boldsymbol{\theta}) = (\mathbf{y} - \mathbf{S}_\theta \boldsymbol{\theta})^{\mathsf{T}} \mathbf{W} (\mathbf{y} - \mathbf{S}_\theta \boldsymbol{\theta})$ is the weighted least-squares objective of Chapter 30.

Chapter 7 of Volume I measured the precision of a maximum likelihood estimate by the curvature of the log-likelihood: for one parameter the Fisher information is $\mathcal{I}(\theta) = \mathbb{E}[-\mathrm{d}^2\ell/\mathrm{d}\theta^2]$, and the large-sample standard error of the estimate is $1/\sqrt{\mathcal{I}}$. For a vector of parameters the second derivative becomes the Hessian matrix. Differentiating the Gaussian log-likelihood once,

$$\nabla_{\boldsymbol{\theta}}\, \ell = \mathbf{S}_\theta^{\mathsf{T}} \mathbf{W} (\mathbf{y} - \mathbf{S}_\theta \boldsymbol{\theta}),$$

and once more,

$$-\nabla^2_{\boldsymbol{\theta}}\,\ell = \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta \quad\Longrightarrow\quad \mathcal{I}(\boldsymbol{\theta}) = \mathbb{E}\big[-\nabla^2_{\boldsymbol{\theta}}\,\ell\big] = \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta. \tag{31.12}$$

For a nonlinear model the Hessian contains a further term, the second derivatives of h weighted by the residuals; its expectation at the true parameters is zero because the errors have zero mean, and the Fisher information is again $\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta$ with $\mathbf{S}_\theta$ evaluated there. For Experiment B with $\sigma_Q = 1\,\mathrm{kW}$,

$$\mathcal{I}(\boldsymbol{\theta}) = \begin{bmatrix} 0 & 10 & 20 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 10 & 1 \\ 20 & 1 \end{bmatrix} = \begin{bmatrix} \sum_i \Delta T_i^2 & \sum_i \Delta T_i \\ \sum_i \Delta T_i & m \end{bmatrix} = \begin{bmatrix} 500 & 30 \\ 30 & 3 \end{bmatrix},$$

in units of $\mathrm{K^2/kW^2}$, $\mathrm{K/kW^2}$ and $\mathrm{kW^{-2}}$ for the entries GG, GP and PP. Its determinant is $500 \times 3 - 30^2 = 1500 - 900 = 600$; a positive determinant confirms full information rank. For Experiment A the same computation gives a singular matrix, determinant zero.

The *Cramér–Rao bound* is the standard result that connects the Fisher information to precision: for any unbiased estimator $\tilde{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$,

$$\mathrm{Cov}(\tilde{\boldsymbol{\theta}}) \succeq \mathcal{I}(\boldsymbol{\theta})^{-1}, \tag{31.13}$$

where $\mathbf{A} \succeq \mathbf{B}$ means that $\mathbf{A} - \mathbf{B}$ is positive semidefinite, so that every linear combination $\mathbf{a}^{\mathsf{T}}\boldsymbol{\theta}$ has a variance of at least $\mathbf{a}^{\mathsf{T}}\mathcal{I}^{-1}\mathbf{a}$. No estimation method can do better than the information the experiment contains.

31.3.2 The parameter covariance and correlation

For the linear Gaussian model the maximum likelihood estimate minimizes $J(\boldsymbol{\theta})$. Setting the gradient to zero gives the normal equations $\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta\,\hat{\boldsymbol{\theta}} = \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{y}$, so that $\hat{\boldsymbol{\theta}} = \mathcal{I}^{-1}\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{y}$, a linear function of $\mathbf{y}$. Its covariance follows from $\mathrm{Cov}(\mathbf{A}\mathbf{y}) = \mathbf{A}\mathbf{R}\mathbf{A}^{\mathsf{T}}$ with $\mathbf{A} = \mathcal{I}^{-1}\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}$:

$$\boldsymbol{\Sigma}_\theta = \mathrm{Cov}(\hat{\boldsymbol{\theta}}) = \mathcal{I}^{-1}\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{R}\mathbf{W}\mathbf{S}_\theta\mathcal{I}^{-1} = \mathcal{I}^{-1}\left(\mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta\right)\mathcal{I}^{-1} = \mathcal{I}(\boldsymbol{\theta})^{-1}, \tag{31.14}$$

using $\mathbf{W}\mathbf{R}\mathbf{W} = \mathbf{W}$. The estimate is unbiased and attains the Cramér–Rao bound; for nonlinear models the same expression holds approximately for large samples. With the inverse of a 2×2 matrix, $[a\ b;\ b\ d]^{-1} = [d\ \ -b;\ -b\ a]/(ad - b^2)$, Experiment B gives

$$\boldsymbol{\Sigma}_\theta = \frac{1}{600}\begin{bmatrix} 3 & -30 \\ -30 & 500 \end{bmatrix} = \begin{bmatrix} 0.005 & -0.05 \\ -0.05 & 0.83333 \end{bmatrix}.$$

The parameter standard deviations are the square roots of the diagonal,

$$\sigma_G = \sqrt{0.005} = 0.0707\,\mathrm{kW/K}, \qquad \sigma_P = \sqrt{0.83333} = 0.9129\,\mathrm{kW},$$

so that, under the model,

$$G = 2.000 \pm 0.071\,\mathrm{kW/K}, \qquad P_{\mathrm{int}} = 40.000 \pm 0.913\,\mathrm{kW},$$

where the $\pm$ values are one standard deviation. They are estimator uncertainties conditional on the proposed model, the experimental inputs and the assumed measurement errors. They do not include any structural discrepancy between the model and the room.

Correlation. The off-diagonal entry is the covariance, $\mathrm{Cov}(\hat{G}, \hat{P}_{\mathrm{int}}) = -0.05\,\mathrm{kW^2/K}$, and the correlation coefficient is

$$\rho_{G,P} = \frac{\mathrm{Cov}(\hat{G}, \hat{P}_{\mathrm{int}})}{\sigma_G\,\sigma_P} = \frac{-0.05}{0.070711 \times 0.912871} = -0.775. \tag{31.15}$$

In general $\rho_{ij} = \Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$. The negative correlation has a physical interpretation: if the estimated conductance is too high, a lower internal heat input partially compensates, and the predicted cooling loads are nearly preserved. The model still distinguishes the parameters, because different temperature differences were measured, but the remaining uncertainty is correlated.

31.3.3 Posterior geometry

Under the Gaussian model the estimation error $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}$ is normal with covariance $\mathcal{I}^{-1}$, so the quadratic form $(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^{\mathsf{T}}\mathcal{I}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$ follows a chi-square distribution with n_θ degrees of freedom. The set

$$\left\{\boldsymbol{\theta} : (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^{\mathsf{T}}\mathcal{I}\,(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \le \chi^2_{n_\theta,0.95}\right\} \tag{31.16}$$

is therefore a 95 % joint confidence region, an ellipse for two parameters with $\chi^2_{2,0.95} = 5.991$. Under a flat prior it is also the 95 % highest-posterior-density region, which is why its shape is called the posterior geometry. For drawing, the ellipse is traced by $G = \hat{G} + c\,\sigma_G \cos t$ and $P_{\mathrm{int}} = \hat{P}_{\mathrm{int}} + c\,\sigma_P(\rho\cos t + \sqrt{1-\rho^2}\sin t)$ for $0 \le t < 2\pi$, with $c = \sqrt{5.991} = 2.448$, the Cholesky factor of $\boldsymbol{\Sigma}_\theta$ applied to a circle.

Figure 31.4 draws these regions in the parameter plane. The horizontal axis is the conductance G and the vertical axis the internal heat P_{int}; the dot is the estimate $(2, 40)$. The small green ellipse is the 95 % region of Experiment B. The large dashed blue ellipse is the region of a narrow design with measurements at 9, 10 and 11 K, analyzed in Section 31.3.5, assumed to return the same estimate. The gray line is the set of parameter vectors consistent with Experiment A, $10\,G + P_{\mathrm{int}} = 60$, along which that experiment has no information at all. Both ellipses slope downward, the compensation of equation (31.15). The narrow design's ellipse is elongated almost exactly along the Experiment A line: measurements near one temperature difference approach the single-condition experiment, and the direction that experiment could not resolve becomes the direction the narrow design resolves poorly.

The correlation gives the direction of the compensation; it does not measure the size of the uncertainty. Two tightly estimated parameters can be highly correlated, and two very uncertain parameters can have a modest correlation. If the measurement error of Experiment B were 0.1 kW instead of 1 kW, both standard deviations would fall tenfold and the correlation would remain -0.775; the ellipse would shrink without changing shape. The covariance magnitudes and the information geometry must be examined together, and the singular values of the next subsection do that.

31.3.4 Singular values and weakly constrained directions

Every real matrix has a singular-value decomposition. For the whitened, scaled sensitivity matrix of equation (31.8),

$$\mathbf{S}_{\mathrm{w}} = \mathbf{U}\,\mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_{n_\theta})\,\mathbf{V}^{\mathsf{T}}, \qquad \sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_{n_\theta} \ge 0, \tag{31.17}$$

where the columns $\mathbf{u}_i$ of $\mathbf{U}$ are orthonormal directions in the space of measurements, the columns $\mathbf{v}_i$ of $\mathbf{V}$ are orthonormal directions in the space of scaled parameters $\boldsymbol{\xi}$, and the σ_i are the singular values. (Singular values carry numerical subscripts; a standard deviation carries the name of its variable, as in σ_G.) The decomposition says that a parameter change along $\mathbf{v}_i$ of unit size changes the whitened predictions by σ_i along $\mathbf{u}_i$. Large singular values are directions

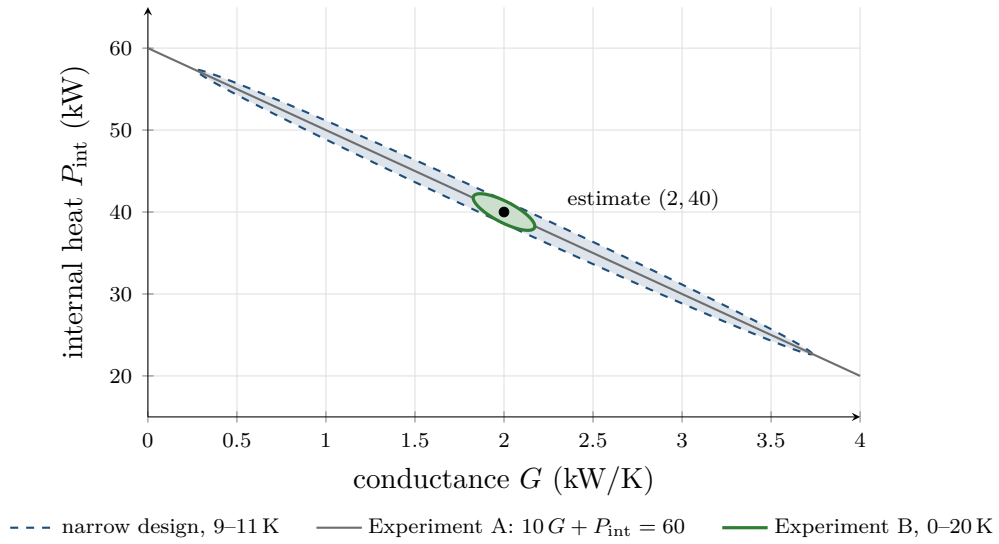


Figure 31.4: Joint 95 % confidence regions of the thermal parameters, equation (31.16) with $\chi^2_{2,0.95} = 5.991$ and $\sigma_Q = 1\,\mathrm{kW}$ (illustrative). Experiment B ($\Delta T = 0, 10, 20\,\mathrm{K}$): $\sigma_G = 0.0707\,\mathrm{kW/K}$, $\sigma_P = 0.913\,\mathrm{kW}$, $\rho = -0.775$. Narrow design ($\Delta T = 9, 10, 11\,\mathrm{K}$): $\sigma_G = 0.707\,\mathrm{kW/K}$, $\sigma_P = 7.09\,\mathrm{kW}$, $\rho = -0.9967$. Gray: the line of parameter vectors that Experiment A cannot distinguish.

that produce strong measurable responses; small ones are directions whose changes produce weak responses; a zero singular value is a locally unidentifiable direction, the null space of equation (31.6).

The connection to the Fisher information follows by substitution. In scaled coordinates the information is $\mathcal{I}_\xi = \mathbf{S}_{\mathrm{w}}^{\mathsf{T}}\mathbf{S}_{\mathrm{w}} = \mathbf{D}_\theta \mathcal{I}(\boldsymbol{\theta})\mathbf{D}_\theta$, and with $\mathbf{U}^{\mathsf{T}}\mathbf{U} = \mathbf{I}$,

$$\mathcal{I}_\xi = \mathbf{V}\,\mathrm{diag}(\sigma_1^2, \ldots, \sigma_{n_\theta}^2)\,\mathbf{V}^{\mathsf{T}}, \qquad \mathrm{Cov}(\hat{\boldsymbol{\xi}}) = \mathcal{I}_\xi^{-1} = \mathbf{V}\,\mathrm{diag}(\sigma_1^{-2}, \ldots, \sigma_{n_\theta}^{-2})\,\mathbf{V}^{\mathsf{T}}. \tag{31.18}$$

The eigenvalues of the Fisher information are the squared singular values, $\lambda_i(\mathcal{I}_\xi) = \sigma_i^2$, its eigenvectors are the $\mathbf{v}_i$, and the standard deviation of the estimate along $\mathbf{v}_i$ is $1/\sigma_i$: the axes of the uncertainty ellipse of Figure 31.4, in scaled coordinates, point along the $\mathbf{v}_i$ with half-lengths proportional to $1/\sigma_i$. The *condition number*

$$\kappa = \frac{\sigma_1}{\sigma_{n_\theta}} \tag{31.19}$$

is the ratio of the longest to the shortest axis. The directions along the smallest singular values are the *weakly constrained directions*, often called sloppy directions: combinations of parameters that the data hardly determine although each parameter separately may seem to matter. A large κ signals poor practical identifiability. For the fragile example of Section 31.1.4, whose parameters are already on a common scale, the singular values are 2.0005 and 4.999×10^{-4} and $\kappa = 4002$.

31.3.5 Two experimental designs compared

Experiment B used $\Delta T = 0$, 10 and 20 K. Consider a second experiment that collects the same number of measurements only near the usual operating condition, at $\Delta T = 9$, 10 and 11 K. Both designs have three measurements and two parameters, and both sensitivity matrices have full rank:

$$\mathbf{S}_\theta^{\mathrm{wide}} = \begin{bmatrix} 0 & 1 \\ 10 & 1 \\ 20 & 1 \end{bmatrix}, \qquad \mathbf{S}_\theta^{\mathrm{narrow}} = \begin{bmatrix} 9 & 1 \\ 10 & 1 \\ 11 & 1 \end{bmatrix}.$$

Their information differs greatly. For the narrow design, $\sum_i \Delta T_i^2 = 81 + 100 + 121 = 302$, so that

$$\mathcal{I}_{\text{narrow}} = \begin{bmatrix} 302 & 30 \\ 30 & 3 \end{bmatrix}, \quad \det = 906 - 900 = 6, \quad \mathcal{I}_{\text{narrow}}^{-1} = \frac{1}{6}\begin{bmatrix} 3 & -30 \\ -30 & 302 \end{bmatrix} = \begin{bmatrix} 0.5 & -5 \\ -5 & 50.3333 \end{bmatrix},$$

with $\sigma_G = \sqrt{0.5} = 0.7071\,\text{kW/K}$, $\sigma_P = \sqrt{50.3333} = 7.0946\,\text{kW}$ and $\rho_{G,P} = -5/(0.70711 \times 7.09460) = -0.9967$. The determinant, a hundred times smaller than the 600 of the wide design, makes the standard deviation of G ten times larger.

Table 31.2 compares the two designs in full. The rows are the structural rank, the determinant of the information matrix, the standard deviations and correlation, and the singular values and condition number of the sensitivity matrix in two forms: raw, in the units kW/K and kW with $\sigma_Q = 1\,\text{kW}$, and scaled by the nominal values, $\mathbf{D}_\theta = \text{diag}(2\,\text{kW/K},\ 40\,\text{kW})$, so that the parameters are relative changes. The last row is the standard deviation of the estimate along the weakest direction in relative units, $1/\sigma_2$ of the scaled matrix.

Table 31.2: Two identifiable designs for the thermal balance, $\sigma_Q = 1\,\text{kW}$ (illustrative). Scaled: $\mathbf{S}_{\text{w}} = \mathbf{R}^{-1/2}\mathbf{S}_\theta\mathbf{D}_\theta$ with $\mathbf{D}_\theta = \text{diag}(2, 40)$; the eigenvalues of the corresponding Fisher information are the squared singular values.

Quantity	Wide: 0, 10, 20 K	Narrow: 9, 10, 11 K
Structural rank	2	2
$\det \mathcal{I}(\boldsymbol{\theta})$	600	6
Conductance standard deviation σ_G	0.0707 kW/K	0.7071 kW/K
Internal heat standard deviation σ_P	0.913 kW	7.095 kW
Parameter correlation $\rho_{G,P}$	-0.775	-0.9967
Raw singular values σ_1, σ_2	22.40, 1.093	17.46, 0.1403
Raw condition number κ	20.5	124.5
Scaled singular values σ_1, σ_2	78.60, 24.93	77.47, 2.529
Scaled condition number κ	3.15	30.6
Weakest-direction standard deviation $1/\sigma_2$ (scaled)	0.040	0.395

Figure 31.5 draws the two standard deviations of each design as bars. Panel (a) is the conductance, panel (b) the internal heat; in each, the left bar is the wide design and the right bar the narrow one, labeled with its value. Both bars grow by about the same factor, ten for G and eight for P_{int}, because the narrow design loses information along the compensating direction, which involves both parameters.

The narrow experiment is identifiable and far less informative: a practical identifiability problem, signaled by a scaled condition number of 30.6 against 3.15 and by a weakest-direction standard deviation ten times larger. Since the variance of $\hat{G}$ falls as the inverse of the number of repetitions, the narrow design would need $0.5/0.005 = 100$ repetitions of its three measurements, 300 readings, to match the precision that the wide design reaches with three. A dataset containing thousands of nearly identical operating conditions may therefore carry less parameter information than a much smaller dataset spanning informative conditions.

Singular vectors and units. Suppose a singular-vector analysis reports the weakest direction as $\mathbf{v} = (0.8,\ -0.6)^{\mathsf{T}}$. It is tempting to read this directly as a physical tradeoff between the two parameters, but the reading depends on the coordinate system. For the narrow design, the weakest right singular vector of the raw matrix is $(0.099,\ -0.995)^{\mathsf{T}}$ in (G, P_{int}), which appears to say that only P_{int} is uncertain; that of the scaled matrix is $(0.894,\ -0.448)^{\mathsf{T}}$ in relative changes,

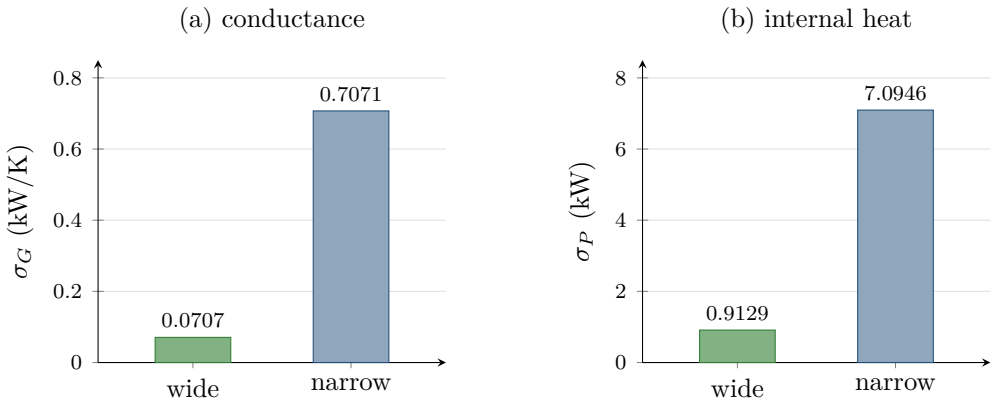


Figure 31.5: Parameter standard deviations of the thermal balance under the wide design ($\Delta T = 0, 10, 20\,\mathrm{K}$) and the narrow design ($\Delta T = 9, 10, 11\,\mathrm{K}$), three measurements each with $\sigma_Q = 1\,\mathrm{kW}$: square roots of the diagonal of $\mathcal{I}^{-1}$, equation (31.14).

which appears to say that G dominates. Both describe nearly the same physical compensation, $\Delta P_{\mathrm{int}} \approx -10\,\Delta G$, the Experiment A line. The numbers themselves depend on the units: for the wide design the raw condition number is 20.5 with P_{int} in kilowatts, 122 with P_{int} in watts, and 3.15 in relative units. Good practice is to report directions in standardized coordinates, for example changes measured in prior standard deviations,

$$\xi_j = \frac{\theta_j - \theta_{j,0}}{\sigma_{\theta_j,\mathrm{prior}}},$$

which makes the components interpretable relative to the prior uncertainty. The conclusion that a direction is or is not identifiable is invariant under any nonsingular reparameterization; the numerical singular values, singular vectors and condition numbers are not.

31.3.6 Prior information and the Gaussian linear posterior

Suppose the measurements constrain only G_1+G_2, as in the two-conductance room. Engineering knowledge of plausible values of G_1 and G_2 separately may still exist, from the wall and roof construction, and Bayesian estimation, Chapter 8 of Volume I, incorporates it through a prior $p(\boldsymbol{\theta})$:

$$p(\boldsymbol{\theta} \mid \mathbf{y}) \propto p(\mathbf{y} \mid \boldsymbol{\theta})\, p(\boldsymbol{\theta}),$$

posterior proportional to likelihood times prior, as in that chapter. A prior can regularize an underdetermined model and produce a finite posterior covariance. That does not mean the data identified the parameters: the posterior combines partial information from the data with prior assumptions, and the distinction must be kept when the result is reported.

With a Gaussian prior $\boldsymbol{\theta} \sim \mathcal{N}(\bar{\boldsymbol{\theta}}_0, \boldsymbol{\Sigma}_0)$ and the linear Gaussian likelihood $\mathbf{y} \mid \boldsymbol{\theta} \sim \mathcal{N}(\mathbf{S}_\theta\boldsymbol{\theta}, \mathbf{R})$, the log posterior is, up to a constant,

$$\ln p(\boldsymbol{\theta} \mid \mathbf{y}) = -\tfrac{1}{2}(\mathbf{y} - \mathbf{S}_\theta\boldsymbol{\theta})^{\mathsf{T}}\mathbf{W}(\mathbf{y} - \mathbf{S}_\theta\boldsymbol{\theta}) - \tfrac{1}{2}(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_0)^{\mathsf{T}}\boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_0).$$

Expanding and collecting the terms quadratic and linear in $\boldsymbol{\theta}$,

$$\ln p(\boldsymbol{\theta} \mid \mathbf{y}) = -\tfrac{1}{2}\boldsymbol{\theta}^{\mathsf{T}}(\boldsymbol{\Sigma}_0^{-1} + \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta)\boldsymbol{\theta} + \boldsymbol{\theta}^{\mathsf{T}}(\boldsymbol{\Sigma}_0^{-1}\bar{\boldsymbol{\theta}}_0 + \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{y}) + \text{const.}$$

A Gaussian density with mean $\mathbf{m}$ and covariance $\mathbf{C}$ has the log form $-\tfrac{1}{2}\boldsymbol{\theta}^{\mathsf{T}}\mathbf{C}^{-1}\boldsymbol{\theta} + \boldsymbol{\theta}^{\mathsf{T}}\mathbf{C}^{-1}\mathbf{m} +$ const; matching the two terms (completing the square) shows that the posterior is Gaussian with

$$\boldsymbol{\Sigma}_{\mathrm{post}} = (\boldsymbol{\Sigma}_0^{-1} + \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta)^{-1}, \qquad \bar{\boldsymbol{\theta}}_{\mathrm{post}} = \boldsymbol{\Sigma}_{\mathrm{post}}(\boldsymbol{\Sigma}_0^{-1}\bar{\boldsymbol{\theta}}_0 + \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{y}). \tag{31.20}$$

The posterior information is the sum of a prior and a data term,

$$\mathcal{I}_{\text{post}} = \boldsymbol{\Sigma}_{\text{post}}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathcal{I}(\boldsymbol{\theta}). \tag{31.21}$$

The decomposition is useful for reporting. Along a direction that the data constrain strongly, $\mathcal{I}(\boldsymbol{\theta})$ dominates and the prior hardly matters; along a weakly constrained direction, the prior dominates the posterior uncertainty. An honest report identifies those prior-dominated directions.

31.3.7 Posterior correlation of two conductances

Return to the two conductances with $Q = (G_1 + G_2)\,\Delta T$. Give both independent priors $G_1, G_2 \sim \mathcal{N}(2, 1^2)$ in kW/K, and suppose one precise observation, $Q = 50\,\text{kW}$ at $\Delta T = 10\,\text{K}$ with $\sigma_Q = 1\,\text{kW}$, which indicates $G_1 + G_2 \approx 5\,\text{kW/K}$ (illustrative values). Then $\mathbf{S}_\theta = (10,\ 10)$, $\mathbf{W} = 1$, and equation (31.21) gives

$$\mathcal{I}_{\text{post}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 100 & 100 \\ 100 & 100 \end{bmatrix} = \begin{bmatrix} 101 & 100 \\ 100 & 101 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_{\text{post}} = \frac{1}{201}\begin{bmatrix} 101 & -100 \\ -100 & 101 \end{bmatrix},$$

the data term alone being singular. The posterior standard deviation of each conductance is $\sqrt{101/201} = 0.709\,\text{kW/K}$, and the posterior correlation is $-100/101 = -0.990$: if G_1 is larger, G_2 must be smaller to preserve the measured total. By equation (31.20) the posterior means are $\bar{\boldsymbol{\theta}}_{\text{post}} = \boldsymbol{\Sigma}_{\text{post}}(2 + 500,\ 2 + 500)^{\mathsf{T}} = (2.498,\ 2.498)^{\mathsf{T}}\,\text{kW/K}$, which add to 4.995.

The two marginal standard deviations hide the structure. The variance of a linear combination $\mathbf{a}^{\mathsf{T}}\boldsymbol{\theta}$ is $\mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}_{\text{post}}\mathbf{a}$; for the sum, $\mathbf{a} = (1, 1)^{\mathsf{T}}$, it is $(101 + 101 - 200)/201 = 0.00995$, a standard deviation of $0.0998\,\text{kW/K}$, and for the difference, $\mathbf{a} = (1, -1)^{\mathsf{T}}$, it is $(101{+}101{+}200)/201 = 2.000$, a standard deviation of $1.414\,\text{kW/K}$. The prior variance of the difference was also $1 + 1 = 2$: the data have constrained the sum very tightly and left the difference exactly where the prior put it. The posterior information matrix has eigenvalues 201 along $(1, 1)$ and 1 along $(1, -1)$, the second equal to the prior information. Reporting the sum with its small uncertainty and the difference as prior dominated is far more informative than reporting two point estimates.

31.4 Model adequacy and model discrepancy

The preceding sections assumed that the model is correct and asked how well its parameters are determined. This section asks whether the model fits at all: it fits an identifiable model to data it cannot explain, tests the residuals, separates model error from measurement error, and shows that a flexible discrepancy term can destroy identifiability.

31.4.1 An identifiable model that does not fit

Suppose the thermal model $Q_{\text{cool}} = G\,\Delta T + P_{\text{int}}$ is fitted to four measurements,

ΔT (K)	0	10	20	30
Q_{cool} (kW)	40	60	80	112

(illustrative). The first three follow $Q_{\text{cool}} = 2\,\Delta T + 40$ exactly. At $\Delta T = 30\,\text{K}$ that line predicts $100\,\text{kW}$, while $112\,\text{kW}$ is measured, an additional $12\,\text{kW}$. Possible explanations include a missing temperature-dependent heat-transfer mechanism, cooling-system performance that changes with the operating point, unmodeled internal heat generation, a sensor error, and a nonrepresentative operating period. Adjusting G and P_{int} cannot make all four measurements agree: the model has two parameters, and no straight line passes through all four points. The best-fit parameters can reduce the residuals but cannot remove their systematic pattern. The model is identifiable, since four distinct temperature differences give a sensitivity matrix of rank two, and the question is now one of model adequacy, not of identifiability.

31.4.2 The least-squares fit

The model is fitted to

$$\mathbf{y} = \begin{bmatrix} 40 \\ 60 \\ 80 \\ 112 \end{bmatrix}, \qquad \mathbf{S}_\theta = \begin{bmatrix} 0 & 1 \\ 10 & 1 \\ 20 & 1 \\ 30 & 1 \end{bmatrix}.$$

For a straight line the normal equations reduce to the familiar form: with the means $\overline{\Delta T} = 15\,\mathrm{K}$ and $\bar{y} = 73\,\mathrm{kW}$, the slope is $\hat{G} = \sum_i (\Delta T_i - \overline{\Delta T})(y_i - \bar{y}) / \sum_i (\Delta T_i - \overline{\Delta T})^2$ and the intercept is $\hat{P}_{\mathrm{int}} = \bar{y} - \hat{G}\,\overline{\Delta T}$. The deviations of ΔT are -15, -5, 5, 15 and those of y are -33, -13, 7, 39, so

$$\hat{G} = \frac{495 + 65 + 35 + 585}{225 + 25 + 25 + 225} = \frac{1180}{500} = 2.36\,\mathrm{kW/K}, \qquad \hat{P}_{\mathrm{int}} = 73 - 2.36 \times 15 = 37.6\,\mathrm{kW}.$$

The predicted observations and the residuals $\mathbf{e} = \mathbf{y} - \mathbf{S}_\theta \hat{\boldsymbol{\theta}}$ are

$$\hat{\mathbf{y}} = \begin{bmatrix} 37.6 \\ 61.2 \\ 84.8 \\ 108.4 \end{bmatrix} \mathrm{kW}, \qquad \mathbf{e} = \begin{bmatrix} 2.4 \\ -1.2 \\ -4.8 \\ 3.6 \end{bmatrix} \mathrm{kW}.$$

Even the optimal parameters cannot reproduce every measurement, and the fit has also moved the first three points off the line they lay on.

Figure 31.6 shows the fit. In panel (a) the horizontal axis is ΔT and the vertical axis the cooling load; the dots are the four measurements, the solid line is the least-squares fit, the dashed line is the line through the first three points, and the dotted curve is the quadratic model M_2 of Section 31.5.1. In panel (b) the vertical axis is the residual of the straight-line fit, drawn as a stem at each measurement, with dotted lines at $\pm 2\sigma_Q = \pm 2\,\mathrm{kW}$. Three of the four residuals lie outside that band, and their signs, positive, negative, negative, positive, form a U-shaped pattern that a curved relationship would produce and random errors would rarely produce.

The four-point fit also has formal standard errors. By equation (31.14), with $\sum \Delta T_i^2 = 1400$, $\sum \Delta T_i = 60$ and $m = 4$, $\mathcal{I} = [1400\ 60;\ 60\ 4]$, $\det \mathcal{I} = 2000$, and the standard deviations are $\sqrt{4/2000} = 0.045\,\mathrm{kW/K}$ for G and $\sqrt{1400/2000} = 0.84\,\mathrm{kW}$ for P_{int}. The estimate $G = 2.36 \pm 0.045\,\mathrm{kW/K}$ looks precise. It excludes the value 2 that three of the four points support by eight standard deviations, and both numbers are conditional on a model the data reject.

Common mistake

Reporting the standard errors of a fit without testing the fit. The covariance $\mathcal{I}^{-1}$ describes the scatter of the estimate if the model is correct and the errors are as assumed. When the model is inadequate, the standard errors measure nothing physical, and a small one gives false confidence.

31.4.3 The residual goodness-of-fit test

Assume independent Gaussian measurement errors with the known standard deviation $\sigma_Q = 1\,\mathrm{kW}$. The weighted residual statistic is the weighted least-squares objective at the estimate,

$$J(\hat{\boldsymbol{\theta}}) = \mathbf{e}^{\mathsf{T}} \mathbf{W} \mathbf{e} = \sum_{i=1}^{m} \frac{e_i^2}{\sigma_Q^2}. \tag{31.22}$$

Substituting,

$$J(\hat{\boldsymbol{\theta}}) = 2.4^2 + (-1.2)^2 + (-4.8)^2 + 3.6^2 = 5.76 + 1.44 + 23.04 + 12.96 = 43.2.$$

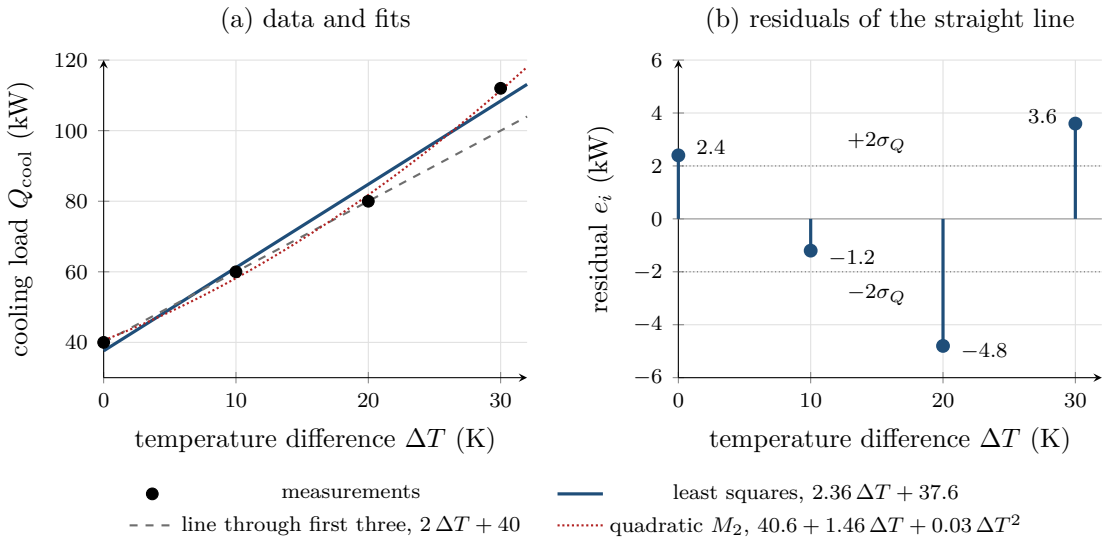


Figure 31.6: An identifiable model that does not fit (illustrative data, $\sigma_Q = 1\,\mathrm{kW}$). (a) Four measurements, the least-squares straight line, the line through the first three points, and the least-squares quadratic of model M_2. (b) Residuals $\mathbf{e} = \mathbf{y} - \mathbf{S}_\theta\hat{\boldsymbol{\theta}}$ of the straight line; their weighted sum of squares is 43.2 against a 99 % point of 9.21 for χ^2_2 (equation (31.23)).

Chapter 30 showed for the state estimator that, when the model is correct and the errors are Gaussian with the assumed covariance, this statistic has a chi-square distribution with $m - n$ degrees of freedom, the number of measurements less the number of fitted unknowns: the fit absorbs n of the m independent error directions and the rest remain in the residual. Here $m = 4$ observations and $n_\theta = 2$ fitted parameters, so

$$J(\hat{\boldsymbol{\theta}}) \sim \chi^2_{m-n_\theta} = \chi^2_2 \quad \text{under the correct linear Gaussian model.} \tag{31.23}$$

The 99th percentile of χ^2_2 is $\chi^2_{2,0.99} = 9.21$ (and the 95th is 5.99). The observed 43.2 is far above it. Because the chi-square distribution with two degrees of freedom is the exponential distribution with mean two, its survival function is $\mathrm{e}^{-x/2}$, and the p-value, the probability of a statistic at least this large under the model, is

$$\Pr(\chi^2_2 \geq 43.2) = \mathrm{e}^{-43.2/2} = \mathrm{e}^{-21.6} = 4.16 \times 10^{-10}.$$

This is overwhelming evidence against the joint assumptions of the linear model and the specified 1 kW Gaussian noise. It does not say which assumption is wrong. The model could be missing physics, the noise could be underestimated, or one or more measurements could be in error. The correct conclusion is model inadequacy under the stated assumptions, not the identification of one particular missing mechanism; the residual pattern of Figure 31.6(b) suggests curvature, and Section 31.5 weighs that suggestion against the alternative of a faulty sensor.

> **Principle 31.3.** *A precise estimate is not necessarily an identifiable or physically credible estimate. Precision describes the scatter of the estimate under the model; credibility requires, in addition, that the model is identifiable under the experiment and that it passes a test of its fit to the data.*

31.4.4 Measurement error and model discrepancy

The model of equation (31.1) assumed that all unexplained variation comes from measurement error. A more general model separates two sources,

$$\mathbf{y} = h(\boldsymbol{\theta}) + \boldsymbol{\delta} + \boldsymbol{\varepsilon}, \tag{31.24}$$

where $\boldsymbol{\varepsilon}$ is the measurement error and $\boldsymbol{\delta}$ is the *model discrepancy*, written as a function $\delta(\cdot)$ of the operating conditions when it is modeled, as in $\delta(\Delta T)$. Measurement error describes the imperfections of the observations: it would vanish with a perfect instrument. Model discrepancy describes the difference between the real physical system and the proposed model even if the parameters were known perfectly: it would remain with a perfect instrument. Examples include missing heat-transfer paths, an incorrect turbulence parameterization, omitted electrical losses, simplified battery aging mechanisms, an incomplete network topology and incorrect control logic. The distinction matters because the two behave differently: independent measurement errors average out over repeated readings, while a discrepancy repeats wherever the operating condition repeats and grows or shrinks with it.

31.4.5 Discrepancy and identifiability

A discrepancy term can itself create identifiability problems. Write the thermal model with a discrepancy function,

$$Q_{\text{cool}} = G\,\Delta T + P_{\text{int}} + \delta(\Delta T).$$

If $\delta(\Delta T)$ may be an arbitrary function, it can explain almost any residual pattern. A discrepancy $\delta(\Delta T) = a\,\Delta T$ is indistinguishable from increasing G by a, and a constant discrepancy is indistinguishable from changing P_{int}: the columns of the sensitivity matrix that represent such discrepancies repeat the columns of the physical parameters, and the rank does not grow. In the extreme, a separate free discrepancy value at each of the four measurements adds four columns, the identity matrix, to the two of $\mathbf{S}_\theta$; the resulting 4×6 matrix has rank four, and two directions are unidentifiable whatever the data. Greater model flexibility can reduce parameter identifiability.

The discrepancy model must therefore be constrained, by physics, by independent evidence, or by a defensible statistical structure; a common choice is a Gaussian-process prior on $\delta(\cdot)$ with a stated amplitude and smoothness, kept orthogonal to the effects of the physical parameters or at least small compared with them. Adding enough free parameters to remove every residual does not establish that the enlarged model is correct.

31.4.6 Model adequacy and predictive usefulness

A model may be inadequate for estimating individual physical parameters and still useful for a particular prediction. In the two-conductance room the measurements constrain $G_{\text{tot}} = G_1 + G_2$ without separating G_1 and G_2. If the only required output is the total heat flow under similar conditions, $Q = G_{\text{tot}}\,\Delta T$, the prediction may be accurate, as the posterior standard deviation of $0.1\,\text{kW/K}$ on the sum showed. If the engineering question is which part of the envelope should be insulated first, the individual conductances become essential, and the same data are insufficient for that decision. Chapter 29 reached the same conclusion for states: what matters is whether the decision-relevant quantity is sufficiently constrained. The converse caution comes from Chapter 6 of Volume I: a model that fits well within the observed range of conditions is not thereby credible for a prediction outside it, where its form, not the data, decides the answer.

> **Principle 31.4.** *Model sufficiency depends on the decision-relevant quantity. Not every state or parameter needs to be identifiable; the evidence must be enough to answer the specific engineering question with appropriate uncertainty.*

31.5 Bayesian model comparison

When an adequacy test fails, several physical explanations usually compete. This section states three candidate models for the unexpected cooling load, shows why the smallest residual cannot

choose among them, derives the Bayesian evidence and its Laplace approximation, and compares the models by their Bayes factors.

31.5.1 Competing physical hypotheses

Three candidate models are considered for the four measurements of Section 31.4.1, with $i = 1, \dots, 4$ indexing the measurements:

$$M_1: \quad Q_i = G\,\Delta T_i + P_{\text{int}}, \tag{31.25}$$

$$M_2: \quad Q_i = G\,\Delta T_i + P_{\text{int}} + k\,\Delta T_i^2, \tag{31.26}$$

$$M_3: \quad Q_i = G\,\Delta T_i + P_{\text{int}} + b_{\text{s}}\,\mathbf{1}_{\{i=4\}}, \tag{31.27}$$

where k in kW/K^2 is the coefficient of an additional temperature-dependent heat transfer, b_{s} in kW is a bias of the sensor at the fourth measurement, and $\mathbf{1}_{\{i=4\}}$ is one for the fourth measurement and zero otherwise. All three are possible explanations, with different physical interpretations: M_1 states that the existing physics and unbiased measurements are sufficient; M_2 that an additional temperature-dependent mechanism exists; M_3 that the fourth measurement has a systematic error. A diagnostic procedure should compare these explanations rather than force one model to fit. Each is a linear model in its parameters, so every computation of the previous sections applies to it.

31.5.2 Fit alone is insufficient

Models M_2 and M_3 each add one parameter to M_1, and M_1 is the special case $k = 0$, respectively $b_{\text{s}} = 0$. A model that contains another as a special case, a nested model, can always reproduce the smaller model's best fit, so its minimized objective is at least as small:

$$\min_{\boldsymbol{\theta}} J_{M_2}(\boldsymbol{\theta}) \le \min_{\boldsymbol{\theta}} J_{M_1}(\boldsymbol{\theta}), \qquad \min_{\boldsymbol{\theta}} J_{M_3}(\boldsymbol{\theta}) \le \min_{\boldsymbol{\theta}} J_{M_1}(\boldsymbol{\theta}).$$

For the four measurements, the least-squares fits give $J = 43.2$ for M_1; $J = 7.2$ for M_2, with $\hat{G} = 1.46\,\text{kW/K}$, $\hat{P}_{\text{int}} = 40.6\,\text{kW}$, $\hat{k} = 0.03\,\text{kW/K}^2$ and residuals $(-0.6,\ 1.8,\ -1.8,\ 0.6)\,\text{kW}$; and $J = 0$ for M_3, with $\hat{G} = 2$, $\hat{P}_{\text{int}} = 40$ and $\hat{b}_{\text{s}} = 12\,\text{kW}$, since three parameters can pass exactly through four points of which three are collinear. A better fit alone does not prove that the added mechanism is real: a more flexible model can fit noise, and M_3 was moreover chosen after seeing which measurement deviates. This motivates comparison methods that account for complexity, predictive performance and uncertainty.

Common mistake

Accepting the model with the lowest residual. A model with more parameters fits at least as well as any model nested in it, whether or not its extra mechanism exists. Complexity, the plausibility of the parameter values, validation on new conditions and the alternative explanations all bear on the choice.

31.5.3 The evidence and its Laplace approximation

Chapter 8 of Volume I outlined Bayesian model comparison: the posterior probability of model M_k in a candidate set is proportional to its prior probability times its marginal likelihood, or *evidence*,

$$p(\mathbf{y} \mid M_k) = \int p(\mathbf{y} \mid \boldsymbol{\theta}_k, M_k)\, p(\boldsymbol{\theta}_k \mid M_k)\, \mathrm{d}\boldsymbol{\theta}_k, \qquad \Pr(M_k \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid M_k)\,\Pr(M_k)}{\sum_j p(\mathbf{y} \mid M_j)\,\Pr(M_j)}, \tag{31.28}$$

where $\boldsymbol{\theta}_k$ is the parameter vector of model M_k. The evidence averages the likelihood over the model's prior parameter distribution; it is not the likelihood at the best-fitting parameters.

A highly flexible model may fit the data well near one parameter setting, but its evidence is reduced if most of its prior parameter space fits poorly.

The integral rarely has a closed form. When the posterior is well behaved near an interior mode $\tilde{\boldsymbol{\theta}}$, write $g(\boldsymbol{\theta}) = \ln[p(\mathbf{y} \mid \boldsymbol{\theta}, M)\, p(\boldsymbol{\theta} \mid M)]$ for the logarithm of the integrand. Its gradient vanishes at the mode, so a second-order Taylor expansion there is

$$g(\boldsymbol{\theta}) \approx g(\tilde{\boldsymbol{\theta}}) - \tfrac{1}{2}(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}})^{\mathsf{T}} \mathcal{I}_{\text{post}} (\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}), \qquad \mathcal{I}_{\text{post}} = -\nabla^2_{\boldsymbol{\theta}}\, g(\tilde{\boldsymbol{\theta}}),$$

where $\mathcal{I}_{\text{post}}$, the negative Hessian of the log posterior density at the mode, is the posterior information of equation (31.21) evaluated there. Exponentiating and integrating with the Gaussian integral $\int \exp(-\tfrac{1}{2}\mathbf{z}^{\mathsf{T}}\mathbf{A}\mathbf{z})\,\mathrm{d}\mathbf{z} = (2\pi)^{n/2}(\det \mathbf{A})^{-1/2}$ for a positive definite $n \times n$ matrix $\mathbf{A}$ gives the *Laplace approximation*

$$p(\mathbf{y} \mid M) \approx p(\mathbf{y} \mid \tilde{\boldsymbol{\theta}}, M)\, p(\tilde{\boldsymbol{\theta}} \mid M)\, (2\pi)^{n_\theta/2} \left(\det \mathcal{I}_{\text{post}}\right)^{-1/2}. \tag{31.29}$$

For a linear model with Gaussian prior and errors the log integrand is exactly quadratic and the approximation is exact. It requires a positive definite curvature at the mode; strong degeneracy, several modes, a mode on the boundary of the parameter domain or nearly unidentifiable parameters invalidate the simple form.

In logarithms, equation (31.29) separates the evidence into a fit term and a penalty,

$$\ln p(\mathbf{y} \mid M) \approx \underbrace{\ln p(\mathbf{y} \mid \tilde{\boldsymbol{\theta}}, M)}_{\text{fit}} + \underbrace{\ln p(\tilde{\boldsymbol{\theta}} \mid M) + \tfrac{n_\theta}{2} \ln 2\pi - \tfrac{1}{2} \ln \det \mathcal{I}_{\text{post}}}_{\text{Occam factor}}. \tag{31.30}$$

For a Gaussian prior with covariance $\boldsymbol{\Sigma}_0$ the second term equals $\tfrac{1}{2}\ln(\det \boldsymbol{\Sigma}_{\text{post}} / \det \boldsymbol{\Sigma}_0) - \tfrac{1}{2} d^2$, where $d^2 = (\tilde{\boldsymbol{\theta}} - \bar{\boldsymbol{\theta}}_0)^{\mathsf{T}} \boldsymbol{\Sigma}_0^{-1} (\tilde{\boldsymbol{\theta}} - \bar{\boldsymbol{\theta}}_0)$: the logarithm of the ratio of the posterior to the prior parameter volume, less the cost of a mode far from the prior. Each parameter that the data pin down much more tightly than the prior did costs the model evidence; that is how the evidence penalizes complexity without an explicit count of parameters.

31.5.4 Bayes factors

Two models are compared by the ratio of their evidences, the *Bayes factor*

$$B_{12} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_2)}, \qquad \frac{\Pr(M_1 \mid \mathbf{y})}{\Pr(M_2 \mid \mathbf{y})} = B_{12}\, \frac{\Pr(M_1)}{\Pr(M_2)}, \tag{31.31}$$

which converts prior odds into posterior odds; B_{jk} is defined in the same way for any pair, and $B_{kj} = 1/B_{jk}$. If $B_{21} = 20$, the observations are twenty times as probable under M_2 as under M_1, after averaging over each model's parameter prior. This does not mean a 95.2 % probability that M_2 is true: $20/21 = 0.952$ is the posterior probability of M_2 only if M_1 and M_2 are the only candidates and had equal prior probabilities, and even then it is a probability relative to that candidate set. With prior probability 0.2 on M_2 the posterior odds are $20 \times 0.25 = 5$ and the probability is 0.83.

Table 31.3 gives a scale in common use for reading a Bayes factor. The first column is the factor in favor of one model, the second its value as $2 \ln B$, the scale on which the deviance and chi-square statistics are also read, and the third the conventional description of the strength of the evidence. The scale is a reading aid, not a decision rule; a decision depends also on prior probabilities and on consequences.

Worked comparison of the three models. Give the parameters independent illustrative priors $G \sim \mathcal{N}(2, 1^2)$ in kW/K, $P_{\text{int}} \sim \mathcal{N}(40, 10^2)$ in kW, $k \sim \mathcal{N}(0, 0.02^2)$ in kW/K^2 (an extra 18 kW at 30 K is one prior standard deviation), and $b_{\text{s}} \sim \mathcal{N}(0, 10^2)$ in kW, with $\sigma_Q = 1$ kW.

Table 31.3: A conventional scale for the strength of evidence expressed by a Bayes factor B_{12} in favor of M_1 over M_2.

B_{12}	$2\ln B_{12}$	Evidence for M_1
1 to 3	0 to 2	not worth more than a bare mention
3 to 20	2 to 6	positive
20 to 150	6 to 10	strong
above 150	above 10	very strong

Each model is linear and Gaussian, so equation (31.29) is exact, and the evidence is also the density of $\mathbf{y}$ under $\mathcal{N}(\mathbf{S}_\theta\bar{\boldsymbol{\theta}}_0, \mathbf{R} + \mathbf{S}_\theta\boldsymbol{\Sigma}_0\mathbf{S}_\theta^{\mathsf{T}})$, which gives the same numbers. Table 31.4 lists the results: for each model, the number of parameters, the minimized residual statistic, the two terms of equation (31.30) at the posterior mode, the log evidence, and the Bayes factor against M_1. One row of the arithmetic, for M_1: the posterior mode is $(2.359,\ 37.63)$, where $J = 43.20$ and the log-likelihood is $-\frac{1}{2}(43.20) - 2\ln 2\pi = -25.28$; the log prior density there is -4.23; $(n_\theta/2)\ln 2\pi = \ln 2\pi = 1.84$; and $\det \mathcal{I}_{\text{post}} = 2018$, so $-\frac{1}{2}\ln 2018 = -3.81$. The Occam factor is $-4.23 + 1.84 - 3.81 = -6.20$ and the log evidence $-25.28 - 6.20 = -31.48$.

Table 31.4: Evidence for the three explanations of the four cooling-load measurements, with the illustrative priors stated in the text and $\sigma_Q = 1\,\text{kW}$; natural logarithms. The fit and Occam terms are those of equation (31.30) at the posterior mode.

Model	n_θ	$\min J$	fit	Occam	$\ln p(\mathbf{y} \mid M_k)$	B_{k1}
M_1 straight line	2	43.2	-25.28	-6.20	-31.48	1
M_2 quadratic	3	7.2	-7.37	-8.62	-15.99	5.3×10^6
M_3 sensor bias	3	0	-3.70	-8.50	-12.20	2.4×10^8

Both added mechanisms pay an Occam penalty 2.3 to 2.4 larger than M_1's for their extra parameter, and both gain far more in fit: by Table 31.3 the evidence against the straight line is very strong under either alternative, the Bayesian counterpart of the chi-square rejection. Between the two alternatives, $\ln B_{32} = -12.20 + 15.99 = 3.79$, so $B_{32} = 44$: strong evidence for the sensor bias over the quadratic, and with equal prior probabilities on the three models the posterior probabilities are 0.978 for M_3, 0.022 for M_2 and below 10^{-8} for M_1.

That conclusion is fragile in two ways the chapter's argument predicts. First, it depends on the parameter priors of the added terms, which the four points barely constrain: with prior standard deviations 0.1 for k and 50 kW for b_{s}, $B_{32} = 32$; with 0.005 and 3 kW, $B_{32} = 639$. Second, M_3 placed the bias at the fourth measurement after the data showed that measurement to be the odd one. Had the bias model been stated beforehand as a bias at one of the four sensors, each position equally likely a priori, its evidence would be the average of the four positions' evidences, $\ln p = -13.58$, and the factor against M_2 would fall to 11, positive rather than strong. The data cannot settle the matter, and the two models disagree sharply where no measurement exists: at $\Delta T = 25\,\text{K}$ the fitted M_2 predicts 95.9 kW and M_3 predicts 90 kW, six measurement standard deviations apart. One further measurement there, or a second sensor beside the fourth, would discriminate between them; choosing such measurements is the subject of a later chapter on the value of information.

Principle 31.5. *Bayesian model comparison weighs competing explanations by their evidence, the likelihood averaged over each model's prior. Its conclusions are relative to the candidate set and depend on the parameter priors, so a prior sensitivity analysis and a predictive check on new conditions belong to every such comparison.*

31.6 Applications, design of experiments and the credible modeling workflow

The ideas of the chapter apply wherever a reliability result rests on an estimated physical parameter. This section applies them to transmission line ratings, battery energy storage and reliability predictions, tabulates how the range of operating conditions controls the precision of the thermal balance, and assembles the checks into a workflow with three kinds of inconclusive result.

31.6.1 Transmission line ratings

The conductor heat balance of Chapter 1 of Volume I is

$$mc_p \frac{\mathrm{d}T_\mathrm{c}}{\mathrm{d}t} = I^2 R(T_\mathrm{c}) + q_\mathrm{solar} - q_\mathrm{conv} - q_\mathrm{rad},$$

with conductor temperature T_c, heat capacity mc_p, current I, temperature-dependent resistance $R(T_\mathrm{c})$ (not the reliability function), solar gain q_solar, and convective and radiative losses q_conv and q_rad. Its parameters may include the convective heat-transfer coefficients, the surface emissivity, the solar absorptivity, the conductor resistance, the thermal capacitance and wind-exposure corrections for the span. If current and conductor temperature are measured but the wind and solar conditions vary little during the measurement period, several of these parameters act on the measurements in nearly the same way and are strongly correlated, exactly as G and P_int are under the narrow design. The model may then predict the conductor temperature well under the observed conditions while being poorly constrained under a new weather regime, which is where a weather-dependent rating matters most.

A reliability analysis of the line should keep four uncertainties apart, because they affect the probability of exceeding the thermal limit T_max in different ways: the uncertainty in the current operating temperature, a state; the uncertainty in the model parameters; the uncertainty in future weather, an input; and the structural inadequacy of the thermal model, a discrepancy. The first shrinks with better state estimation, the second with excitation, the third only with a better forecast, and the fourth only with a better model.

31.6.2 Battery energy storage

A lumped thermal model of a battery module is

$$C_\mathrm{th} \frac{\mathrm{d}T_\mathrm{b}}{\mathrm{d}t} = I^2 R_\mathrm{int} - G\,(T_\mathrm{b} - T_\mathrm{a}),$$

with module temperature T_b, ambient temperature T_a, thermal capacitance C_th in kJ/K, internal resistance R_int (an electrical resistance, not a reliability) and conductance to ambient G, with unknown parameters $\boldsymbol{\theta} = (C_\mathrm{th}, R_\mathrm{int}, G)^\mathsf{T}$. With a constant current and slowly changing temperature, the module reaches the steady state

$$T_\mathrm{b} - T_\mathrm{a} = \frac{I^2 R_\mathrm{int}}{G},$$

which constrains only the ratio R_int/G at the fixed current. A transient, such as the response to a step in current, $T_\mathrm{b} - T_\mathrm{a} = (I^2 R_\mathrm{int}/G)(1 - \mathrm{e}^{-t\,G/C_\mathrm{th}})$ from ambient, adds the rate G/C_th, the inverse of the thermal time constant. The three original parameters remain unidentifiable from current and temperature alone under any current profile: dividing the model by C_th gives $\mathrm{d}T_\mathrm{b}/\mathrm{d}t = (R_\mathrm{int}/C_\mathrm{th})I^2 - (G/C_\mathrm{th})(T_\mathrm{b} - T_\mathrm{a})$, so multiplying all three parameters by a common factor leaves every temperature trajectory unchanged. An additional measurement or an independent constraint fixes the scale, for example R_int from the voltage response to a current step. A model may thus reproduce the observed temperature trajectory exactly while leaving the underlying physical parameters undetermined.

31.6.3 Reliability predictions under parameter uncertainty

Suppose a failure criterion depends on uncertain parameters, such as a conductor temperature $T_{\mathrm{c}}(\boldsymbol{\theta}, \mathbf{u})$ compared with a limit $T_{\max}(\boldsymbol{\theta})$ that may itself depend on them. The probability of exceeding the limit given the evidence is the indicator of exceedance averaged over the posterior,

$$\Pr(T_{\mathrm{c}} > T_{\max} \mid \mathbf{y}) = \int \mathbf{1}_{\{T_{\mathrm{c}}(\boldsymbol{\theta},\mathbf{u}) > T_{\max}(\boldsymbol{\theta})\}}\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}, \tag{31.32}$$

the same averaging that gave the posterior predictive reliability in Chapter 8 of Volume I. The result depends on the posterior uncertainty, not merely on the best-fit parameters.

For the thermal balance, take as an illustrative failure criterion that the cooling load at $\Delta T = 30\,\mathrm{K}$ exceeds a cooling capacity of $C = 105\,\mathrm{kW}$ assigned to the enclosure. The predicted load is $\mathbf{s}^{\mathsf{T}}\boldsymbol{\theta}$ with $\mathbf{s} = (30,\ 1)^{\mathsf{T}}$, mean $2 \times 30 + 40 = 100\,\mathrm{kW}$ under both designs, and variance $\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}_{\theta}\mathbf{s}$. For the wide design, $900 \times 0.005 + 2 \times 30 \times (-0.05) + 0.8333 = 2.333$, a standard deviation of $1.53\,\mathrm{kW}$; for the narrow design, $900 \times 0.5 + 2 \times 30 \times (-5) + 50.333 = 200.3$, a standard deviation of $14.2\,\mathrm{kW}$. With a Gaussian posterior, equation (31.32) gives $\Pr(Q_{\mathrm{cool}} > 105 \mid \mathbf{y}) = 1 - \Phi(5/1.53) = 0.0005$ for the wide design and $1 - \Phi(5/14.2) = 0.36$ for the narrow one: the same best fit, and probabilities of exceedance seven hundred times apart. Both numbers are conditional on the model, and Section 31.4.1 showed a measurement of $112\,\mathrm{kW}$ at that very condition. If the model is inadequate, the posterior under it can still produce overconfident predictions, and the wide design's confident 0.0005 is the more misleading of the two. Identifiability and model adequacy thus reach directly into the physical reliability risk.

31.6.4 Design of experiments: the range of temperature variation

The precision of the thermal balance depends on how far the measured conditions spread. Consider three measurements at $\Delta T = 10 - s$, 10 and $10 + s$ kelvin, a spread s around the usual condition, with independent errors of standard deviation σ_Q. Then $\sum_i \Delta T_i = 30$, $\sum_i \Delta T_i^2 = 300 + 2s^2$, and $\det[\mathbf{S}_{\theta}^{\mathsf{T}}\mathbf{S}_{\theta}] = 3(300 + 2s^2) - 30^2 = 6s^2$, so that equation (31.14) gives in closed form

$$\sigma_G = \frac{\sigma_Q}{\sqrt{2}\,s}, \qquad \sigma_P = \sigma_Q \sqrt{\frac{50}{s^2} + \frac{1}{3}}, \qquad \rho_{G,P} = -\frac{1}{\sqrt{1 + s^2/150}}. \tag{31.33}$$

Table 31.5 evaluates the design for $\sigma_Q = 1\,\mathrm{kW}$ and six spreads, from the narrow design ($s = 1\,\mathrm{K}$) to Experiment B ($s = 10\,\mathrm{K}$). The columns are the spread, the three conditions, the two standard deviations, the correlation, the two singular values of the scaled matrix with $\mathbf{D}_{\theta} = \mathrm{diag}(2\,\mathrm{kW/K},\ 40\,\mathrm{kW})$, its condition number, and, for comparison, the condition number of the raw matrix in kW/K and kW. One row of the arithmetic, for $s = 5\,\mathrm{K}$: $\sum \Delta T_i^2 = 350$, $\det = 150$, $\sigma_G^2 = 3/150 = 0.02$, $\sigma_P^2 = 350/150 = 2.333$ and $\mathrm{Cov} = -30/150 = -0.2$, giving $\sigma_G = 0.141\,\mathrm{kW/K}$, $\sigma_P = 1.53\,\mathrm{kW}$ and $\rho = -0.2/(0.1414 \times 1.5275) = -0.926$.

The trend is the one equation (31.33) predicts. The standard deviation of G falls in inverse proportion to the spread, tenfold from $s = 1$ to $s = 10\,\mathrm{K}$; that of P_{int} falls almost as fast and levels off at $\sigma_Q/\sqrt{3} = 0.58\,\mathrm{kW}$, the precision of an average of three readings; and the correlation weakens from -0.997 to -0.775. The largest singular value hardly moves, while the smallest grows in proportion to s: widening the conditions adds information only along the weakly constrained direction, which is where it was missing, and the condition number falls accordingly. Both parameters remain structurally identifiable for every nonzero spread. The measurement precision acts separately: every standard deviation in the table is proportional to σ_Q, so with $\sigma_Q = 3\,\mathrm{kW}$ the wide design gives $\sigma_G = 0.212\,\mathrm{kW/K}$ and the narrow one 2.12, while the correlations and condition numbers are unchanged. Parameter accuracy depends jointly on experimental excitation and measurement precision.

Table 31.5: Precision of the thermal balance against the spread s of three measurement conditions centered at 10 K, $\sigma_Q = 1$ kW (illustrative). Scaled singular values and κ use $\mathbf{D}_\theta = \mathrm{diag}(2, 40)$; raw κ uses the units kW/K and kW. Rows $s = 1$ and $s = 10$ are the narrow design and Experiment B; the others are added to show the trend.

s (K)	ΔT_i (K)	σ_G (kW/K)	σ_P (kW)	$\rho_{G,P}$	σ_1	σ_2	κ scaled	κ raw
1	9, 10, 11	0.707	7.09	−0.997	77.47	2.53	30.6	124.5
2	8, 10, 12	0.354	3.58	−0.987	77.50	5.06	15.3	63.5
3	7, 10, 13	0.236	2.43	−0.971	77.55	7.58	10.2	43.7
5	5, 10, 15	0.141	1.53	−0.926	77.72	12.61	6.17	28.8
7	3, 10, 17	0.101	1.16	−0.868	77.99	17.59	4.43	23.3
10	0, 10, 20	0.071	0.91	−0.775	78.60	24.93	3.15	20.5

31.6.5 A credible modeling workflow

A defensible reliability analysis that rests on estimated parameters passes through nine stages, and Figure 31.7 draws them. Each box names a stage and its content; the arrows give the order, reading along the first row, back along the second and along the third; the dashed arrow from the last stage to the third shows that recommended measurements, once taken, are audited and analyzed again. Stages 4 and 5 are the identifiability checks of Sections 31.1 to 31.3, stage 7 the adequacy checks of Sections 31.4 and 31.5, and stage 8 the propagation of equation (31.32).

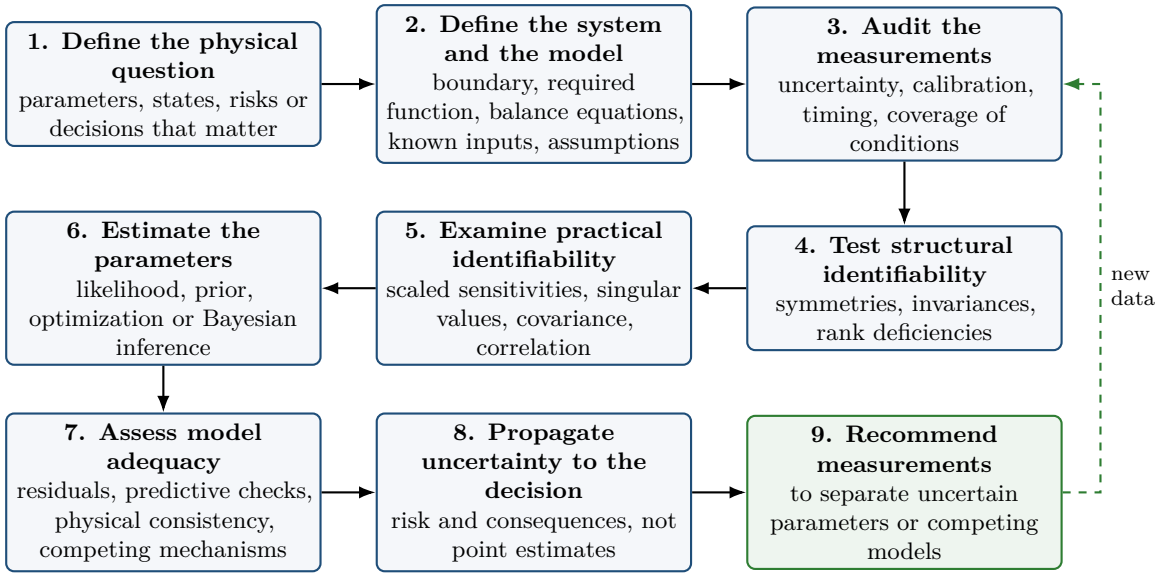


Figure 31.7: A credible modeling workflow for reliability analyses that rest on estimated physical parameters. Stages 4 and 5 test identifiability, stage 7 tests adequacy, and stage 9 closes the loop by targeting the measurements that would resolve what remains uncertain.

31.6.6 Three kinds of inconclusive results

An analysis that does not deliver a usable parameter can fail in three different ways, and Table 31.6 separates them. Each row is a finding; the middle column says what it means and the right column what response it calls for. The three conclusions should not be conflated. Adding more measurements of the same observable combination may reduce noise but cannot resolve a structural nonidentifiability, as Experiment A showed. Adding a parameter may improve the fit and worsen identifiability, as the discrepancy term did. And precise parameter estimates under

an incorrect model do not make its predictions reliable, as the four-point fit and the exceedance probability of Section 31.6.3 showed.

Table 31.6: Three kinds of inconclusive result and the response each calls for.

Finding	Interpretation	Appropriate response
Structurally unidentifiable	distinct parameter combinations produce identical observations	add independent measurements or change the experimental design
Practically weakly identifiable	a unique solution exists but its uncertainty is large	improve excitation, measurement quality or data coverage
Model inadequate	no plausible parameter choice explains the observations at the assumed uncertainty	investigate missing physics, erroneous data or competing model structures

Reliability inference should use only the information that the measurements and credible physical assumptions actually support. The next chapter turns from assessing the information in existing data to choosing the next measurement: Bayesian updating under new evidence and the expected value of the information a measurement would provide.

31.7 What to retain

Definition 31.1 (Credible parameter estimate)**.** A parameter estimate is credible when the parameter is structurally identifiable under the model and the experiment (the sensitivity matrix has full rank), practically identifiable from the actual data (the inverse Fisher information, read through the singular values of the scaled sensitivity matrix, is small enough for the decision), and the model is adequate (its residuals pass a goodness-of-fit test and it is preferred to the competing explanations); a precise estimate is not necessarily an identifiable or physically credible estimate.

Table 31.7 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 31.7: The terms of identifiability, parameter uncertainty and model adequacy, and where each is defined. Units for the thermal example.

Term	Expression	Units	Meaning	Defined in
Forward model	$\mathbf{y} = h(\boldsymbol{\theta}, \mathbf{u}) + \boldsymbol{\varepsilon}$	kW	measurements from parameters	31.1.2
Structural identifiability	$h(\boldsymbol{\theta}_1) = h(\boldsymbol{\theta}_2) \Rightarrow \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2$	–	uniqueness with ideal data	31.1.3
Practical identifiability	small $\mathrm{Cov}(\hat{\boldsymbol{\theta}})$	–	useful precision with real data	31.1.4
Sensitivity matrix	$\mathbf{S}_\theta = \partial h / \partial \boldsymbol{\theta}$	varies	response to each parameter	31.2.1
Local rank condition	$\mathrm{rank}(\mathbf{S}_\theta) = n_\theta$	–	local identifiability	31.2.2
Scaled sensitivity matrix	$\mathbf{S}_\mathrm{w} = \mathbf{R}^{-1/2}\mathbf{S}_\theta\mathbf{D}_\theta$	–	dimensionless form	31.2.3
Fisher information	$\mathcal{I}(\boldsymbol{\theta}) = \mathbf{S}_\theta^\mathsf{T}\mathbf{W}\mathbf{S}_\theta$	varies	curvature of the log-likelihood	31.3.1
Cramér–Rao bound	$\mathrm{Cov}(\tilde{\boldsymbol{\theta}}) \succeq \mathcal{I}^{-1}$	–	best attainable precision	31.3.1
Parameter covariance	$\boldsymbol{\Sigma}_\theta = \mathcal{I}(\boldsymbol{\theta})^{-1}$	varies	estimator uncertainty	31.3.2
Parameter correlation	$\rho_{ij} = \Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$	–	compensation between parameters	31.3.2
Confidence ellipse	$\Delta\boldsymbol{\theta}^\mathsf{T}\mathcal{I}\,\Delta\boldsymbol{\theta} \le \chi^2_{n_\theta,0.95}$	–	posterior geometry	31.3.3
Singular values	$\mathbf{S}_\mathrm{w} = \mathbf{U}\,\mathrm{diag}(\sigma_i)\,\mathbf{V}^\mathsf{T}$	–	response along each direction	31.3.4
Condition number	$\kappa = \sigma_1/\sigma_{n_\theta}$	–	spread of information	31.3.4
Weakly constrained direction	$\mathbf{v}_i$ with small σ_i	–	sloppy parameter combination	31.3.4
Gaussian linear posterior	$\boldsymbol{\Sigma}_\mathrm{post} = (\boldsymbol{\Sigma}_0^{-1} + \mathcal{I})^{-1}$	varies	prior plus data information	31.3.6
Residual statistic	$J(\hat{\boldsymbol{\theta}}) = \mathbf{e}^\mathsf{T}\mathbf{W}\mathbf{e} \sim \chi^2_{m-n_\theta}$	–	goodness-of-fit test	31.4.3
Model discrepancy	$\mathbf{y} = h(\boldsymbol{\theta}) + \boldsymbol{\delta} + \boldsymbol{\varepsilon}$	kW	model error, not sensor error	31.4.4
Evidence	$\int p(\mathbf{y} \mid \boldsymbol{\theta}, M)\, p(\boldsymbol{\theta} \mid M)\, \mathrm{d}\boldsymbol{\theta}$	–	likelihood averaged over prior	31.5.3
Laplace approximation	$p(\mathbf{y} \mid \tilde{\boldsymbol{\theta}})\, p(\tilde{\boldsymbol{\theta}}) \times (2\pi)^{n_\theta/2} (\det \mathcal{I}_\mathrm{post})^{-1/2}$	–	evidence from posterior curvature	31.5.3
Bayes factor	$B_{12} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_2)$	–	relative support of two models	31.5.4
Design of experiments	choice of conditions $\mathbf{u}$	–	excitation for identifiability	31.6.4

Further reading. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*: inverse problems and uncertainty. Kaipio and Somersalo, *Statistical and Computational Inverse Problems*: Bayesian inverse problems and regularization. Gelman and coauthors, *Bayesian Data Analysis*: posterior inference, model checking and uncertainty. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*: parameter inference, sensitivity and predictive uncertainty.

31.8 Exercises

Exercise 31.1. (a) State what structural identifiability is: whether distinct parameters can produce identical ideal observations, whether the optimizer converges quickly, or whether the sensors have low noise. (Whether distinct admissible parameter vectors can produce identical ideal observations under the specified model and experiment.) (b) What does a rank-deficient parameter sensitivity matrix indicate? (At least one locally unidentifiable parameter direction, a perturbation that changes no prediction to first order.) (c) Explain why a model whose sensitivity matrix has full rank can still have poor practical identifiability. (Some directions of the scaled sensitivity matrix may respond extremely weakly relative to the measurement noise, a small singular value; full rank does not imply precise estimation.) (d) What does a negative posterior correlation between two parameters often indicate? (That one parameter can partially compensate for the other in the predictions, as G and P_{int} do; not that the parameters are negative or the observations wrong.)

Exercise 31.2. (a) What does a very large residual chi-square statistic suggest under an assumed Gaussian model? (That the model and the stated measurement uncertainty jointly fit the data poorly; it does not say which assumption fails.) (b) Can adding a flexible model-discrepancy term create identifiability problems? (Yes: the discrepancy can absorb effects otherwise attributed to the physical parameters, such as a discrepancy proportional to ΔT absorbing part of G.) (c) Does a lower best-fit residual prove that a more complex physical model is correct? (No: complexity, priors, validation and alternative explanations matter, and a nested model always fits at least as well.) (d) The thermal model $Q_{\text{cool}} = G\,\Delta T + P_{\text{int}}$ is fitted to 41, 59, 81 and 99 kW measured at $\Delta T = 0$, 10, 20 and 30 K with $\sigma_Q = 1$ kW. Compute the least-squares estimates, the residuals and $J(\hat{\boldsymbol{\theta}})$, and test the fit at the 5 % level. ($\hat{G} = 980/500 = 1.96$ kW/K, $\hat{P}_{\text{int}} = 40.6$ kW; residuals 0.4, -1.2, 1.2, -0.4 kW; $J = 3.2 < \chi^2_{2,0.95} = 5.99$, p-value $\mathrm{e}^{-1.6} = 0.20$; the fit is not rejected.)

Exercise 31.3. The thermal balance is measured with four readings, two at $\Delta T = 0$ and two at $\Delta T = 20$ K, with $\sigma_Q = 1$ kW. (a) Compute $\mathcal{I}(\boldsymbol{\theta})$, its determinant, $\boldsymbol{\Sigma}_\theta$, σ_G, σ_P and $\rho_{G,P}$. ([800 40; 40 4]; 1600; [0.0025 -0.025; -0.025 0.5]; 0.050 kW/K, 0.707 kW, -0.707.) (b) Compare with four readings at 0, 10, 10 and 20 K. ($\mathcal{I}$ = [600 40; 40 4], determinant 800; $\sigma_G = 0.0707$ kW/K, $\sigma_P = 0.866$ kW; placing the readings at the ends of the range halves the variance of $\hat{G}$.) (c) Compute the singular values and the condition number of the scaled matrix with $\mathbf{D}_\theta = \mathrm{diag}(2, 40)$ for the design of (a). (91.53 and 34.96; $\kappa = 2.62$, against 3.15 for Experiment B.) (d) How many repetitions of the narrow design 9, 10, 11 K would give the σ_G of (a)? ($0.5/n = 0.0025$, so $n = 200$ repetitions, 600 readings, against four.)

Exercise 31.4. Two envelope conductances have independent priors $G_1, G_2 \sim \mathcal{N}(2, 0.5^2)$ in kW/K, and one measurement gives $Q = 50$ kW at $\Delta T = 10$ K with $\sigma_Q = 1$ kW, under the model $Q = (G_1 + G_2)\,\Delta T$. (a) Compute $\mathcal{I}_{\text{post}}$, $\boldsymbol{\Sigma}_{\text{post}}$, the posterior means and the posterior correlation. ([104 100; 100 104]; [104 -100; -100 104]/816; 2.490 kW/K each; $-100/104 = -0.962$.) (b) Compute the posterior standard deviations of $G_1 + G_2$ and of $G_1 - G_2$, and state which is prior dominated. ($\sqrt{8/816} = 0.099$ and $\sqrt{408/816} = 0.707$ kW/K; the difference, whose prior standard deviation was also $0.5\sqrt{2} = 0.707$.) (c) A Bayes factor $B_{21} = 20$ is reported for a model M_2 against M_1. Give the posterior probability of M_2 with equal prior probabilities, with $\Pr(M_2) = 0.2$, and with a third model M_3 of the same evidence as M_2 and the three equally probable a priori. ($20/21 = 0.952$; odds $20 \times 0.25 = 5$, probability 0.833; $20/(1+20+20) = 0.488$.)

Chapter 32

Bayesian Updating and Value of Information

32.1 Bayesian updating of competing hypotheses

Chapters 29 to 31 established three limits of engineering inference. A system may not be observable from its existing sensors (Chapter 29); its physical parameters may be unidentifiable or poorly constrained by the available data (Chapter 31); and a model may fit the observations while failing to represent the physical mechanism that produced them (Section 31.4). Each limit ends in the same practical situation. A data center cooling system behaves abnormally; several physical explanations are possible; the instrumentation is limited; and the maintenance budget is finite. The engineer can install another temperature sensor, calibrate a flow meter, inspect a piece of equipment, or proceed directly to repair. The question of this chapter is

when several physical explanations are possible, which measurement should be obtained next, and how much the information it provides is worth.

The answer requires more than selecting the measurement with the smallest uncertainty. A measurement is worth obtaining when it is expected to improve the understanding of the system and when that improvement is expected to change what should be done. The chapter follows the chain

$$\text{current evidence} \rightarrow \text{posterior over hypotheses} \rightarrow \text{candidate observations} \rightarrow \text{updated decisions},$$

and develops two complementary quantities along it: the expected information gain, which measures how much a measurement is expected to reduce the uncertainty about the hypotheses, and the expected value of sample information, which measures how much it is expected to improve a decision, in the units of the decision objective, here reais. The two are not the same quantity, and the chapter is organized around the difference. It updates the probabilities of competing explanations of an abnormal cooling reading; measures uncertainty by entropy and the information of a measurement by the expected entropy reduction; defines the decision value of information and its bound; works a pump replacement decision with one and with three candidate inspections; extends the calculations to continuous parameters and to competing physical models; and ends with sequential decisions, the economics of inspection over a range of inputs, and the credibility of the recommendation.

Jargon: prior, likelihood, posterior, posterior odds, Bayes factor, entropy, information gain, mutual information, Kullback–Leibler divergence, loss function, expected value of sample and perfect information, preposterior analysis

The *prior* probability $\Pr(H_i)$ of a hypothesis H_i is its probability before a new observation; the *likelihood* $\Pr(y \mid H_i)$ is the probability of the observation y if H_i is true; the *posterior* $\Pr(H_i \mid y)$ is its probability after the observation. The *posterior odds* of two hypotheses are the ratio of their posteriors, and the *Bayes factor* is the ratio of their likelihoods, the factor by which an observation multiplies the prior odds. The Shannon *entropy* $\mathcal{H}$ of a probability distribution, in bits, measures its dispersion. The *information gain* of an observation is the reduction of entropy it produces; the *expected information gain* EIG is its average over the possible outcomes, equal to the *mutual information* $I(H;Y)$ between the hypothesis and the outcome, and to the average *Kullback–Leibler divergence* D_{KL} of the posterior from the prior. A *loss function* $L(a, H)$ gives the loss, in reais, of action a when hypothesis H is true. The *expected value of sample information* (EVSI) is the expected reduction in the loss of the best decision that an imperfect measurement brings, before its cost; the *expected value of perfect information* (EVPI) is the same reduction for a measurement that reveals the true hypothesis. *Preposterior analysis* is the evaluation of a measurement before it is made, by averaging the decision that would follow each possible outcome over the predictive probabilities of the outcomes.

Key idea

A measurement has two different values. Its information value is the expected reduction of uncertainty about the hypotheses; its decision value is the expected reduction of the loss of the best action, and it exists only when some outcome of the measurement would change that action. The best next measurement is the one that most improves the decision net of its cost, not necessarily the one that most reduces uncertainty.

This section restates Bayes' theorem for a discrete set of competing hypotheses, works the update of an abnormal cooling reading for both outcomes of a test, and writes the update in odds form.

32.1.1 Competing hypotheses for an abnormal cooling reading

An industrial cooling system reports an abnormal temperature reading. Three explanations are considered,

$$H_0 : \text{the system is healthy (the reading is a normal fluctuation)},$$
$$H_{\mathrm{C}} : \text{the cooling capacity is insufficient},$$
$$H_{\mathrm{S}} : \text{the temperature sensor is biased},$$

the same three kinds of explanation that Section 22.5.1 showed a single temperature sensor cannot separate. Before collecting further measurements, the engineer assigns the prior probabilities

$$\Pr(H_0) = 0.50, \qquad \Pr(H_{\mathrm{C}}) = 0.30, \qquad \Pr(H_{\mathrm{S}}) = 0.20.$$

These numbers are hypothetical. They express the current degree of belief within the stated hypothesis set, conditional on the evidence so far; they are not verified frequencies of real faults. The set is treated as exhaustive and its members as mutually exclusive, so the priors add to one.

An independent cooling-performance test is now performed and returns the observation $Y = y$. Bayes' theorem, stated for a continuous parameter in Chapter 8 of Volume I as $p(\lambda \mid D) = p(D \mid \lambda)\, p(\lambda)/p(D)$, takes for a discrete set of hypotheses the form

$$\Pr(H_i \mid y) = \frac{\Pr(y \mid H_i)\, \Pr(H_i)}{\sum_j \Pr(y \mid H_j)\, \Pr(H_j)} = \frac{\Pr(y \mid H_i)\, \Pr(H_i)}{\Pr(y)}. \tag{32.1}$$

It follows from writing the joint probability of hypothesis and observation in two ways, $\Pr(H_i, y) = \Pr(y \mid H_i)\Pr(H_i) = \Pr(H_i \mid y)\Pr(y)$. The four terms are the prior $\Pr(H_i)$; the likelihood $\Pr(y \mid H_i)$ of the observation under hypothesis H_i; the posterior $\Pr(H_i \mid y)$; and the denominator $\Pr(y) = \sum_j \Pr(y \mid H_j)\Pr(H_j)$, the *predictive probability* of the observation, obtained by the law of total probability over the hypotheses. The predictive probability does not depend on i; it normalizes the posteriors so that they add to one, and it is also the probability with which the engineer should expect to see y before the test is made, which the later sections use.

32.1.2 The update after an abnormal test

Suppose the test result is abnormal, $Y = Y^+$, and that the test has the likelihoods of the third column of Table 32.1: an abnormal result is rare under a healthy system (0.05, the false-alarm probability of the test), very probable under a capacity shortfall (0.90), and unlikely under a sensor bias (0.10), because the test measures the cooling performance independently of the suspect sensor. The predictive probability of an abnormal result is

$$\Pr(Y^+) = 0.50 \times 0.05 + 0.30 \times 0.90 + 0.20 \times 0.10 = 0.025 + 0.270 + 0.020 = 0.315.$$

The three products are the unnormalized posterior weights, $w_0 = 0.025$, $w_{\mathrm{C}} = 0.270$ and $w_{\mathrm{S}} = 0.020$, and dividing each by their sum gives the posteriors

$$\Pr(H_0 \mid Y^+) = \frac{0.025}{0.315} = 0.0794,$$
$$\Pr(H_{\mathrm{C}} \mid Y^+) = \frac{0.270}{0.315} = 0.8571,$$
$$\Pr(H_{\mathrm{S}} \mid Y^+) = \frac{0.020}{0.315} = 0.0635.$$

Table 32.1 collects the calculation; its last three columns anticipate the normal outcome of Section 32.1.4.

Table 32.1: Bayesian update of three explanations of an abnormal cooling reading after a cooling-performance test, for both outcomes of the test (illustrative probabilities).

Hypothesis	Prior	$\Pr(Y^+ \mid H_i)$	Weight	$\Pr(H_i \mid Y^+)$	$\Pr(Y^- \mid H_i)$	$\Pr(H_i \mid Y^-)$
Healthy, H_0	0.50	0.05	0.025	0.0794	0.95	0.6934
Cooling shortfall, H_{C}	0.30	0.90	0.270	0.8571	0.10	0.0438
Sensor bias, H_{S}	0.20	0.10	0.020	0.0635	0.90	0.2628
Sum	1	–	0.315	1	–	1

Figure 32.1 draws the three distributions as bars. The horizontal axis lists the hypotheses; the vertical axis is the probability in percent; for each hypothesis the left bar is the prior, the middle bar the posterior after an abnormal test and the right bar the posterior after a normal test, each labeled with its value. After the abnormal result the cooling-shortfall hypothesis becomes dominant, from 30 % to 85.7 %, because an abnormal result is much more probable under it than under either alternative. The healthy hypothesis falls from 50 % to 7.9 % although it had the largest prior.

Bayesian inference does not ask which hypothesis best fits the observation alone. The observation is most probable under H_{C}, but it is also possible under the other two, and the posterior combines the likelihood with the prior. Had the prior of the healthy hypothesis been 0.95, with 0.03 and 0.02 for the others, the same abnormal result would have left $\Pr(H_0 \mid Y^+) = 0.0475/0.0765 = 0.62$: the healthy hypothesis would still be the most probable.

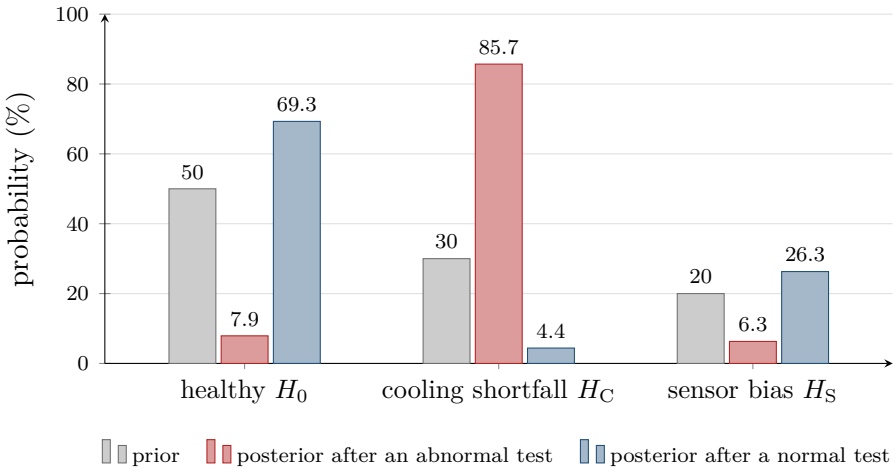


Figure 32.1: Probabilities of the three explanations of the abnormal cooling reading before the cooling-performance test and after each of its outcomes, equation (32.1) with the likelihoods of Table 32.1 (illustrative values).

32.1.3 Posterior odds and Bayes factors

Dividing equation (32.1) for H_i by the same equation for H_j cancels the predictive probability:

$$\frac{\Pr(H_i \mid y)}{\Pr(H_j \mid y)} = \frac{\Pr(H_i)}{\Pr(H_j)} \, \frac{\Pr(y \mid H_i)}{\Pr(y \mid H_j)}, \tag{32.2}$$

that is,

$$\text{posterior odds} = \text{prior odds} \times \text{Bayes factor}.$$

The Bayes factor of two simple hypotheses is the likelihood ratio of the observation; Section 31.5.4 defined it for two models with parameters, $B_{12} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_2)$, where each likelihood is averaged over the parameter prior of its model. The odds form needs only the two hypotheses being compared, and it shows the evidence of the observation separately from the prior.

For cooling shortfall against sensor bias the prior odds are $0.30/0.20 = 1.5$ and the Bayes factor of an abnormal result is $0.90/0.10 = 9$, so

$$\frac{\Pr(H_C \mid Y^+)}{\Pr(H_S \mid Y^+)} = 1.5 \times 9 = 13.5,$$

which agrees with $0.8571/0.0635 = 13.5$. The abnormal test shifts the evidence strongly toward an actual cooling deficiency. Against the healthy hypothesis the prior odds of a shortfall are 0.6 and the Bayes factor is $0.90/0.05 = 18$, giving posterior odds of 10.8. The Bayes factor of a sensor bias against health is only $0.10/0.05 = 2$: an abnormal result from an independent test says little about whether the suspect sensor is biased, because a biased sensor does not change what the test measures.

32.1.4 The update after a normal test

A normal result, $Y = Y^-$, is also informative. Its likelihoods are the complements of the abnormal ones,

$$\Pr(Y^- \mid H_0) = 0.95, \qquad \Pr(Y^- \mid H_C) = 0.10, \qquad \Pr(Y^- \mid H_S) = 0.90,$$

and its predictive probability is

$$\Pr(Y^-) = 0.50 \times 0.95 + 0.30 \times 0.10 + 0.20 \times 0.90 = 0.475 + 0.030 + 0.180 = 0.685 = 1 - \Pr(Y^+).$$

The posteriors, in the last column of Table 32.1 and the right bars of Figure 32.1, are $0.475/0.685 = 0.6934$ for health, $0.030/0.685 = 0.0438$ for a shortfall and $0.180/0.685 = 0.2628$ for a sensor bias. A normal test makes a cooling deficiency unlikely relative to the alternatives: its odds against a bias fall from 1.5 to $1.5 \times (0.10/0.90) = 0.167$. It also raises the probability of a sensor bias, from 0.20 to 0.26, because a normal test leaves an abnormal reading of the sensor in need of explanation and the bias is the explanation that predicts both.

The test has two outcomes, with probabilities 0.315 and 0.685, and each leads to a different posterior. Before the test is made, its value must account for every possible outcome, weighted by its predictive probability, not merely for the outcome that is expected or hoped for. That averaging is the subject of the rest of the chapter.

32.2 Entropy and expected information gain

To compare measurements by how much they are expected to teach, the uncertainty of a probability distribution over hypotheses must be measured by one number. This section defines that number, the Shannon entropy, distinguishes the entropy reduction produced by one outcome from its expectation, identifies the expectation with the mutual information and with an average Kullback–Leibler divergence, and characterizes a measurement that carries no information.

32.2.1 Shannon entropy of a set of hypotheses

Let the current probabilities of n mutually exclusive and exhaustive hypotheses be $p_1, p_2, \ldots, p_n$, and let H denote the unknown true hypothesis, a discrete random variable with these probabilities. The *Shannon entropy* of the distribution is

$$\mathcal{H}(H) = -\sum_{i=1}^{n} p_i \log_2 p_i, \tag{32.3}$$

in bits when the logarithm has base 2, with the convention $0 \log_2 0 = 0$; Section 22.5.5 used the same measure for the diagnosis of a pump. For two equally probable hypotheses, $\Pr(H_1) = \Pr(H_2) = 0.5$, the entropy is 1 bit; if one hypothesis has probability one and the others zero, the entropy is zero; and for n hypotheses it is largest, $\log_2 n$ bits, when they are equally probable. For the three explanations of Section 32.1.1,

$$\mathcal{H}(H) = -(0.50 \log_2 0.50 + 0.30 \log_2 0.30 + 0.20 \log_2 0.20) = 0.500 + 0.521 + 0.464 = 1.485 \text{ bits},$$

against a maximum of $\log_2 3 = 1.585$ bits.

Entropy measures the dispersion of a probability distribution. It does not measure physical risk or financial loss. A highly uncertain diagnosis may have little operational consequence, if every explanation leads to the same harmless action; a nearly certain diagnosis may indicate a catastrophic condition. The distinction between uncertainty and its consequences returns in Section 32.3.

Remark 32.1. In this chapter H is the unknown hypothesis and $\mathcal{H}$ an entropy. Neither is the operating horizon H of the earlier chapters, which does not appear here, and the mutual information $I(H;Y)$ defined below is not the Fisher information $\mathcal{I}(\boldsymbol{\theta})$ of Chapter 31.

32.2.2 Realized and expected entropy reduction

After the outcome $Y = y$ is observed, the uncertainty becomes the entropy of the posterior, $\mathcal{H}(H \mid Y = y) = -\sum_i \Pr(H_i \mid y) \log_2 \Pr(H_i \mid y)$, and the *realized entropy reduction* is

$$\Delta\mathcal{H}(y) = \mathcal{H}(H) - \mathcal{H}(H \mid Y = y). \tag{32.4}$$

For the cooling test the posterior entropies are 0.733 bits after an abnormal result and 1.071 bits after a normal one, so the realized reductions are $1.485 - 0.733 = 0.752$ bits and $1.485 - 1.071 = 0.415$ bits. Both are positive here, but the realized reduction can be negative for a particular surprising outcome: a test can make the posterior more diffuse than the prior.

A binary example shows it. Let a failure hypothesis have prior probability 0.20, with entropy $-0.2\log_2 0.2 - 0.8\log_2 0.8 = 0.722$ bits, and let a test detect the failure with probability 0.90 (its sensitivity) and correctly clear its absence with probability 0.80 (its specificity); these are the numbers of the pump inspection of Section 32.4. A positive result, with predictive probability $0.90 \times 0.20 + 0.20 \times 0.80 = 0.34$, raises the failure probability to $0.18/0.34 = 0.529$, whose entropy is 0.998 bits: the realized reduction is $0.722 - 0.998 = -0.276$ bits. A negative result, with probability 0.66, lowers it to $0.02/0.66 = 0.030$, whose entropy is 0.196 bits, a reduction of 0.526 bits. What is guaranteed is that the average,

$$0.34 \times (-0.276) + 0.66 \times 0.526 = 0.254 \text{ bits},$$

is not negative.

Figure 32.2 shows why. The horizontal axis is the probability p of the failure hypothesis and the vertical axis the binary entropy $-p\log_2 p - (1-p)\log_2(1-p)$ in bits, drawn as the solid curve. The black dot is the prior, the red and green dots the posteriors after a positive and a negative outcome, and the dashed chord joins the posteriors. Because the prior is the average of the posteriors weighted by the outcome probabilities, $0.34 \times 0.529 + 0.66 \times 0.030 = 0.20$, the expected posterior entropy lies on the chord directly below the prior, at 0.468 bits (open circle). The entropy curve is concave, so every chord lies below it, and the vertical gap between the prior and the open circle, the expected reduction, cannot be negative. A single posterior, such as the red one, can lie higher on the curve than the prior.

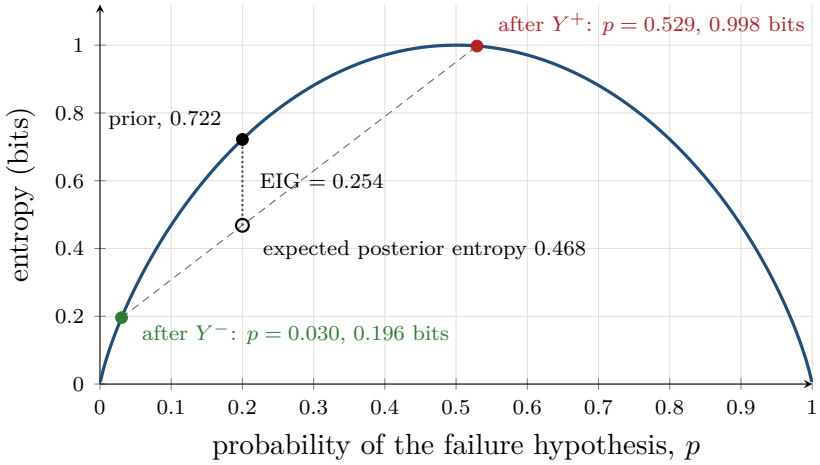


Figure 32.2: Binary entropy $-p\log_2 p - (1-p)\log_2(1-p)$ of a failure hypothesis, with the prior $p = 0.20$ and the two posteriors of a test with sensitivity 0.90 and specificity 0.80 (illustrative values). The expected posterior entropy lies on the chord between the posteriors, below the prior; the dotted gap is the expected information gain, although the positive outcome alone raises the entropy.

32.2.3 Expected information gain as mutual information

Before the measurement is made, its outcome is unknown, but its predictive distribution $\Pr(y)$ is known. Averaging the posterior entropy over the outcomes of a discrete measurement,

$$\mathbb{E}_Y\left[\mathcal{H}(H \mid Y)\right] = \sum_y \Pr(y)\,\mathcal{H}(H \mid Y = y),$$

defines the *expected information gain*

$$\mathrm{EIG}(Y) = \mathcal{H}(H) - \mathbb{E}_Y\big[\mathcal{H}(H \mid Y)\big], \tag{32.5}$$

in bits, the definition of equation (22.29). For the cooling test of Table 32.1,

$$\mathrm{EIG} = 1.485 - (0.315 \times 0.733 + 0.685 \times 1.071) = 1.485 - 0.964 = 0.521 \text{ bits},$$

about a third of the prior uncertainty.

The expected information gain is the *mutual information* $I(H;Y)$ between the hypothesis and the outcome. Two further forms follow by rewriting the entropies. Write $\mathcal{H}(H)$ with each prior expanded by the law of total probability, $\Pr(H_i) = \sum_y \Pr(y)\ \Pr(H_i \mid y)$:

$$\mathcal{H}(H) = -\sum_i \Pr(H_i)\log_2 \Pr(H_i) = -\sum_y \Pr(y) \sum_i \Pr(H_i \mid y)\log_2 \Pr(H_i).$$

Subtracting the expected posterior entropy, which has the same double sum with $\log_2 \Pr(H_i \mid y)$ in place of $\log_2 \Pr(H_i)$, gives

$$\mathrm{EIG}(Y) = I(H;Y) = \sum_y \Pr(y) \sum_i \Pr(H_i \mid y)\log_2 \frac{\Pr(H_i \mid y)}{\Pr(H_i)} = \mathbb{E}_Y\Big[D_{\mathrm{KL}}\big(\Pr(H \mid Y) \parallel \Pr(H)\big)\Big], \tag{32.6}$$

where the *Kullback–Leibler divergence* of a distribution q from a distribution p over the same hypotheses is $D_{\mathrm{KL}}(q\|p) = \sum_i q_i \log_2(q_i/p_i)$. The divergence is never negative, and is zero only when $q = p$: this is Gibbs' inequality, which follows from Jensen's inequality applied to the concave logarithm, $-D_{\mathrm{KL}}(q\|p) = \sum_i q_i \log_2(p_i/q_i) \le \log_2 \sum_i q_i(p_i/q_i) = \log_2 1 = 0$. Every term of the average in equation (32.6) is therefore nonnegative, and so is the expected information gain; it is zero only when every posterior equals the prior. For the cooling test the divergence is 0.982 bits after an abnormal result and 0.309 bits after a normal one, and $0.315 \times 0.982 + 0.685 \times 0.309 = 0.521$ bits, as before. A third form, derived in Section 22.5.5, is the symmetric one,

$$I(H;Y) = \mathcal{H}(Y) - \mathbb{E}_H\big[\mathcal{H}(Y \mid H)\big],$$

the entropy of the predicted outcome minus the average entropy of the outcome under each hypothesis: here $\mathcal{H}(Y) = 0.899$ bits for the outcome probabilities 0.315 and 0.685, and the average conditional entropy is $0.50 \times 0.286 + 0.30 \times 0.469 + 0.20 \times 0.469 = 0.378$ bits, again a difference of 0.521 bits.

The divergence form of equation (32.6) is the one that generalizes. It needs no entropy of the hypotheses themselves, only the ratio of posterior to prior, so it applies unchanged to continuous parameters and continuous measurements, with densities in place of probabilities and integrals in place of sums; Section 32.5.2 uses it in that form.

32.2.4 A measurement with zero information

Suppose a temperature measurement has exactly the same predictive distribution under all three hypotheses,

$$p(y \mid H_0) = p(y \mid H_{\mathrm{C}}) = p(y \mid H_{\mathrm{S}}).$$

Then the likelihoods cancel from the numerator and the denominator of equation (32.1), $\Pr(H_i \mid y) = \Pr(H_i)$ for every possible outcome, every divergence in equation (32.6) is zero, and

$$\mathrm{EIG}(Y) = 0.$$

Conversely, a zero expected information gain implies that every posterior equals the prior, so that Y and H are independent. Collecting such a measurement does not distinguish the

hypotheses, however precise the instrument; the flow-meter calibration of Table 22.9, whose outcome probability was the same under every hypothesis, was an example. The measurement might still be useful for another purpose, for example to verify compliance with a temperature limit or to test a hypothesis that the set omits. The value of information is always relative to a stated unknown quantity and a stated question.

Principle 32.1. *The expected information gain of a measurement is the mutual information between the unknown hypothesis and the outcome, the average Kullback–Leibler divergence of the posterior from the prior. It is never negative, although a single outcome can increase the uncertainty, and it is zero exactly when the outcome has the same distribution under every hypothesis.*

32.3 Expected value of information for decisions

Information gain measures learning, not usefulness. This section explains why the two differ, introduces the loss function and the expected loss of an action, and defines the expected value of sample and of perfect information with the inequalities that bound them.

32.3.1 Information gain and decision value

Suppose two instruments are available. Sensor A substantially reduces the uncertainty about a minor equipment parameter. Sensor B provides less information overall, but it distinguishes whether an urgent intervention on the cooling system is necessary. Selecting measurements by entropy reduction alone would choose Sensor A. If Sensor B improves the maintenance decision and avoids significant losses, it is the more valuable instrument. The entropy of Section 32.2.1 weighs every hypothesis equally by its probability; it does not know which distinctions matter for what will be done.

To formalize the difference, the decision must be stated. Let $a \in \mathcal{A}$ denote an action from the set $\mathcal{A}$ of available actions, for example continue operating, inspect, repair or replace, and let

$$L(a, H)$$

be the *loss* incurred, in reais, by taking action a when hypothesis H is true: the cost of the action plus the cost of what follows from it over the decision period. The loss function plays for decisions the role that the likelihood plays for inference. It encodes the consequences, which the probabilities do not.

32.3.2 Expected loss and the Bayes action

The current expected loss of action a is its loss averaged over the current probabilities of the hypotheses,

$$\rho(a) = \sum_{i} \Pr(H_i)\, L(a, H_i), \tag{32.7}$$

in reais; decision theory calls it the posterior expected loss when the probabilities are a posterior. The best action, the *Bayes action*, minimizes it,

$$a^{*} = \arg\min_{a \in \mathcal{A}} \rho(a). \tag{32.8}$$

Choosing the action of least expected loss is the rule that Chapter 24 applied to the replacement of a pump, where the two expected losses J_{rep} and J_{wait} of equation (24.12) were compared, and that Section 22.5.6 applied to four maintenance actions under three diagnostic hypotheses.

Remark 32.2. The loss $L(a, H)$ is not the likelihood $L(\theta)$ of Chapter 7 of Volume I, and the expected loss $\rho(a)$ is not the ratio λ/μ of Chapter 19 of Volume I; both letters follow the standard notation of statistical decision theory and are used in this sense from this chapter on.

32.3.3 The expected value of sample information

With the current information, the minimum expected loss is

$$C_{\text{now}} = \min_a \sum_i \Pr(H_i)\, L(a, H_i) = \rho(a^*). \tag{32.9}$$

Suppose a measurement Y can be obtained before deciding. After its outcome y is observed, the action is chosen again with the posterior probabilities,

$$a^*(y) = \arg\min_a \mathbb{E}\big[L(a, H) \mid Y = y\big] = \arg\min_a \sum_i \Pr(H_i \mid y)\, L(a, H_i),$$

and it may differ from a^*. Before the measurement is made, the outcome is unknown; averaging the minimum posterior expected loss over the predictive distribution of the outcome is the preposterior analysis of the measurement, and its result is the expected loss with the information, before the cost of the measurement:

$$C_{\text{sample}} = \mathbb{E}_Y\Big[\min_a \mathbb{E}\big[L(a, H) \mid Y\big]\Big] = \sum_y \Pr(y)\, \min_a \sum_i \Pr(H_i \mid y)\, L(a, H_i). \tag{32.10}$$

The *expected value of sample information* is the difference,

$$\text{EVSI} = C_{\text{now}} - C_{\text{sample}}, \tag{32.11}$$

in reais. It is the quantity that Chapter 24 computed for one pump inspection, equation (24.18), and Chapter 22 for several diagnostic measurements, equation (22.32), written here in general form.

Before the cost of the measurement,

$$\text{EVSI} \geq 0. \tag{32.12}$$

The reason is that after observing the outcome the decision maker can always ignore it and take the original action a^*. Formally, for every outcome y the minimum over actions is at most the posterior expected loss of a^*,

$$\min_a \sum_i \Pr(H_i \mid y)\, L(a, H_i) \leq \sum_i \Pr(H_i \mid y)\, L(a^*, H_i).$$

Averaging both sides over $\Pr(y)$, and using $\sum_y \Pr(y)\ \Pr(H_i \mid y) = \Pr(H_i)$ (the posterior averaged over the predictive distribution is the prior), the left side becomes C_{sample} and the right side becomes $\rho(a^*) = C_{\text{now}}$. Equality holds when $a^*(y) = a^*$ for every outcome: a measurement whose outcomes cannot change the action has zero decision value, however much it reduces the entropy.

32.3.4 The expected value of perfect information

Suppose, hypothetically, that the true hypothesis could be learned with certainty before choosing the action. The best action would then be chosen separately for each hypothesis, and the expected loss would be

$$C_{\text{perfect}} = \sum_i \Pr(H_i)\ \min_a L(a, H_i). \tag{32.13}$$

The *expected value of perfect information* is

$$\text{EVPI} = C_{\text{now}} - C_{\text{perfect}}. \tag{32.14}$$

For the same decision problem, with the same actions, losses and current probabilities,

$$0 \le \mathrm{EVSI} \le \mathrm{EVPI}. \tag{32.15}$$

The left inequality is equation (32.12). For the right one, fix an outcome y. For every action a and every hypothesis, $L(a, H_i) \ge \min_{a'} L(a', H_i)$; averaging over the posterior, $\sum_i \Pr(H_i \mid y)\, L(a, H_i) \ge \sum_i \Pr(H_i \mid y) \min_{a'} L(a', H_i)$ for every a, and therefore also for the minimizing a. Averaging over $\Pr(y)$ and using again $\sum_y \Pr(y)\, \Pr(H_i \mid y) = \Pr(H_i)$ gives $C_{\text{sample}} \ge C_{\text{perfect}}$, hence EVSI $\le$ EVPI. The chain is $C_{\text{perfect}} \le C_{\text{sample}} \le C_{\text{now}}$: perfect information is at least as useful as any imperfect measurement, which is at least as useful as none.

The EVPI is an upper bound on the economic value, for the stated decision, of learning the uncertain state, and so an upper bound on what any measurement of that state is worth to that decision. It is not necessarily the largest worthwhile instrumentation budget: a sensor installed for this decision may also serve operation, compliance, fleet statistics or later diagnoses, which the bound does not count. The bound applies to the specified uncertainty and decision objective.

Principle 32.2. *The expected value of sample information of a measurement is the expected loss of the best decision without it minus the expected loss of the best decision chosen after its outcome. It is never negative, never exceeds the expected value of perfect information, and is zero when no outcome of the measurement would change the action.*

32.4 The pump replacement decision worked in full

The definitions of Section 32.3 are now applied to a two-hypothesis maintenance decision for the cooling-water pump: the loss matrix, the decision without inspection, the inspection and its posteriors, the expected cost with inspection and the value of the information, and finally a choice among three candidate inspections. The numbers are hypothetical; they differ from those of Chapter 24 and are chosen to make every step visible.

32.4.1 Hypotheses, actions and the loss matrix

A cooling-water pump of the data center cooling plant has an uncertain bearing condition. Two hypotheses are considered,

$$H_{\mathrm{F}} : \text{the pump will fail within the next 30 days,}$$
$$H_{\mathrm{N}} : \text{the pump will not fail within the next 30 days,}$$

and the current prognostic model, built as in Chapter 23, assigns

$$\Pr(H_{\mathrm{F}}) = 0.20, \qquad \Pr(H_{\mathrm{N}}) = 0.80.$$

Two actions are available: replace the pump now, or continue operating for 30 days. Replacing now costs $C_{\mathrm{R}} = \mathrm{R\$}\,8000$ and prevents the failure event within the planning horizon. If the pump is kept in operation and fails, the total loss is $C_{\mathrm{F}} = \mathrm{R\$}\,50\,000$; if it does not fail, there is no incremental loss within the horizon. Table 32.2 is the loss matrix $L(a, H)$: one row per action, one column per hypothesis.

The assumptions are deliberately simple. A complete maintenance model would also include the risk of the maintenance window while the replacement is performed, the remaining life of the equipment, and the costs after the 30-day horizon, as Chapter 24 did. The matrix has the structure of the replace-or-wait comparison of Section 24.3.3, with the cost of acting early C_{R} in place of the gain from waiting G_{w} and the loss of a failure C_{F} in place of the incremental failure penalty Δc, so its results carry over with that substitution. The symbols C_{R} and C_{F} are costs in reais, not the available capacity C nor the consequence C of the earlier chapters.

Table 32.2: Loss matrix $L(a, H)$ of the pump decision over a 30-day horizon, in reais (illustrative values).

Action a	Pump fails within 30 days, H_{F}	Pump survives 30 days, H_{N}
Replace now	$C_{\mathrm{R}} = 8000$	$C_{\mathrm{R}} = 8000$
Continue operating	$C_{\mathrm{F}} = 50\,000$	0

32.4.2 The decision without inspection

By equation (32.7), the expected losses of the two actions are

$$\rho(\text{replace}) = 0.20 \times 8000 + 0.80 \times 8000 = \text{R\$}\,8000,$$
$$\rho(\text{continue}) = 0.20 \times 50\,000 + 0.80 \times 0 = \text{R\$}\,10\,000,$$

so the Bayes action is to replace now and

$$C_{\text{now}} = \min\{8000,\ 10\,000\} = \text{R\$}\,8000.$$

The two expected losses are equal when $C_{\mathrm{R}} = p\,C_{\mathrm{F}}$, with $p = \Pr(H_{\mathrm{F}})$, which gives the economic replacement threshold

$$p^{*} = \frac{C_{\mathrm{R}}}{C_{\mathrm{F}}} = \frac{8000}{50\,000} = 0.16, \tag{32.16}$$

the counterpart of equation (24.13). If the predicted probability of failure within 30 days exceeds 16 %, replacement is cheaper in expectation under these assumptions; the prognosis of 20 % exceeds the threshold, and the decision is to replace.

32.4.3 The inspection and its posteriors

A condition inspection can be performed first. It returns a positive outcome Y^{+}, indicating elevated risk, or a negative outcome Y^{-}, with

$$\Pr(Y^{+} \mid H_{\mathrm{F}}) = 0.90, \qquad \Pr(Y^{-} \mid H_{\mathrm{N}}) = 0.80,$$

a sensitivity of 90 % and a specificity of 80 %, defined against the 30-day failure event as in Section 24.4.1. The test is informative but imperfect: its false-negative probability is $\Pr(Y^{-} \mid H_{\mathrm{F}}) = 0.10$ and its false-positive probability $\Pr(Y^{+} \mid H_{\mathrm{N}}) = 0.20$. By the law of total probability,

$$\Pr(Y^{+}) = 0.90 \times 0.20 + 0.20 \times 0.80 = 0.18 + 0.16 = 0.34, \qquad \Pr(Y^{-}) = 1 - 0.34 = 0.66.$$

Nearly half of the positive outcomes (0.16 of 0.34) are false alarms from the larger group of pumps that will not fail. By equation (32.1), the posterior failure probabilities are

$$\Pr(H_{\mathrm{F}} \mid Y^{+}) = \frac{0.90 \times 0.20}{0.34} = \frac{0.18}{0.34} = 0.5294,$$
$$\Pr(H_{\mathrm{F}} \mid Y^{-}) = \frac{0.10 \times 0.20}{0.66} = \frac{0.02}{0.66} = 0.0303.$$

In odds, the prior odds of failure $0.20/0.80 = 0.25$ are multiplied by the Bayes factor $0.90/0.20 = 4.5$ after a positive outcome, to 1.125, and by $0.10/0.80 = 0.125$ after a negative one, to 0.03125; the corresponding probabilities are $1.125/2.125 = 0.5294$ and $0.03125/1.03125 = 0.0303$.

Table 32.3 gives the resulting decision policy. Each row is an outcome; the columns are its probability, the updated failure probability, the action that equation (32.16) selects and its expected loss. The posterior after a positive outcome lies far above the threshold 0.16 and the one after a negative outcome far below it, so the two outcomes lead to different actions. The inspection is valuable because of that, and only because of that.

Table 32.3: Decision policy after the pump inspection (sensitivity 0.90, specificity 0.80, prior 0.20; illustrative values).

Test result	Probability	Updated failure probability	Optimal action	Expected loss (R$)
Positive, Y^+	0.34	0.5294	replace now	8000
Negative, Y^-	0.66	0.0303	continue operating	1515.15

32.4.4 Expected cost with inspection, EVSI and EVPI

After a positive outcome the pump is replaced, with loss $C_+ = \text{R\$}\,8000$. After a negative outcome it is kept in operation, with expected loss

$$C_- = \Pr(H_\text{F} \mid Y^-)\, C_\text{F} = 0.030303 \times 50\,000 = \text{R\$}\,1515.15.$$

By equation (32.10), averaging over the outcomes,

$$C_\text{sample} = 0.34 \times 8000 + 0.66 \times 1515.15 = 2720 + 1000 = \text{R\$}\,3720,$$

where $0.66 \times 0.030303 = 0.02 = \Pr(H_\text{F})\,\Pr(Y^- \mid H_\text{F})$ makes the second term exactly $0.02 \times 50\,000 = \text{R\$}\,1000$: it is the expected loss of the missed failures. The expected value of sample information is

$$\text{EVSI} = C_\text{now} - C_\text{sample} = 8000 - 3720 = \text{R\$}\,4280.$$

If the inspection costs $c_\text{I} = \text{R\$}\,1000$, the expected loss of inspecting and then deciding is $3720 + 1000 = \text{R\$}\,4720$, and the net expected value of the inspection is $\text{EVSI} - c_\text{I} = 4280 - 1000 = \text{R\$}\,3280$. The inspection is economically worthwhile under this decision model.

Figure 32.3 draws the decision as a tree, in the style of Figure 24.4. Squares are decision nodes and circles chance nodes, with the probability of each chance branch on it; the numbers at the right are the losses at the end of each path in reais, and the boxed numbers above the nodes are the expected losses obtained by rolling the tree back from the right, the probability-weighted average at a chance node and the minimum at a decision node. Thick branches are the ones chosen: inspect, then replace after Y^+ and continue after Y^-.

Figure 32.4 compares the three strategies. The horizontal axis lists them; the height of each bar is its expected loss in reais, with the value above it, and the bar of the inspection strategy is split into the expected loss after the outcome-dependent decision (lower part) and the inspection cost (upper part). Inspecting and then deciding costs R$ 4720 in expectation, against R$ 8000 for replacing now and R$ 10 000 for continuing. The model assumes that the result is available at once and that the inspection adds no outage risk; an inspection that takes the pump offline would add the risk of its maintenance window, as in Section 24.5.2.

With perfect information the operator would replace the pump exactly when it would fail, at R$ 8000, and continue at no incremental cost otherwise. By equation (32.13),

$$C_\text{perfect} = 0.20 \times 8000 + 0.80 \times 0 = \text{R\$}\,1600, \qquad \text{EVPI} = C_\text{now} - C_\text{perfect} = 8000 - 1600 = \text{R\$}\,6400.$$

The imperfect inspection captures

$$\frac{\text{EVSI}}{\text{EVPI}} = \frac{4280}{6400} = 0.669,$$

or 66.9 %, of the decision value available from perfect information. The remaining $\text{EVPI} - \text{EVSI} = C_\text{sample} - C_\text{perfect} = \text{R\$}\,2120$ is lost because the inspection sometimes gives a wrong answer. Splitting C_sample by the joint outcomes,

$$C_\text{sample} = \underbrace{0.18 \times 8000}_{\text{detected failures}} + \underbrace{0.16 \times 8000}_{\text{false alarms}} + \underbrace{0.02 \times 50\,000}_{\text{missed failures}} = 1440 + 1280 + 1000 = \text{R\$}\,3720,$$

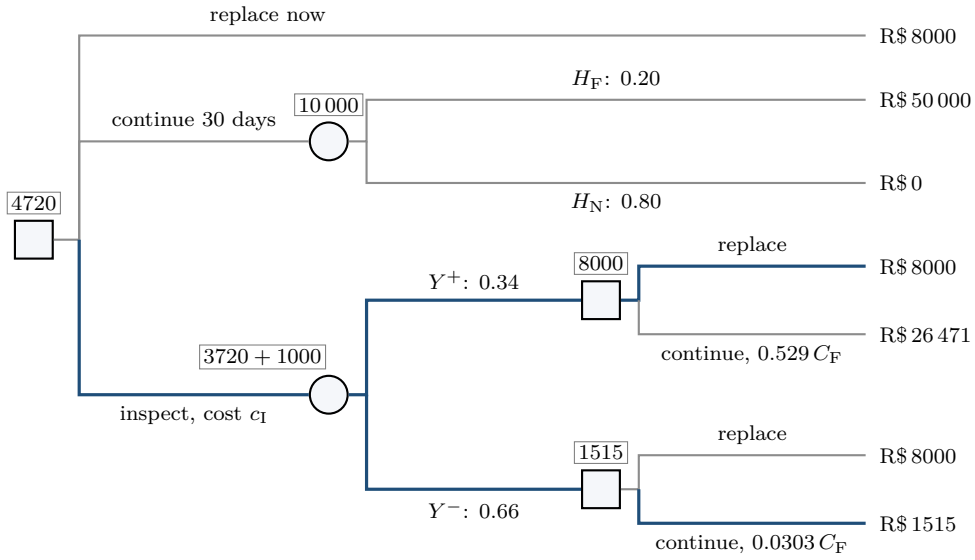


Figure 32.3: Decision tree of the pump decision with the inspection. Squares are decisions, circles are chance outcomes with their probabilities; losses in reais at the right; boxed values are the rolled-back expected losses. Thick branches are chosen: inspect, then replace after Y^+ and continue after Y^- ($C_R = \text{R\$}\,8000$, $C_F = \text{R\$}\,50\,000$, $c_I = \text{R\$}\,1000$; illustrative values).



Figure 32.4: Expected losses of the three strategies for the pump: replace now (C_R), continue 30 days ($\Pr(H_F)\,C_F$), and inspect then decide ($C_{sample} + c_I$), in reais (illustrative values).

while perfect information replaces exactly the failing pumps, $C_{\text{perfect}} = (0.18 + 0.02) \times 8000 =$ R\$ 1600. The difference consists of the unnecessary replacements triggered by false alarms, $0.16 \times 8000 =$ R\$ 1280, and of the missed failures, which cost C_{F} where perfect information would have paid C_{R}, $0.02 \times (50\,000 - 8000) =$ R\$ 840. For this inspection the four values are EVSI = R\$ 4280, EVPI = R\$ 6400, $c_{\text{I}} =$ R\$ 1000 and a net inspection benefit of R\$ 3280.

32.4.5 Three candidate inspections

Suppose three tests are available for the same pump, all used from the same prior $\Pr(H_{\text{F}}) = 0.20$ and with the same costs C_{R} and C_{F}: Test A, the standard inspection of the preceding subsections (sensitivity 0.90, specificity 0.80, cost R\$ 1000); Test B, a targeted diagnostic check (sensitivity 0.65, specificity 0.95, cost R\$ 400); and Test C, a comprehensive laboratory evaluation (sensitivity and specificity 0.98, cost R\$ 6000). For each test, the posterior failure probabilities after each outcome, the optimal action after each, the expected loss C_{sample}, the EVSI and the expected information gain are computed as above. Table 32.4 lists the results; its columns are the test characteristics, the probability of a positive outcome, the two posteriors, the expected information gain in bits, the EVSI, its share of the EVPI of R\$ 6400, and the net value after the test cost.

Table 32.4: Three candidate inspections of the pump, prior $\Pr(H_{\text{F}}) = 0.20$, $C_{\text{R}} =$ R\$ 8000, $C_{\text{F}} =$ R\$ 50 000, EVPI = R\$ 6400, prior entropy 0.722 bits (illustrative values; money in reais). In every case the action is to replace after Y^+ and continue after Y^-.

Test	Sens., spec.	Cost	$\Pr(Y^+)$	$\Pr(H_{\text{F}} \mid Y^+)$	$\Pr(H_{\text{F}} \mid Y^-)$	EIG (bits)	EVSI	Share of EVPI	Net
A, standard inspection	0.90, 0.80	1000	0.340	0.5294	0.0303	0.253	4280	66.9 %	3280
B, targeted check	0.65, 0.95	400	0.170	0.7647	0.0843	0.242	3140	49.1 %	2740
C, laboratory evaluation	0.98, 0.98	6000	0.212	0.9245	0.0051	0.604	6104	95.4 %	104

One row of the arithmetic, Test B: $\Pr(Y^+) = 0.65 \times 0.20 + 0.05 \times 0.80 = 0.13 + 0.04 = 0.17$; after Y^+ the failure probability is $0.13/0.17 = 0.7647 > p^*$, so the pump is replaced; after Y^- it is $0.07/0.83 = 0.0843 < p^*$, so it is kept, with expected loss $0.07 \times 50\,000 =$ R\$ 3500 over the negative branch. Hence $C_{\text{sample}} = 0.17 \times 8000 + 3500 = 1360 + 3500 =$ R\$ 4860 and EVSI $= 8000 - 4860 =$ R\$ 3140. Its expected information gain is $\mathcal{H}(Y) - \mathbb{E}_H[\mathcal{H}(Y \mid H)] = 0.658 - (0.20 \times 0.934 + 0.80 \times 0.286) = 0.658 - 0.416 = 0.242$ bits.

Test C provides the largest expected information gain, 0.604 bits, removing 84 % of the prior entropy, and the largest gross decision benefit, R\$ 6104, almost the whole EVPI. After its cost of R\$ 6000, its net value is only R\$ 104. Test A provides the greatest net benefit. Under this decision model, therefore,

$$\text{optimal test} = \text{Test A}, \qquad \text{most informative test} = \text{Test C}.$$

Tests A and B also illustrate the difference between the two measures: their information gains are nearly equal, 0.253 and 0.242 bits, but Test A is worth R\$ 1140 more before costs. Test B misses 35 % of the failures, and at a prior above the threshold the costly error is the missed failure, whereas Test A's errors are mostly false alarms, which cost only an early replacement. The entropy weighs the two kinds of error alike; the loss matrix does not. In expectation, Test A produces false alarms with probability 0.16 and missed failures with probability 0.02; Test B produces them with probabilities 0.04 and 0.07, and each missed failure costs R\$ 42 000 more than a replacement while each false alarm costs R\$ 8000; the difference of their EVSIs is exactly $(0.07 - 0.02) \times 42\,000 - (0.16 - 0.04) \times 8000 = 2100 - 960 =$ R\$ 1140.

Low-cost targeted measurements. A measurement does not need to characterize the whole physical system to have high operational value; it may only need to distinguish between actions. When the condition of a pump could justify either an immediate replacement or continued operation, an inexpensive check may suffice to move the posterior failure probability decisively above or below the replacement threshold, as Test B does here. A comprehensive characterization of the bearing geometry, the lubrication state and the microscopic damage could provide much more scientific information and add little to the immediate maintenance decision. This does not make comprehensive diagnostics wasteful. They may have long-term value for root-cause analysis, for fleet reliability models and for future prognostics, and that value would appear in a calculation whose objectives and time horizon include those uses. The value of a test depends on the objectives and the horizon that the calculation includes.

> **Principle 32.3.** *The most informative measurement is not necessarily the best one. A test is preferred when its expected value of sample information, net of its cost, is the largest among the candidates; information gain ranks measurements by what they teach, and the net decision value ranks them by what they are worth.*

32.5 Measurement selection for parameters and competing models

So far the unknown was a discrete event. Many engineering models carry continuous parameters, and several physical models may compete to explain the same data. This section derives the expected information gain of a linear measurement under a Gaussian posterior, applies it to the thermal balance of Chapter 31, contrasts information-optimal with decision-optimal measurement placement, and then treats measurement selection under model uncertainty, a sequential instrumentation strategy and the conditional information of a measurement.

32.5.1 From discrete hypotheses to continuous parameters

In many engineering models the uncertain quantities are continuous parameters rather than a binary event such as H_{F}. The running example is the data center thermal balance of Section 31.2.4, $Q_{\mathrm{cool}} = G\,\Delta T + P_{\mathrm{int}}$, with the parameter vector

$$\boldsymbol{\theta} = \begin{bmatrix} G \\ P_{\mathrm{int}} \end{bmatrix},$$

where G is the effective envelope conductance in kW/K, P_{int} the internal heat generation in kW, and $\Delta T = T_{\mathrm{out}} - T_{\mathrm{in}}$ the temperature difference across the envelope in kelvin. Given the evidence $\mathcal{Y}$ collected so far, the state of knowledge is the posterior density $p(\boldsymbol{\theta} \mid \mathcal{Y})$. A candidate measurement of the cooling load at a chosen operating condition is

$$Y = G\,\Delta T + P_{\mathrm{int}} + \varepsilon, \tag{32.17}$$

with measurement error ε. The engineer can choose the temperature difference at which it is taken, or the type of sensor, and the question is which choice maximizes the expected information about $\boldsymbol{\theta}$. Here the hypotheses form a continuum, and the information gain is the divergence form of equation (32.6) with densities and integrals.

32.5.2 Gaussian posterior update and information gain

Suppose the current posterior is Gaussian, $\boldsymbol{\theta} \mid \mathcal{Y} \sim \mathcal{N}(\bar{\boldsymbol{\theta}}, \boldsymbol{\Sigma})$, with mean $\bar{\boldsymbol{\theta}}$ and covariance $\boldsymbol{\Sigma}$ of dimension $n_\theta \times n_\theta$, and that the candidate measurement is linear in the parameters,

$$Y = \mathbf{s}^{\mathsf{T}}\boldsymbol{\theta} + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_y^2), \tag{32.18}$$

with noise independent of $\boldsymbol{\theta}$ and of the earlier measurements. The vector $\mathbf{s} = \partial Y / \partial \boldsymbol{\theta}$ is the *sensitivity vector* of the measurement, one row of the sensitivity matrix $\mathbf{S}_\theta$ of equation (31.5) written as a column; it is not the residual sensitivity matrix $\mathbf{S}$ of Chapter 30. Equation (31.21) showed that a Gaussian linear update adds the information of the measurement to the current information, so the posterior covariance after the measurement is

$$\boldsymbol{\Sigma}_{\text{new}} = \Big(\boldsymbol{\Sigma}^{-1} + \frac{\mathbf{s}\,\mathbf{s}^{\mathsf{T}}}{\sigma_y^2} \Big)^{-1}, \tag{32.19}$$

and for several measurements with sensitivity matrix $\mathbf{S}_\theta$ and weight matrix $\mathbf{W} = \mathbf{R}^{-1}$ it is $(\boldsymbol{\Sigma}^{-1} + \mathbf{S}_\theta^{\mathsf{T}}\mathbf{W}\mathbf{S}_\theta)^{-1}$. Under the linear Gaussian assumptions the new covariance does not depend on the numerical value of Y: how much will be learned is known before the measurement is made, although what will be learned is not.

The *differential entropy* of an n_θ-dimensional Gaussian density with covariance $\boldsymbol{\Sigma}$ is the standard result

$$\mathcal{H}(\boldsymbol{\theta}) = \tfrac{1}{2} \ln\big[(2\pi \mathrm{e})^{n_\theta} \det \boldsymbol{\Sigma} \big], \tag{32.20}$$

in nats, the unit of information when the natural logarithm is used (1 nat $= 1/\ln 2 = 1.443$ bits). It depends on the covariance only. Since the posterior covariance after the measurement is the same for every outcome, the expected posterior entropy equals the entropy at $\boldsymbol{\Sigma}_{\text{new}}$, and equation (32.5) gives

$$\mathrm{EIG} = \mathcal{H}(\boldsymbol{\theta}) - \mathcal{H}(\boldsymbol{\theta} \mid Y) = \frac{1}{2} \ln \frac{\det \boldsymbol{\Sigma}}{\det \boldsymbol{\Sigma}_{\text{new}}}, \tag{32.21}$$

the constants $(2\pi\mathrm{e})^{n_\theta}$ cancelling. The ratio of determinants is the factor by which the volume of the uncertainty ellipsoid shrinks, squared.

For one scalar measurement the determinant has a closed form. The *matrix determinant lemma* states that for an invertible matrix $\mathbf{A}$ and vectors $\mathbf{u}$, $\mathbf{v}$, $\det(\mathbf{A} + \mathbf{u}\mathbf{v}^{\mathsf{T}}) = (1 + \mathbf{v}^{\mathsf{T}}\mathbf{A}^{-1}\mathbf{u}) \det \mathbf{A}$. It follows from the block factorization

$$\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{v}^{\mathsf{T}} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} + \mathbf{u}\mathbf{v}^{\mathsf{T}} & \mathbf{u} \\ \mathbf{0}^{\mathsf{T}} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{v}^{\mathsf{T}} & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{u} \\ \mathbf{0}^{\mathsf{T}} & 1 + \mathbf{v}^{\mathsf{T}}\mathbf{u} \end{bmatrix},$$

whose outer factors have determinant one, so that $\det(\mathbf{I} + \mathbf{u}\mathbf{v}^{\mathsf{T}}) = 1 + \mathbf{v}^{\mathsf{T}}\mathbf{u}$; writing $\mathbf{A} + \mathbf{u}\mathbf{v}^{\mathsf{T}} = \mathbf{A}(\mathbf{I} + \mathbf{A}^{-1}\mathbf{u}\mathbf{v}^{\mathsf{T}})$ gives the lemma. With $\mathbf{A} = \boldsymbol{\Sigma}^{-1}$ and $\mathbf{u} = \mathbf{v} = \mathbf{s}/\sigma_y$,

$$\frac{\det \boldsymbol{\Sigma}}{\det \boldsymbol{\Sigma}_{\text{new}}} = \frac{\det \boldsymbol{\Sigma}_{\text{new}}^{-1}}{\det \boldsymbol{\Sigma}^{-1}} = 1 + \frac{\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{s}}{\sigma_y^2},$$

and therefore

$$\mathrm{EIG} = \frac{1}{2} \ln\Big(1 + \frac{\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{s}}{\sigma_y^2} \Big). \tag{32.22}$$

The same result follows from the symmetric form of the mutual information: the predicted measurement has variance $\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{s} + \sigma_y^2$ and, given $\boldsymbol{\theta}$, variance σ_y^2, and the difference of the two Gaussian entropies is $\tfrac{1}{2}\ln[(\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{s} + \sigma_y^2)/\sigma_y^2]$. For several measurements the same argument with Sylvester's determinant identity gives $\tfrac{1}{2}\ln\det(\mathbf{I} + \mathbf{R}^{-1/2}\mathbf{S}_\theta\boldsymbol{\Sigma}\mathbf{S}_\theta^{\mathsf{T}}\mathbf{R}^{-1/2})$.

Equation (32.22) has a useful reading. The quantity $\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{s}$ is the variance of the noise-free prediction $\mathbf{s}^{\mathsf{T}}\boldsymbol{\theta}$ under the current uncertainty, so the information gain grows with the ratio of

the predicted signal variance to the noise variance: a measurement is informative when its predicted value varies strongly over the parameters that are still plausible, relative to its noise. A large raw signal is not enough. The measurement must respond along the directions in which the parameters are still uncertain; a measurement whose sensitivity vector lies along a direction the data have already fixed has a small $\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{s}$ however large $\mathbf{s}$ is.

32.5.3 Sensitivity of a thermal measurement

For the thermal measurement of equation (32.17), the sensitivity vector is

$$\mathbf{s} = \begin{bmatrix} \Delta T \\ 1 \end{bmatrix}, \qquad \mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{s} = \Delta T^2\,\Sigma_{GG} + 2\,\Delta T\,\Sigma_{GP} + \Sigma_{PP},$$

where Σ_{GG}, Σ_{GP} and Σ_{PP} are the entries of the current covariance. A measurement at $\Delta T = 0$ constrains P_{int} directly; a measurement at a large $|\Delta T|$ is more sensitive to G; several measurements at varied ΔT separate the two parameters, which is the excitation principle of Section 31.2.6 restated as the design of the next measurement.

As an illustration, give the parameters the independent prior $G \sim \mathcal{N}(2, 1^2)$ in kW/K and $P_{\text{int}} \sim \mathcal{N}(40, 10^2)$ in kW, so that $\boldsymbol{\Sigma}_0 = \operatorname{diag}(1, 100)$, and let every cooling-load measurement have $\sigma_y = 1\,\text{kW}$ (illustrative values). A first measurement is taken at the usual operating condition, $\Delta T = 10\,\text{K}$, as in Experiment A of Section 31.2.5. Its sensitivity vector is $(10, 1)^{\mathsf{T}}$, $\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}_0\mathbf{s} = 100 \times 1 + 100 = 200$, and its expected information gain is $\frac{1}{2}\ln 201 = 2.652$ nats. By equation (32.19) the posterior information and covariance after it are

$$\boldsymbol{\Sigma}_1^{-1} = \begin{bmatrix} 1 & 0 \\ 0 & 0.01 \end{bmatrix} + \begin{bmatrix} 100 & 10 \\ 10 & 1 \end{bmatrix} = \begin{bmatrix} 101 & 10 \\ 10 & 1.01 \end{bmatrix},$$

$$\boldsymbol{\Sigma}_1 = \frac{1}{2.01}\begin{bmatrix} 1.01 & -10 \\ -10 & 101 \end{bmatrix} = \begin{bmatrix} 0.5025 & -4.975 \\ -4.975 & 50.25 \end{bmatrix},$$

with standard deviations 0.709 kW/K and 7.09 kW and a correlation of -0.990: the measurement fixed the combination $10\,G + P_{\text{int}}$ and left the compensating direction of Section 31.2.5 almost as uncertain as before.

Which second measurement is best? For a candidate at temperature difference ΔT,

$$\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}_1\mathbf{s} = 0.5025\,\Delta T^2 - 9.950\,\Delta T + 50.25,$$

a parabola with its minimum, 0.990, at $\Delta T = 9.95/1.005 = 9.90\,\text{K}$, essentially the condition already measured. Table 32.5 evaluates equation (32.22) at five candidate conditions. Its columns are the temperature difference, the predicted signal variance, the expected information gain in nats and in bits, and the posterior standard deviations of the two parameters after the second measurement. One row of the arithmetic, $\Delta T = 20\,\text{K}$: $0.5025 \times 400 - 9.950 \times 20 + 50.25 = 201.0 - 199.0 + 50.25 = 52.24$, so EIG $= \frac{1}{2}\ln 53.24 = 1.987$ nats, or 2.867 bits.

Figure 32.5 draws the expected information gain of the next measurement against the temperature difference at which it is taken. The horizontal axis is ΔT in kelvin and the vertical axis the information gain in nats. The dashed gray curve is the gain of a first measurement under the prior, $\frac{1}{2}\ln(1 + \Delta T^2 + 100)$; the solid blue curve is the gain of the second measurement after the first at 10 K; the dots mark the conditions of Table 32.5, and the dotted vertical lines bound an assumed feasible operating range from 0 to 20 K. Before any measurement every condition is informative and the gain rises with $|\Delta T|$. After the first measurement the curve has a deep minimum at the condition already measured: repeating it is worth 0.345 nats, a sixth of a measurement at either end of the range. The two ends are nearly equivalent, 1.968 nats at 0 and 1.987 nats at 20 K, because either one, combined with the first, fixes both parameters.

Table 32.5: Expected information gain of a second cooling-load measurement after a first one at $\Delta T = 10\,\mathrm{K}$, prior $\boldsymbol{\Sigma}_0 = \mathrm{diag}(1, 100)$, $\sigma_y = 1\,\mathrm{kW}$ (illustrative values).

ΔT (K)	$\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}_1\mathbf{s}$ (kW^2)	EIG (nats)	EIG (bits)	σ_G after (kW/K)	σ_P after (kW)
0	50.25	1.968	2.840	0.140	0.990
5	13.06	1.322	1.907	0.267	2.112
10 (repeat)	0.995	0.345	0.498	0.708	7.080
15	14.05	1.356	1.956	0.258	3.282
20	52.24	1.987	2.867	0.137	2.164

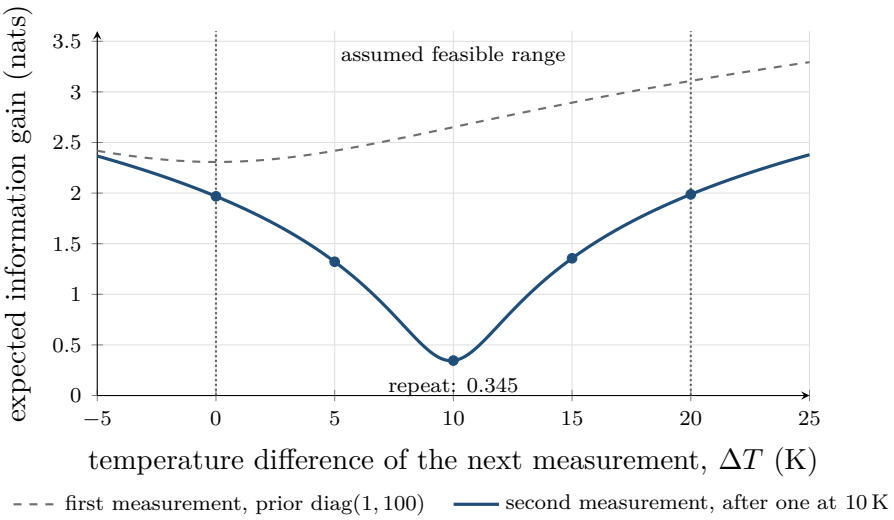


Figure 32.5: Expected information gain of one cooling-load measurement against its temperature difference, equation (32.22) with $\mathbf{s} = (\Delta T, 1)^{\mathsf{T}}$ and $\sigma_y = 1\,\mathrm{kW}$: before any measurement (prior covariance $\mathrm{diag}(1, 100)$) and after a first measurement at $10\,\mathrm{K}$ (covariance $\boldsymbol{\Sigma}_1$), EIG $= \tfrac{1}{2}\ln(1 + 0.5025\,\Delta T^2 - 9.950\,\Delta T + 50.25)$ (illustrative values).

The figure connects the practical identifiability problem of Chapter 31 with the design of experiments: the measurement that adds the most information is the one that responds along the direction the existing data leave unconstrained, and a repetition of the existing condition responds only along the direction already fixed.

Common mistake

Taking the condition of largest information gain without checking that it can be reached. The gain of equation (32.22) keeps increasing as $|\Delta T|$ grows, beyond any range the plant can or should operate in. A measurement at an extremely large temperature difference may have a large theoretical information gain and be unsafe, infeasible, or outside the range in which the physical model has been validated, where its sensitivity vector no longer describes the system. Experimental design must respect the operating constraints and the validity of the model.

32.5.4 Information-optimal and decision-optimal placement

Suppose two candidate measurements are available: Sensor A, which mainly reduces the uncertainty of G, and Sensor B, which mainly reduces that of P_{int}. If the engineering decision concerns the envelope insulation, the uncertainty of G may matter most; if it concerns the internal electrical heat generation, for example the sizing of a power upgrade, P_{int} matters more. A *D-optimal* measurement plan maximizes the determinant of the posterior information, equivalently minimizes $\det \boldsymbol{\Sigma}_{\text{new}}$, the squared volume of the posterior uncertainty ellipsoid; under the linear Gaussian model of Section 32.5.2 this is the plan of largest expected information gain, by equation (32.21). A *decision-optimal* plan instead minimizes the expected operational loss after the measurements are used.

The two can be computed side by side for the prior of Section 32.5.3, before any measurement, with one measurement to place at either $\Delta T = 0$ or $\Delta T = 20\,\text{K}$. Table 32.6 gives, for each, the expected information gain and the posterior variances of the two parameters. The measurement at 20 K has the larger information gain, 3.108 nats against 2.308, and is the D-optimal choice. It reduces the variance of G from 1 to 0.202 but leaves that of P_{int} at $80.0\,\text{kW}^2$; the measurement at 0 K leaves G untouched and reduces the variance of P_{int} from 100 to $0.990\,\text{kW}^2$.

Table 32.6: One measurement placed at $\Delta T = 0$ or 20 K under the prior $\text{diag}(1, 100)$ with $\sigma_y = 1\,\text{kW}$: information gain and posterior variances (illustrative values).

Measurement	EIG (nats)	$\text{Var}(G)$ after (kW^2/K^2)	$\text{Var}(P_{\text{int}})$ after (kW^2)
None (prior)	–	1.000	100.0
At $\Delta T = 0$	2.308	1.000	0.990
At $\Delta T = 20\,\text{K}$	3.108	0.202	80.04

Suppose the decision is to set a quantity a, for example the electrical capacity reserved for the internal load, with a loss that grows with the square of the error, $L(a, \boldsymbol{\theta}) = k\,(a - P_{\text{int}})^2$, with k in reais per kW^2. The expected loss $k\,\mathbb{E}[(a - P_{\text{int}})^2]$ is minimized by the posterior mean of P_{int}, and its minimum is $k\ \text{Var}(P_{\text{int}})$. Since the posterior variance does not depend on the outcome, the EVSI of a measurement is k times the reduction of $\text{Var}(P_{\text{int}})$: $99.0\,k$ for the measurement at 0 K and $20.0\,k$ for the one at 20 K. For a decision about the envelope, with loss $k\,(a - G)^2$, the ranking reverses: $0.798\,k$ against zero. The information-optimal and the decision-optimal choices differ,

$$a^*_{\text{EIG}} \neq a^*_{\text{EVSI}},$$

in general, where a^{*}_{EIG} and a^{*}_{EVSI} denote the measurement choices that maximize the information gain and the decision value. The difference is not a defect of either objective; they answer different questions, the first about the volume of what is unknown and the second about the part of it the decision depends on.

32.5.5 Posterior predictive distributions under competing models

Uncertainty may concern the structure of the model rather than its parameters. Consider a cooling system with three competing explanations of an anomaly,

$$M_1 : \text{insufficient heat-transfer capacity},$$
$$M_2 : \text{an unmetered internal heat source},$$
$$M_3 : \text{a bias of the temperature sensor},$$

a candidate set different from the three models of Section 31.5.1, and suppose that the current measurements cannot distinguish them. A new sensor is valuable if its possible outputs have sufficiently different predictive distributions under the models. For a candidate measurement Y, the *posterior predictive distribution* under model M_k averages the measurement model over the posterior of that model's parameters $\boldsymbol{\theta}_k$,

$$p(y \mid \mathcal{Y}, M_k) = \int p(y \mid \boldsymbol{\theta}_k, M_k)\, p(\boldsymbol{\theta}_k \mid \mathcal{Y}, M_k)\, \mathrm{d}\boldsymbol{\theta}_k, \tag{32.23}$$

so that it accounts for the uncertainty of the parameters as well as for the measurement noise; Chapter 8 of Volume I used the same construction for the reliability of a new unit. The overall predictive distribution mixes the models with their current posterior probabilities, which Section 31.5.3 computed from the evidence,

$$p(y \mid \mathcal{Y}) = \sum_{k} \Pr(M_k \mid \mathcal{Y})\, p(y \mid \mathcal{Y}, M_k). \tag{32.24}$$

These two distributions are all that the preposterior analysis needs. For each possible outcome y, Bayes' theorem with the predictive densities as likelihoods gives the posterior model probabilities $\Pr(M_k \mid \mathcal{Y}, y) = \Pr(M_k \mid \mathcal{Y})\, p(y \mid \mathcal{Y}, M_k)/p(y \mid \mathcal{Y})$; averaging their entropy over $p(y \mid \mathcal{Y})$ gives the expected information gain about the model, and averaging the minimum expected loss gives the EVSI. Measurement selection is therefore based on what the competing models predict the instrument will read, not on the quantity the instrument nominally measures.

32.5.6 An instrumentation comparison and a sequential strategy

Table 32.7 lists the qualitative response of four candidate measurements under the three models. Each row is a candidate measurement, each column a model, and each entry the response the model predicts. The entries are illustrative signatures, not universal fault rules.

An independent temperature sensor can distinguish a sensor bias from a real thermal abnormality, but it cannot distinguish insufficient cooling from an additional heat source, since both raise the true return temperature. An electrical-load audit can separate those two. To put numbers on the table, translate each entry into the probability of an "abnormal" outcome under each model, take equal prior probabilities of $1/3$, an entropy of $\log_2 3 = 1.585$ bits, and compute the expected information gain of each measurement with the symmetric form of Section 32.2.3. Table 32.8 gives the result. Its first three numeric columns are the assumed probabilities of an abnormal outcome under each model; the fourth is the expected information gain of the measurement taken first; and the last is its expected information gain after the independent temperature sensor has been read, the conditional information of Section 32.5.7.

Table 32.7: Predicted responses of four candidate measurements under three explanations of a cooling anomaly (illustrative signatures).

Candidate measurement	M_1, insufficient cooling	M_2, unmetered heat source	M_3, biased return sensor
Independent return-temperature sensor	elevated	elevated	normal
Cooling-water flow measurement	may be low	may be normal	may be normal
Independent IT electrical-load audit	may be normal	elevated	may be normal
Cooling-equipment capacity test	may show deficiency	may show adequate performance	may show adequate performance

Table 32.8: Expected information gain of the four candidate measurements about the three models, equal priors (illustrative probabilities of an abnormal outcome). In the last column, the entry of the return-temperature sensor is that of a second, independent sensor of the same kind.

Measurement	M_1	M_2	M_3	EIG first (bits)	EIG after the return sensor (bits)
Independent return-temperature sensor	0.95	0.95	0.05	0.648	0.159
Cooling-water flow measurement	0.60	0.20	0.20	0.113	0.090
Independent IT electrical-load audit	0.10	0.90	0.10	0.479	0.384
Cooling-equipment capacity test	0.85	0.15	0.15	0.351	0.281

One row of the arithmetic, the return-temperature sensor taken first: the probability of an elevated reading is $(0.95 + 0.95 + 0.05)/3 = 0.65$, whose binary entropy is 0.934 bits; the outcome entropy under each model is 0.286 bits; so EIG $= 0.934 - 0.286 = 0.648$ bits. The temperature sensor is the most informative first measurement, because it isolates the bias, the hypothesis that would make every other thermal reading misleading. After an elevated reading, the bias has probability 0.026 and the two physical explanations share the rest equally; after a normal one, the bias has probability 0.905.

A useful measurement strategy is therefore sequential:

1. check whether the temperature anomaly is real (independent temperature sensor);
2. if it is real, distinguish a change of heat generation from a deficiency of cooling (electrical-load audit or capacity test);
3. if necessary, estimate the specific degraded parameters, as in Section 32.5.3.

Each step is chosen with the posterior left by the previous one, and a later step is skipped when the earlier result already settles the decision. This can be considerably more economical than installing all possible sensors at once.

32.5.7 Conditional information gain

Two candidate sensors may provide nearly identical information. The first may be valuable; after it is installed, the second may add little. For two measurements Y_1 and Y_2, the information that Y_2 adds after Y_1 has been observed is the *conditional mutual information*

$$I(H; Y_2 \mid Y_1) = \mathbb{E}_{Y_1}\big[\mathcal{H}(H \mid Y_1)\big] - \mathbb{E}_{Y_1,Y_2}\big[\mathcal{H}(H \mid Y_1, Y_2)\big], \tag{32.25}$$

the expected entropy reduction by Y_2 from the posterior that Y_1 leaves, averaged over both outcomes. It satisfies the chain rule

$$I(H; Y_1, Y_2) = I(H; Y_1) + I(H; Y_2 \mid Y_1),$$

which follows by adding and subtracting $\mathbb{E}_{Y_1}[\mathcal{H}(H \mid Y_1)]$ in the definition of $I(H; Y_1, Y_2)$, and it need not equal $I(H; Y_2)$. In the last column of Table 32.8, a second, independent return-temperature sensor, with the same characteristics and independent errors given the model, adds 0.159 bits after the first, against the 0.648 bits the first one brought; the two together are worth $0.648 + 0.159 = 0.807$ bits. The electrical-load audit, worth 0.479 bits first, still adds 0.384 bits after the temperature sensor, and the pair is worth $0.648 + 0.384 = 1.031$ bits, the best pair of the table. The ranking of the second measurement is not the ranking of the first: after the temperature sensor, the audit is worth more than twice a duplicate of it.

Information value is conditional on the evidence already acquired, and the same holds for the economic value of information: a sensor that is valuable before an inspection may be redundant after it, because the action it could change has already been chosen. Optimal instrumentation planning is therefore sequential rather than a static ranking of sensors, each choice being made with the posterior left by the previous ones.

32.6 Sequential decisions, inspection economics and model credibility

Measurements and actions alternate over time, and every calculation of the previous sections rests on a model. This section states the sequential decision problem on the belief state, prices the option of waiting for evidence, tabulates the economics of the pump inspection over a range of inputs, and examines the conditions under which a value of information is credible, ending with a workflow for recommending a measurement.

32.6.1 The Bellman equation for the belief state

Suppose the actions at each decision epoch are $a_t \in \{\text{operate}, \text{inspect}, \text{repair}, \text{replace}\}$. The state of knowledge at time t is the belief state of equation (24.1), $b_t = p(\mathbf{x}_t, \boldsymbol{\theta} \mid \mathcal{Y}_{0:t})$, the posterior over the physical condition $\mathbf{x}_t$ and the uncertain model parameters $\boldsymbol{\theta}$ given the evidence up to t. A new measurement changes the belief; an action changes the physical future. A replacement may reset the degradation, a load reduction may slow it, an inspection changes the knowledge but not necessarily the physical condition, and continued operation preserves revenue but creates failure exposure. The decision is therefore sequential. Section 24.6.2 stated its optimality condition, the Bellman equation of the partially observable Markov decision process, which for a finite horizon and without discounting reads

$$V_t(b) = \min_{a} \Big[\, c(b,a) + \mathbb{E}_{Y_{t+1}}\, V_{t+1}\big(\mathcal{B}(b,a,Y_{t+1})\big) \Big], \tag{32.26}$$

where $V_t(b)$ is the minimum expected future cost from epoch t with belief b, $c(b,a)$ the immediate expected cost of action a under that belief, $\mathcal{B}(b,a,y)$ the Bayesian belief update of equation (24.30) after action a and observation y, and Y_{t+1} the future observation, whose distribution depends on b and a. The optimal policy accounts both for the immediate consequence of an action and for how it changes the future state and the future information.

The single inspection of Section 32.4 is the two-epoch case. At the last epoch, $V_1(b) = \min\{C_{\mathrm{R}},\ p(b)\,C_{\mathrm{F}}\}$, with $p(b)$ the failure probability under belief b; at the first,

$$\begin{aligned} V_0(b) &= \min\Big\{C_{\mathrm{R}},\ p\,C_{\mathrm{F}},\ c_{\mathrm{I}} + \Pr(Y^{+})\,V_1(b^{+}) + \Pr(Y^{-})\,V_1(b^{-})\Big\} \\ &= \min\{8000,\ 10\,000,\ 1000 + 2720 + 1000\} = \text{R\$}\,4720, \end{aligned}$$

where b^{+} and b^{-} are the beliefs after each outcome; the third term is $c_{\mathrm{I}} + C_{\mathrm{sample}}$, and equation (32.26) reproduces the preposterior analysis.

32.6.2 The value of waiting

Sometimes the best decision is to delay an irreversible intervention while more evidence is collected. Waiting has a cost: a degrading component may fail before the next observation. The value of waiting is therefore not the expected information gain of the coming observation; it is

$$\begin{aligned} \text{value of waiting} = \text{expected decision improvement} &- \text{failure exposure during the delay} \\ &- \text{measurement and operational cost}, \end{aligned} \tag{32.27}$$

all in reais and relative to the best action available now.

A worked case makes the terms concrete. Instead of the paid inspection, the operator of the pump of Section 32.4 can wait 7 days for a reading of an installed condition-monitoring sensor, which costs nothing and has the same sensitivity 0.90 and specificity 0.80 against H_{F}. Assume, illustratively, that a fraction $f = 0.25$ of the failures under H_{F} occur within the first 7 days. A failure during the wait has probability $0.20 \times 0.25 = 0.05$ and costs $C_{\mathrm{F}} = \text{R\$}\,50\,000$. If the pump survives the week, the failure probability is updated by the survival itself, $0.20 \times 0.75/0.95 = 0.158$, and then by the reading: a positive reading, with probability 0.311, raises it to 0.458 and the pump is replaced; a negative one lowers it to 0.023 and the pump is kept, which adds an expected loss of $0.1 \times 0.158 \times 50\,000 = \text{R\$}\,789$ from the failures the reading misses. The expected loss after the week is $0.311 \times 8000 + 789 = \text{R\$}\,3274$, and the expected loss of the waiting strategy is

$$0.05 \times 50\,000 + 0.95 \times 3274 = 2500 + 3110 = \text{R\$}\,5610,$$

an instance of equation (32.26) with the immediate cost $c(b, \text{wait}) = \text{R\$}\,2500$.

Against the best immediate action without information, replacing now at R$ 8000, waiting is worth $8000 - 5610 =$ R$ 2390. The terms of equation (32.27) are: the expected decision improvement of the reading, had it been available at once, $8000 - 3720 =$ R$ 4280 (the EVSI of Section 32.4.4); the failure exposure during the delay, $5610 - 3720 =$ R$ 1890; and no measurement cost. The exposure has a closed form. The only paths on which the delay changes the loss are the failures in the first week, with probability $\Pr(H_{\mathrm{F}})\, f$; on them the immediate reading would have led to a replacement with probability 0.90, saving $C_{\mathrm{F}} - C_{\mathrm{R}}$, so the exposure is $\Pr(H_{\mathrm{F}})\, f \times 0.90 \times (C_{\mathrm{F}} - C_{\mathrm{R}}) = 0.20 \times 0.25 \times 0.90 \times 42\,000 =$ R$ 1890. Waiting is better than replacing now but worse than the paid inspection, R$ 4720, because the inspection's fee of R$ 1000 is smaller than the exposure of R$ 1890. The exposure equals $7560\, f$ reais, so waiting beats inspecting when $f < 1000/7560 = 0.13$; with $f = 0.10$ the waiting strategy costs $3720 + 756 =$ R$ 4476. If an inspection requires taking a critical pump offline, its effect on the redundancy of the plant must be included in the same way; if a thermal test requires operating equipment near its limits, its operational risk must be included. The ideal mathematical experiment may not be the safest physical experiment.

32.6.3 Inspection economics over a range of inputs

The value of the pump inspection depends on the prior failure probability, on the sensitivity and specificity of the test and on the costs. Table 32.9 evaluates the single-period model of Section 32.4 for several inputs, changed one group at a time from the base case (prior 0.20, sensitivity 0.90, specificity 0.80, $C_{\mathrm{R}} =$ R$ 8000, $C_{\mathrm{F}} =$ R$ 50 000, $c_{\mathrm{I}} =$ R$ 1000). The model assumes that replacement prevents the failure event, that the inspection is instantaneous and does not change the condition of the equipment, and that there is no inspection downtime, action delay, intervention risk or later operating cost. The columns are the two posteriors, the minimum expected loss without information and the action that attains it (R, replace now; C, continue), the expected loss with the inspection before its cost, the EVSI, the EVPI, and the strategy of least expected loss with that loss. One row of the arithmetic, prior 0.10: $\Pr(Y^{+}) = 0.09 + 0.18 = 0.27$, posteriors $0.09/0.27 = 0.333$ (replace) and $0.01/0.73 = 0.0137$ (continue); $C_{\mathrm{now}} = \min\{8000,\ 5000\} =$ R$ 5000; $C_{\mathrm{sample}} = 0.27 \times 8000 + 0.01 \times 50\,000 = 2160 + 500 =$ R$ 2660; EVSI = R$ 2340; EVPI $= 5000 - 0.10 \times 8000 =$ R$ 4200; inspecting costs $2660 + 1000 =$ R$ 3660 < 5000.

Three trends follow. First, raising the inspection cost leaves the EVSI unchanged at R$ 4280, because the information content and the consequences are unchanged, but it lowers the net value: at R$ 3000 inspecting still pays, at R$ 5000 it no longer does, and the break-even cost is the EVSI itself. Second, reducing the sensitivity and specificity brings the two posteriors together and reduces the expected improvement of the decision, from R$ 4280 to R$ 1960 at 0.70 and R$ 480 at 0.60; at 0.55 both posteriors lie above the threshold 0.16, both outcomes lead to replacement, and the EVSI is zero although the test is not entirely uninformative. Third, moving the prior toward zero makes replacement uneconomic whatever the result: at 0.02 even a positive outcome leaves the failure probability at 0.084, below the threshold, the preferred actions after both outcomes coincide, and the inspection is worth nothing. The same happens at the other end, where at 0.60 even a negative outcome leaves the probability at 0.158, a hair below the threshold, and the EVSI is only R$ 40.

Figure 32.6 draws the EVSI and the EVPI against the prior failure probability for the base test and costs. The horizontal axis is the prior $\Pr(H_{\mathrm{F}})$ and the vertical axis the value in reais; the solid blue line is the EVSI, the dashed red line the EVPI, and the dotted horizontal line the inspection cost R$ 1000, above which the inspection pays. Both curves are piecewise linear, because each expected loss is linear in the prior between the points where a best action changes, and both peak at the threshold $p^{*} = 0.16$, where the operator is indifferent without information: the EVPI at R$ 6720, the EVSI at R$ 4704. The EVSI is zero below 0.0406, where even a

Table 32.9: The pump inspection decision for several priors, test qualities and costs, changed from the base case (prior 0.20, sensitivity 0.90, specificity 0.80, $C_{\mathrm{R}} = \mathrm{R\$}\,8000$, $C_{\mathrm{F}} = \mathrm{R\$}\,50\,000$, $c_{\mathrm{I}} = \mathrm{R\$}\,1000$); illustrative values, money in reais. R: replace now; C: continue; tie: both actions have the same expected loss.

Case	$\Pr(H_{\mathrm{F}} \mid Y^{+})$	$\Pr(H_{\mathrm{F}} \mid Y^{-})$	C_{now}	C_{sample}	EVSI	EVPI	Preferred (expected loss)
Base case	0.529	0.0303	8000 (R)	3720	4280	6400	inspect (4720)
$c_{\mathrm{I}} = 3000$	0.529	0.0303	8000 (R)	3720	4280	6400	inspect (6720)
$c_{\mathrm{I}} = 5000$	0.529	0.0303	8000 (R)	3720	4280	6400	replace (8000)
Sens., spec. 0.70, 0.70	0.368	0.0968	8000 (R)	6040	1960	6400	inspect (7040)
Sens., spec. 0.60, 0.60	0.273	0.143	8000 (R)	7520	480	6400	replace (8000)
Sens., spec. 0.55, 0.55	0.234	0.170	8000 (R)	8000	0	6400	replace (8000)
Prior 0.02	0.084	0.0025	1000 (C)	1000	0	840	continue (1000)
Prior 0.05	0.191	0.0065	2500 (C)	2130	370	2100	continue (2500)
Prior 0.10	0.333	0.0137	5000 (C)	2660	2340	4200	inspect (3660)
Prior 0.16	0.462	0.0233	8000 (tie)	3296	4704	6720	inspect (4296)
Prior 0.40	0.750	0.0769	8000 (R)	5840	2160	4800	inspect (6840)
Prior 0.60	0.871	0.158	8000 (R)	7960	40	3200	replace (8000)
$C_{\mathrm{F}} = \mathrm{R\$}\,100\,000$	0.529	0.0303	8000 (R)	4720	3280	6400	inspect (5720)
$C_{\mathrm{R}} = \mathrm{R\$}\,12\,000$	0.529	0.0303	10 000 (C)	5080	4920	7600	inspect (6080)

positive outcome leaves the posterior below p^{*}, and above 0.604, where even a negative outcome leaves it above. With the R\$ 1000 fee the inspection pays for priors between 0.066 and 0.509. The EVSI curve lies everywhere below the EVPI curve, as equation (32.15) requires.

32.6.4 Value of information depends on model adequacy

Every expected information gain and every EVSI of this chapter was computed from predictive distributions supplied by models. If the models are wrong, the calculated value of information may also be wrong. Suppose a temperature sensor is believed to have a random error of standard deviation $\sigma_y = 0.05$°C, but it has an unmodeled systematic bias of 0.5°C. Measuring a temperature whose prior standard deviation is 1°C, the design calculation of equation (32.22) promises $\frac{1}{2}\log_2(1 + 1/0.05^2) = 4.32$ bits, and the posterior reports the temperature with a standard deviation of 0.050°C. The reported value is in fact displaced by about 0.50°C, ten reported standard deviations: the experiment delivers a confident and wrong conclusion. Had the bias been represented as an unknown offset of standard deviation 0.5°C, the error variance would have been $0.05^2 + 0.5^2$ and the honest information gain $\frac{1}{2}\log_2[1 + 1/(0.05^2 + 0.5^2)] = 1.16$ bits. In the same way, a test may have an impressive sensitivity under laboratory conditions and perform differently under the actual loading of the equipment, and the posteriors of Table 32.3 would then be wrong.

Instrumentation selection should therefore settle the following questions before the value of information is trusted:

- the calibration and possible bias of each sensor;
- the independence between measurements, including shared calibrations and common causes;
- the correctness of the physical response model used to predict the outcomes;
- the representativeness of the tested operating conditions;
- competing hypotheses that the set omits;
- whether the test outcome will be available in time to affect the decision.

An information-theoretic optimum computed under an inadequate model is not necessarily a

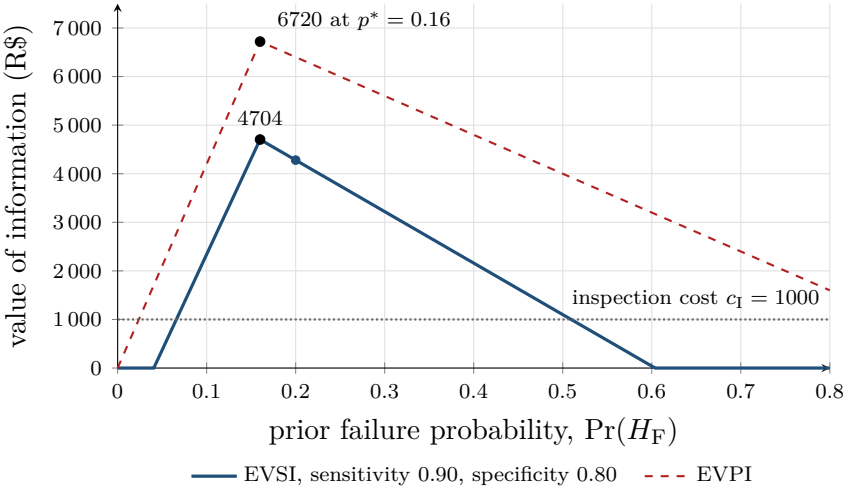


Figure 32.6: Expected values of sample and perfect information of the pump inspection against the prior failure probability, C_R = R$ 8000, C_F = R$ 50 000 (illustrative values). EVPI $= 42\,000\,p$ for $p \le 0.16$ and $8000(1-p)$ above; EVSI $= 39\,400\,p - 1600$ on $[0.0406, 0.16]$ and $6400 - 10\,600\,p$ on $[0.16, 0.604]$, zero elsewhere. The blue dot at 0.20 is the base case, EVSI = R$ 4280.

credible engineering recommendation; the adequacy checks of Section 31.4 apply to the models that predict a measurement as much as to the models fitted to it.

32.6.5 Ambiguities that no measurement resolves

Additional measurements cannot resolve every ambiguity. If two physical hypotheses produce identical predictions for every feasible candidate measurement of a family, no measurement of that family can distinguish them: their predictive distributions coincide, and by Section 32.2.4 the information gain is zero for every member of the family. The room with two envelope paths of Section 31.1.3 is the example: from repeated total-heat-flow measurements,

$$Q = (G_1 + G_2)\,\Delta T,$$

the conductances G_1 and G_2 cannot be identified separately at any temperature difference. The sensitivity vector of every such measurement is $\Delta T\,(1,1)^{\mathsf{T}}$, which has no component along the difference direction $(1,-1)^{\mathsf{T}}$; with a prior covariance that does not correlate the sum and the difference, as the independent priors of Section 31.3.7 do, the posterior variance of $G_1 - G_2$ stays at its prior value of $2\,\mathrm{kW^2/K^2}$ after any number of them. A different kind of measurement is needed, such as a heat-flux sensor on one path, or additional spatial temperature measurements interpreted with a model that separates the paths.

Two rules summarize the design of measurements. An experimental-design rule: choose measurements that respond differently to the uncertain explanations. A decision rule: choose the test whose expected benefit justifies its full cost and risk.

32.6.6 A measurement-recommendation workflow

The steps of the chapter assemble into a workflow for recommending the next measurement. Figure 32.7 draws it. Each box is one step, numbered in the order in which it is performed and read row by row, left to right; arrows connect consecutive steps, and the arrow from the last box back to the second returns the updated posterior to the start of the next cycle.

1. *Define the engineering decision*: the operational actions, the constraints and the consequences, as a loss function.

2. *Represent the current uncertainty*: posterior probabilities over states, parameters or competing models.
3. *Identify the candidate measurements*: the feasible sensors, tests, inspections and operating experiments.
4. *Predict the possible observations*: each test's outcome distribution under the competing hypotheses, the posterior predictive distributions of equation (32.23).
5. *Compute the information gain*: which tests are expected to reduce the uncertainty, equation (32.6).
6. *Compute the decision value*: the optimal action under each possible outcome and the averaged loss, the EVSI of equation (32.11).
7. *Account for the cost and risk of the test*: procurement, delay, equipment downtime and operational exposure.
8. *Check the adequacy of the model*: whether the predicted outcomes and the uncertainty assumptions are credible.
9. *Select, measure and update*: obtain the preferred measurement, update the posterior, and reconsider the decision.

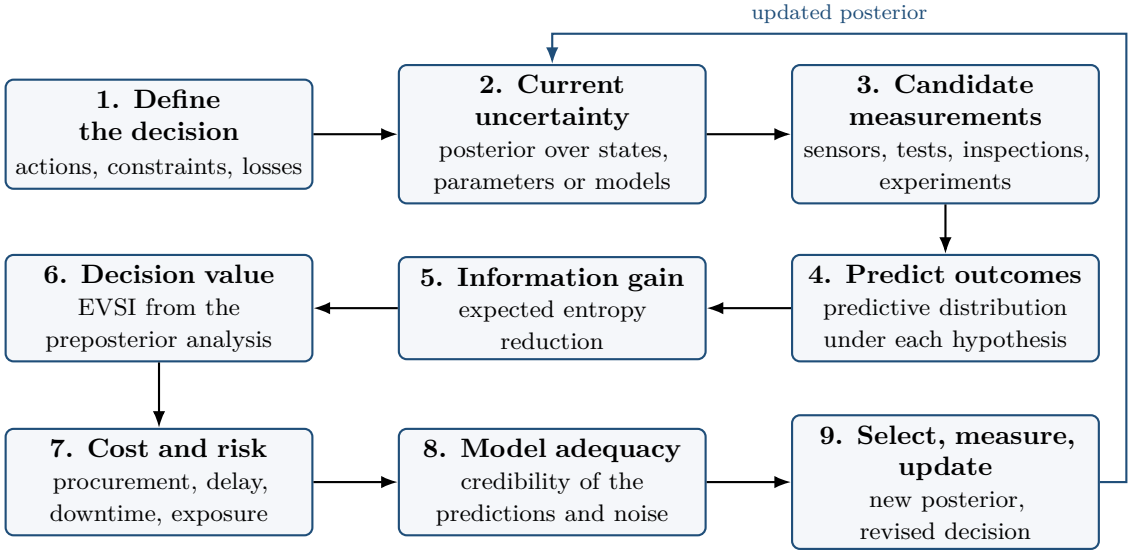


Figure 32.7: The measurement-recommendation workflow. Steps 1–3 state the decision, the uncertainty and the candidates; steps 4–6 evaluate each candidate by its predicted outcomes, its information gain and its decision value; steps 7–9 subtract its cost and risk, check the model and act. The loop returns the updated posterior for the next measurement.

The workflow makes the order of the questions explicit. The decision is defined before any measurement is evaluated, because without it only information gain can be computed, and information gain cannot say which uncertainty matters. The information gain is computed before the decision value because it is cheaper and identifies the measurements that cannot be worth anything, those with zero gain. Model adequacy is checked before the measurement is made, because a value of information computed with a wrong predictive model is itself wrong. The loop makes the procedure sequential, as Section 32.5.7 requires.

Principle 32.4. *The best next measurement is the one that most improves the decision, net of its cost and risk, under a model whose predictions are credible, not necessarily the one that most reduces uncertainty.*

The methods of this part of the book establish what the measurements of a system can and cannot determine, and what the next measurement is worth. The next chapter, Chapter 33,

turns to how such probabilistic and physical analyses are embedded in engineering practice: reliability requirements and acceptance criteria, the standards that govern them, and the documentation and review that make an assessment auditable.

32.7 What to retain

Definition 32.1 (Value of information)**.** The value of a measurement for a decision is its expected value of sample information: the minimum expected loss of the decision with the current information, C_{now}, minus the expected value of the minimum loss after the measurement's outcome, C_{sample}. It is never negative, never exceeds the expected value of perfect information, and differs from the expected information gain, which measures the expected reduction of uncertainty rather than of loss.

Bayesian updating converts a measurement into revised probabilities over the physical explanations, by the likelihood of the observation under each and the prior; in odds form the observation contributes a Bayes factor. The Shannon entropy measures the dispersion of the probabilities, and the expected information gain, the mutual information between hypothesis and outcome, measures the expected reduction of that dispersion; a single outcome can increase it, the average cannot. The expected value of sample information measures the expected improvement of the decision, and the expected value of perfect information bounds it. The most accurate or most informative test may not be the economically preferred one, and two tests with the same information gain can have very different decision values. For continuous parameters under a Gaussian posterior, the information of a measurement is $\frac{1}{2}\ln(1+\mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{s}/\sigma_y^2)$, so experimental design should target the directions that existing measurements do not constrain. The value of a measurement depends on the evidence already available and on the actions it can change, which makes instrumentation planning sequential, and it is only as credible as the models that predict its outcomes. Table 32.10 lists the chapter's terms, their meaning, expression and units, and the subsection that defines each.

Table 32.10: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Posterior probability	updated by an observation	equation (32.1)	–	32.1.1
Predictive probability	probability of an outcome before it is seen	$\Pr(y) = \sum_j \Pr(y \mid H_j)\Pr(H_j)$	–	32.1.1
Posterior odds, Bayes factor	odds after; likelihood ratio	equation (32.2)	–	32.1.3
Shannon entropy	dispersion of the probabilities	$-\sum_i p_i \log_2 p_i$	bit	32.2.1
Realized entropy reduction	reduction by one outcome; may be negative	$\mathcal{H}(H) - \mathcal{H}(H \mid Y = y)$	bit	32.2.2
Expected information gain, mutual information	expected entropy reduction	$\mathcal{H}(H) - \mathbb{E}_Y[\mathcal{H}(H \mid Y)] = I(H;Y)$	bit	32.2.3
Kullback–Leibler divergence	divergence of posterior from prior	$\sum_i q_i \log_2(q_i/p_i)$	bit	32.2.3
Loss function	loss of an action under a hypothesis	$L(a, H)$	R$	32.3.1
Expected loss, Bayes action	averaged loss; its minimizer	$\rho(a)$, $a^* = \arg\min_a \rho(a)$	R$	32.3.2
Expected loss with sample information	preposterior expected loss	C_{sample}, equation (32.10)	R$	32.3.3
Expected value of sample information	loss reduction from a test	$C_{\text{now}} - C_{\text{sample}} \geq 0$	R$	32.3.3
Expected value of perfect information	upper bound on EVSI	$C_{\text{now}} - C_{\text{perfect}}$	R$	32.3.4
Replacement threshold	failure probability of indifference	$p^* = C_{\text{R}}/C_{\text{F}}$	–	32.4.2
Sensitivity vector	response of a measurement to the parameters	$\mathbf{s} = \partial Y/\partial \boldsymbol{\theta}$	varies	32.5.2
Gaussian information gain	information of a linear measurement	$\frac{1}{2}\ln(1 + \mathbf{s}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{s}/\sigma_y^2)$	nat	32.5.2
D-optimal plan	largest information determinant	$\min \det \boldsymbol{\Sigma}_{\text{new}}$	–	32.5.4
Posterior predictive distribution	outcome under a model, parameters averaged	equation (32.23)	–	32.5.5
Conditional mutual information	information added after earlier data	$I(H; Y_2 \mid Y_1)$	bit	32.5.7
Bellman equation	sequential optimal policy on beliefs	equation (32.26)	R$	32.6.1
Value of waiting	improvement minus exposure and cost	equation (32.27)	R$	32.6.2

Further reading. Gelman and coauthors, *Bayesian Data Analysis*: posterior inference, predictive distributions and model checking. Chaloner and Verdinelli, *Bayesian Experimental Design: A Review*: Bayesian optimal experimental design. Berger, *Statistical Decision Theory and Bayesian Analysis*: Bayesian decisions and the value of information. Clemen, *Making Hard Decisions*: decision trees and value-of-information methods.

32.8 Exercises

Exercise 32.1. (a) State what Bayes' theorem combines. (b) State what the expected information gain measures. (c) State what the EVSI measures. (d) State the inequality that holds between EVSI and EVPI under a coherent common decision model. (e) State whether a highly informative test can be economically inferior to a less informative one, and why. (f) State why varying the operating conditions can improve parameter identifiability. (g) State what an engineer should do when competing models predict nearly identical outcomes for the candidate measurements. (Answers: (a) prior probabilities and the likelihoods of the observation; (b) the expected reduction of uncertainty about a specified unknown quantity, not a reduction of cost; (c) the expected improvement of the optimal decision's expected loss brought by an imperfect measurement, before its cost; (d) $0 \le \text{EVSI} \le \text{EVPI}$; (e) yes, once the test costs and the way the outcomes change the decisions are counted, as Test C against Test A shows; (f) it creates different parameter sensitivities across the measurements; (g) seek a different measurement or experiment whose predictions separate the models, rather than repeat the same test or accept the model with the most parameters.)

Exercise 32.2. For the three explanations of Section 32.1.1, with priors 0.50, 0.30 and 0.20, a cooling-equipment capacity test gives an abnormal result with probability 0.05 under H_0, 0.70 under H_C and 0.05 under H_S. (a) Compute the probability of an abnormal result and the three posteriors after it. (b) Compute the Bayes factor of H_C against H_S and the posterior odds. (c) Compute the posteriors after a normal result. (d) Compute the expected information gain of the test, and compare it with the 0.521 bits of the test of Section 32.2.3. (Answers: (a) $\Pr(Y^+) = 0.025 + 0.210 + 0.010 = 0.245$; posteriors 0.102, 0.857, 0.041; (b) $0.70/0.05 = 14$, posterior odds $1.5 \times 14 = 21$; (c) $\Pr(Y^-) = 0.755$; posteriors 0.629, 0.119, 0.252; (d) posterior entropies 0.715 and 1.287 bits, so EIG $= 1.485 - (0.245 \times 0.715 + 0.755 \times 1.287) = 0.338$ bits, less than 0.521: the capacity test misses 30 % of the shortfalls, and its Bayes factor of 1 between health and bias leaves that distinction where the prior put it.)

Exercise 32.3. A later prognosis gives the pump of Section 32.4 a failure probability of 0.10 within 30 days, with $C_\text{R} = \text{R\$}\,8000$ and $C_\text{F} = \text{R\$}\,50\,000$. (a) Compute C_now and the best action without inspection. (b) For Test B (sensitivity 0.65, specificity 0.95, cost R$ 400), compute the posteriors, the policy, C_sample, the EVSI and the net value. (c) Compare with Test A (sensitivity 0.90, specificity 0.80, cost R$ 1000) and Test C (sensitivity and specificity 0.98, cost R$ 6000), and state which test to perform. (d) Compute the EVPI. (Answers: (a) $\min\{8000,\ 0.10 \times 50\,000\} = \text{R\$}\,5000$, continue; (b) $\Pr(Y^+) = 0.065 + 0.045 = 0.11$, $\Pr(H_\text{F} \mid Y^+) = 0.591$, replace; $\Pr(H_\text{F} \mid Y^-) = 0.035/0.89 = 0.0393$, continue; $C_\text{sample} = 0.11 \times 8000 + 0.035 \times 50\,000 = 880 + 1750 = \text{R\$}\,2630$, EVSI $= \text{R\$}\,2370$, net R$ 1970; (c) Test A: $C_\text{sample} = \text{R\$}\,2660$, EVSI $= \text{R\$}\,2340$, net R$ 1340; Test C: $C_\text{sample} = \text{R\$}\,1028$, EVSI $= \text{R\$}\,3972$, net $-\text{R\$}\,2028$; perform Test B, which at this prior has the largest EVSI of the two cheaper tests although its information gain, 0.149 bits, barely exceeds Test A's 0.145; (d) EVPI $= 5000 - 0.10 \times 8000 = \text{R\$}\,4200$.)

Exercise 32.4. The thermal parameters have the prior covariance $\text{diag}(1, 100)$ (G in kW/K, P_int in kW), and one cooling-load measurement with a noisier meter, $\sigma_y = 2\,\text{kW}$, can be taken at $\Delta T = 0$ or at $\Delta T = 20\,\text{K}$. (a) Compute the expected information gain of each, in nats and bits, and state the D-optimal choice. (b) Compute the posterior variances of G and of P_int after each. (c) With the loss $k\,(a - P_\text{int})^2$ of Section 32.5.4, state the decision-optimal choice and its EVSI. (d) A colleague proposes to repeat the measurement at 20 K a hundred times instead; state, with the reason, whether this resolves P_int as well as one measurement at $\Delta T = 0$. (Answers: (a) at 0: $\tfrac{1}{2}\ln(1 + 100/4) = 1.629$ nats, 2.350 bits; at 20 K: $\tfrac{1}{2}\ln(1 + 500/4) = 2.418$ nats, 3.489 bits; D-optimal at 20 K; (b) at 0: $\text{Var}(G) = 1$, $\text{Var}(P_\text{int}) = 100 - 100^2/104 = 3.85\,\text{kW}^2$; at 20 K: $\text{Var}(G) = 1 - 400/504 = 0.206$, $\text{Var}(P_\text{int}) = 100 - 100^2/504 = 80.2\,\text{kW}^2$; (c) the measurement

at $\Delta T = 0$, with EVSI $= 96.2\,k$ against $19.8\,k$; (d) no: repetitions at one condition constrain only the combination $20\,G + P_{\text{int}}$, so P_{int} remains limited by the prior uncertainty of G, and its variance can fall no lower than $400 \times 100/(400 + 100) = 80\,\text{kW}^2$ however many readings are taken.)

Part IX

Engineering Applications and Decisions

Chapter 33

Reliability Standards and Engineering Practice

33.1 Engineering assessments and the role of standards

The preceding chapters developed methods for estimating failure probabilities, availability, expected unserved energy, degradation, parameter uncertainty and engineering risk. A technically sophisticated model does not, by that fact alone, constitute an acceptable engineering assessment. An organization that receives a reliability study must approve it, act on it and keep it current as the system and the data change, and it can do so only when the meaning of every number, the evidence behind it and the authority that accepts the remaining risk are established. The question of this chapter is

how a reliability model is turned into an auditable engineering assessment that an organization can trust, approve, implement and maintain.

The chapter answers it in six steps. It states why engineering standards exist and separates compliance requirements from additional engineering analysis. It describes three families of standards that bear on reliability work: the transmission planning standards of the North American Electric Reliability Corporation, the functional safety standard IEC 61508 and the asset management standards ISO 55000 and 55001. It states what makes a reliability requirement useful and how a target differs from the evidence that a system meets it. It carries out a complete assessment of a two-pump cooling system against a stated requirement, including a design improvement and its economics. It separates verification from validation and builds the traceability from requirement to decision. Finally, it treats risk registers, acceptance authority and the effect of parameter uncertainty on an acceptance decision.

Jargon: engineering standard, compliance, requirement, acceptance criterion, verification, validation, traceability, configuration management, assurance, risk register, acceptance authority

An *engineering standard* is a document, issued by a recognized body, that sets common requirements, processes, methods, terminology or evidence expectations. *Compliance* is the demonstrated satisfaction of the obligations that a standard, a regulation or a contract places on a specified entity. A *requirement* is a statement of what a system must do or achieve, written so that it can be checked; an *acceptance criterion* is the measurable condition, with a threshold, that decides whether the requirement is met. *Verification* is the demonstration that a model or a design has been implemented and solved correctly against its specification; *validation* is the demonstration that it represents the real system adequately for its intended use. *Traceability* is the documented chain that connects each requirement to the models, inputs, tests, results and decisions that address it. *Configuration management* is the control of the versions of models, data, software and assumptions from which a result was produced. *Assurance* is the body of evidence and review that justifies confidence in a claim. A *risk register* is a structured record of identified risks, their causes, consequences, controls, owners and status. The *acceptance authority* is the person or body formally responsible for accepting a result, or the residual risk it describes, on behalf of the organization.

Key idea

A numerical result becomes engineering evidence only when its meaning, its assumptions, its validity for the intended use and its relevance to an approved requirement are established and recorded, and when a named authority accepts the decision it supports.

This section states the problem with an example, explains why engineering standards exist, separates compliance from additional analysis, and compares the three families of standards that the next three sections examine.

33.1.1 A probability is not yet an assessment

Suppose a report delivered to a utility or a data center operator states that the probability of a critical equipment outage is 0.00038. The number is the steady-state unavailability of the two redundant pumps of Chapter 19 of Volume I with two repair crews, $\pi_0 = 0.00038447$, and it is computed correctly. Before acting on it, the organization must still establish:

1. what precisely counts as an outage (both pumps down, flow below the design value, or a temperature excursion in the served equipment);
2. over what time period, or at what instant, the probability applies;
3. which components, dependencies and operating conditions the model includes;
4. where the failure and repair parameters come from;
5. whether the implementation of the model has been verified;
6. whether the model has been validated against relevant evidence;
7. whether the result satisfies an approved reliability requirement;
8. who is authorized to accept the remaining risk.

The second question alone changes the reading of the number. The value 0.00038 is a point unavailability, the probability that the system is down at a randomly chosen instant in the stationary regime. The probability of at least one complete outage during a year that starts with both pumps operating is 0.281 (Section 19.4.4), more than seven hundred times larger. Both numbers are correct; they answer different questions, and a report that does not say which question it answers does not support a decision. The eight questions move the work from mathematical reliability analysis into engineering governance and assurance, which is the subject of this chapter.

33.1.2 Why engineering standards exist

Engineering standards provide common requirements, processes, methods, terminology and evidence expectations. Their first function is coordination: when every party to a project uses the same definitions of a contingency, a safety function or an asset-management objective, the results of different analysts can be compared and audited. Their second function is protection against inconsistent assumptions and unsubstantiated claims. A standard that specifies which events a planning study must examine, or which failure measure a safety function must satisfy, removes the freedom to choose the events or the measure that make a design look acceptable.

Different standards serve different purposes. Some establish mandatory performance or process requirements for the entities to which they apply, so that failing to meet them is a violation with formal consequences. Others describe a management system, a set of recommended practices, or methods for achieving a specific engineering objective; an organization adopts them, and may be certified against them, but their content is a framework within which the organization sets its own objectives. Which kind a standard is, which edition is in force, and which entity and system it applies to must be established before any claim of conformity is made.

33.1.3 Compliance requirements and additional engineering analysis

An essential distinction runs through the chapter:

compliance requirements against additional engineering analysis.

A compliance requirement is an obligation defined by an applicable standard or regulation, assessed by the method that the standard prescribes. Additional engineering analysis is any study that an organization performs to understand and manage its risk beyond those obligations: a probabilistic risk assessment, an expected-unserved-energy calculation, a value-of-information study.

The two do not substitute for each other in either direction. A probabilistic risk assessment may improve decision-making considerably without satisfying a particular regulatory requirement, because the requirement prescribes a specific deterministic test that the assessment did not perform. Conversely, meeting a prescribed minimum planning requirement does not prove that every relevant physical risk has been eliminated, because the prescribed events are a selection and the physical system can fail in ways outside it.

Principle 33.1. *A compliance claim and a risk estimate are different deliverables. Each must state its own basis, and neither may be presented as the other.*

33.1.4 Three families of standards

Table 33.1 compares the three families of standards examined in this chapter. Each row is one family; the columns give its principal purpose and the methods of this book to which it connects. The families are not interchangeable. The NERC transmission planning (TPL) standards define planning obligations for the bulk transmission system. IEC 61508 addresses how the required integrity of safety-related functions is achieved and maintained. ISO 55000 establishes the principles of asset management, and ISO 55001 specifies the requirements of an asset-management system. Their applicability depends on the organization, the system, the jurisdiction and the intended function: a data center operator is not a registered entity under the NERC standards, and a transmission planner's contingency study is not a functional safety assessment.

Table 33.1: Three families of standards that bear on reliability work and their connection to the methods of this book.

Standard	Principal purpose	Connection to the methods of this book
NERC TPL standards	Planning performance of the Bulk Electric System transmission network	Contingency analysis, planning criteria, corrective action (Chapter 27)
IEC 61508	Functional safety of electrical, electronic and programmable electronic safety-related systems	Safety functions, safety integrity, lifecycle assurance
ISO 55000 family	Management of assets throughout their lifecycle	Reliability, maintenance, risk, performance and cost

33.2 Transmission planning standards

The first family applies to the transmission networks of Chapters 27 and 28. This section states what the NERC planning standard TPL-001-5.1 addresses, summarizes its planning-event categories, contrasts its deterministic criteria with the probabilistic contingency analysis of Chapter 27, lists the evidence a planning assessment retains, and describes the companion standard for extreme-temperature events.

Jargon: NERC, registered entity, Bulk Electric System, Transmission Planner, Planning Coordinator, planning assessment, planning event, corrective action plan

The *North American Electric Reliability Corporation* (NERC) develops and enforces reliability standards for the bulk power system of North America. A *registered entity* is an organization registered with NERC in one or more functional roles, to which the standards for those roles apply. The *Bulk Electric System* is the part of the power system, defined by voltage level and function, to which the standards apply. The *Transmission Planner* and the *Planning Coordinator* are the functional roles responsible for planning the transmission system of an area and for coordinating the plans of several areas. A *planning assessment* is the documented evaluation of the future transmission system against the performance requirements. A *planning event* is a contingency, or a defined sequence of contingencies, that the assessment must simulate. A *corrective action plan* is the documented set of actions that removes a performance deficiency found by the assessment.

33.2.1 Scope of TPL-001-5.1

NERC establishes reliability standards that apply to specified registered entities. TPL-001-5.1, *Transmission System Planning Performance Requirements*, establishes planning-performance requirements for the Bulk Electric System across a range of system conditions and contingencies. It applies to Transmission Planners and Planning Coordinators. Its requirements address the maintenance of the system models used for planning, the annual planning assessment, the steady-state and stability studies that the assessment contains, and the development of corrective action plans when the studies show that the performance requirements are not met.

The performance obligations are tied to defined planning-event categories. The standard is not a generic instruction to calculate every imaginable $N-k$ outage, nor an instruction to compute a probability of any kind. It specifies which events must be simulated, under which system conditions, and what performance is required after each: which facility ratings

and voltage limits must be respected, whether the system must remain stable, and whether interruption of firm transmission service or loss of load is permitted as part of the response. In the vocabulary of Chapter 27, the standard fixes the contingency list and the response model of the deterministic analysis, and makes both obligations.

33.2.2 Planning-event categories

Table 33.2 gives a conceptual summary of the broad planning-event categories of TPL-001. Each row is one category; the columns give its label, whether it is a single or a multiple contingency, and a simplified description. The summary is not a reproduction of the standard's event definitions. The actual requirements are those of Table 1 of the standard and its footnotes, which define each event precisely (which elements, which fault types, which system adjustments are allowed between the events of a sequence) and state the performance required after each, and they govern any assessment.

Table 33.2: Planning-event categories of TPL-001: a conceptual summary, not a reproduction of the standard's definitions, which are those of its Table 1 and footnotes.

Category	Contingency class	Simplified description
P0	none	Normal system, no contingency
P1	single	Single contingencies: loss of one generator, line, transformer or similar element
P2	single	Certain bus, breaker or line-section events
P3	multiple	Loss of a generator followed, after system adjustments, by the loss of another element
P4	multiple	Stuck-breaker events: a fault whose clearing fails because a breaker does not open
P5	multiple	A fault combined with the failure of a nonredundant component of the protection system
P6	multiple	Sequential loss of two elements, with system adjustments between them
P7	multiple	Loss of circuits on a common structure, or specified DC-line events

The categories cover the vocabulary of Part III of Volume I. Category P1 is the $N-1$ criterion of Chapter 27. Categories P4 and P5 are failures of the protection that should clear a fault, the sequences that an event tree (Chapter 15 of Volume I) would branch on. Category P7 is a common-cause failure in the sense of Chapter 13 of Volume I: two circuits on one tower share a coupling factor, the structure, and are prescribed as a joint event rather than as two independent outages.

33.2.3 Deterministic planning and probabilistic risk

The three-bus network of Chapter 27 illustrates the difference between the two kinds of analysis. With all three lines in service and a transfer of 100 MW from bus 1 to bus 3, the DC flows of Section 27.2.3 are $P_{12} = P_{23} = P_{13} = 50\,\text{MW}$. After the outage of the direct line L13 the whole transfer crosses the series path (Section 27.2.4),

$$P_{12} = P_{23} = 100\,\text{MW},$$

against assumed ratings $P_{12}^{\max} = P_{23}^{\max} = 70\,\text{MW}$. The deterministic analysis identifies a post-contingency violation: an overload of 30 MW on each line of the indirect path. In planning

terms this is a single-contingency event that does not meet its performance criterion, and the planner must document a corrective action, such as an upgrade of L12 and L23, a new line, or an operating limit on the transfer, that restores compliant performance. Whether curtailing load is an admissible response to a given event is itself fixed by the event category and the footnotes of the standard, not by its expected cost.

A probabilistic model asks something else. With the illustrative unavailability $q_{13} = 0.01$ of Section 27.4.1, and the other two lines in service with probabilities 0.98 and 0.97, the network is in the state "L13 alone out" with probability $0.98 \times 0.97 \times 0.01 = 0.009506$. Under the curtailment response model of Chapter 27, that state requires 30 MW of curtailment, and its contribution to the expected curtailment is

$$\mathcal{R}_{\mathrm{L13}} = 0.009506 \times 30\,\mathrm{MW} = 0.28518\,\mathrm{MW},$$

or $0.28518 \times 8760 = 2498$ MWh of expected unserved energy per year if the same conditions held all year (Table 27.4). The outage itself is expected about once a year, with a mean duration of 87.6 h (Table 27.5).

The two calculations answer different questions. Table 33.3 pairs them: each row places a question of the deterministic assessment beside the corresponding question of the probabilistic one. The left column is about a prescribed event and a pass-or-fail criterion; the right column is about frequencies, consequence distributions and the allocation of improvement effort.

Table 33.3: Questions answered by a deterministic planning assessment and by a probabilistic risk assessment of the same network.

Deterministic assessment	Probabilistic assessment
Does the modeled planning event satisfy the applicable performance criteria?	How frequently might a specified loss occur?
Which limits are exceeded?	What is the distribution of consequences?
What corrective actions are required?	Which improvements yield the greatest expected risk reduction?
What evidence demonstrates that the study was performed?	How uncertain are the probability and consequence estimates?

A low modeled probability does not exempt a required planning event from assessment: the outage of L13, unavailable only 1 % of the time, must still be simulated and its violation corrected. Similarly, a favorable expected-risk result does not establish compliance with TPL-001: an expected unserved energy judged acceptable by a cost-benefit criterion says nothing about whether each prescribed event meets its performance requirement. The probabilistic analysis remains valuable for choosing among corrective actions and for risks the prescribed events do not cover, but it is additional analysis in the sense of Section 33.1.3.

33.2.4 Planning models and documented evidence

A defensible transmission planning assessment may need to document the study-area topology; the generator and load assumptions; the reactive-power capability of the generators and compensating devices; the facility ratings; the protection systems and automatic controls; the expected operating configurations, including planned outages; the contingency definitions; the steady-state and stability results; and the identified deficiencies with their corrective action plans. Each item is an input or an output of the models of Chapters 27 and 28, and the assessment is only as credible as the weakest of them.

TPL-001-5.1 includes explicit evidence obligations for the models, the planning assessments and the studies on which they rest. Reliability engineering under such a standard therefore

consists not only of performing calculations, but of retaining evidence that the models and analyses correspond to the system conditions the standard requires: the model version that represented the network, the data that populated it, the contingency list that was simulated and the results that were obtained. Section 33.5.5 develops the same idea as a general traceability chain.

33.2.5 Extreme-temperature planning

A companion standard, TPL-008-1, addresses transmission planning for extreme-temperature events. It introduces planning requirements for extreme heat and extreme cold conditions, under which demand, generator availability and facility ratings all depart from their normal values at once, and its requirements take effect in phases under an implementation plan. The obligations that apply to a given entity on a given date are therefore read from the approved standard and its implementation plan, requirement by requirement, and are not assumed.

Probabilistic weather modeling of the kind developed in Chapter 28 can complement such assessments by estimating, with their uncertainty, the conditional probabilities

$$\Pr(\text{contingency} \mid \text{weather}) \qquad \text{and} \qquad \Pr(\text{overload} \mid \text{contingency, weather}),$$

the first from the weather-dependent hazard $\lambda(t \mid Z)$ of the equipment and the second from the weather-dependent line ratings. The probabilistic model and the compliance assessment should nevertheless remain distinguishable deliverables, each with its own basis: the compliance assessment simulates the extreme-temperature scenarios the standard defines, and the probabilistic model estimates how likely and how costly the outcomes are across the full range of weather.

33.3 Functional safety and asset management standards

The second family of standards concerns protective functions rather than the equipment that performs the required function, and the third places reliability within the management of an organization's assets. This section separates reliability from functional safety, states how a safety function is specified, defines the safety integrity levels and their failure measures, derives the average probability of failure on demand of a periodically tested channel, and describes the safety lifecycle; it then describes the ISO 55000 family and a maintenance-strategy decision in its terms.

Jargon: functional safety, safety function, safety-related system, dangerous and safe failures, dangerous undetected failure, demand, low-demand and high-demand mode, proof test, SIL, $\mathrm{PFD}_{\mathrm{avg}}$, PFH, risk reduction factor

Functional safety is the part of the overall safety of a system that depends on safety-related systems operating correctly in response to their inputs. A *safety function* is a function, implemented by a *safety-related system*, that is intended to achieve or maintain a safe state with respect to a specific hazardous event. A *dangerous failure* is one that impedes or disables the safety function; a *safe failure* is one that does not (a spurious trip is safe in this sense, although it is a failure of the equipment). A *dangerous undetected failure* is a dangerous failure that the system's diagnostics do not reveal, so that it remains latent until a test or a demand reveals it; its rate is λ_{DU}. A *demand* is an occurrence of the hazardous condition that requires the safety function to act. A function is in *low-demand mode* when demands occur no more than once per year, and in *high-demand or continuous mode* otherwise. A *proof test* is a periodic test that reveals dangerous undetected failures and restores the channel to its as-new condition; the time between proof tests is the *proof-test interval* T_{p}. The *safety integrity level* (SIL) is one of four discrete levels of required safety integrity. $\mathrm{PFD}_{\mathrm{avg}}$ is the average probability of dangerous failure on demand, the failure measure of low-demand functions; PFH is the average frequency of dangerous failure per hour, the failure measure of high-demand and continuous functions. The *risk reduction factor* of a low-demand function is $1/\mathrm{PFD}_{\mathrm{avg}}$.

33.3.1 Reliability and functional safety

A cooling pump can be highly reliable and still be part of an unsafe system. The pump may rarely fail, but if a dangerous overheating condition arises, from a blocked heat exchanger or a valve closed in error, the emergency shutdown function that should de-energize the equipment may fail to operate. The reliability of the pump and the integrity of the shutdown function are different properties of different items, and a high value of the first says nothing about the second.

IEC 61508 addresses the electrical, electronic and programmable electronic safety-related systems used to implement safety functions: the sensors, the logic solvers and the final elements, together with their software. It provides a lifecycle approach to functional safety and a framework for specifying the safety integrity that each function requires and for demonstrating that the implemented system achieves it. The edition referred to in this chapter is IEC 61508-1:2010, the second edition of its first part, which states the general requirements.

33.3.2 Specification of a safety function

A safety function is specified by five elements:

1. the hazardous condition to detect;
2. the required protective action;
3. the time available to complete that action;
4. the conditions under which the function must operate;
5. the required reliability, expressed as a safety integrity.

For the cooling-water system an example is:

> When cooling-water flow falls below the safe minimum while critical equipment is energized, the safety system shall initiate the specified protective shutdown within the required response time.

This is much more precise than the statement "the cooling system should be safe." It identifies a physical condition (flow below a stated minimum, with the equipment energized), a required

response (the specified shutdown) and a timing constraint (the response time, which must be shorter than the time the equipment takes to reach a damaging temperature after the flow is lost, the time to the thermal limit t_{crit} of Chapter 34). The fifth element, the required integrity, follows from the hazard and risk analysis and is stated as a SIL with its failure measure. A specification that omits any of the five cannot be verified, because there is nothing against which to test the implementation.

33.3.3 Safety integrity levels and their failure measures

IEC 61508 uses four safety integrity levels, SIL 1 to SIL 4. A higher SIL corresponds to more demanding requirements for safety integrity, which include a quantitative target failure measure and requirements on systematic capability, the measures that guard against design, software and procedural errors, which no failure rate represents.

For low-demand safety functions the target failure measure is $\mathrm{PFD}_{\mathrm{avg}}$, the average probability of dangerous failure on demand: the probability, averaged over time, that the function is unable to act when a demand arrives. For high-demand and continuous functions it is PFH, the average frequency of dangerous failure per hour, in h^{-1}. A low-demand function fails dangerously only if a latent failure coincides with a demand, so a probability is the relevant measure; a continuous function is exercised all the time, so every dangerous failure leads toward the hazard and a frequency is the relevant measure. Table 33.4 gives the ranges of the two measures that IEC 61508-1 assigns to each SIL. Each row is one level; the columns give the range of $\mathrm{PFD}_{\mathrm{avg}}$, the corresponding risk reduction factor and the range of PFH. Each level spans one decade.

Table 33.4: Target failure measures of the four safety integrity levels for a safety function, as IEC 61508-1 assigns them.

Level	$\mathrm{PFD}_{\mathrm{avg}}$ (low demand)	Risk reduction factor	PFH (high demand, continuous)
SIL 4	10^{-5} to below 10^{-4}	above 10^{4} to 10^{5}	10^{-9} to below $10^{-8}\,\mathrm{h}^{-1}$
SIL 3	10^{-4} to below 10^{-3}	above 10^{3} to 10^{4}	10^{-8} to below $10^{-7}\,\mathrm{h}^{-1}$
SIL 2	10^{-3} to below 10^{-2}	above 10^{2} to 10^{3}	10^{-7} to below $10^{-6}\,\mathrm{h}^{-1}$
SIL 1	10^{-2} to below 10^{-1}	above 10 to 10^{2}	10^{-6} to below $10^{-5}\,\mathrm{h}^{-1}$

These measures are not interchangeable with ordinary equipment availability or with the probability that any component fails. The same pump, or the same flow transmitter, may have several benign failure modes and one dangerous failure mode relevant to a particular safety function: a transmitter that fails to a high reading hides a loss of flow, whereas one that fails to a low reading causes a spurious trip. Only the first enters λ_{DU} of the low-flow shutdown, and the failure mode and effects analysis of Chapter 4 of Volume I is the tool that separates them. The functional safety framework distinguishes the target failure measures from the risk reduction that the hazard analysis requires of the function; the second determines which SIL is required, and the first is what the implemented function must achieve.

33.3.4 Average probability of failure on demand of one channel

Consider a protective function that is demanded only occasionally, implemented by a single channel (one sensor, one logic solver, one final element treated together). Assume that its dangerous undetected failures occur at a constant rate λ_{DU}, that a proof test is performed every T_{p} hours, that each proof test reveals every dangerous undetected failure and restores the channel to the as-new condition, that the test and any repair take negligible time, and that demands arrive at times unrelated to the test schedule.

Let $\tau \in [0, T_{\mathrm{p}})$ be the time since the last proof test. The channel is unable to act at τ if a dangerous undetected failure has occurred since the test, so its unavailability for the safety function is the cumulative failure probability of the exponential model,

$$\mathrm{PFD}(\tau) = 1 - \mathrm{e}^{-\lambda_{\mathrm{DU}}\tau}, \qquad 0 \le \tau < T_{\mathrm{p}}, \tag{33.1}$$

which returns to zero at every proof test. A demand arriving at a random time finds the channel at an age τ uniformly distributed over the interval, so the probability that it finds the channel failed is the time average of equation (33.1),

$$\mathrm{PFD}_{\mathrm{avg}} = \frac{1}{T_{\mathrm{p}}} \int_0^{T_{\mathrm{p}}} \left(1 - \mathrm{e}^{-\lambda_{\mathrm{DU}}\tau}\right) \mathrm{d}\tau = 1 - \frac{1 - \mathrm{e}^{-\lambda_{\mathrm{DU}} T_{\mathrm{p}}}}{\lambda_{\mathrm{DU}} T_{\mathrm{p}}}. \tag{33.2}$$

When $\lambda_{\mathrm{DU}} T_{\mathrm{p}}$ is small the unavailability is nearly linear within the interval, $\mathrm{PFD}(\tau) \approx \lambda_{\mathrm{DU}}\tau$, and the average of a linear function that rises from 0 to $\lambda_{\mathrm{DU}} T_{\mathrm{p}}$ is half its final value:

$$\mathrm{PFD}_{\mathrm{avg}} \approx \frac{1}{T_{\mathrm{p}}} \int_0^{T_{\mathrm{p}}} \lambda_{\mathrm{DU}}\tau \, \mathrm{d}\tau = \frac{\lambda_{\mathrm{DU}} T_{\mathrm{p}}}{2}. \tag{33.3}$$

The error of the approximation follows from the series $\mathrm{e}^{-x} = 1 - x + x^2/2 - x^3/6 + \dots$, which gives $(1 - \mathrm{e}^{-x})/x = 1 - x/2 + x^2/6 - \dots$ and therefore $\mathrm{PFD}_{\mathrm{avg}} = x/2 - x^2/6 + \dots$ with $x = \lambda_{\mathrm{DU}} T_{\mathrm{p}}$. The approximation overstates the exact value by a relative amount of about $x/3$, so it is conservative.

Remark 33.1. In this chapter T_{p} is the proof-test interval of a safety function. It is not the preventive replacement age T_p of the age-replacement policy of Chapter 18 of Volume I, although both are fixed maintenance intervals.

As an illustrative example, let $\lambda_{\mathrm{DU}} = 2 \times 10^{-6}\,\mathrm{h}^{-1}$ and $T_{\mathrm{p}} = 4000\,\mathrm{h}$. Then $x = 2 \times 10^{-6} \times 4000 = 0.008$ and

$$\mathrm{PFD}_{\mathrm{avg}} \approx \frac{(2 \times 10^{-6}\,\mathrm{h}^{-1})(4000\,\mathrm{h})}{2} = 0.004,$$

or 0.4 %, a risk reduction factor of $1/0.004 = 250$. The exact value of equation (33.2) is $1 - (1 - \mathrm{e}^{-0.008})/0.008 = 1 - 0.007968085/0.008 = 0.0039894$, which the approximation overstates by 0.27 %, close to $x/3 = 0.27\,\%$.

Figure 33.1 shows the calculation. The horizontal axis is time over four proof-test intervals; the vertical axis is the instantaneous probability that the channel is in a dangerous failed state, in units of 10^{-3}. The solid blue sawtooth is equation (33.1): it rises almost linearly to 0.00797 at the end of each interval and drops to zero at each proof test, marked by the dotted verticals. The dashed red line is its time average, the $\mathrm{PFD}_{\mathrm{avg}}$ of equation (33.2). The average sits at half the peak because the sawtooth is nearly a triangle; a demand arriving just before a proof test meets a channel twice as likely to be failed as the average suggests.

Table 33.5 evaluates the model for the same failure rate and six proof-test intervals, from three weeks to three years. The columns are the interval, the product $x = \lambda_{\mathrm{DU}} T_{\mathrm{p}}$, the approximation of equation (33.3), the exact value of equation (33.2), the relative amount by which the approximation overstates it, and the SIL band of Table 33.4 in which the exact value falls. The row for $T_{\mathrm{p}} = 8760\,\mathrm{h}$ shows the arithmetic: $x = 0.01752$, the approximation is 0.00876, the exact value is $1 - (1 - \mathrm{e}^{-0.01752})/0.01752 = 0.0087091$, and the overstatement is $0.00876/0.0087091 - 1 = 0.58\,\%$. Doubling the interval doubles $\mathrm{PFD}_{\mathrm{avg}}$, and the approximation degrades in proportion to x, from 0.03 % to 1.8 % over the table. A proof test every 500 h places the target measure in the SIL 3 band, intervals from 2000 to 8760 h in the SIL 2 band, and intervals of two and three years in the SIL 1 band.

This is a model-based calculation, not proof that the safety function meets a particular SIL. Falling in a band of Table 33.4 is necessary for the quantitative part of the requirement

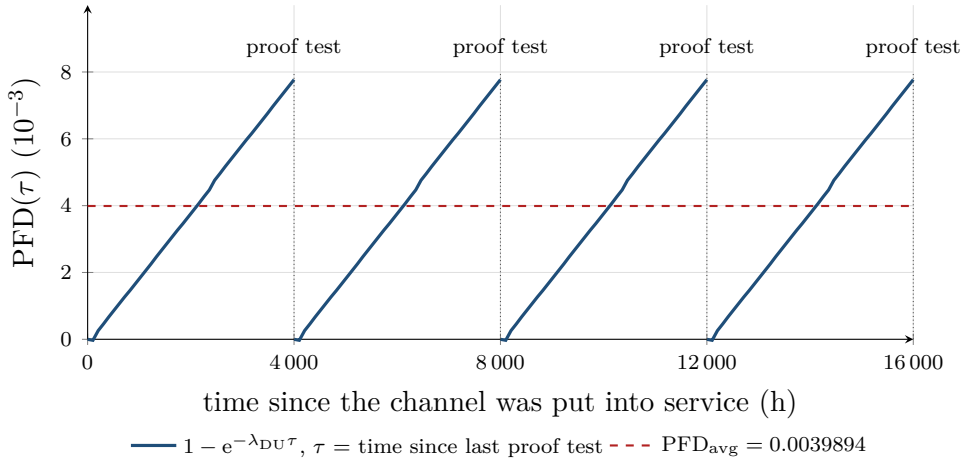


Figure 33.1: Instantaneous probability that a single periodically tested channel is in a dangerous undetected failed state, $\mathrm{PFD}(\tau) = 1 - \mathrm{e}^{-\lambda_{\mathrm{DU}}\tau}$ (equation (33.1)), with $\lambda_{\mathrm{DU}} = 2 \times 10^{-6}\,\mathrm{h}^{-1}$ and a perfect proof test every $T_{\mathrm{p}} = 4000\,\mathrm{h}$ (illustrative values). The dashed line is the time average, 0.0039894 exactly and $\lambda_{\mathrm{DU}} T_{\mathrm{p}}/2 = 0.004$ in the linear approximation.

Table 33.5: Average probability of failure on demand of one periodically tested channel, $\lambda_{\mathrm{DU}} = 2 \times 10^{-6}\,\mathrm{h}^{-1}$ (illustrative), against the proof-test interval: linear approximation, exact value and the SIL band of the target measure.

T_{p} (h)	$x = \lambda_{\mathrm{DU}} T_{\mathrm{p}}$	$x/2$	Exact, eq. (33.2)	Overstatement	SIL band
500	0.001	0.000500	0.00049983	0.03 %	SIL 3
2000	0.004	0.002000	0.0019973	0.13 %	SIL 2
4000	0.008	0.004000	0.0039894	0.27 %	SIL 2
8760	0.01752	0.008760	0.0087091	0.58 %	SIL 2
17 520	0.03504	0.017520	0.017317	1.17 %	SIL 1
26 280	0.05256	0.026280	0.025826	1.76 %	SIL 1

and is not sufficient for the requirement as a whole. A complete assessment accounts for the architecture of the channels and its constraints, the diagnostic coverage, the proof-test coverage (a test that reveals only part of the dangerous failures leaves the rest to accumulate over the life of the equipment), common-cause failures between redundant channels (Chapter 14 of Volume I), systematic failures, the repair time after a revealed failure, the configuration actually installed and the detailed requirements of the standard. If the same channel were used in continuous mode, every dangerous failure would lead toward the hazard, and to first order its PFH would equal its dangerous failure rate, $2 \times 10^{-6}\,\mathrm{h}^{-1}$, a value in the SIL 1 band: the same hardware supports different integrity levels in different modes of operation.

33.3.5 The safety lifecycle

Functional safety is organized as a lifecycle. A simplified version has six phases:

1. *hazard and risk analysis*: identify the hazardous events and the risk reduction each requires;
2. *safety function specification*: define the required actions, response times and integrity targets;
3. *design and implementation*: develop the sensors, logic, actuators, software and interfaces;
4. *verification and validation*: demonstrate conformity with the specification and with the safety needs;
5. *operation and maintenance*: preserve the safety integrity through proof testing, repair and monitoring;
6. *modification and reassessment*: reevaluate the function when its design or its operating conditions change.

Figure 33.2 draws the six phases. The boxes are the phases, read along the upper row from left to right and along the lower row from right to left; each box names the phase and its main output. The dashed arrow from the last phase to the first shows that a modification, or a change in the operating conditions, sends the function back to the hazard and risk analysis. The actual IEC 61508 lifecycle is more detailed, with separate phases for the overall safety requirements, their allocation to safety-related systems, installation and commissioning, and decommissioning, but the figure shows its organizing principle: safety assurance does not end with the initial design, and the proof tests of Section 33.3.4 are part of the lifecycle, not an addition to it.

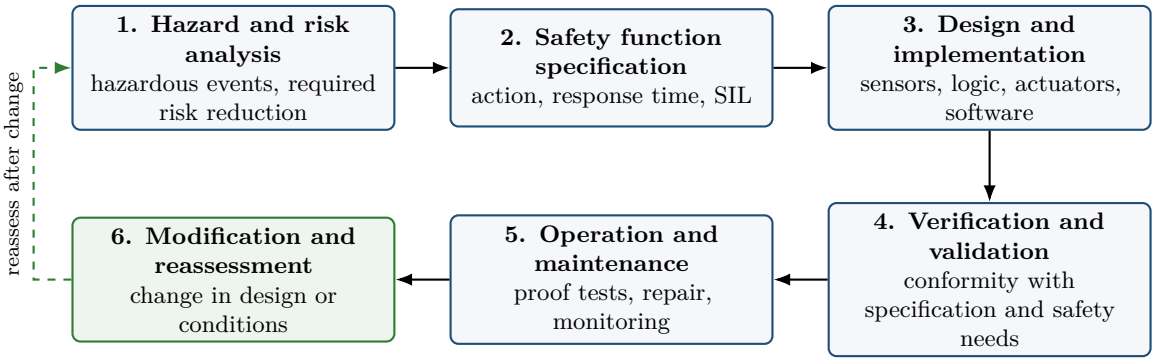


Figure 33.2: A simplified functional-safety lifecycle. The six phases run from the hazard analysis to operation; a modification or a change of operating conditions returns the safety function to the hazard and risk analysis (dashed arrow). The lifecycle of IEC 61508 itself has more phases.

33.3.6 Asset management and the ISO 55000 family

Reliability engineering concentrates on probabilities and failure mechanisms. Asset management asks a broader question: how an organization should manage its assets over their lifecycles to

achieve its objectives, while balancing performance, risk and expenditure. For a utility the assets include transmission lines, substations, transformers, sensors and the supporting infrastructure; for a data center they include UPS systems, chillers, pumps, switchgear, generators and control systems. The reliability of each asset is one input to that balance, together with its cost, its remaining life and the organization's objectives for the service it supports.

The editions that apply to this discussion are:

- ISO 55000:2024, *Vocabulary, overview and principles*;
- ISO 55001:2024, *Asset-management-system requirements*;
- ISO 55002, *Guidance for applying ISO 55001.*

ISO 55000 provides the underlying vocabulary and principles, while ISO 55001 specifies the requirements of a management system, the policies, objectives, plans, processes and reviews through which an organization manages its assets. An organization can be audited against ISO 55001; ISO 55000 and ISO 55002 explain and guide. None of the three prescribes a reliability model, a failure rate or a maintenance interval: they require that such decisions be made through a documented system aligned with the organization's objectives.

33.3.7 A maintenance-strategy decision

Consider two maintenance strategies for a pump. Under strategy A the pump is replaced on a fixed calendar schedule, the age-replacement policy of Chapter 18 of Volume I. Under strategy B its condition is monitored and it is replaced when the predicted failure risk exceeds an approved threshold, the condition-based policy of Chapter 24. The two can be compared by their expected long-run cost rate,

$$C(a) = C_{\text{maint}}(a) + \mathbb{E}\big[C_{\text{down}} \mid a\big] + C_{\text{mon}}(a), \qquad a \in \{\text{A}, \text{B}\}, \tag{33.4}$$

in R$/yr, where $C_{\text{maint}}(a)$ is the cost of the preventive and corrective replacements under strategy a, $\mathbb{E}[C_{\text{down}} \mid a]$ the expected cost of the downtime it leaves, and $C_{\text{mon}}(a)$ the cost of the monitoring it requires (zero for strategy A). The comparison is made subject to acceptable reliability and safety constraints: a strategy with the lower cost rate is admissible only if it also meets the reliability requirement of the service and leaves every safety function intact. Equation (33.4) generalizes the cost rate $C(T_p)$ of age replacement to any policy a.

The comparison is consistent with an asset-management perspective because it weighs lifecycle value, the cost of maintenance, monitoring and failures together, rather than availability alone. It is nevertheless a decision-analysis tool chosen by the analyst: ISO 55001 does not prescribe this or any other mathematical objective. The organization is responsible for determining its objectives, its decision criteria, the requirements that apply to it and the risks it accepts, and the management system is the means by which it records and reviews those choices.

33.4 Reliability requirements and a two-pump cooling assessment

Every assessment, under any standard, is judged against requirements. This section states what makes a reliability requirement useful and separates a target from the evidence that it is met, then carries out a complete assessment of a two-pump cooling system against such a requirement: it defines the system and two candidate configurations, takes their Markov models from Chapter 19 of Volume I, translates the results into the acceptance metric, evaluates an improvement and its economics, and lists the assumptions on which the conclusion rests.

33.4.1 What makes a reliability requirement useful

Consider the statement:

> The cooling system shall have high availability.

It is not a sufficiently precise engineering requirement. It does not say what “high” means, what constitutes a failure of cooling, or over what period the availability is assessed, so no analysis can show that it is met or that it is not. A better requirement is:

> The cooling-water pumping subsystem shall provide at least 100 % of the specified design flow under the defined demand and environmental envelope, with expected functional unavailability not exceeding four hours per 8760-hour assessment year.

This requirement specifies five elements: the system function (cooling-water flow), the required operating envelope (the defined demand and environmental conditions), the reliability metric (expected functional unavailability, in hours), the acceptance threshold (four hours) and the assessment period (a year of 8760 h). It uses the vocabulary of Chapter 2 of Volume I: a required function stated as a measurable performance criterion, with failure defined as the loss of that function. A complete requirement also states which events are included (for example, whether a flow reduction below 100 % that does not stop the flow counts) and how planned maintenance is treated (whether hours of scheduled outage count toward the four hours, or are excluded and governed separately).

Types of requirements. Table 33.6 lists seven types of requirement that a reliability assessment may have to address, each with an example for the cooling-water system. The rows address different properties, and each is checked by different evidence: a functional requirement by a performance test, a reliability or availability requirement by a model and its inputs, a maintainability requirement by repair records and drills, a monitoring requirement by inspection of the monitoring program, and the model-quality and assurance requirements by review of the analysis itself.

Table 33.6: Types of requirement for the cooling-water pumping subsystem, with an example of each (illustrative values).

Requirement type	Example
Functional	Deliver at least 500 $\mathrm{m^3/h}$ at the required differential pressure
Reliability	The probability of losing the required pumping capacity during a defined mission shall not exceed a stated threshold
Availability	Expected annual functional downtime shall be no more than four hours
Maintainability	A specified repair shall be completed within eight hours with 90 % probability
Monitoring	Bearing condition shall be measured and checked at an approved interval
Model quality	Predictions shall satisfy documented calibration and adequacy criteria
Assurance	Every acceptance claim shall reference approved evidence and model versions

Because the requirements address different properties, a subsystem may satisfy one and fail another. It may satisfy the availability requirement and fail the maintainability requirement: with exponential repairs of mean 20 h, the probability of completing a repair within 8 h is

$1-\mathrm{e}^{-0.05\times 8}=0.33$, far from 0.90, although the redundant pair may still meet a four-hour annual downtime target. Likewise it may have a favorable expected downtime and fail a particular single-contingency criterion, as a network with a small expected unserved energy can still violate a rating after a prescribed outage (Section 33.2.3).

33.4.2 Targets and evidence

Suppose the accepted requirement is

$$\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}, \tag{33.5}$$

where D_{yr} is the functional downtime of the subsystem in an assessment year, in hours. The target does not determine whether the system complies. To decide that, the assessment still needs:

1. a clearly defined failure event;
2. a suitable reliability model;
3. justified inputs;
4. verified calculations;
5. validation evidence;
6. an uncertainty assessment;
7. an authorized acceptance decision.

The rest of the chapter supplies each item for the two-pump example. A numerical threshold of this kind may be contractual (a service-level agreement with a customer), internal (a design standard of the organization) or regulatory. The four-hour requirement of this chapter is hypothetical; it is not attributed to NERC, IEC or ISO.

Remark 33.2. The downtime D_{yr} and the configuration downtimes D_{A} and D_{B} below are times in hours per year, as in Section 19.5.4. They are not the degradation state $D(t)$ of Chapter 21 nor the FMEA detection rating D.

33.4.3 The system, two configurations and the Markov model

A data center cooling plant has two identical active pumps. Each pump can supply the required cooling flow by itself, so the subsystem is functionally available when at least one pump operates. Each operating pump fails at the constant rate

$$\lambda = 0.001\,\mathrm{h}^{-1},$$

and each repair process in progress is completed at the constant rate

$$\mu = 0.05\,\mathrm{h}^{-1},$$

so that $\mathrm{MTTF} = 1/\lambda = 1000\,\mathrm{h}$ and $\mathrm{MTTR} = 1/\mu = 20\,\mathrm{h}$. These values are illustrative; they are chosen so that the effects of redundancy and of repair resources are both measurable. Two configurations are considered:

- *Configuration A*: two active pumps and one shared repair crew;
- *Configuration B*: two active pumps and independent repair capacity for both pumps.

The pumps and their rates are the same in both; only the repair resources differ.

Let the state $i \in \{2, 1, 0\}$ be the number of working pumps. The model is the three-state continuous-time Markov chain of Figure 19.2, Chapter 19 of Volume I. In configuration A the transitions are

$$2 \xrightarrow{2\lambda} 1, \qquad 1 \xrightarrow{\mu} 2, \qquad 1 \xrightarrow{\lambda} 0, \qquad 0 \xrightarrow{\mu} 1,$$

with the generator $\mathbf{Q}_{\text{one crew}}$ of equation (19.40); in configuration B the last transition has rate 2μ, because two repairs proceed at once, and the generator is that of equation (19.26). With

$\rho = \lambda/\mu = 0.02$, the stationary probabilities of the state with no working pump, derived in Section 19.5.3 and Section 19.5.1, are

$$\pi_0^{\mathrm{A}} = \frac{2\rho^2}{1 + 2\rho + 2\rho^2} \qquad \text{and} \qquad \pi_0^{\mathrm{B}} = \frac{\rho^2}{(1+\rho)^2} = \left(\frac{\lambda}{\lambda+\mu}\right)^2, \tag{33.6}$$

equations (19.41) and (19.39). Numerically,

$$\pi_0^{\mathrm{A}} = \frac{2 \times 0.0004}{1 + 0.04 + 0.0008} = \frac{0.0008}{1.0408} = 0.00076864, \qquad \pi_0^{\mathrm{B}} = \left(\frac{0.001}{0.051}\right)^2 = 0.00038447,$$

and the steady-state functional availabilities are $A_{\mathrm{A}} = 1 - \pi_0^{\mathrm{A}} = 0.99923136$, or 99.9231 %, and $A_{\mathrm{B}} = 1 - \pi_0^{\mathrm{B}} = 0.99961553$, or 99.9616 %. Chapter 19 also gives the failure frequency $f_{\mathrm{F}} = \pi_1 \lambda$ and the mean down time $\mathrm{MDT} = 1/q_{01}$ of each configuration (Section 19.5.5), which the assessment will need for the distribution of annual downtime.

33.4.4 The acceptance metric

Under stationary operation the expected functional downtime in an assessment year is the steady-state unavailability times the hours in the year, equation (19.42):

$$\mathbb{E}[D_{\mathrm{A}}] = 8760\,\mathrm{h} \times \pi_0^{\mathrm{A}}, \qquad \mathbb{E}[D_{\mathrm{B}}] = 8760\,\mathrm{h} \times \pi_0^{\mathrm{B}}. \tag{33.7}$$

For the baseline, configuration A,

$$\mathbb{E}[D_{\mathrm{A}}] = 8760 \times 0.00076864 = 6.73\,\mathrm{h/yr}.$$

Against the requirement $\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}$ of equation (33.5), the baseline configuration does not satisfy the target; it exceeds it by 2.73 h/yr, or 68 %. This is the first engineering finding of the assessment. An assessment that stopped at this statement would leave the organization with a failure and no course of action, so the next step evaluates feasible corrective actions.

33.4.5 An improvement: independent repair capacity

If both failed pumps can be repaired simultaneously by independent crews, the transition out of state 0 becomes

$$0 \xrightarrow{2\mu} 1,$$

and the model is configuration B. The probability of no working pump halves, to $\pi_0^{\mathrm{B}} = 0.00038447$, the availability rises to 99.9616 %, and

$$\mathbb{E}[D_{\mathrm{B}}] = 8760 \times 0.00038447 = 3.37\,\mathrm{h/yr}.$$

The improved configuration meets the numerical requirement under the assumed Markov model. The pumps fail no less often in configuration B; the failure frequency is 0.3368 outages per year in B against 0.3367 in A. What changes is the duration of each complete outage, a mean of 10 h instead of 20 h, because two repairs progress at once (Section 19.5.3). Table 33.7 collects the results: each row is one quantity of the model and the two columns are the configurations.

Figure 33.3 shows the comparison against the criterion. The horizontal axis lists the two configurations; the vertical axis is the expected annual functional downtime in hours per year. Each bar is the value of equation (33.7) for its configuration, with the value printed above it, and the dashed red line is the acceptance criterion of four hours per year. The bar of configuration A crosses the line and that of configuration B stays below it, by a margin of 0.63 h/yr that Section 33.6.3 examines.

Table 33.7: Stationary results of the two configurations of the cooling pumps, $\lambda = 0.001\,\mathrm{h}^{-1}$ and $\mu = 0.05\,\mathrm{h}^{-1}$ (illustrative), against the hypothetical requirement $\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}$.

Quantity	A: one shared crew	B: independent crews
Repair rate out of state 0, q_{01}	$\mu = 0.05\,\mathrm{h}^{-1}$	$2\mu = 0.10\,\mathrm{h}^{-1}$
Both pumps working, π_2	0.960799	0.961169
One pump working, π_1	0.038432	0.038447
No pump working, π_0	0.00076864	0.00038447
Functional availability	99.9231 %	99.9616 %
Failure frequency $f_{\mathrm{F}} = 8760\,\pi_1\lambda$	$0.3367\,\mathrm{yr}^{-1}$	$0.3368\,\mathrm{yr}^{-1}$
Mean down time MDT $= 1/q_{01}$	20 h	10 h
Expected annual downtime $8760\,\pi_0$	6.73 h/yr	3.37 h/yr
Meets $\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}$	no	yes

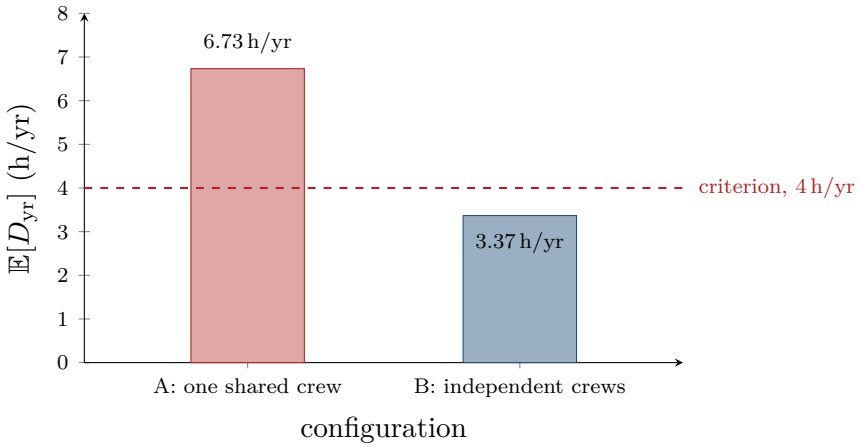


Figure 33.3: Expected annual functional downtime $\mathbb{E}[D_{\mathrm{yr}}] = 8760\,\pi_0$ of the two pump configurations against the hypothetical acceptance criterion of 4 h/yr (dashed), with $\lambda = 0.001\,\mathrm{h}^{-1}$ and $\mu = 0.05\,\mathrm{h}^{-1}$ (illustrative). One shared crew: $\pi_0 = 2\rho^2/(1 + 2\rho + 2\rho^2)$; independent crews: $\pi_0 = [\lambda/(\lambda+\mu)]^2$; $\rho = \lambda/\mu = 0.02$.

33.4.6 Economics of the improvement

Suppose the modeled consequence of lost pumping function is a cost of downtime

$$c_D = \text{R\$ } 25\,000 \text{ per hour}$$

(illustrative). The expected annual cost of downtime of each configuration is its risk in the sense of Chapter 1 of Volume I,

$$\mathcal{R}(a) = c_D \, \mathbb{E}[D_a], \qquad a \in \{\text{A}, \text{B}\}, \tag{33.8}$$

which gives $\mathcal{R}(\text{A}) = 25\,000 \times 6.7333 = \text{R\$ } 168\,332$ per year and $\mathcal{R}(\text{B}) = 25\,000 \times 3.3679 = \text{R\$ } 84\,198$ per year. The expected reduction in downtime is

$$\Delta D = \mathbb{E}[D_{\text{A}}] - \mathbb{E}[D_{\text{B}}] = 6.7333 - 3.3679 = 3.365 \,\text{h/yr},$$

and the expected avoided interruption cost, the risk reduction of the improvement, is

$$\Delta\mathcal{R} = c_D \, \Delta D = 25\,000 \times 3.365 = \text{R\$ } 84\,134 \text{ per year},$$

about R\$ 84 000 per year. Suppose the additional repair capacity, the second crew on call with its training and tools, has an annualized cost

$$c_{\text{crew}} = \text{R\$ } 45\,000 \text{ per year}.$$

The modeled net expected annual benefit is then

$$\Delta\mathcal{R} - c_{\text{crew}} = 84\,134 - 45\,000 = \text{R\$ } 39\,134 \text{ per year},$$

about R\$ 39 000. Equivalently, the total expected annual cost, downtime plus crew, is R\$ 168 332 for configuration A and $84\,198 + 45\,000 = \text{R\$ } 129\,198$ for configuration B. The decision is not close on these numbers: the second crew pays for itself as long as an hour of lost pumping costs more than $45\,000/3.365 = \text{R\$ } 13\,372$.

Under these assumptions the additional repair capability both satisfies the proposed availability target and has a positive expected economic value. That is not yet a final engineering approval. The result is a model output under nominal parameters, and it still requires the assurance evidence of Sections 33.5 and 33.6.

33.4.7 Critical assumptions

The numerical conclusion relies on several assumptions. Table 33.8 lists them. Each row is an assumption; the second column says why it matters, and the third names the evidence that would test it, which Section 33.5.3 develops into a validation checklist.

Some assumptions can be tested directly: a single-pump flow test settles the first, a failover test the seventh. Others require additional field evidence or a sensitivity analysis: independence, constant rates and the form of the repair distribution can be checked against operating records only as those records accumulate, and until then their effect is bounded by computing the result under alternative assumptions. A credible report identifies which of its conclusions are strongly supported and which depend on uncertain assumptions. Several of the assumptions have quantified effects in Chapter 19: a common-cause fraction of $\beta = 0.1$ multiplies the steady-state unavailability of the pair by 3.45 (Section 19.6.2), which would raise $\mathbb{E}[D_{\text{B}}]$ from 3.37 to about 11.6 h/yr, and a fivefold load-sharing increase of the survivor's failure rate multiplies it by 4.99 (Section 19.6.1).

Table 33.8: Critical assumptions of the two-pump assessment, why each matters, and the evidence that tests it.

Assumption	Why it matters	Evidence that tests it
Either pump can carry the full required flow	Determines what counts as functional availability	Single-pump flow and pressure test over the envelope
Failures are independent	Shared electrical or control failures could defeat the redundancy	Review of shared supplies; failure records of both pumps
Failure rates are constant	Aging or operating stress may change the hazard	Trend test of failure times (Chapter 17 of Volume I)
Repairs are exponential	Actual logistics can have long-tailed restoration times	Distribution of recorded repair durations
Independent crews are always available	Resource sharing and staffing may reduce the effective repair capacity	Staffing roster; maintenance drill
No scheduled maintenance impact	Planned interventions temporarily remove the redundancy	Maintenance plan and outage records
Switching and controls work	A healthy standby path may be unusable if the control system fails	Failover test
Cooling demand is within the design envelope	A single pump may not suffice at extreme load	Worst-case demand analysis and test

33.5 Verification, validation and traceability

A model result enters an engineering decision only after it has been checked twice, once against its own specification and once against the physical system, and after its path from requirement to decision has been recorded. This section separates verification from validation, gives a checklist for each in the two-pump example, ties credibility to a context of use, and builds the traceability chain, the requirements traceability matrix and the configuration management that make the assessment auditable.

33.5.1 Verification and validation

The two terms are often used together, but they answer different questions. *Verification* asks whether the model has been implemented and solved correctly: whether the code computes what the equations say, and the equations say what the analyst intended. *Validation* asks whether the model adequately represents the real system for its intended use. The first compares the implementation with the model; the second compares the model with the physical system and the evidence about it.

A model can pass verification and fail validation. A perfectly implemented Markov model of the two pumps, whose generator rows sum to zero and whose stationary distribution agrees with the hand calculation to every digit, can still produce wrong outage-risk estimates if the real pumps exhibit correlated failures through a shared electrical supply, or if the repair delays are not represented adequately because spare parts take days to arrive. Verification establishes that the answer is the right answer to the stated model; only validation establishes that the stated model is the right one. The distinction is the one between model adequacy and parameter precision in Chapter 31 (Section 31.4.6): a precise answer is not thereby a credible one.

33.5.2 Verification of the two-pump model

For the two-pump model, verification includes the following checks:

- every row of the generator matrix sums to zero;
- the off-diagonal transition rates are nonnegative;
- the state probabilities remain between zero and one;
- the state probabilities sum to one;
- the analytical steady-state solutions agree with numerical solutions;
- both repair-resource configurations reduce correctly to their limiting cases;
- the units are consistent;
- the model reproduces hand-calculated reference examples.

In particular, the calculated stationary distribution must satisfy

$$\boldsymbol{\pi}\,\mathbf{Q} = \mathbf{0} \qquad \text{and} \qquad \sum_i \pi_i = 1, \tag{33.9}$$

equation (19.36) of Chapter 19 of Volume I, and the solver should be tested to respect these relations within justified numerical tolerances. For the example, a numerical solution of equation (33.9) by linear algebra reproduces $\pi_0^{\mathrm{A}} = 0.00076864$ and $\pi_0^{\mathrm{B}} = 0.00038447$ to every digit shown, with residuals $\boldsymbol{\pi}\mathbf{Q}$ below $10^{-18}\,\mathrm{h}^{-1}$, the round-off level of double-precision arithmetic. The limiting cases are equally specific. With $\lambda = 0$ both configurations must stay in state 2, $\boldsymbol{\pi} = [1,\ 0,\ 0]$. Configuration B must reproduce the binomial probabilities of two independent pumps, each down with probability $\lambda/(\lambda{+}\mu)$. As $\rho \to 0$ the ratio $\pi_0^{\mathrm{A}}/\pi_0^{\mathrm{B}} = 2(1{+}\rho)^2/(1{+}2\rho{+}2\rho^2)$ must tend to 2, the ratio of the mean down times; for $\rho = 0.02$ it is 1.9992. Unit consistency means that both rates are in h^{-1} and the year in hours; a failure rate taken from a database in failures per 10^6 h and a repair time in days is the common source of error. A second, independent calculation of the same quantity, such as $1 - A = f_{\mathrm{F}} \cdot \mathrm{MDT}$ (equation (19.45)), checks the first: $0.33679\,\mathrm{yr}^{-1} \times 10\,\mathrm{h}/8760\,\mathrm{h} = 0.00038447$.

33.5.3 Validation against the real plant

Validation requires evidence about the real cooling plant. The possible tests include:

- verifying that either pump alone meets the required flow and pressure;
- comparing the observed failure frequencies with the modeled hazards;
- analyzing the repair histories and the actual logistics delays;
- testing switchover and the performance of the automatic controls;
- investigating shared power, control and environmental dependencies;
- comparing the predicted and the observed frequency and duration of functional outages;
- assessing whether failures occur under conditions missing from the model.

A particularly important validation question is whether, when one pump fails, the other actually preserves the required cooling function over the full operating envelope. A redundant equipment count is not the same as verified functional redundancy: two pumps installed in parallel are redundant in the sense of the model only if the survivor can carry the full flow at the worst-case demand, the check valves prevent backflow through the failed pump, and the controls start or load the survivor in time. Each of these is a property of the physical system that the Markov model assumes and does not test.

The comparison of predicted and observed outages has a statistical limit that the validation report must state. The model predicts 0.34 complete outages per year; a plant observed for three years is expected to show about one, and observing none or two is entirely consistent with the model. The two pumps, by contrast, accumulate $2 \times 8760 = 17\,520$ pump-hours and about 17.5 pump failures a year at $\lambda = 0.001\,\mathrm{h}^{-1}$. Operating experience at this level of redundancy therefore validates the component rates far better than it validates the system outage frequency, which is the quantity of interest, and the system frequency is validated indirectly, through the component rates and the structure of the model.

33.5.4 Credibility for a context of use

Suppose the model has been validated for routine operation. That does not establish its credibility for an extreme-heat event. At elevated cooling demand one pump may no longer deliver the required flow, so the system is no longer a 1-out-of-2 system but a 2-out-of-2 one; the failure intensity of the pumps may increase with load and temperature; repair access may become more difficult; and other cooling equipment, the chillers and the cooling towers, may approach their limits at the same time, introducing a dependence the model does not contain. The assumptions of the original model may no longer hold.

Model credibility should therefore be tied to an explicitly defined context of use: the question the model answers, the operating envelope over which it has been validated, and the consequence of an error in the decision it supports. A model adequate for annual maintenance budgeting may be inadequate for a safety-critical emergency-response decision, because the second operates at the edge of the envelope and the cost of an error is larger. This is the workflow of Figure 31.7 applied to the assessment as a whole: stage 1 defines the question, stage 7 tests adequacy for it, and a change in the question sends the model back through the stages.

> **Principle 33.2.** *Validation establishes credibility for a stated context of use and operating envelope. A model validated for one decision is not thereby credible for another, and the context of use is part of every acceptance claim.*

33.5.5 The traceability chain

A defensible engineering assessment connects requirements to models, assumptions, tests, results and decisions. For the two-pump example the chain has eight links:

1. *requirement*: expected functional downtime of at most 4 h/yr;

2. *system definition*: two pumps, either sufficient to meet the specified cooling load;
3. *model*: a repairable-system Markov model with explicit repair-resource constraints;
4. *input evidence*: failure histories, repair records, capacity tests and switching tests;
5. *verification evidence*: state-balance checks and independent analytical reference calculations;
6. *validation evidence*: comparison with observed operation and failure behavior;
7. *acceptance assessment*: comparison of the uncertainty-qualified model results with the approved threshold;
8. *engineering decision*: approve the repair-capacity improvement, require further evidence, or reject the proposal.

Figure 33.4 draws the chain. Each box is a link, read from top to bottom, with its content for the example; the braces on the right group the links into the five stages of an assessment, requirements, models, evidence, assurance and decision; the dashed arrow on the left is the audit trail, along which a reviewer traces the decision back to the requirement. A useful assessment allows a reviewer to follow that arrow without reconstructing the project from memory or from the analyst's files.

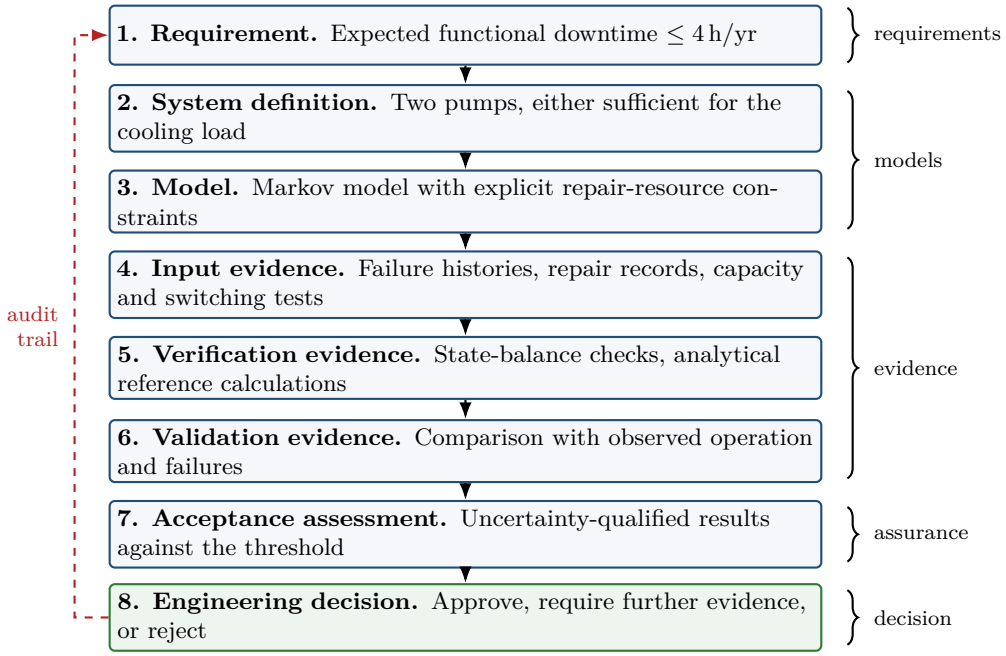


Figure 33.4: The traceability chain of the two-pump assessment, from the requirement to the engineering decision, grouped into the stages requirements, models, evidence, assurance and decision. The dashed arrow is the audit trail along which a reviewer traces the decision back to the requirement.

33.5.6 The requirements traceability matrix

The chain of one requirement generalizes to a table for all of them, the requirements traceability matrix. Table 33.9 is an example for the cooling assessment. Each row is one requirement with an identifier; the columns give the requirement, the evidence that addresses it and the method by which acceptance is judged. The methods differ by row: a quantitative evaluation for the availability requirement, a demonstration for the duty of one pump, an inspection and a drill for the repair capacity, an audit for the data, a test for the numerical model and a technical review for its validity.

The evidence column should ultimately reference controlled records, not descriptions of

Table 33.9: Example requirements traceability matrix for the cooling assessment (illustrative identifiers).

ID	Requirement	Evidence	Acceptance method
REL-01	Expected functional downtime $\leq 4\,\mathrm{h/yr}$	Markov availability calculation	Quantitative evaluation
FUNC-01	One pump meets the design duty	Commissioning performance test	Demonstration
MAINT-01	Independent repair capacity available for two simultaneous failures	Staffing and maintenance procedure	Inspection and drill
DATA-01	Failure and repair inputs are documented	Maintenance records	Data audit
MOD-01	Numerical and analytical models agree	Verification test results	Test
VAL-01	Model is credible for the specified operating envelope	Validation report	Technical review

them. The entry for VAL-01, for example, becomes a pointer to a specific document, such as "Validation Report VR-2026-014, revision 2" (an illustrative identifier), and that document identifies the dataset, the model version, the test procedure and the acceptance result. A matrix whose evidence column says only "validation report" records an intention, not evidence.

33.5.7 Configuration management

Suppose a revised model changes the pump failure rate from $\lambda = 0.001\,\mathrm{h}^{-1}$ to $\lambda = 0.0015\,\mathrm{h}^{-1}$, on the basis of new failure data. The change alters the availability calculation enough to reverse the acceptance decision, as Section 33.6.3 shows. Two reports, one produced before the change and one after, would then reach opposite conclusions about the same plant, and a reader who cannot tell which inputs each used cannot tell which to believe. Configuration management prevents this by recording, for every result:

- the model version;
- the input dataset version;
- the parameter-estimation method;
- the software version;
- the scenario definition;
- the assumptions;
- the numerical outputs;
- the reviewer and approval history.

A reliability report should be reproducible from its retained inputs and approved computational artifacts: another analyst, given the recorded versions, obtains the same numbers. Without configuration control, two analysts may report different answers while believing that they are evaluating the same system.

33.6 Risk registers, uncertainty and acceptance decisions

The last step is the acceptance decision itself. This section places the quantitative model beside the risk register, states acceptance criteria and the authority that applies them, shows how a narrow margin and an uncertain parameter change the decision, separates confidence in the

criterion from the predictive distribution of a year's downtime, and ends with independent review and an example of a qualified engineering conclusion.

33.6.1 Risk registers and quantitative models

A conventional risk register records, for each entry, the hazard or failure scenario, its causes, its consequences, its likelihood, the existing controls, the mitigation actions, the responsible owner, the residual risk and the review status. A quantitative model can support the likelihood and consequence estimates of an entry, but the register is not a substitute for it: an entry labeled "high risk" does not, by itself, establish a failure probability, because the label summarizes a judgment and not a calculation. Conversely, a calculated expected annual loss does not determine whether a risk is acceptable. Acceptance depends on approved organizational criteria, applicable law and standards, safety obligations and the authority responsible for the decision.

Table 33.10 is an example register for the cooling system. Each row is one risk with an identifier; the columns give the risk, the current mitigation, the further action proposed and the status. The proposed repair-capacity improvement addresses R-01, both pumps unavailable, and the assessment of Section 33.4 quantifies it. It does not necessarily resolve R-02 to R-05: a shared electrical supply, a failed switchover, a long repair and an extreme demand each defeat the redundancy by a mechanism outside the Markov model, the critical assumptions of Table 33.8 written as risks. A narrow numerical result about R-01 should therefore not be described as proving the overall reliability of the cooling system.

Table 33.10: Example risk register for the cooling-water system (illustrative).

ID	Risk	Current mitigation	Further action	Status
R-01	Both pumps unavailable	Two-pump redundancy	Add independent repair capability	Proposed
R-02	Shared electrical supply disables both pumps	Electrical protection	Assess power-supply diversity	Open
R-03	Switchover control fails	Automated controls	Conduct failover test	Open
R-04	Repair time exceeds the assumed MTTR	Spare inventory	Review the restoration-time distribution	Open
R-05	Extreme cooling demand exceeds one-pump capability	Capacity margin	Validate the worst-case operating envelope	Open

33.6.2 Risk acceptance criteria and authority

Organizations may apply several acceptance criteria at once. For the cooling system they might be

$$\begin{gathered}\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}, \qquad \mathbb{E}[C_{\mathrm{annual}}] \le B, \\ \Pr(\text{complete cooling loss during a critical mission}) \le p^{*},\end{gathered} \tag{33.10}$$

together with the mandatory safety and regulatory requirements, where p^{*} is an approved bound on the mission failure probability, C_{annual} the annual cost of downtime and mitigation and B the approved budget. The criteria are not interchangeable. An option that minimizes expected cost may violate the bound on the mission failure probability, because the expected cost averages rare long outages with many years without any. An option with a favorable average availability

may still carry an unacceptable common-cause risk, which appears in the tail of the consequence distribution rather than in its mean. The decision process should keep the criteria separate, record which ones each option meets, and refer the trade-offs to the acceptance authority rather than resolve them silently in a weighted sum.

The acceptance authority is part of the criteria. A calculation can show that a criterion is met; it cannot accept the residual risk, which is a decision on behalf of the organization and belongs to a named person or body with the authority to make it. A report that ends with "requirement passed" and no named authority leaves the decision unmade.

33.6.3 The narrow margin and a revised failure rate

The improved configuration gives $\mathbb{E}[D_{\mathrm{B}}] = 3.37\,\mathrm{h/yr}$ against the threshold of $4\,\mathrm{h/yr}$, a margin of

$$4 - 3.37 = 0.63\,\mathrm{h/yr},$$

16 % of the criterion. The margin is narrow, and the assessment must ask what happens if the actual failure rate is higher. Suppose new field data suggest $\lambda = 0.0015\,\mathrm{h^{-1}}$ instead of $0.001\,\mathrm{h^{-1}}$. With independent repairs, by equation (33.6),

$$\pi_0^{\mathrm{B}} = \left(\frac{\lambda}{\lambda+\mu}\right)^2 = \left(\frac{0.0015}{0.0515}\right)^2 = 0.029126^2 = 0.00084834,$$

and the expected annual downtime becomes

$$\mathbb{E}[D_{\mathrm{B}}] = 8760 \times 0.00084834 = 7.43\,\mathrm{h/yr}.$$

The configuration would fail the hypothetical four-hour target by 86 %. A 50 % increase in one parameter more than doubles the downtime, because the unavailability of a redundant pair grows roughly with the square of the failure rate, $(1.5)^2 = 2.25$, and the acceptance decision reverses.

The rate at which configuration B just meets the criterion follows from setting $8760\,[\lambda/(\lambda+\mu)]^2 = 4$:

$$\frac{\lambda^*}{\lambda^*+\mu} = \sqrt{\frac{4}{8760}} = 0.021369 \quad\Longrightarrow\quad \lambda^* = \mu\,\frac{0.021369}{1-0.021369} = 0.0010918\,\mathrm{h^{-1}}. \tag{33.11}$$

The nominal failure rate is only 9.2 % below the value at which the decision reverses. The same calculation for the repair rate, at $\lambda = 0.001\,\mathrm{h^{-1}}$, gives $\mu^* = 0.04580\,\mathrm{h^{-1}}$: a mean repair time of 21.8 h instead of 20 h is enough to fail the criterion.

Figure 33.5 shows the dependence on the failure rate. The horizontal axis is the pump failure rate in units of $10^{-3}\,\mathrm{h^{-1}}$; the vertical axis is the expected annual downtime. The red curve is configuration A and the blue curve configuration B, both with $\mu = 0.05\,\mathrm{h^{-1}}$; the dashed horizontal line is the criterion. The dots mark the nominal rate $0.001\,\mathrm{h^{-1}}$ and the revised rate $0.0015\,\mathrm{h^{-1}}$ on each curve, and the vertical dotted lines mark the rates at which each curve crosses the criterion: $0.000767\,\mathrm{h^{-1}}$ for A and $0.001092\,\mathrm{h^{-1}}$ for B. Both curves rise faster than linearly. The nominal point of B lies just left of its crossing, which is the narrow margin seen as a distance along the axis.

A complete acceptance assessment should therefore evaluate the conclusion over credible ranges or posterior distributions of the relevant inputs, rather than at the nominal values alone. Nominal compliance with a small margin may warrant further testing, a more conservative design or both.

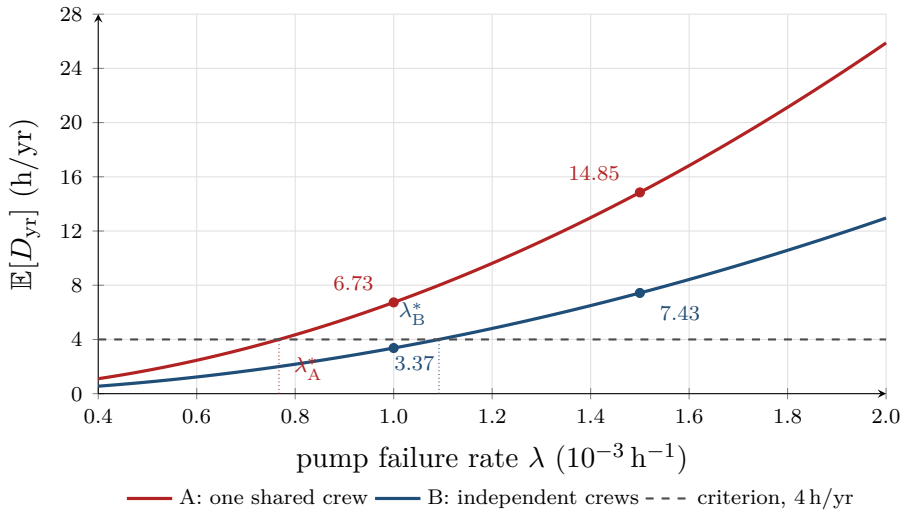


Figure 33.5: Expected annual functional downtime of the two configurations against the pump failure rate, $\mu = 0.05\,\mathrm{h}^{-1}$ (illustrative). A: $8760 \cdot 2\rho^2/(1+2\rho+2\rho^2)$; B: $8760\,[\lambda/(\lambda+\mu)]^2$; $\rho = \lambda/\mu$. Dots: nominal ($\lambda = 0.001\,\mathrm{h}^{-1}$) and revised ($\lambda = 0.0015\,\mathrm{h}^{-1}$) rates. Dotted verticals: the rates at which each configuration just meets the criterion, $\lambda^*_{\mathrm{A}} = 0.000767\,\mathrm{h}^{-1}$ and $\lambda^*_{\mathrm{B}} = 0.001092\,\mathrm{h}^{-1}$.

33.6.4 Posterior confidence in the criterion

Let $\boldsymbol{\theta} = (\lambda, \mu)$ be the parameters of the model and $\mathbf{y}$ the data from which they are estimated, the failure and repair records. The quantity the criterion constrains, $\mathbb{E}[D_{\mathrm{yr}} \mid \boldsymbol{\theta}]$, the long-run expected downtime given the parameters, is then uncertain, and the question becomes

$$\Pr\Big(\mathbb{E}[D_{\mathrm{yr}} \mid \boldsymbol{\theta}] \le 4\,\mathrm{h/yr} \;\Big|\; \mathbf{y}\Big), \tag{33.12}$$

the posterior probability, given the evidence, that the modeled long-run expected downtime satisfies the criterion. For configuration B with the repair rate taken as known, the expected downtime increases with λ, so the event in equation (33.12) is $\lambda \le \lambda^* = 0.0010918\,\mathrm{h}^{-1}$, and its probability is the posterior distribution function of λ at λ^*.

As an illustrative example, suppose the pump failure rate is estimated from r failures observed in a total exposure of T_{tot} pump-hours, with $r/T_{\mathrm{tot}} = 0.001\,\mathrm{h}^{-1}$, and a noninformative prior (the limit of a gamma prior whose shape and rate tend to zero), so that the posterior is the gamma distribution $\mathrm{Gamma}(r, T_{\mathrm{tot}})$ of Chapter 8 of Volume I, with shape r, rate T_{tot} and mean $r/T_{\mathrm{tot}} = 0.001\,\mathrm{h}^{-1}$. For $r = 10$ and $T_{\mathrm{tot}} = 10\,000$ pump-hours, about seven months of the two pumps' operation, the posterior probability that $\lambda \le 0.0010918\,\mathrm{h}^{-1}$ equals the probability that a Poisson variable of mean $0.0010918 \times 10\,000 = 10.918$ is at least 10, which is 0.650. The criterion is met at the nominal value, and the evidence gives only a 65 % posterior probability that it is met at all.

Table 33.11 evaluates the same calculation for five amounts of evidence with the same point estimate. The columns are the number of failures, the exposure, the 90 % upper credible bound of λ, the expected downtime of configuration B at that bound, the posterior mean of $\mathbb{E}[D_{\mathrm{B}} \mid \boldsymbol{\theta}]$, and the posterior probability, equation (33.12), that each configuration meets the criterion. More evidence raises the confidence that B meets the criterion, from 0.64 to 0.82, and lowers the confidence that A does, from 0.34 to 0.006: data sharpen the conclusion in whichever direction the nominal value points, and they sharpen it slowly when the nominal value lies close to the threshold. Even with 100 failures the 90 % upper bound of λ gives 4.28 h/yr for B. The posterior mean of the expected downtime exceeds its value at the posterior mean of λ, 3.37 h/yr, because the downtime is a convex function of λ; with five failures it reaches 3.98 h/yr, nearly the criterion.

Table 33.11: Posterior confidence that the expected annual downtime meets 4 h/yr, for gamma posteriors Gamma(r, T_{tot}) of the pump failure rate with the same mean $0.001\,\mathrm{h}^{-1}$; $\mu = 0.05\,\mathrm{h}^{-1}$ known (illustrative).

r	T_{tot} (h)	$\lambda_{0.90}$ (h^{-1})	$\mathbb{E}[D_{\mathrm{B}}]$ at $\lambda_{0.90}$ (h/yr)	Posterior mean of $\mathbb{E}[D_{\mathrm{B}} \mid \boldsymbol{\theta}]$ (h/yr)	Pr(B meets)	Pr(A meets)
5	5000	0.001599	8.41	3.98	0.636	0.339
10	10 000	0.001421	6.69	3.68	0.650	0.244
20	20 000	0.001295	5.58	3.52	0.682	0.145
40	40 000	0.001207	4.87	3.45	0.731	0.060
100	100 000	0.001130	4.28	3.40	0.822	0.006

33.6.5 The predictive distribution of annual downtime

The posterior confidence of equation (33.12) concerns the long-run expected downtime. It is different from the probability that the realized downtime in one particular year is below four hours, $\Pr(D_{\mathrm{yr}} \le 4\,\mathrm{h} \mid \mathbf{y})$. The second requires the predictive distribution of the annual downtime, which includes the random occurrence and duration of the outages in that year even when the parameters are known.

For configuration B with known parameters, complete outages begin at the failure frequency $f_{\mathrm{F}} = 0.3368$ per year and last an exponential time of mean MDT $= 10\,\mathrm{h}$. Because outages are rare and short against the mean time between them, 26 000 h (Section 19.5.5 of Volume I), their number N in a year is approximately Poisson with mean $m = 0.3368$, and the annual downtime is the sum of N independent exponential durations. Given $N = n \ge 1$ the sum has an Erlang distribution, so

$$\Pr(D_{\mathrm{yr}} > d) = \sum_{n=1}^{\infty} \frac{\mathrm{e}^{-m} m^{n}}{n!}\, \mathrm{e}^{-d/\mathrm{MDT}} \sum_{k=0}^{n-1} \frac{(d/\mathrm{MDT})^{k}}{k!}, \qquad d \ge 0, \tag{33.13}$$

whose mean is $m \cdot \mathrm{MDT} = 3.37\,\mathrm{h}$, the expected downtime of equation (33.7). At $d = 4\,\mathrm{h}$, with $d/\mathrm{MDT} = 0.4$ and $\mathrm{e}^{-0.4} = 0.67032$, the probabilities of one, two and three outages are 0.2405, 0.0405 and 0.0045, the probabilities that their total duration exceeds 4 h are 0.67032, $0.67032 \times 1.4 = 0.93845$ and $0.67032 \times 1.48 = 0.99207$, and the terms are $0.1612 + 0.0380 + 0.0045$, with 0.0004 from four outages, so $\Pr(D_{\mathrm{yr}} > 4\,\mathrm{h}) = 0.204$. A direct simulation of the Markov model over 100 000 years, started from the stationary distribution, gives 0.204 as well, and 0.238 for configuration A against the 0.242 of the approximation.

Figure 33.6 plots equation (33.13) for both configurations. The horizontal axis is a downtime level d in hours; the vertical axis is the probability that the downtime of one year exceeds it. Both curves start at about 0.286, the probability of at least one outage in the year, and fall toward zero; configuration A falls more slowly because its outages last twice as long. The dashed vertical line is the four-hour criterion. In configuration B, a configuration that meets the criterion on expectation, about one year in five still exceeds four hours of downtime, and one year in twenty exceeds twenty hours; 71 % of years have no outage at all. The expected value of 3.37 h is not the downtime of a typical year, which is zero, and it is not a bound on the downtime of any year.

The two probabilities answer different acceptance questions. A requirement stated on the expected downtime, as equation (33.5) is, is judged by the posterior confidence of equation (33.12); a service-level agreement that pays a penalty in any year with more than four hours of downtime is judged by the predictive distribution, which also averages equation (33.13) over the posterior

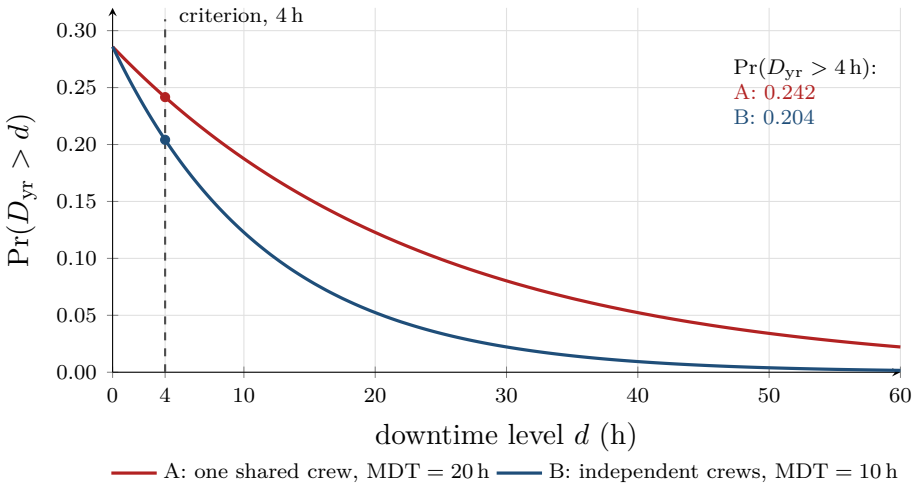


Figure 33.6: Predictive probability that the functional downtime of one year exceeds d hours, equation (33.13) truncated at four outages, with known parameters $\lambda = 0.001\,\mathrm{h}^{-1}$ and $\mu = 0.05\,\mathrm{h}^{-1}$ (illustrative): Poisson number of outages with mean $m = 0.3367$ (A) and 0.3368 (B) per year, exponential durations of mean MDT. Dots: $\Pr(D_{\mathrm{yr}} > 4\,\mathrm{h})$.

of the parameters. A design that meets the first with confidence may still breach the second in one year out of five.

Common mistake

Reading the posterior probability that the long-run expected downtime meets a criterion as the probability that the downtime of next year will meet it. The first concerns a parameter-dependent average and shrinks toward certainty as data accumulate; the second includes the randomness of outages in a single year, which no amount of data removes.

33.6.6 Independent technical review

Before a consequential reliability assessment is accepted, an independent reviewer should be able to examine:

- the correctness and completeness of the requirements;
- the modeling context and the system boundaries;
- the data quality and the provenance of the parameters;
- the structural assumptions and the common-cause mechanisms;
- the numerical verification and the physical validation;
- the sensitivity to parameters and scenarios;
- the compliance claims and their exact normative basis, that is, which standard, edition and requirement each claim invokes;
- the conclusions, limitations and recommended actions.

Independent review should be proportionate to the consequences of an erroneous decision. A maintenance-budgeting model does not necessarily require the same assurance rigor as a protective function whose failure could cause serious injury or widespread disruption of the system. The proportionality is the same graded approach that the SIL scale applies to safety functions: the more risk reduction a decision or a function carries, the more independent the evidence that it is correct must be.

33.6.7 A qualified engineering conclusion

An appropriately qualified conclusion for the numerical case of this chapter reads:

> The two-pump system with independent repair resources has a modeled steady-state functional availability of approximately 99.9616 %, corresponding to 3.37 expected unavailable hours per 8760-hour year under the baseline failure and restoration assumptions.
>
> This satisfies the project's illustrative four-hour annual expected-downtime target at the nominal parameter values. However, the result is sensitive to the assumed pump failure intensity, excludes several identified common-cause and control-system failure mechanisms, and has not been established as valid across the full operating envelope.
>
> The proposed repair-capacity improvement is supported for further implementation planning, subject to validation of the input assumptions, completion of functional redundancy testing, and resolution of the identified model gaps. Final acceptance remains with the designated engineering authority.

This conclusion is more useful than the declaration "reliability requirement passed." It separates the calculated result, the assumptions on which it rests, the evidence still missing and the proposed engineering decision, and it names who makes the decision. It could be strengthened by quantifying its own caveats: the failure rate at which the conclusion reverses, $0.0010918\,\mathrm{h}^{-1}$, 9.2 % above the nominal value; the posterior confidence that the criterion is met, 0.65 for ten observed failures; and the one-in-five probability that a single year exceeds four hours. The next chapter applies the methods of the book to the integrated power and cooling infrastructure of a data center, in which electrical and cooling failures are coupled through shared equipment and every one of those caveats becomes a modeled dependency.

33.7 What to retain

Definition 33.1 (Engineering reliability assessment)**.** An engineering reliability assessment is a traceable argument from an approved requirement, through a stated system definition, a verified model, justified and validated inputs and an uncertainty analysis, to a decision accepted by a named authority: requirements, models, evidence, assurance, decision. A numerical result is evidence within it only when its meaning, assumptions, validity for the context of use and relevance to the requirement are established.

Standards define different forms of engineering assurance, and applying one does not satisfy another: the NERC planning standards prescribe transmission planning events and performance, IEC 61508 prescribes the lifecycle and integrity of safety functions, and the ISO 55000 family prescribes a management system for assets. A reliability requirement must state the function, the operating envelope, the metric, the threshold and the period, and a mathematically correct model does not establish compliance without an assessment of the applicable requirements. Verification and validation are separate activities, and both are needed. Results must be traceable to physical assumptions, data, model versions and tests. Risk registers organize mitigation but do not replace quantitative modeling or physical validation. Uncertainty can change an acceptance decision, and nominal compliance with a small margin may warrant further testing or a more conservative design. The final engineering decision requires accountable judgment, not only a numerical output. Table 33.12 lists the chapter's terms, their meaning and units, and the subsection that defines each.

Table 33.12: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Compliance requirement	obligation of an applicable standard	assessed by its own method	–	33.1.3
Planning event	contingency a planning study must simulate	categories P0–P7	–	33.2.2
Safety function	action that achieves a safe state on a hazard	five specified elements	–	33.3.2
Safety integrity level	one of four levels of required integrity	SIL 1 to 4	–	33.3.3
$\mathrm{PFD}_{\mathrm{avg}}$	average probability of dangerous failure on demand	$\approx \lambda_{\mathrm{DU}} T_{\mathrm{p}}/2$	–	33.3.4
PFH	average frequency of dangerous failure	per hour	h^{-1}	33.3.3
Proof-test interval	time between tests revealing latent failures	T_{p}	h	33.3.4
Safety lifecycle	phases from hazard analysis to modification	Figure 33.2	–	33.3.5
Long-run cost rate of a policy	maintenance, downtime and monitoring cost	$C(a)$, eq. (33.4)	R$/yr	33.3.7
Reliability requirement	function, envelope, metric, threshold, period	e.g. $\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}$	h/yr	33.4.1
Expected annual downtime	hours unavailable per year	$8760\,\pi_0$	h/yr	33.4.4
Risk reduction of an action	avoided expected downtime cost	$c_D\,\Delta D$	R$/yr	33.4.6
Verification	model implemented and solved correctly	$\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$, $\sum_i \pi_i = 1$	–	33.5.1
Validation	model adequate for its intended use	tests against the plant	–	33.5.1
Context of use	question, envelope and consequence of error	–	–	33.5.4
Traceability	chain from requirement to decision	Figure 33.4	–	33.5.5
Configuration management	control of versions behind a result	–	–	33.5.7
Risk register	record of risks, controls, owners, status	Table 33.10	–	33.6.1
Threshold failure rate	rate at which the criterion is just met	λ^{*}, eq. (33.11)	h^{-1}	33.6.3
Posterior confidence	probability that the criterion holds	eq. (33.12)	–	33.6.4
Predictive downtime	distribution of one year's downtime	eq. (33.13)	h	33.6.5

Further reading. NERC, TPL-001-5.1, *Transmission System Planning Performance Requirements*; IEC 61508-1:2010, *Functional Safety General Requirements*; ISO 55001:2024, *Asset Management System Requirements*; ISO 55000:2024, *Vocabulary, Overview and Principles*.

33.8 Exercises

Exercise 33.1. (a) State the primary purpose of NERC TPL-001. (b) State what IEC 61508 principally addresses. (c) Name the document of the ISO 55000 family that specifies asset-

management-system requirements, and state what ISO 55000 provides instead. (d) A study shows that the outage of one transmission line has an unavailability of 0.01 and contributes only 0.285 MW to the expected curtailment. Explain whether this exempts the outage from a planning assessment. (Answers: (a) transmission-system planning performance assessment for the Bulk Electric System, through defined planning events and performance requirements. (b) The functional safety of electrical, electronic and programmable electronic safety-related systems. (c) ISO 55001; ISO 55000 provides the vocabulary, overview and principles. (d) No: a low modeled probability does not exempt a required planning event, which must be simulated and, if its performance criterion is violated, corrected; the probabilistic result is additional analysis.)

Exercise 33.2. (a) State the distinction between verification and validation. (b) Explain whether meeting a modeled expected-downtime target proves compliance with every reliability standard. (c) State what makes a requirement traceable. (d) Explain why an assessment can pass its nominal criteria and still be unsuitable for final acceptance. (Answers: (a) verification checks that the model has been implemented and solved correctly; validation assesses whether it represents the real system adequately for its intended physical use. (b) No: applicability and the specific normative requirements of each standard must be assessed separately. (c) A documented connection to models, test evidence, results and approval. (d) Parameter uncertainty, common causes or missing validation may undermine the result.)

Exercise 33.3. A low-demand safety channel has $\lambda_{\mathrm{DU}} = 5\times10^{-7}\,\mathrm{h}^{-1}$ (illustrative). (a) Compute $\mathrm{PFD}_{\mathrm{avg}}$ for an annual proof test, $T_{\mathrm{p}} = 8760\,\mathrm{h}$, with the approximation of equation (33.3) and exactly with equation (33.2), and state the SIL band of the result. (b) Find the longest proof-test interval for which the approximate $\mathrm{PFD}_{\mathrm{avg}}$ does not exceed 10^{-3}. (c) For the channel of Section 33.3.4, $\lambda_{\mathrm{DU}} = 2\times10^{-6}\,\mathrm{h}^{-1}$, compute $\mathrm{PFD}_{\mathrm{avg}}$ and the risk reduction factor for $T_{\mathrm{p}} = 2000\,\mathrm{h}$. (d) State two reasons why the result of (a) does not demonstrate that the function meets that SIL. (Answers: (a) $x = 0.00438$; approximation 0.00219; exact $1-(1-\mathrm{e}^{-0.00438})/0.00438 = 0.0021868$; SIL 2 band. (b) $T_{\mathrm{p}} = 2\times10^{-3}/(5\times10^{-7}) = 4000\,\mathrm{h}$. (c) 0.002 (exact 0.0019973), a risk reduction factor of 500. (d) Any two of: the architectural constraints, diagnostic and proof-test coverage, common-cause failures between channels, systematic failures and systematic capability, repair time and the installed configuration are not represented.)

Exercise 33.4. For the two-pump assessment with $\mu = 0.05\,\mathrm{h}^{-1}$ and the criterion $\mathbb{E}[D_{\mathrm{yr}}] \le 4\,\mathrm{h/yr}$: (a) verify the threshold failure rate $\lambda^{*} = 0.0010918\,\mathrm{h}^{-1}$ of configuration B; (b) find the failure rate at which configuration A, one shared crew, just meets the criterion; (c) at $\lambda = 0.001\,\mathrm{h}^{-1}$, find the repair rate and the mean repair time that configuration A would need to meet the criterion; (d) at the revised rate $\lambda = 0.0015\,\mathrm{h}^{-1}$ compute $\mathbb{E}[D_{\mathrm{A}}]$, $\mathbb{E}[D_{\mathrm{B}}]$, the risk reduction of the second crew at $c_D = \mathrm{R\$}\,25\,000$ per hour and its net value at R\$ 45 000 per year, and state what the result implies for the decision. (Answers: (a) $8760\,[0.0010918/0.0510918]^2 = 8760\times0.00045662 = 4.000\,\mathrm{h/yr}$. (b) Solving $8760\cdot2\rho^2/(1+2\rho+2\rho^2) = 4$ gives $\rho = 0.015344$, $\lambda = 0.000767\,\mathrm{h}^{-1}$. (c) $\mu = 0.0652\,\mathrm{h}^{-1}$, a mean repair time of 15.3 h instead of 20 h. (d) $\rho = 0.03$: $\mathbb{E}[D_{\mathrm{A}}] = 8760\times0.0018/1.0618 = 14.85\,\mathrm{h/yr}$ and $\mathbb{E}[D_{\mathrm{B}}] = 7.43\,\mathrm{h/yr}$; $\Delta\mathcal{R} = 25\,000\times7.419 = \mathrm{R\$}\,185\,471$ per year, net R\$ 140 471 per year. The second crew becomes more valuable as the failure rate rises, but neither configuration meets the criterion, so the decision needs a further action, such as a third pump or faster repairs, and the requirement cannot be accepted on the crew alone.)

Chapter 34

Data Center Power and Cooling Reliability

34.1 Data center reliability as a service-level property

Chapter 20 of Volume I built the availability model of a pair of cooling-water pumps, and Chapter 15 of Volume I built the event tree of a loss of utility supply at a data center, with an uninterruptible power supply and a standby generator as its headings. Each of those models treated one infrastructure and assumed the other adequate. A data center does not fail that way. Its IT equipment needs electrical power and an acceptable thermal environment at the same time; the cooling equipment that maintains the thermal environment itself needs electrical power; and almost all the electrical energy that the IT equipment consumes becomes heat that the cooling must remove. A failure in one infrastructure can therefore disable the other, a shared control system or switchboard can defeat the redundancy of both, and a loss of cooling becomes an interruption of IT service only if it lasts longer than the thermal inertia of the room allows. The question of this chapter is

how to quantify the probability that a data center loses its required IT service when electrical and cooling failures are coupled through shared physical infrastructure.

The chapter answers it in steps. It defines the reliability of a data center at the level of the IT service it supports; models the reliability of the electrical infrastructure (utility supply, standby generator, uninterruptible power supply, switchgear and distribution) and of the cooling infrastructure (chillers, pumps, air handlers and the chilled-water loop); compares N+1 and 2N redundancy; represents shared dependencies and common-cause failures explicitly; computes the availability and the expected deficiency time of a complete illustrative facility; models the thermal ride-through after a loss of cooling and the event tree of cooling restoration that follows; and closes with importance measures, design and maintenance decisions, and a unified physical–probabilistic formulation of the problem.

Jargon: UPS, standby generator, transfer controls, switchgear, PDU, CRAH and CRAC units, chiller, chilled-water loop, heat rejection, N+1 and 2N redundancy, thermal ride-through

An *uninterruptible power supply* (UPS) carries the critical load from stored energy, usually batteries, during disturbances of the normal supply and during the transition to an alternate source; its *usable autonomy* is the time for which it can carry the load. A *standby generator* is an engine-driven generator that starts after a loss of utility supply and provides sustained alternate power. *Automatic transfer controls* (an automatic transfer switch and its logic) move the load from the failed source to the alternate one. *Switchgear* is the assembly of buses, breakers and protective devices that routes and protects electrical power. A *power distribution unit* (PDU), with the busways, feeders and branch breakers around it, delivers power from the UPS output to the IT racks; one such chain is a *distribution path.* A *computer-room air handler* (CRAH) moves room air across a chilled-water coil; a *computer-room air conditioner* (CRAC) has its own refrigeration circuit. A *chiller* is a refrigeration machine that cools the water of the *chilled-water loop*, the pumps, pipes and headers that carry heat from the air handlers to the chillers. *Heat rejection* equipment (cooling towers, dry coolers) transfers the collected heat to the outdoor environment. *N+1 redundancy* installs the N capacity units that the demand requires plus one; *2N redundancy* installs two complete systems, each able to carry the whole demand. The *thermal ride-through* time is the time for which the IT equipment stays within its temperature limits after active cooling is lost.

Key idea

The reliability of a data center is the probability that its required IT service is supported, not the reliability of any one piece of equipment. It follows from the equipment reliabilities combined with the physical dependencies between power and cooling and with the timing of failure, thermal response and restoration: a cooling failure that is restored before the thermal limit is reached is an equipment failure, not a service interruption.

This section defines the required function and the service indicator of a data center, draws the physical chain that couples its power and cooling, and writes the thermal balance that decides how much time a cooling failure leaves.

34.1.1 Equipment failure and IT-service interruption

A data center contains many systems that can fail without interrupting its IT operations. A chiller can trip while the redundant chillers keep operating; a utility feeder can lose power while the UPS batteries carry the load; a cooling-water pump can fail while another pump maintains the required flow. The failure of an item and the loss of the service that the facility exists to provide are therefore different events,

$$\text{equipment failure} \neq \text{IT-service interruption.}$$

The analysis begins, as Chapter 2 of Volume I prescribed for every reliability study, with the required function. For this chapter it is

maintain electrical power and an acceptable thermal environment for the specified IT workload throughout the assessment period.

The definition is more demanding than requiring power at the UPS output. IT equipment also needs adequate cooling, heat rejection, air and water distribution and the controls that coordinate them, and a failure of any of these can defeat the function while every electrical component operates.

The state of the function is described by the *service indicator*

$$Y(t) = \begin{cases} 1, & \text{the required IT service is supported at time } t, \\ 0, & \text{the required IT service cannot be supported at time } t, \end{cases} \tag{34.1}$$

which is the binary system state of equation (20.2) in Chapter 20 applied to the IT service. Its probability of being one is the *IT-service availability*,

$$A_{\mathrm{IT}}(t) = \Pr[Y(t) = 1] = \mathbb{E}[Y(t)], \tag{34.2}$$

dimensionless. For a stationary system the probability does not depend on t, and the steady-state value is $A_{\mathrm{IT},\infty} = \lim_{t\to\infty} A_{\mathrm{IT}}(t) = \Pr[Y = 1]$. The functional downtime over an assessment period $[0, H]$ is the time integral of the indicator of interruption, $D_H = \int_0^H [1 - Y(t)]\,\mathrm{d}t$, and by linearity of the expectation its expected value is

$$\mathbb{E}[D_H] = \int_0^H [1 - A_{\mathrm{IT}}(t)]\,\mathrm{d}t, \tag{34.3}$$

in hours, which is equation (20.5) for the IT service. If $A_{\mathrm{IT}}(t)$ is constant over an 8760-hour year, the integral is a product,

$$\mathbb{E}[D_{\mathrm{yr}}] = 8760\,\mathrm{h} \times (1 - A_{\mathrm{IT}}). \tag{34.4}$$

This is an expected amount of time, not the prediction of one uninterrupted outage. An availability of 0.9999, for example, gives 0.876 h or 52.6 minutes per year; that total may arrive as one interruption, as several short ones, or, in most years of a facility with rare failures, not at all, with a long interruption in a rare year carrying the average.

34.1.2 The coupled physical chain

Figure 34.1 draws the physical chain on which the required function rests. Each box is a part of the facility, and the arrows are what passes between the parts: blue arrows carry electric power, red arrows carry heat, and the green arrow is the cooled air or liquid that returns to the IT equipment. The electrical supply (utility, generators, UPS, switchgear and distribution) feeds the IT equipment, which consumes the electricity and releases it as heat. The heat is carried by air, water or refrigerant distribution to the cooling equipment, which also consumes electricity and rejects the heat to the environment. The box on the right is the required IT service, which needs both the power and an acceptable inlet temperature.

Two dependencies appear in the figure. First, the cooling equipment requires electrical power: the blue arrow into the cooling box means that an electrical failure upstream of it removes cooling as well as IT power. Second, nearly all the electrical energy consumed by the IT equipment becomes heat that must be removed from the IT environment: the red arrows close a loop, and a facility that loses cooling while it keeps its IT load powered keeps generating the heat it can no longer reject. The electrical and cooling subsystems are therefore physically coupled, and their reliabilities cannot always be calculated independently and multiplied.

34.1.3 The central thermal balance

The time that a cooling failure leaves before it becomes a service interruption follows from an energy balance. Let $\dot{Q}_{\mathrm{IT}}$ be the rate of heat generation of the IT equipment, $\dot{Q}_{\mathrm{other}}$ the other heat gains of the IT environment (lighting, people, the building envelope, fan and electrical losses), and $\dot{Q}_{\mathrm{cool}}$ the rate at which heat is removed, all in kilowatts. Represent the IT environment by one effective thermal storage state with a representative temperature $T_{\mathrm{IT}}(t)$, in degrees Celsius, and an effective thermal capacitance C_{eff}, in MJ/K: the heat that the air, the racks,

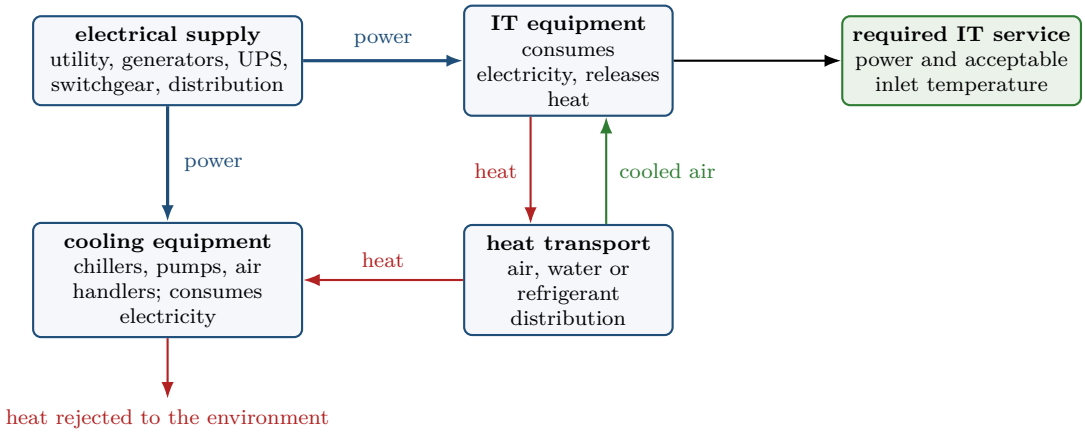


Figure 34.1: The coupled physical chain of a data center. Electric power (blue) feeds both the IT equipment and the cooling equipment; heat (red) flows from the IT equipment through the heat-transport system to the cooling equipment and out to the environment; cooled air or liquid (green) returns to the IT equipment. The required IT service needs both power and an acceptable inlet temperature.

the structure and any water in thermal contact with them absorb per kelvin of temperature rise. The first law of thermodynamics for this control volume, with no work done and the stored energy written as $C_{\text{eff}}T_{\text{IT}}$ above a reference, states that the rate of change of stored energy equals the heat gained minus the heat removed,

$$C_{\text{eff}}\,\frac{\mathrm{d}T_{\text{IT}}}{\mathrm{d}t} = \dot{Q}_{\text{IT}} + \dot{Q}_{\text{other}} - \dot{Q}_{\text{cool}}. \tag{34.5}$$

The units check: a power in MJ/s divided by a capacitance in MJ/K is a rate in K/s, and a temperature difference of one kelvin is one degree Celsius. In normal operation the cooling removes what the room gains, $\dot{Q}_{\text{cool}} = \dot{Q}_{\text{IT}} + \dot{Q}_{\text{other}}$, and the temperature is constant.

Remark 34.1. T_{IT} is a temperature, not the time to failure T; C_{eff} is a thermal capacitance, not the available capacity C, the consequence C or a minimal cut set C_j; and the heat rates $\dot{Q}$ are not the generator matrix $\mathbf{Q}$ of a Markov model.

Equation (34.5) explains why losing cooling capacity does not necessarily cause an immediate interruption. When $\dot{Q}_{\text{cool}}$ falls, the temperature rises at the finite rate $(\dot{Q}_{\text{IT}}+\dot{Q}_{\text{other}}-\dot{Q}_{\text{cool}})/C_{\text{eff}}$, and if the stored thermal headroom and the residual heat removal are large enough, there is time to restore cooling before a limit is reached. The equation has the same form as the battery energy balance of the UPS in Section 15.5.5, with stored heat in place of stored electrical energy, and as the conductor heat balance of Chapter 28.

The aggregate balance cannot by itself guarantee acceptable temperatures at every server inlet. The airflow distribution, local hot spots, the heat capacity of individual racks and the operating limits of the equipment matter as well, and a room whose average temperature is acceptable can have racks above their limits. Published thermal guidelines for data processing environments define equipment environmental classes and their allowable operating conditions; they provide the context for choosing a thermal limit, but the temperature actually permitted depends on the equipment class and on the manufacturers' specifications. The limits used in this chapter are illustrative.

34.2 Electrical infrastructure reliability

This section describes the critical electrical chain of a data center, derives the probability that a standby generator and a UPS together carry the IT load through a utility outage, computes

the unavailability of a simplified source arrangement, and treats redundant distribution paths.

34.2.1 The critical electrical chain

A simplified electrical system consists of the chain

$$\text{utility} \longrightarrow \text{switchgear} \longrightarrow \text{UPS} \longrightarrow \text{power distribution} \longrightarrow \text{IT},$$

with a standby generator providing an alternate source during a utility failure. Figure 34.2 draws one such arrangement. The utility supply and the standby generator, on the left, both feed the switchgear and its transfer controls; the switchgear feeds the UPS, whose batteries are drawn below it; the UPS output divides into two distribution paths, A and B, which both feed the IT racks on the right; and a dashed branch from the switchgear supplies the cooling plant. In this arrangement, which is common but not universal, the mechanical loads are fed upstream of the UPS: they ride on the generator but not on the batteries, so they stop during the generator start while the IT load does not.

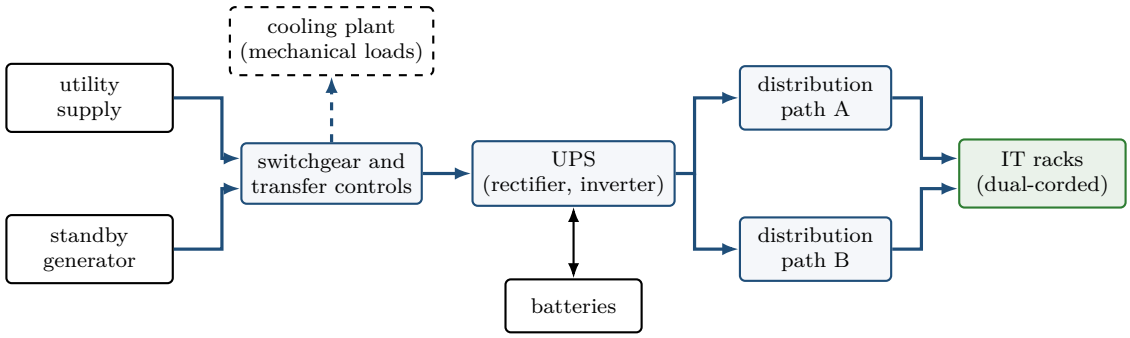


Figure 34.2: A simplified critical electrical chain (illustrative). The utility supply and the standby generator feed the switchgear and transfer controls; the UPS, with its batteries, bridges the transfer; two distribution paths, either able to carry the full load, feed dual-corded IT racks. The cooling plant is fed from the switchgear, upstream of the UPS.

Table 34.1 lists the main components, with the function each performs and the issues that dominate its reliability. The precise arrangement of generators, UPS modules, busways, transfer equipment and distribution paths determines which failures are consequential: the same breaker failure is a single point of failure in one arrangement and a tolerated fault in another. An equipment availability calculation must therefore reflect the actual electrical architecture, as recorded in the facility's single-line diagram, and not an equipment list.

34.2.2 UPS ride-through and the backup probability

Suppose the utility supply fails at $t = 0$. The standby generator needs several seconds to start, to synchronize where the arrangement requires it, and to be connected to the load. The UPS must preserve the critical load while this transition occurs. Let T_{G} be the time from the utility failure until the generator is usable, and T_{UPS} the usable UPS autonomy. The generator preserves the load through the transition if

$$T_{\mathrm{G}} < T_{\mathrm{UPS}}$$

and the UPS and the transfer equipment operate correctly in the meantime.

Let G be the event that the generator starts and becomes usable, with probability $p_G = \Pr(G)$, and let B be the event that the UPS bridges the transition, operating correctly until the generator takes the load ($T_{\mathrm{UPS}} > T_{\mathrm{G}}$). Bridging is defined only when the generator becomes usable,

Table 34.1: Components of the critical electrical chain.

Component	Principal function	Reliability issue
Utility connection	Normal electrical supply	External outages, voltage disturbances
Standby generator	Sustained alternate power	Start failure, fuel, engine cooling, mechanical faults
UPS	Maintain supply during disturbances and transitions	Inverter, batteries, bypass, common controls
Automatic transfer controls	Coordinate the alternate supply	Failed or delayed transfer
Switchgear	Route and protect electrical power	Breaker failure, bus faults, maintenance
Power distribution	Deliver power to the IT racks	Feeder, busway, PDU and breaker faults

since otherwise T_{G} does not exist, so its probability is stated conditionally, $p_B = \Pr(B \mid G)$. The backup succeeds when both events occur, and by the definition of conditional probability, $\Pr(G \cap B) = \Pr(G)\,\Pr(B \mid G)$, or

$$p_{\mathrm{backup}} = p_G\, p_B. \tag{34.6}$$

The product is exact provided p_B is the probability conditional on generator success; an unconditional bridging probability, estimated over all utility outages including those in which the generator failed, would mix the two events and count the generator failures twice. The model extends directly to transfer switches, UPS faults and several standby generators, each adding a conditional factor.

Chapter 15 of Volume I quantified the same race in an event tree with the UPS, the generator and timely recovery as headings, and derived the timely-recovery branch from the restoration-time distribution evaluated at the battery autonomy, $\Pr(\mathrm{REC}) = \Pr(T_{\mathrm{R}} \le T_{\mathrm{UPS}})$ (Section 15.5.4). The bridging probability p_B is of the same kind: it combines the probability that the UPS operates on demand with the probability that the time needed is shorter than the time available.

34.2.3 Utility-to-generator backup: a numerical example

Assume that at a randomly selected assessment instant the utility supply is unavailable with probability $q_U = 0.01$, which corresponds to 87.6 hours per year. Conditional on a utility outage, suppose $p_G = 0.95$ and $p_B = 0.98$. By equation (34.6),

$$p_{\mathrm{backup}} = 0.95 \times 0.98 = 0.931.$$

The simplified upstream supply is unavailable when the utility is out and the backup has failed. When the utility is available the source is available, so by the law of total probability

$$q_{\mathrm{source}} = q_U\,(1 - p_{\mathrm{backup}}), \tag{34.7}$$

and

$$q_{\mathrm{source}} = 0.01 \times (1 - 0.931) = 0.01 \times 0.069 = 0.00069,$$

or 0.069 %, about 6.04 hours per year.

This is a hypothetical *snapshot state-mixture model*: its probabilities describe the conditions found at a randomly sampled instant, including whether the backup is supporting the load at that instant. It does not follow from multiplying an annual frequency of utility outages

by a generator start probability. The product $f_{\mathrm{IE}}(1 - p_{\mathrm{backup}})$, with the initiating-event frequency f_{IE} of Chapter 15, is a frequency of unsupported outages per year, for example $0.40 \times 0.069 = 0.0276\,\mathrm{yr}^{-1}$; turning it into a fraction of time requires the duration of each unsupported outage, and a backup that starts but fails to run, or a battery that is exhausted before the generator connects, makes that duration depend on time. Duration-dependent UPS autonomy and generator restoration should normally be modeled chronologically, as Sections 34.5.4 and 34.6.6 do for cooling.

34.2.4 Redundant distribution paths

Suppose the IT equipment can be supplied through either of two distribution paths, A and B, for example through the two power supplies of dual-corded servers, and that each path is rated to carry the full IT load. Let each path have the stationary unavailability $q_A = q_B = 0.005$. The IT supply is lost only when both paths are unavailable, the parallel structure of Chapter 9 of Volume I, and if their failures are independent

$$q_{\mathrm{dist}} = q_A\, q_B = 0.005^2 = 0.000025, \tag{34.8}$$

so the distribution availability is $A_{\mathrm{dist}} = 1 - 0.000025 = 99.9975\,\%$. One path alone would leave the load without supply $8760 \times 0.005 = 43.8$ hours per year on average; two independent paths reduce that to 0.219 hours, a factor of 200.

The improvement rests on the independence assumption. If both paths share the same switchgear room, cable route, control power or upstream bus, their failures are dependent, and the shared element must be represented explicitly rather than assumed away. With a shared element of unavailability q_{s} whose failure disables both paths, the common-cause event model of Chapter 13 of Volume I gives $q_{\mathrm{s}} + (1 - q_{\mathrm{s}})\, q_A q_B$; for $q_{\mathrm{s}} = 0.0002$ this is $0.0002 + 0.9998 \times 0.000025 = 0.000225$, nine times the independent value. Section 34.4.2 develops the model.

34.3 Cooling infrastructure reliability

This section describes the physical cooling chain and its required functions, derives the availability of N+1 and 2N cooling arrangements, and compares the two.

34.3.1 The physical cooling chain

A chilled-water cooling system carries heat along the chain drawn in Figure 34.3. The boxes, from left to right, are the IT equipment that produces the heat, the computer-room air handlers that move air across the IT equipment and transfer its heat to chilled water, the chilled-water loop with its pumps and headers, the chillers that produce the cooling capacity, and the heat-rejection equipment that releases the collected heat outdoors. The red arrows are the heat flow. Beneath each box is the required function it performs, and the bar across the bottom is the function that spans the whole chain, carried out by the controls.

The cooling function therefore consists of five required functions: moving sufficient air or liquid across the IT equipment; transporting heat away from the IT environment; producing sufficient cooling capacity; rejecting the collected heat to the environment; and maintaining the required temperatures and flow rates. A failure of any essential function can compromise cooling, so at the level of functions the chain is a series system, and redundancy can only be provided within each function. Healthy chillers cannot maintain server inlet temperatures if a common chilled-water distribution header fails: the header is one element in series with every chiller, a minimal cut set of order one that no number of chillers removes.

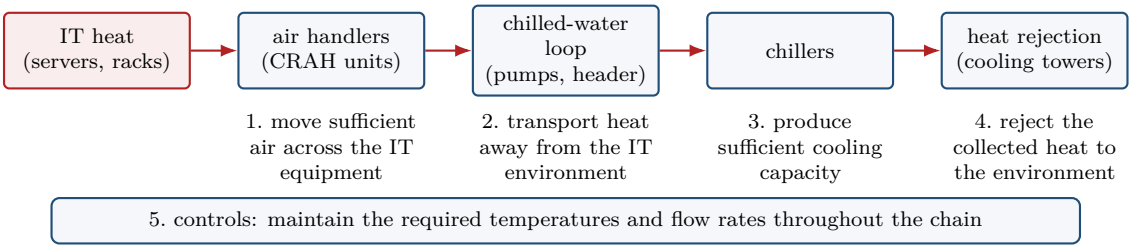


Figure 34.3: The physical cooling chain of a chilled-water data center and the required function of each stage. Heat (red arrows) passes from the IT equipment through the air handlers and the chilled-water loop to the chillers and the heat-rejection equipment; the controls span the chain.

34.3.2 N+1 capacity redundancy

Suppose each chiller has the capacity $K = 500\,\mathrm{kW}$ and the cooling demand is $L = 900\,\mathrm{kW}$. Two chillers meet the demand, since $2 \times 500 = 1000\,\mathrm{kW} \geq 900\,\mathrm{kW}$, and one does not. A plant of three chillers is then N+1 capacity redundancy at this duty point, with $N = 2$: two units are required and one additional unit is installed. The qualifier matters, because redundancy is a property of the plant and the load together; at a demand above 1000 kW the same three chillers would be N+0. This is the running cooling plant of this book, three chillers of which two are required, introduced in Chapter 1 of Volume I. Panel (a) of Figure 34.4 draws it as a reliability block diagram: three chiller blocks in parallel feed a voting element that requires two of them.

The system is adequate when at least two of the three chillers are available. Let each chiller be available with probability a, independently of the others. The number of available chillers is then binomial with parameters 3 and a, and the cooling-capacity availability is the probability that it is two or three,

$$A_{\mathrm{chiller}} = \binom{3}{2} a^2(1-a) + \binom{3}{3} a^3 = 3a^2 - 3a^3 + a^3,$$

that is,

$$A_{\mathrm{chiller}} = 3a^2 - 2a^3. \tag{34.9}$$

This is the 2-out-of-3 formula of equation (1.19) with an availability in place of a mission reliability; the algebra is the same because only the joint state of the chillers at one instant enters. For $a = 0.97$,

$$A_{\mathrm{chiller}} = 3 \times 0.9409 - 2 \times 0.912673 = 2.8227 - 1.825346 = 0.997354,$$

or 99.7354 %. The availability of one chiller, 97 %, is not the availability of the 2-out-of-3 capacity system: the system unavailability, 0.002646, is 11.3 times smaller than that of one chiller, 0.03. For a general plant of n identical units of which k are required, the same argument gives the k-out-of-n availability $\sum_{j=k}^{n} \binom{n}{j} a^j (1-a)^{n-j}$ of equation (9.19) in Chapter 9 of Volume I.

The calculation assumes that every available chiller delivers its full modeled capacity, and that partial deratings, load sharing between the units and common causes are absent. A chiller derated to 400 kW on a hot afternoon turns two available units into 800 kW, below the demand, and the binary availability model of equation (34.9) would not see it; Chapter 26 treats such multistate units.

34.3.3 2N redundancy

In a 2N arrangement two separate sets of capacity are each sufficient to satisfy the entire required function. For example, two independent cooling trains may each be rated to meet

$L = 900\,\text{kW}$, each train containing its own chillers, pumps and air handlers, as in panel (b) of Figure 34.4. If trains A and B each have unavailability q, either one can support the full function, and their failures are independent, the function is lost only when both are unavailable, and

$$A_{2\text{N}} = 1 - q^2. \tag{34.10}$$

For $q = 0.02$, $A_{2\text{N}} = 1 - 0.0004 = 0.9996$, or $99.96\,\%$.

This is not automatically the availability of a physically installed 2N data center. The formula omits shared distribution elements, the switching between trains, common controls, maintenance configurations and the other dependencies that can defeat the separation of the trains, and each train's unavailability q is itself the result of a series of chillers, pumps, valves and air handlers that the formula takes as given.

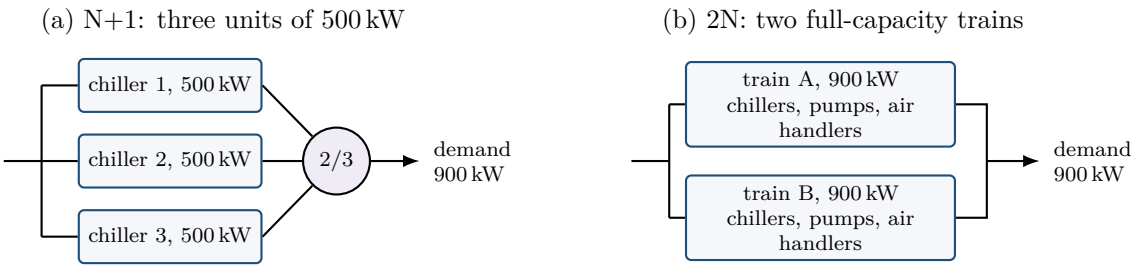


Figure 34.4: Reliability block diagrams of two cooling architectures. (a) N+1: three chillers of 500 kW, any two of which meet the 900 kW demand, availability $3a^2 - 2a^3$. (b) 2N: two trains, each able to meet the whole demand, availability $1 - q^2$ under independence.

Table 34.2 compares the two formulas for several values of the availability of one unit, taken equal to the availability of one 2N train so that only the architecture differs. Its columns are the unit availability a, the unavailability of one unit, the availability and unavailability of the 2-out-of-3 plant from equation (34.9), the availability and unavailability of the 2N arrangement from equation (34.10) with $q = 1 - a$, and the ratio of the two unavailabilities. For $a = 0.98$ the row reads $3 \times 0.9604 - 2 \times 0.941192 = 0.998816$ and $1 - 0.02^2 = 0.9996$. The ratio has a closed form: the 2-out-of-3 unavailability is $1 - 3a^2 + 2a^3 = 3q^2 - 2q^3$, so the ratio is $(3q^2 - 2q^3)/q^2 = 3 - 2q$, close to three for every reliable unit. The 2-out-of-3 plant has three minimal cut sets of order two, any two chillers, and the 2N arrangement has one. At equal unit availability the 2N arrangement therefore has about one third of the unavailability, but it installs 1800 kW of capacity against 1500 kW, and a full train contains more equipment than one chiller, so in practice its q is larger than $1 - a$.

Table 34.2: Availability of N+1 (2-out-of-3) and 2N cooling for several unit availabilities a, independent units, 2N train unavailability $q = 1 - a$ (illustrative).

a	$1 - a$	2-out-of-3 availability	2-out-of-3 unavail.	2N availability	2N unavail.	ratio $3 - 2q$
0.90	0.10	0.972000	0.028000	0.990000	0.010000	2.80
0.95	0.05	0.992750	0.007250	0.997500	0.002500	2.90
0.97	0.03	0.997354	0.002646	0.999100	0.000900	2.94
0.98	0.02	0.998816	0.001184	0.999600	0.000400	2.96
0.99	0.01	0.999702	0.000298	0.999900	0.000100	2.98

34.3.4 N+1 against 2N

Table 34.3 sets the two architectures side by side, one feature per row: the capacity architecture, the typical design goal, the role of shared infrastructure, the capital cost, the physical separation, and what the availability result depends on. N+1 tolerates the loss of one capacity unit, but the units commonly share headers, controls and electrical supply; 2N aims to preserve service through the loss of a complete train, and its benefit over N+1 comes as much from the separation of the trains as from the extra capacity.

Table 34.3: N+1 and 2N redundancy compared.

Feature	N+1	2N
Capacity architecture	Enough units for the demand plus one extra	Two full-capacity systems
Typical goal	Tolerate the loss of one capacity unit	Preserve service through the loss of one complete train
Shared infrastructure	May be significant	Should be minimized, but depends on the design
Capital cost	Usually lower	Usually higher
Physical separation	Not guaranteed	Often a major design objective
Availability result	Depends on the k-out-of-n capacity model	Depends on independence and on the distribution architecture

Two further properties are distinguished from redundant capacity in data center practice. *Concurrent maintainability* is the ability to take any capacity component or distribution path out of service for planned maintenance without interrupting the IT load. *Fault tolerance* is the ability to sustain any single unplanned failure or event, in any component or distribution path, without affecting the IT load. A widely used commercial classification of data center site infrastructure is built on this distinction: its third level is characterized by concurrent maintainability, and its fourth level additionally requires fault tolerance. A claim of N+1 or 2N equipment does not establish either property, nor a certification under such a classification, because both properties concern the distribution paths, the controls and the physical separation of the systems, not only the count of capacity units.

34.4 Coupled electrical and cooling failures and the integrated case study

This section models the dependence of cooling on the electrical infrastructure and a common-cause failure of redundant pumps, and then computes the availability and the expected deficiency time of a complete illustrative data center.

34.4.1 Shared dependencies

Suppose the cooling equipment is powered by the same electrical infrastructure that supplies the IT equipment. One electrical failure can then cause at the same time the loss of IT power, the loss of chiller power, the loss of pump power and the loss of the cooling controls. These consequences are correlated through one common initiating event. A fault tree that treats "IT power fails" and "cooling power fails" as independent basic events would double-count or misrepresent the shared mechanism. The model must instead be built around the underlying causes, with each shared cause entered once as a basic event that appears in both functions, the treatment of repeated basic events in Section 15.5.2 of Chapter 15 of Volume I.

Define the events

$$E = \{\text{the electrical infrastructure supports the required service}\},$$
$$C = \{\text{the cooling infrastructure supports the required heat removal}\}.$$

A simplified instantaneous service condition is that both hold,

$$Y = E \wedge C, \qquad \text{that is,} \qquad Y = \mathbf{1}_E \, \mathbf{1}_C. \tag{34.11}$$

If E and C were independent, $\Pr(Y = 1) = \Pr(E)\,\Pr(C)$. Where power failures also disable cooling they are not, and the general expression is the definition of conditional probability,

$$\Pr(E \wedge C) = \Pr(E)\,\Pr(C \mid E), \tag{34.12}$$

which holds whatever the dependence. In this section E and C are events, not the available capacity C or the consequence C.

The conditional formulation is more general, and it separates two quantities that must not be confused: the cooling-capacity availability conditional on a powered cooling system, $A_{C|E} = \Pr(C \mid E)$, and the unconditional availability of the entire cooling function, $\Pr(C)$. If the cooling cannot operate without the electrical infrastructure, $\Pr(C \mid \bar{E}) = 0$, and by the law of total probability $\Pr(C) = \Pr(C \mid E)\,\Pr(E)$. Multiplying the two unconditional availabilities then gives $\Pr(E)\,\Pr(C) = \Pr(E)^2\,\Pr(C \mid E)$, which counts the electrical unavailability twice; Section 34.4.6 shows the size of the error for the case study.

Common mistake

Multiplying an electrical availability by a cooling availability that already includes the loss of cooling from electrical failures. The product counts the shared electrical failures twice; the coupled availability is $\Pr(E)\,\Pr(C \mid E)$, with the cooling availability computed given electrical support.

34.4.2 Common-cause failure of redundant pumps

Suppose two cooling pumps, either of which suffices, are each unavailable with probability $q = 0.02$. If their failures are independent, both are unavailable with probability $q^2 = 0.02^2 = 0.0004$. Now introduce a shared control-system failure with probability $q_{\text{ctrl}} = 0.005$, assumed to disable both pumps; when it does not occur, the individual pump failures remain independent with unavailability 0.02. This is the explicit common-cause event model of Chapter 13 of Volume I (Section 13.3.3): with a common cause of probability c and independent failures of probability q, the redundant pair fails with probability $c + (1 - c)\,q^2$, equation (13.14), where the factor $1 - c$ removes the cases in which the common cause has already failed both units. Here

$$q_{\text{pumping}} = q_{\text{ctrl}} + (1 - q_{\text{ctrl}})\,q^2 = 0.005 + 0.995 \times 0.0004 = 0.005398. \tag{34.13}$$

The functional unavailability of the pumping system rises from 0.04 % to about 0.5398 %, a factor of 13.5, and the common-control failure carries $0.005/0.005398 = 92.6\,\%$ of it.

Adding a third pump would not substantially solve the problem unless the shared dependency were addressed. Table 34.4 evaluates the model for two and three pumps, any one sufficient, and for three values of the common-control probability; its last column is the relative reduction obtained from the third pump. Without the common cause, the third pump divides the unavailability by fifty; with $q_{\text{ctrl}} = 0.005$ it reduces it by 7.2 %, from 0.005398 to 0.005008, while improving the control system to $q_{\text{ctrl}} = 0.001$ with only two pumps reduces it by 74 %, to 0.001400. Redundancy without independence can provide much less reliability improvement than its count of units suggests.

Table 34.4: Unavailability of the pumping function, $q_{\text{ctrl}} + (1 - q_{\text{ctrl}})\, q^n$, for n pumps any one sufficient, pump unavailability $q = 0.02$ (illustrative).

q_{ctrl}	two pumps	three pumps	reduction from third pump
0	0.000400	0.000008	98.0 %
0.001	0.001400	0.001008	28.0 %
0.005	0.005398	0.005008	7.2 %

Principle 34.1. *A shared dependency of redundant units sets a floor under their joint unavailability. Once the shared cause dominates, reliability is improved by reducing the shared cause, through separation, diverse controls or independent supplies, and not by adding identical units.*

34.4.3 The facility

The rest of this section analyzes a hypothetical data center with an IT load of

$$P_{\text{IT}} = 800\,\text{kW},$$

constant over the assessment. All of it becomes heat, so $\dot{Q}_{\text{IT}} = 800\,\text{kW}$, and the total heat entering the modeled IT thermal environment is $\dot{Q}_{\text{IT}} + \dot{Q}_{\text{other}} = 900\,\text{kW}$, the additional 100 kW representing the other applicable heat gains. The facility contains a utility supply with generator and UPS backup; two independent IT distribution paths; two parallel chillers, each sufficient for the modeled cooling load; two parallel chilled-water pumps, each sufficient for the modeled flow; and common electrical and chilled-water infrastructure. The reliability parameters are illustrative and are not representative fleet averages.

Figure 34.5 draws the infrastructure as a reliability block diagram with the parameters of the next two subsections. The upper row is the electrical support E: the utility in parallel with the generator-and-UPS backup, in series with the shared electrical infrastructure, in series with the two distribution paths in parallel. The lower row is the cooling support C: the two chillers in parallel, in series with the two pumps in parallel, in series with the shared chilled-water header. The dashed arrow is the electric power that the cooling equipment draws from the electrical infrastructure, the dependence of Section 34.4.1; the AND gate on the right combines the two supports into the service condition $Y = E \wedge C$.

34.4.4 Electrical availability

Table 34.5 lists the electrical assumptions, with the symbol under which each enters the model. The first three rows are the source arrangement of Section 34.2.3, which gave $q_{\text{source}} = 0.00069$. The two distribution paths give, by equation (34.8), $q_{\text{dist}} = 0.005^2 = 0.000025$. The shared electrical infrastructure, the switchgear bus and the control power common to both paths, has $q_{\text{shared},E} = 0.0002$.

Assume that the source, the distribution-path subsystem and the shared electrical failure mechanism are mutually independent in this simplified model. The electrical support holds when none of the three mechanisms is active, which is a series structure, so its availability is the product of the three availabilities,

$$A_E = (1 - q_{\text{source}})\,(1 - q_{\text{dist}})\,(1 - q_{\text{shared},E}), \tag{34.14}$$

and

$$A_E = 0.99931 \times 0.999975 \times 0.9998 = 0.99908516,$$

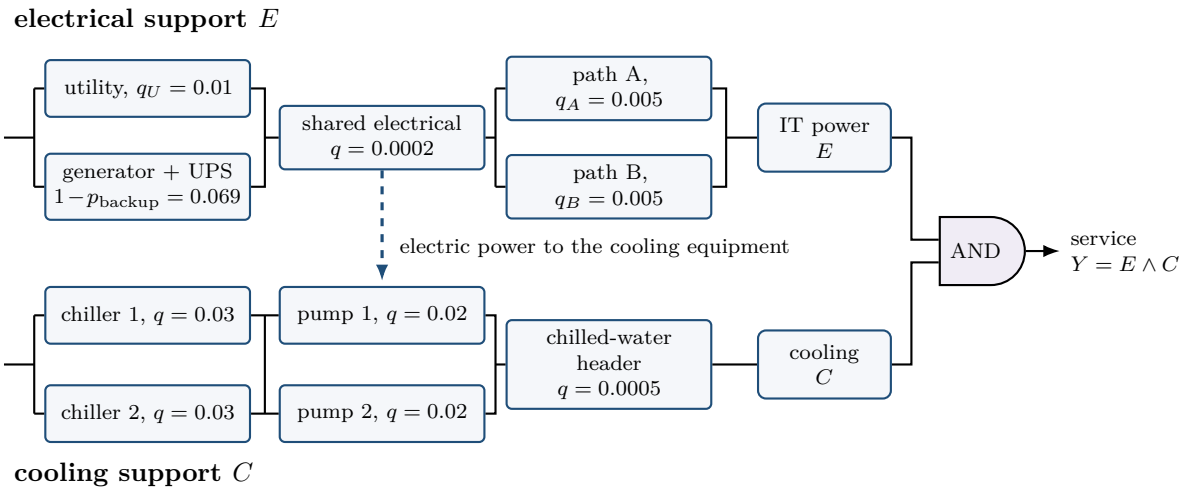


Figure 34.5: Reliability block diagram of the illustrative facility. Upper row: electrical support, the utility in parallel with the backup, the shared electrical infrastructure, and two parallel distribution paths. Lower row: cooling support given power, two parallel chillers, two parallel pumps and the shared chilled-water header. The dashed arrow is the cooling equipment's electrical dependence; the AND gate forms the service condition. Unavailabilities are illustrative.

Table 34.5: Electrical assumptions of the illustrative facility.

Element	Assumed value	Symbol
Utility source	1 % snapshot unavailability	$q_U = 0.01$
Generator successfully becomes usable on utility loss	95 %	$p_G = 0.95$
UPS bridges the transition, conditional on generator usability	98 %	$p_B = 0.98$
Distribution path A	0.5 % unavailability	$q_A = 0.005$
Distribution path B	0.5 % unavailability	$q_B = 0.005$
Shared electrical infrastructure	0.02 % unavailability	$q_{\text{shared},E} = 0.0002$

or 99.9085 %, an expected $8760 \times 0.00091484 = 8.01$ hours per year without electrical support. The electrical availability is limited principally by the assumed source-backup reliability: the source contributes 0.00069 of the 0.00091484 unavailability, 75 %, while the redundant distribution paths contribute less than 3 %.

34.4.5 Cooling availability given electrical support

Table 34.6 lists the cooling assumptions. Its columns are the element, its configuration, the unavailability of one unit and the unavailability of the resulting subsystem. For two independent chillers, either sufficient, $q_{\text{chillers}} = 0.03^2 = 0.0009$; for two independent pumps, either sufficient, $q_{\text{pumps}} = 0.02^2 = 0.0004$; and the shared header has $q_{\text{header}} = 0.0005$.

Table 34.6: Cooling assumptions of the illustrative facility.

Element	Configuration	Unavailability	Subsystem
Chillers	Two independent, either sufficient	3 % per chiller	$q_{\text{chillers}} = 0.0009$
Pumps	Two independent, either sufficient	2 % per pump	$q_{\text{pumps}} = 0.0004$
Common chilled-water header	Shared element	0.05 %	$q_{\text{header}} = 0.0005$

Assuming that these three subsystem failure mechanisms are independent, and that the cooling equipment has a usable electrical supply, the cooling support is again a series of three subsystems, and

$$A_{C|E} = (1 - q_{\text{chillers}})\,(1 - q_{\text{pumps}})\,(1 - q_{\text{header}}), \tag{34.15}$$

so that

$$A_{C|E} = 0.9991 \times 0.9996 \times 0.9995 = 0.99820101,$$

or 99.8201 %. This is the modeled probability that the cooling equipment and its distribution are functionally adequate given the necessary electrical support.

Three features of these assumptions connect them with the rest of the book. First, the case study's chillers are each sized for the full load, a one-out-of-two arrangement, unlike the running plant of Section 34.3.2, whose 500 kW chillers are two out of three; with that plant at $a = 0.97$, q_{chillers} would be 0.002646 instead of 0.0009, and $A_{C|E}$ would fall to 0.99645658. Second, the pump unavailability 0.02 is, rounded, the steady-state unavailability $1 - 0.980392 = 0.019608$ of one pump of Chapter 20 of Volume I, with failure rate $0.001\,\text{h}^{-1}$ and a mean time to restore of 20 h. Treating the two pump states as independent corresponds to that chapter's model with two repair crews, whose unavailability $(0.019608)^2 = 0.00038447$ rounds to the 0.0004 used here; with one crew the two repairs queue, the pump states are dependent, and the unavailability is 0.00076864, equation (20.12). Third, the case study keeps the pump controls independent; the shared control failure of Section 34.4.2 would raise q_{pumps} to 0.005398, more than all the rest of the cooling unavailability.

34.4.6 Integrated availability and deficiency time

Under these assumptions the instantaneous infrastructure-support condition is $Y = E \wedge C$, and by equation (34.12)

$$A_{\text{infrastructure}} = A_E\, A_{C|E}, \tag{34.16}$$

so that

$$A_{\text{infrastructure}} = 0.99908516 \times 0.99820101 = 0.99728782,$$

or 99.7288 %: integrated snapshot infrastructure support under the illustrative component probabilities and independence assumptions. The result is lower than the availability of either subsystem alone, because both electricity and active cooling must be available to satisfy this strict instantaneous definition.

Under a stationary interpretation, equation (34.4) gives the expected annual infrastructure-deficiency time,

$$\mathbb{E}[D_{\text{infra}}] = 8760 \times (1 - A_{\text{infrastructure}}) = 8760 \times 0.00271218 = 23.76\,\text{h per year.}$$

It divides into the time without electrical support, $8760 \times (1 - A_E) = 8.01\,\text{h}$, and the time with electrical support but without adequate cooling, $8760 \times A_E\,(1 - A_{C|E})$. The conditional cooling deficiency alone, $8760 \times (1 - A_{C|E}) = 8760 \times 0.00179899 = 15.76\,\text{h}$ per year, exceeds the second part by the overlap $8760 \times (1 - A_E)(1 - A_{C|E}) = 0.014\,\text{h}$ of hours in which both would be deficient, so that $8.014 + 15.759 - 0.014 = 23.76\,\text{h}$. The cooling side carries two thirds of the deficiency time. Had the cooling availability been taken unconditionally, as $\Pr(C) = A_E\,A_{C|E}$, and multiplied by A_E, the result would have been $A_E^2 A_{C|E} = 0.99637546$, or 31.75 h per year: the double counting of Section 34.4.1 adds eight hours.

The 23.76 h is not yet expected IT downtime. It is the expected time during which the simplified electrical-and-active-cooling infrastructure condition is not satisfied. An active-cooling deficiency may be short enough for IT temperatures to remain acceptable, and the electrical systems may offer ride-through that the snapshot model does not represent. Those dynamic effects require another layer of analysis, which the next section supplies.

34.5 Thermal ride-through and the coupled event tree with restoration

This section computes how long the IT environment of the case study stays within its thermal limit after a complete loss of active cooling, builds the event tree of cooling restoration, and derives from it the probability, the expected duration and the expected unserved energy of IT-service interruptions.

34.5.1 Cooling loss and the rate of temperature rise

Suppose active cooling is lost while electrical power to the IT equipment remains available. Heat generation continues at $\dot{Q}_{\text{IT}} + \dot{Q}_{\text{other}} = 900\,\text{kW}$. Assume that residual passive heat removal or thermal exchange (conduction through the building envelope, mixing with adjacent spaces, the remaining air circulation) continues at $\dot{Q}_{\text{res}} = 300\,\text{kW}$, so that $\dot{Q}_{\text{cool}} = \dot{Q}_{\text{res}}$ in equation (34.5). The net rate of heat accumulation is

$$\dot{Q}_{\text{net}} = 900 - 300 = 600\,\text{kW},$$

and the aggregate temperature evolution is $C_{\text{eff}}\,\mathrm{d}T_{\text{IT}}/\mathrm{d}t = 600\,\text{kW}$. With an effective thermal capacitance $C_{\text{eff}} = 120\,\text{MJ/K}$, and since $600\,\text{kW} = 0.6\,\text{MJ/s}$,

$$\frac{\mathrm{d}T_{\text{IT}}}{\mathrm{d}t} = \frac{\dot{Q}_{\text{net}}}{C_{\text{eff}}} = \frac{0.6\,\text{MJ/s}}{120\,\text{MJ/K}} = 0.005\,\text{K/s}, \tag{34.17}$$

or 0.30°C per minute.

34.5.2 Time to the thermal limit

Suppose the initial representative temperature is $T_{\text{IT}}(0) = 25°\text{C}$ and the operational limit is $T_{\max} = 33°\text{C}$. These values are chosen for the example and are not universal limits of any

thermal guideline. The available temperature rise is $\Delta T_{\max} = T_{\max} - T_{\mathrm{IT}}(0) = 8\,\mathrm{K}$. Under the constant net-heat assumption, integrating equation (34.17) from 0 to t gives the straight line $T_{\mathrm{IT}}(t) = T_{\mathrm{IT}}(0) + (\dot{Q}_{\mathrm{net}}/C_{\mathrm{eff}})\,t$, and setting $T_{\mathrm{IT}}(t) = T_{\max}$ gives the time to the thermal limit,

$$t_{\mathrm{crit}} = \frac{C_{\mathrm{eff}}\,\Delta T_{\max}}{\dot{Q}_{\mathrm{net}}}. \tag{34.18}$$

For the case study,

$$t_{\mathrm{crit}} = \frac{120\,\mathrm{MJ/K} \times 8\,\mathrm{K}}{0.6\,\mathrm{MJ/s}} = 1600\,\mathrm{s} = 26.67\,\mathrm{min}.$$

The numerator is the thermal headroom, 960 MJ or 266.7 kWh of heat that the room can absorb before reaching its limit, and the denominator is the rate at which it is used. The expression has the same structure as the depletion time of the UPS batteries in Chapter 15 of Volume I, $T_{\mathrm{dep}} = (E_{\mathrm{batt}}(0) - E_{\min})/P_{\mathrm{load}}$, equation (15.20): a stored energy margin divided by a net drain.

Figure 34.6 plots the temperature rise. The horizontal axis is the time since the complete loss of active cooling, in minutes; the vertical axis is the representative IT temperature. The solid blue line is the case study, rising at 0.30°C per minute from 25°C; the dashed red line is the operational limit of 33°C; and the dot where they meet marks $t_{\mathrm{crit}} = 26.67$ minutes. The dashed green line, which rises at 0.20°C per minute and meets the limit at 40 minutes, is a larger thermal capacitance studied in Section 34.6.3. A cooling outage of 10 minutes ends at 28°C and is tolerable in this simplified model; an outage of 40 minutes would carry the temperature to 37°C and exceed the limit unless the load or the thermal conditions changed.

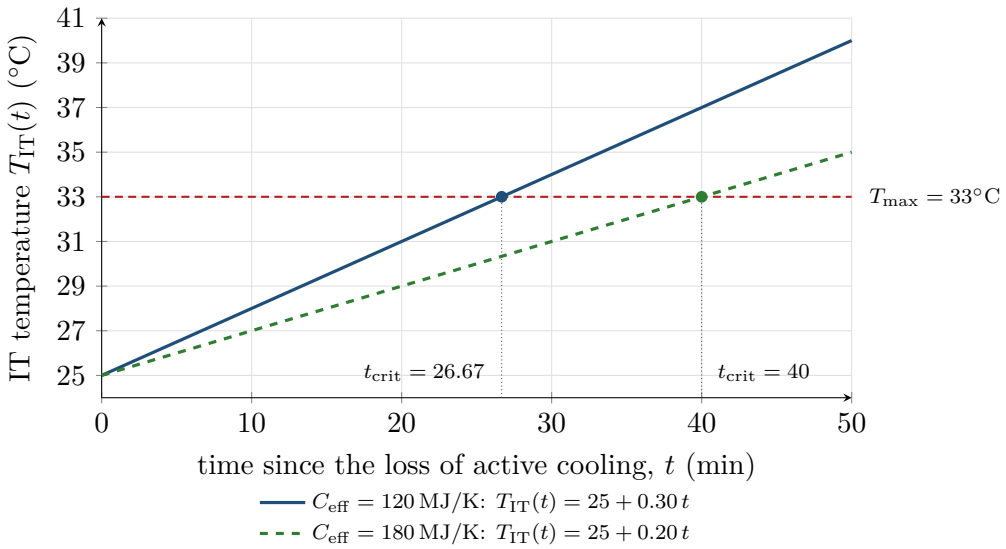


Figure 34.6: Illustrative temperature rise after a complete loss of active cooling, $T_{\mathrm{IT}}(t) = T_{\mathrm{IT}}(0) + (\dot{Q}_{\mathrm{net}}/C_{\mathrm{eff}})\,t$ with $T_{\mathrm{IT}}(0) = 25$°C and $\dot{Q}_{\mathrm{net}} = 600\,\mathrm{kW}$, for $C_{\mathrm{eff}} = 120\,\mathrm{MJ/K}$ (0.30°C per minute) and 180 MJ/K (0.20°C per minute), against the operational limit $T_{\max} = 33$°C. Constant net heat accumulation and a uniform effective temperature are assumed; real rack-inlet temperatures may rise nonuniformly and more rapidly.

34.5.3 Why thermal ride-through must be modeled dynamically

Under the simplified model in which a threshold crossing directly triggers an IT-service interruption, the interruption occurs when the restoration of cooling takes longer than the time to the limit. With T_{R} the restoration time, measured from the loss of cooling,

$$\Pr(\text{IT thermal interruption} \mid \text{cooling loss}) = \Pr(T_{\mathrm{R}} > t_{\mathrm{crit}}) = 1 - \Pr(T_{\mathrm{R}} \le t_{\mathrm{crit}}), \tag{34.19}$$

the complement of the maintainability function of Chapter 3 of Volume I, equation (3.13), evaluated at the time available. This is the most important equation of the chapter. The probability depends on two things at once: the restoration-time distribution, which is a property of the repair organization, the spares and the controls; and the time-dependent thermal state, which fixes t_{crit} through the heat load, the residual removal and the thermal capacitance. For example, if T_{R} were exponential with a mean of 30 minutes, the probability would be $\mathrm{e}^{-26.67/30} = 0.411$.

Two data centers with identical chiller availability can therefore have different IT-service availability, because their thermal inertia, airflow, heat density and restoration capabilities differ. This is why simply multiplying power and cooling equipment availabilities, as in Section 34.4.6, is not always sufficient: the product counts every hour of cooling deficiency as service lost, whatever the time each deficiency leaves.

The constant net-heat model is itself the simplest dynamic model. If the residual removal grows with the temperature, as conduction to the surroundings does, linearize it as $\dot{Q}_{\mathrm{res}}(T_{\mathrm{IT}}) = \dot{Q}_{\mathrm{res}} + k\,(T_{\mathrm{IT}} - T_{\mathrm{IT}}(0))$ with a coefficient k in kW/K. Equation (34.5) then becomes the first-order model of Chapter 28, equation (28.10), with time constant $\tau_{\mathrm{th}} = C_{\mathrm{eff}}/k$ and steady temperature $T_\infty = T_{\mathrm{IT}}(0) + \dot{Q}_{\mathrm{net}}/k$, and the time to the limit is equation (28.12),

$$t_{\mathrm{crit}} = \tau_{\mathrm{th}}\, \ln \frac{T_\infty - T_{\mathrm{IT}}(0)}{T_\infty - T_{\max}}. \tag{34.20}$$

As $k \to 0$, $\ln[1/(1-x)] \to x$ with $x = k\,\Delta T_{\max}/\dot{Q}_{\mathrm{net}}$, and equation (34.20) tends to equation (34.18). With an illustrative $k = 25\,\mathrm{kW/K}$, $\tau_{\mathrm{th}} = 120\,000/25 = 4800\,\mathrm{s} = 80\,\mathrm{min}$ and $T_\infty - T_{\mathrm{IT}}(0) = 600/25 = 24\,\mathrm{K}$, so $t_{\mathrm{crit}} = 80\ln(24/16) = 32.44$ minutes, against 26.67 for constant net heat. The constant net-heat model is conservative when the residual removal grows with temperature, and it is not conservative when local hot spots heat faster than the aggregate state; only measured transients can tell which applies.

34.5.4 The cooling-loss event tree

Let the initiating event I_C be a complete loss of active cooling, occurring at the assumed frequency

$$\nu_C = 0.5\,\mathrm{yr}^{-1},$$

the initiating-event frequency f_{IE} of Chapter 15 for this event. That chapter defined it as the expected number of occurrences per year and showed that a sequence's frequency is the initiating-event frequency times the sequence's conditional probability, $f_{S_k} = f_{\mathrm{IE}}\,\Pr(S_k \mid I)$, equation (15.7). This is an event-frequency model, separate from the snapshot unavailability model of Section 34.4. Assume that the electrical supply to the IT equipment remains available during these cooling-loss events, and that the restoration time takes three possible values: 10 minutes with probability 0.60, 40 minutes with probability 0.30, and 120 minutes with probability 0.10. The probabilities sum to one. The thermal model's threshold-crossing time is taken as exactly $t_{\mathrm{crit}} = 26.67$ minutes.

Figure 34.7 draws the event tree. The initiating event is the box on the left. The one heading is the restoration of cooling, with three branches, one per restoration time, each labeled with its time and conditional probability. The next column states whether the restoration time exceeds t_{crit}. To the right of the tree, each of the three sequences carries its number k, its end state (green when the IT service is maintained, red when it is interrupted), its conditional probability $\Pr(S_k \mid I_C)$, its frequency f_{S_k} and the IT interruption duration $D_{\mathrm{IT},k}$ it causes.

Cooling restored within 10 minutes does not cross the modeled limit; restoration at 40 or 120 minutes does. Since the three sequences are mutually exclusive, by equation (34.19)

$$\Pr(\text{thermal interruption} \mid I_C) = \Pr(T_{\mathrm{R}} > 26.67\,\mathrm{min}) = 0.30 + 0.10 = 0.40.$$

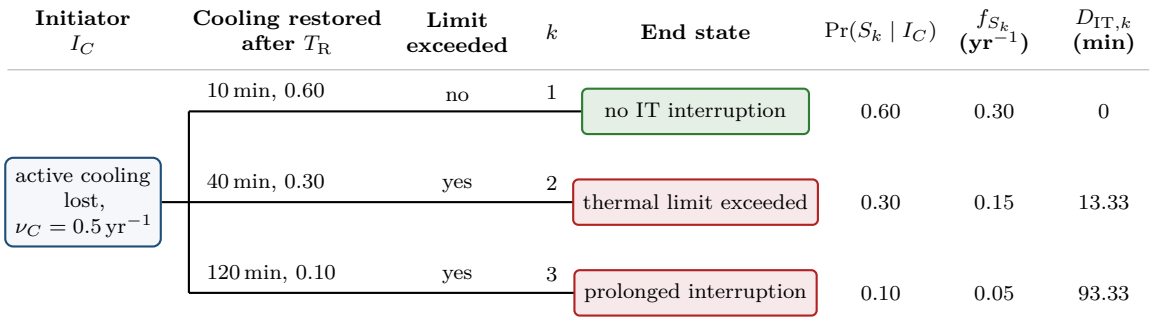


Figure 34.7: Event tree for a complete loss of active cooling with the IT load powered, $\nu_C = 0.5\,\mathrm{yr}^{-1}$ and $t_{\mathrm{crit}} = 26.67\,\mathrm{min}$ (illustrative). The restoration heading has three branches; the IT interruption lasts from the threshold crossing until cooling returns, $D_{\mathrm{IT},k} = (T_{\mathrm{R},k} - t_{\mathrm{crit}})^+$. The sequence probabilities sum to 1 and the frequencies to $0.5\,\mathrm{yr}^{-1}$.

Conditional on a complete active-cooling-loss event, the probability of crossing the threshold is 40 %. The expected annual frequency of thermal-interruption events is

$$\nu_{\mathrm{thermal}} = \nu_C \times 0.40 = 0.5 \times 0.40 = 0.20\,\mathrm{yr}^{-1},$$

under a stationary marked-event model, in which the cooling losses form a Poisson process and each carries an independent restoration time as its mark, and on the assumption that events do not overlap. The thinning property used in Section 15.2.1 makes the thermal-interruption events themselves a Poisson process with rate $0.20\,\mathrm{yr}^{-1}$.

34.5.5 Expected interruption duration and unserved IT energy

Suppose the IT service is interrupted immediately upon the threshold crossing and is restored when active cooling returns, ignoring for now any additional restart and thermal recovery time. The interruption of sequence k then lasts $D_{\mathrm{IT},k} = (T_{\mathrm{R},k} - t_{\mathrm{crit}})^+$, where $(x)^+ = \max(x, 0)$. For a 40-minute cooling outage, $D_{\mathrm{IT},2} = 40 - 26.67 = 13.33$ minutes, and for a 120-minute outage $D_{\mathrm{IT},3} = 120 - 26.67 = 93.33$ minutes. By the law of total expectation over the three sequences, the expected IT interruption duration per cooling-loss event is

$$\mathbb{E}[D_{\mathrm{IT}} \mid I_C] = \sum_k \Pr(S_k \mid I_C)\, D_{\mathrm{IT},k} = 0.60 \times 0 + 0.30 \times 13.33 + 0.10 \times 93.33 = 13.33\,\mathrm{min}, \quad (34.21)$$

the sum of 4.00 and 9.33 minutes; the equality with $D_{\mathrm{IT},2}$ is a coincidence of the numbers. Given that an interruption occurs, its expected duration is $13.33/0.40 = 33.33$ minutes. The expected annual interruption duration is the frequency times the expected duration per event, by the random-sum argument of Section 15.3.5,

$$\mathbb{E}[D_{\mathrm{IT,yr}}] = \nu_C\, \mathbb{E}[D_{\mathrm{IT}} \mid I_C] = 0.5\,\mathrm{yr}^{-1} \times \frac{13.33}{60}\,\mathrm{h} = 0.1111\,\mathrm{h\ per\ year},$$

approximately 6.67 expected minutes of IT-service interruption per year from the modeled cooling-loss events.

Assume the full IT load $P_{\mathrm{IT}} = 800\,\mathrm{kW}$ is interrupted after the threshold crossing. The unserved IT energy of sequence k is $E_k^{\mathrm{unserved}} = P_{\mathrm{IT}}\, D_{\mathrm{IT},k}$, as in equation (15.13) with a constant interrupted load, and the expected unserved IT energy per cooling-loss event is

$$\mathbb{E}[E^{\mathrm{unserved}} \mid I_C] = P_{\mathrm{IT}}\, \mathbb{E}[D_{\mathrm{IT}} \mid I_C] = 800\,\mathrm{kW} \times \frac{13.333}{60}\,\mathrm{h} = 177.78\,\mathrm{kWh}.$$

The annual expected unserved IT energy is, as equation (15.14), the initiating frequency times the expected energy per event,

$$\mathrm{EUE}_{\mathrm{IT}} = \nu_C \, \mathbb{E}[E^{\mathrm{unserved}} \mid I_C] = 0.5 \times 177.78 = 88.89\,\mathrm{kWh\ per\ year}. \tag{34.22}$$

Table 34.7 collects the sequence-by-sequence numbers. Its columns are the sequence, the restoration time, the conditional probability and the frequency of the sequence, whether the limit is exceeded, the IT interruption duration and its probability-weighted contribution, and the unserved energy and its contribution. The prolonged interruption, one tenth of the events, carries 70 % of the expected duration and of the expected unserved energy. This small value is not a general data center reliability estimate: it follows from the selected event frequency, restoration distribution, thermal storage and immediate-return model.

Table 34.7: Sequences of the cooling-loss event tree, $\nu_C = 0.5\,\mathrm{yr}^{-1}$, $t_{\mathrm{crit}} = 26.67\,\mathrm{min}$ and $P_{\mathrm{IT}} = 800\,\mathrm{kW}$ (illustrative).

k	$T_{\mathrm{R},k}$ (min)	$\Pr(S_k \mid I_C)$	f_{S_k} (yr^{-1})	Limit exceeded	$D_{\mathrm{IT},k}$ (min)	$\Pr \cdot D$ (min)	E_k^{unserved} (kWh)	$\Pr \cdot E$ (kWh)
1	10	0.60	0.30	no	0	0	0	0
2	40	0.30	0.15	yes	13.33	4.00	177.78	53.33
3	120	0.10	0.05	yes	93.33	9.33	1244.44	124.44
	Total	1.00	0.50			13.33		177.78

34.5.6 Cooling outage against IT interruption

The expected duration of a cooling-loss event is

$$\mathbb{E}[T_{\mathrm{R}}] = 0.60 \times 10 + 0.30 \times 40 + 0.10 \times 120 = 30\,\mathrm{min},$$

and with $\nu_C = 0.5\,\mathrm{yr}^{-1}$ the expected active-cooling downtime from this event model is $0.5 \times 0.5\,\mathrm{h} = 0.25\,\mathrm{h}$ per year. The resulting expected IT interruption is only 0.1111 h per year, 44 % of it, because some cooling outages are shorter than the modeled thermal ride-through, and the longer ones are shortened by t_{crit}. This is the central distinction of the chapter,

$$\text{cooling equipment unavailable} \neq \text{IT service unavailable}.$$

A further qualification is essential. The snapshot model of Section 34.4.6 implies 15.76 h per year of cooling-infrastructure deficiency conditional on electrical support, $8760 \times (1 - A_{C|E})$, within 23.76 h per year of integrated infrastructure deficiency; the event model of this section implies 0.25 h per year of complete active-cooling loss. The two are different hypothetical reliability models, sixty-three times apart, and they must not be combined or assumed mutually consistent without calibrating outage frequencies and durations to the same underlying stochastic process. A consistent model would obtain each snapshot unavailability as a frequency times a mean down time, the relation $\mathbb{E}[D_{\mathrm{yr}}] = f_{\mathrm{F}} \cdot \mathrm{MDT}$ of equation (20.17) in Chapter 20 of Volume I: 15.76 h per year of cooling deficiency made of 30-minute events would require 31.5 events per year, not 0.5. Only then can the deficiency hours be passed through the thermal ride-through to give IT downtime.

Common mistake

Reading the expected hours of equipment deficiency as expected IT downtime, or adding hours from a snapshot model to hours from an event model. Deficiency becomes interruption only after the thermal and electrical ride-through is applied to each deficiency's duration, and that requires one model that gives both the frequency and the duration of the deficiencies.

34.6 Reliability importance, engineering decisions and a unified physical–probabilistic model

This section ranks the failure mechanisms of the case study, distinguishes failure importance from risk importance, evaluates two decisions (more thermal ride-through, and a maintenance configuration), and states the unified model that the earlier sections approximate, with the evidence needed to trust it.

34.6.1 Contributions of the failure mechanisms

The integrated infrastructure model contains six failure mechanisms: source-backup failure, loss of both distribution paths, shared electrical infrastructure failure, loss of both chillers, loss of both pumps, and shared chilled-water header failure. They are not equally important. Table 34.8 lists them. Its columns are the mechanism, its unavailability q_i, its share of the sum of the six, the exact reduction of the infrastructure unavailability if the mechanism were eliminated, ΔQ_i^- (defined in Section 34.6.2), the corresponding hours per year, and the effect of the mechanism on the IT service.

Table 34.8: Contributions of the six failure mechanisms to the infrastructure unavailability of the case study (illustrative).

Failure mechanism	q_i	Share	ΔQ_i^-	h/yr	Effect on IT service
Source backup unavailable	0.000690	25.4 %	0.00068860	6.03	immediate interruption
Both electrical paths unavailable	0.000025	0.9 %	0.00002493	0.22	immediate interruption
Shared electrical infrastructure	0.000200	7.4 %	0.00019950	1.75	immediate interruption
Both chillers unavailable	0.000900	33.1 %	0.00089837	7.87	after the ride-through
Both pumps unavailable	0.000400	14.7 %	0.00039907	3.50	after the ride-through
Shared chilled-water header	0.000500	18.4 %	0.00049889	4.37	after the ride-through
Sum (rare-event approximation)	0.002715	100 %			

For the rare-event approximation of Chapter 11 of Volume I, the sum of the six probabilities, 0.002715, approximates the infrastructure unavailability 0.00271218; the difference, 2.8×10^{-6}, is the overlap of mechanisms that are active at the same time. The entries are individual basic or subsystem probabilities, not probabilities of mutually exclusive states, and their sum approximates the total only when the overlap terms are small. The greatest single contribution comes from the dual-chiller subsystem, 33 %.

Choosing an improvement by that ranking alone can still mislead. A chiller failure may be restored quickly and covered by thermal storage, while a shared electrical failure may interrupt the IT service immediately, as the last column of the table records. Interventions must therefore be compared by their consequences for the IT service, not only by the probabilities of the failed equipment. Whether thermal storage helps a mechanism depends on the duration of its deficiencies relative to t_{crit}: a ride-through of 27 minutes covers a control trip that is reset in ten minutes, but not a pump repair of the 20 h assumed in Chapter 20 of Volume I.

34.6.2 Failure importance and risk importance

A component importance measure based on availability is

$$I_i^{(A)} = A_{\mathrm{system}}(q_i = 0) - A_{\mathrm{system}}(q_i), \tag{34.23}$$

the increase in modeled system availability from eliminating the unavailability of mechanism i. Since $A_{\mathrm{system}} = 1 - Q_{\mathrm{sys}}$, it equals the reduction $\Delta Q_i^{-} = Q_{\mathrm{sys}} - Q(0_i)$ of equation (12.17) in Chapter 12 of Volume I, which that chapter showed to be $q_i\, I_i^{\mathrm{B}}$, the probability of the mechanism times its Birnbaum importance. For the series-of-mechanisms model of the case study, $A_{\mathrm{system}} = \prod_j (1 - q_j)$, so $Q(0_i) = 1 - A_{\mathrm{system}}/(1 - q_i)$ and

$$\Delta Q_i^{-} = A_{\mathrm{system}}\, \frac{q_i}{1 - q_i};$$

for the chillers, $0.99728782 \times 0.0009/0.9991 = 0.00089837$, the value in Table 34.8.

An economically useful measure compares expected losses instead. With $\mathcal{R}$ the expected annual loss of the IT service (R$ per year), and $\mathcal{R}(a_i)$ the same quantity after improvement a_i, the risk importance of the improvement is the risk reduction

$$\Delta\mathcal{R}(a_i) = \mathcal{R} - \mathcal{R}(a_i). \tag{34.24}$$

It can account for what the availability measure ignores: the different downtime durations of different mechanisms, the different amounts of interrupted IT load, the thermal ride-through, repair costs, operational disruption, and common-cause dependencies. Improving the reliability of a shared electrical bus may therefore be more valuable than adding another cooling pump, even though the pump subsystem has the larger failure probability in Table 34.8 (0.0004 against 0.0002): the bus failure interrupts the whole IT load at once, possibly with a restart time after power returns, while a short pumping deficiency may be absorbed by the thermal ride-through.

34.6.3 Improving thermal ride-through

Return to the event tree of Section 34.5.4, with $t_{\mathrm{crit}} = 26.67$ minutes. Suppose additional thermal buffering raises the effective capacitance from 120 to 180 MJ/K while the net heat accumulation remains 600 kW. By equation (34.18),

$$t_{\mathrm{crit}} = \frac{180 \times 8}{0.6} = 2400\,\mathrm{s} = 40\,\mathrm{min}.$$

In the discrete restoration model, a 40-minute outage now reaches the limit exactly at restoration. Under the idealized convention that restoration at the threshold avoids interruption, only the 120-minute branch creates IT downtime, and the expected interruption duration per cooling event becomes

$$\mathbb{E}[D_{\mathrm{IT}} \mid I_C] = 0.10 \times (120 - 40) = 8\,\mathrm{min},$$

the annual expected interruption $0.5 \times 8/60 = 0.0667$ h per year, against the original 0.1111 h, a reduction of $(0.1111 - 0.0667)/0.1111 = 40\,\%$, without any change in the chiller or pump failure

probabilities. Under the same convention the frequency of thermal interruptions falls from 0.20 to $0.05\,\mathrm{yr}^{-1}$, by 75 %. The two conventions differ only in whether the 40-minute branch counts as an interruption of zero duration; the duration result does not depend on the convention, the frequency result does.

Table 34.9 evaluates the event tree for four effective capacitances, with the restoration distribution, ν_C, the 8-kelvin headroom, the 600 kW net heat and the 800 kW IT load unchanged. Its columns are the capacitance, the time to the limit, the probability of a thermal interruption per event, the frequency of thermal interruptions, the expected interruption duration per event, the expected annual interruption and the annual expected unserved IT energy. For $C_{\mathrm{eff}} = 60\,\mathrm{MJ/K}$ the row reads $t_{\mathrm{crit}} = 60 \times 8/0.6 = 800\,\mathrm{s} = 13.33$ minutes, $\mathbb{E}[D_{\mathrm{IT}} \mid I_C] = 0.30 \times 26.67 + 0.10 \times 106.67 = 18.67$ minutes, $0.5 \times 18.67/60 = 0.1556\,\mathrm{h}$ per year and $0.1556 \times 800 = 124.44\,\mathrm{kWh}$ per year. The trend is stepwise: the probability of interruption changes only when t_{crit} passes one of the restoration times, while the expected duration falls continuously, by $\nu_C \Pr(T_{\mathrm{R}} > t_{\mathrm{crit}})$ hours per year for each hour added to t_{crit}.

Table 34.9: The cooling-loss event tree for four effective thermal capacitances, $\dot{Q}_{\mathrm{net}} = 600\,\mathrm{kW}$, $\Delta T_{\max} = 8\,\mathrm{K}$, $\nu_C = 0.5\,\mathrm{yr}^{-1}$, restoration after 10, 40 or 120 minutes with probabilities 0.6, 0.3, 0.1, and restoration at the threshold counted as no interruption (illustrative).

C_{eff} (MJ/K)	t_{crit} (min)	$\Pr(T_{\mathrm{R}} > t_{\mathrm{crit}})$	ν_{thermal} (yr^{-1})	$\mathbb{E}[D_{\mathrm{IT}} \mid I_C]$ (min)	$\mathbb{E}[D_{\mathrm{IT,yr}}]$ (h/yr)	$\mathrm{EUE}_{\mathrm{IT}}$ (kWh/yr)
60	13.33	0.40	0.20	18.67	0.1556	124.44
120	26.67	0.40	0.20	13.33	0.1111	88.89
180	40.00	0.10	0.05	8.00	0.0667	53.33
240	53.33	0.10	0.05	6.67	0.0556	44.44

This is an important form of reliability improvement: it changes the physical consequences of an equipment failure rather than the likelihood of the failure. In practice, a design that targets exactly the predicted threshold, as the 180 MJ/K row does for the 40-minute restoration, provides no margin for model or sensor uncertainty, and a real design would require a more conservative assessment, with the uncertainty in C_{eff}, $\dot{Q}_{\mathrm{net}}$ and the restoration times propagated to the result.

Principle 34.2. *A reliability improvement can act on the probability of failure, on the duration of restoration, or on the physical consequence of the failure. Raising the thermal ride-through acts on the consequence, and it can reduce IT interruption as much as an equipment upgrade without changing any equipment failure probability.*

34.6.4 Maintenance configurations and common-cause risk

Suppose one chiller is deliberately taken offline for maintenance. If the second chiller is the only remaining unit able to meet the cooling demand, redundancy is temporarily lost. Let the remaining chiller have the failure rate $\lambda = 0.001\,\mathrm{h}^{-1}$, and let the maintenance last $T_{\mathrm{M}} = 8\,\mathrm{h}$. Assuming an exponential lifetime and an initially healthy unit, the probability that the operating chiller fails before the maintenance ends is the exponential distribution function at T_{M},

$$\Pr(F) = 1 - \mathrm{e}^{-\lambda T_{\mathrm{M}}} = 1 - \mathrm{e}^{-0.008} = 0.0079681,$$

about 0.797 %, close to the first-order value $\lambda T_{\mathrm{M}} = 0.008$. In the running plant of three 500 kW chillers, maintenance of one leaves two units that are both required, a series system, and the probability of losing capacity during the window is $1 - \mathrm{e}^{-2 \times 0.008} = 0.0159$, twice as large.

This probability concerns the loss of active chiller capacity during maintenance. It is not automatically the probability of an IT interruption: whether the IT service is lost also depends on the thermal buffering, the restoration, any alternative cooling and the duration of the resulting deficiency, through equation (34.19). Maintenance also creates its own common causes. The maintenance dependency of Chapter 13 of Volume I, the same procedure, crew or error applied to redundant units, is a coupling factor that can disable the unit left in service, for example through an isolation valve left in the wrong position on the shared header. Planned maintenance should therefore be evaluated as a temporary configuration with its own reliability block diagram, capacity constraints and dependencies; a concurrently maintainable design is one in which that configuration still meets the demand.

34.6.5 The hybrid system state

The models of this chapter are snapshots of a richer process. A more advanced model defines the system state

$$\mathbf{x}(t) = \begin{bmatrix} \mathbf{x}_E(t) \\ \mathbf{x}_C(t) \\ T_{\mathrm{IT}}(t) \\ T_{\mathrm{water}}(t) \\ E_{\mathrm{UPS}}(t) \end{bmatrix},$$

where $\mathbf{x}_E$ collects the electrical equipment states and operating variables, $\mathbf{x}_C$ the cooling equipment states and operating variables, T_{IT} is the representative IT environmental temperature, T_{water} the chilled-water temperature and E_{UPS} the energy remaining in the UPS batteries. Discrete events change the configuration: a generator trip, a pump failure, a breaker operation, a chiller restoration. Continuous dynamics act between them: thermal energy accumulates in the room, the water temperatures drift, the batteries discharge.

The system is naturally modeled as a hybrid stochastic system, the hybrid dynamical system of Chapter 28 with random events. Between events the continuous state follows

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}, s), \tag{34.25}$$

where s is the discrete equipment configuration (which units are available), $\mathbf{u}$ the control inputs (setpoints, staging, transfers) and $\boldsymbol{\theta}$ the physical parameters, such as C_{eff} and the heat-transfer coefficients; the configuration changes by stochastic transitions whose rates may depend on the state. Equation (34.5) is one row of $\mathbf{f}$, and the battery balance $\mathrm{d}E_{\mathrm{batt}}/\mathrm{d}t = -P_{\mathrm{load}}$ of Chapter 15 of Volume I, equation (15.19), another. The vector field $\mathbf{f}$ is not the failure density f, and $\boldsymbol{\theta}$ is the parameter vector of Chapter 31.

34.6.6 The reliability event and chronological measures

The required IT service is defined by conditions on the hybrid state: the delivered power meets the required power,

$$P_{\mathrm{delivered}}(t) \geq P_{\mathrm{IT,required}}(t),$$

the inlet temperature at each critical IT location j stays within its limit,

$$T_{\mathrm{inlet},j}(t) \leq T_{\mathrm{max},j},$$

and any other required operating conditions hold. The service indicator of equation (34.1) is then the indicator of all of them,

$$Y(t) = \mathbf{1}\{\text{all required IT-service conditions are satisfied at } t\}, \qquad A_{\mathrm{IT}}(t) = \mathbb{E}[Y(t)]. \tag{34.26}$$

For chronological scenarios k with probabilities p_k, each a trajectory of the hybrid system over $[0, H]$, the expected IT-service interruption duration is

$$\mathbb{E}[D_H] = \sum_k p_k \int_0^H [1 - Y_k(t)]\,\mathrm{d}t, \tag{34.27}$$

and if each scenario specifies its interrupted IT power $P_{\text{unsupplied},k}(t)$,

$$\mathrm{EUE}_{\mathrm{IT}} = \sum_k p_k \int_0^H P_{\text{unsupplied},k}(t)\,\mathrm{d}t. \tag{34.28}$$

In a Monte Carlo simulation the scenarios are sampled and $p_k = 1/N$.

The earlier models are special cases. The event tree of Section 34.5.4 has three scenarios per cooling-loss event, with $p_k = 0.6$, 0.3 and 0.1 and $Y_k(t) = 0$ between t_{crit} and $T_{\mathrm{R},k}$; equation (34.27) then returns $\int [1 - Y_k]\,\mathrm{d}t = (T_{\mathrm{R},k} - t_{\text{crit}})^+$ and the 13.33 minutes per event of equation (34.21). The snapshot model of Section 34.4 is the case in which Y depends only on the instantaneous equipment configuration, with no thermal or electrical storage. Equations (34.27) and (34.28) are a more defensible system-level formulation than multiplying equipment availability factors whenever physical coupling and event timing matter.

34.6.7 Model credibility

A numerical result is useful only if the measurements and the model support it. Table 34.10 lists the questions a credibility assessment of a data center reliability model should examine, one per row, with the evidence each requires. The first three concern the structure of the model (capacity, independence, ride-through), the next two its parameters (outage and restoration probabilities), and the last three the instruments, the thermal model and the definition of the consequence.

Table 34.10: Credibility questions for a data center reliability model and the evidence each requires.

Question	Required evidence
Can one cooling pump meet the demand?	Pump curves, flow and pressure measurements
Is the redundancy genuinely independent?	Electrical single-line diagrams, cooling schematics, shared-cause analysis
Is the thermal ride-through realistic?	Temperature-response tests, airflow distribution, IT heat inventory
Are the outage probabilities credible?	Failure and maintenance records, appropriate fleet data
Are the restoration times credible?	Maintenance logs and repair-resource assumptions
Are the sensors trustworthy?	Calibration, synchronization, redundancy
Is the thermal model adequate?	Measured transients across representative loads
Are the consequences correctly defined?	IT workload requirements and actual shutdown criteria

The model should not imply precise thermal protection or service availability merely because it produces a numerical probability. A value of C_{eff} inferred from one measured transient, an outage probability taken from fleet data for different equipment, or a thermal limit taken from a guideline rather than from the installed equipment can each move the result by more than the

differences between the design options it is meant to rank. Chapter 35 turns the importance measures and risk reductions of this section into a formal choice among reliability investments under a limited budget.

34.7 What to retain

Definition 34.1 (Data center IT-service reliability)**.** The reliability of a data center is the probability that its required IT service, electrical power and an acceptable thermal environment for the specified workload, is supported; it is obtained from equipment reliability combined with the physical dependencies between power and cooling and with the timing of failure, thermal response and restoration, so that a cooling loss interrupts the service only with probability $\Pr(T_{\mathrm{R}} > t_{\mathrm{crit}})$.

Data center reliability is a service-level property, not a product of individual equipment reliabilities. N+1 and 2N architectures improve capacity redundancy, but their actual benefit depends on shared infrastructure and physical independence. Electrical and cooling infrastructure are coupled: shared power supplies, controls and distribution paths can defeat nominal redundancy, and common-cause failures can dominate system risk, which makes physical separation and control diversity important. Thermal inertia gives time to respond to cooling failures, and its benefit depends on the workload, the heat removal and the restoration dynamics. Snapshot availability and chronological interruption risk are different modeling problems and must not be combined without making them consistent. The best reliability improvement may reduce a failure probability, shorten restoration, or reduce the physical consequence. Model credibility requires physical testing, reliable instrumentation and an explicit operating envelope. In short, equipment reliability, physical dependencies and event timing together determine IT-service reliability. Table 34.11 lists the chapter's terms, their meaning and units, and the subsection that defines each.

Table 34.11: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Service indicator	required IT service supported	$Y(t) \in \{0,1\}$	–	34.1.1
IT-service availability	probability the service is supported	$A_{\mathrm{IT}}(t) = \Pr[Y(t)=1]$	–	34.1.1
Expected downtime	expected time without service	$\int_0^H [1 - A_{\mathrm{IT}}(t)]\,\mathrm{d}t$	h	34.1.1
Thermal balance	stored heat from gains and removal	$C_{\mathrm{eff}}\dot{T}_{\mathrm{IT}} = \dot{Q}_{\mathrm{IT}} + \dot{Q}_{\mathrm{other}} - \dot{Q}_{\mathrm{cool}}$	kW	34.1.3
Effective thermal capacitance	heat stored per kelvin	C_{eff}	MJ/K	34.1.3
Backup success probability	generator usable and UPS bridges	$p_{\mathrm{backup}} = p_G\, p_B$	–	34.2.2
Snapshot source unavailability	utility out and backup failed	$q_U(1 - p_{\mathrm{backup}})$	–	34.2.3
Redundant paths	either path carries the load	$1 - q_A q_B$	–	34.2.4
N+1 capacity (2-out-of-3)	two of three units required	$3a^2 - 2a^3$	–	34.3.2
2N redundancy	two full-capacity trains	$1 - q^2$	–	34.3.3
Coupled availability	both supports present	$\Pr(E)\,\Pr(C \mid E)$	–	34.4.1
Shared-cause unavailability	common control plus independent pair	$q_{\mathrm{ctrl}} + (1 - q_{\mathrm{ctrl}})q^2$	–	34.4.2
Infrastructure availability	electrical times cooling given power	$A_E\, A_{C\mid E}$	–	34.4.6
Time to the thermal limit	thermal headroom over net heat	$t_{\mathrm{crit}} = C_{\mathrm{eff}}\Delta T_{\max}/\dot{Q}_{\mathrm{net}}$	min	34.5.2
Thermal-interruption probability	restoration slower than ride-through	$\Pr(T_{\mathrm{R}} > t_{\mathrm{crit}})$	–	34.5.3
Expected interruption per event	sequence-weighted duration	$\sum_k \Pr(S_k \mid I_C)(T_{\mathrm{R},k} - t_{\mathrm{crit}})^+$	min	34.5.5
Annual unserved IT energy	frequency times energy per event	$\nu_C\, \mathbb{E}[E^{\mathrm{unserved}} \mid I_C]$	kWh/yr	34.5.5
Availability importance	gain from eliminating a mechanism	$\Delta Q_i^- = A_{\mathrm{system}}(q_i = 0) - A_{\mathrm{system}}$	–	34.6.2
Risk importance	reduction of expected annual loss	$\Delta\mathcal{R}(a_i) = \mathcal{R} - \mathcal{R}(a_i)$	R$/yr	34.6.2
Hybrid stochastic model	continuous state between random events	$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}, s)$	varies	34.6.5

Further reading. Uptime Institute, *Tier Classification System*, for redundancy, concurrent maintainability and fault tolerance; Uptime Institute, *Tier Standard: Topology*, for site infrastructure topology and performance confirmation; *ASHRAE Handbook*, *Data Centers and Telecommunication Facilities*, for equipment environmental requirements and cooling considerations; Billinton and Allan, *Reliability Evaluation of Engineering Systems*, for system reliability, dependencies and repairable-system modeling.

34.8 Exercises

Exercise 34.1. (a) State why cooling-equipment unavailability differs from IT-service unavailability, and what should define the system-level reliability metric of a data center. (b) A

smaller IT room has $C_{\text{eff}} = 90\,\text{MJ/K}$, a net heat accumulation of 450 kW after a complete loss of active cooling, and 6 K of headroom. Compute the rate of temperature rise and t_{crit}. (c) With the restoration distribution of Section 34.5.4 (10, 40 or 120 minutes with probabilities 0.6, 0.3, 0.1), $\nu_C = 0.5\,\text{yr}^{-1}$ and an interrupted load of 800 kW, compute the probability of thermal interruption per event, the expected interruption duration per event and per year, and the annual expected unserved IT energy. (Answers: (a) thermal storage and restoration time can prevent the temperatures from exceeding their limits; whether the required IT service is maintained, not whether every component operates. (b) $0.45/90 = 0.005\,\text{K/s}$, or 0.30°C per minute; $t_{\text{crit}} = 90 \times 6/0.45 = 1200\,\text{s} = 20\,\text{min}$. (c) 0.40; $0.30 \times 20 + 0.10 \times 100 = 16\,\text{min}$ per event and $0.5 \times 16/60 = 0.1333\,\text{h}$ per year; $0.1333 \times 800 = 106.67\,\text{kWh}$ per year.)

Exercise 34.2. (a) State the availability of two independent, fully redundant paths, each with unavailability q, and what N+1 redundancy describes. (b) Compute the availability of a 2-out-of-3 chiller plant with unit availability 0.95 and of a 2N arrangement whose trains each have unavailability 0.05. (c) A plant is enlarged to N+1 with $N = 3$: four chillers of which three are required, each with availability 0.97. Compute its availability and compare it with the 2-out-of-3 plant at the same unit availability. (Answers: (a) $1 - q^2$; enough capacity to satisfy the demand plus one additional unit. (b) $3 \times 0.9025 - 2 \times 0.857375 = 0.99275$; $1 - 0.05^2 = 0.9975$. (c) $a^4 + 4a^3(1-a) = 0.88529281 + 0.10952076 = 0.99481357$, below the 0.997354 of the 2-out-of-3 plant: with more units required, more pairs of failures defeat the plant, and the unavailability is close to $\binom{4}{2}(1-a)^2 = 0.0054$ instead of $3(1-a)^2 = 0.0027$.)

Exercise 34.3. (a) State why multiplying electrical and cooling availabilities can be incorrect, and why improving a shared control system can matter more than installing an additional pump. (b) Two pumps, either sufficient, each have unavailability 0.03, and a shared control failure of probability 0.002 disables both. Compute the unavailability of the pumping function with two pumps, with three pumps, and with two pumps after the control failure probability is reduced to 0.0005. (c) Rank the two improvements of part (b). (Answers: (a) the systems may share failure causes and operating dependencies, so that $\Pr(E \wedge C) = \Pr(E)\,\Pr(C \mid E)$, not $\Pr(E)\,\Pr(C)$; a shared control failure disables all the redundant pumps at once. (b) $0.002 + 0.998 \times 0.0009 = 0.0028982$; $0.002 + 0.998 \times 0.000027 = 0.0020269$; $0.0005 + 0.9995 \times 0.0009 = 0.0013996$. (c) The better control reduces the unavailability by 52 %, the third pump by 30 %.)

Exercise 34.4. (a) State what determines the thermal ride-through time in the simplified energy balance. (b) In the case study of Section 34.4, the shared chilled-water header is duplicated so that its unavailability falls from 0.0005 to 0.0001, all else unchanged. Compute $A_{C|E}$, the infrastructure availability and the expected annual deficiency time, and the reduction from 23.76 h per year. (c) A chiller with failure rate $0.002\,\text{h}^{-1}$ is left as the only unit able to meet the demand during a 12-hour maintenance. Compute the probability that it fails before the maintenance ends, exactly and with the first-order approximation. (Answers: (a) the thermal capacitance, the allowable temperature rise and the net heat accumulation, $t_{\text{crit}} = C_{\text{eff}} \Delta T_{\max} / \dot{Q}_{\text{net}}$. (b) $0.9991 \times 0.9996 \times 0.9999 = 0.99860049$; $0.99908516 \times 0.99860049 = 0.99768693$; $8760 \times 0.00231307 = 20.26\,\text{h}$ per year, 3.50 h less. (c) $1 - \text{e}^{-0.024} = 0.023714$, against 0.024.)

Chapter 35

Risk Importance and Decision Optimization

35.1 Component importance and the value of an intervention

Chapter 34 assessed the reliability of an integrated data center, its electrical and cooling infrastructure, its redundant equipment, its common causes and its thermal ride-through. One of its conclusions was that the component with the highest failure probability is not necessarily the component whose improvement reduces the system risk most. This chapter makes that observation quantitative and turns it into a decision method. An engineer with a limited budget may consider four kinds of investment: improving the availability of a frequently failing transmission line; increasing the rating of a transmission corridor; adding redundant pumping equipment; or installing sensors that reduce uncertainty before a larger investment is committed. Which of them is preferred depends on the objective. The option that maximizes system availability need not be the one that minimizes the expected unserved energy, and once investment costs are counted a third option may become optimal. The question of this chapter is

given limited engineering resources, which intervention produces the greatest credible improvement in system reliability or reduction in operational risk, and why the answer is not a ranking of components by importance.

The central distinction of the chapter is

$$\text{component importance} \neq \text{intervention value}. \tag{35.1}$$

A component may be important, and improving it may still be expensive, technically ineffective, or less beneficial than addressing a shared dependency or the consequences of its failure. The chapter develops the importance measures as sensitivities of a system model, the sensitivity of risk to failure and repair rates, the counterfactual evaluation of interventions on the three-bus transmission network of Chapter 27, the selection of an investment portfolio under a budget, the value of information against the value of an upgrade, and the effect of uncertainty on the rankings.

Jargon: importance measure, Birnbaum importance, RRW, RAW, elasticity, intervention, counterfactual, annualized cost, portfolio, EVPI, EVSI, CVaR

An *importance measure* ranks components by their contribution to a stated system risk metric. The *Birnbaum importance* I_i^{B} is the partial derivative of the system reliability with respect to the reliability of component i; the *risk reduction worth* RRW_i is the ratio of the baseline risk to the risk with component i made perfectly reliable; the *risk achievement worth* RAW_i is the ratio of the risk with component i certainly failed to the baseline risk (all three as in Chapter 12 of Volume I). An *elasticity* is a dimensionless sensitivity, the relative change of an output per relative change of an input. An *intervention* is an engineering action that sets some quantities of the physical system, and therefore of its model, to new values; a *counterfactual* calculation evaluates the model as it would be after the intervention, which has not been carried out. An *annualized cost* spreads a project's capital cost over its life as an equivalent yearly payment and adds its recurring incremental expenses. A *portfolio* is a set of projects selected together. The *expected value of perfect information* (EVPI) and the *expected value of sample information* (EVSI) are the expected improvements of a decision obtained by resolving an uncertainty completely, or by an imperfect test, before deciding (Chapter 24). The *conditional value at risk* CVaR_α is the mean of a loss over its worst $100(1-\alpha)\,\%$ of outcomes (Section 35.6.3).

Key idea

An importance measure describes a component's role in the present model. The value of an intervention is the change in a chosen risk metric when the model is modified exactly as the intervention modifies the physical system, set against its cost. Interventions, and combinations of them, are chosen by recomputing the risk for each candidate, not by reading an importance ranking.

This section states the system reliability function, defines the Birnbaum importance on it, shows with two two-component systems that importance depends on the architecture, and lists what an importance value leaves out of an investment decision.

35.1.1 The system reliability function

Consider a system of n components, each of which either functions or does not. Let

$$p_i = \Pr(\text{component } i \text{ functions}) \tag{35.2}$$

be the availability or the mission reliability of component i, consistently defined for the chosen time horizon: either every p_i is a point availability at a randomly selected instant, or every p_i is a reliability $R_i(H)$ over the same mission of length H. With independent components, the reliability of the system is a function of these probabilities alone,

$$R_{\text{system}} = h(p_1, p_2, \ldots, p_n) = h(\mathbf{p}), \tag{35.3}$$

called the *reliability function* of the system structure. It is the expected value of the structure function $\phi(\mathbf{x})$ of Chapter 9 of Volume I, $h(\mathbf{p}) = \mathbb{E}[\phi(\mathbf{X})]$, when the state indicators X_i are independent with $\Pr(X_i = 1) = p_i$, and its form is set by the architecture of the system.

Remark 35.1. In this chapter $h(\mathbf{p})$ is the reliability function of a system structure, the standard notation of system reliability theory; it is not the measurement model $h(\mathbf{x})$ of Chapters 29–31. The component probability p_i is likewise distinct from the network-state probability p_s of Chapter 27, which carries the subscript s of a state.

Two architectures recur in the chapter. In a *series* system every component is required, so the system functions only when all of them function, and with independent components

$$R_{\text{series}} = \prod_{i=1}^{n} p_i. \tag{35.4}$$

In a *parallel* system of two components either one can perform the required function, so the system fails only when both fail. The probability that both fail is $(1-p_1)(1-p_2)$ by independence, and

$$R_{\text{parallel}} = 1 - (1-p_1)(1-p_2). \tag{35.5}$$

The same component reliabilities produce different system reliabilities in the two architectures, and, more important for this chapter, a change of one component's reliability propagates differently to the system. The fault-tree form of Chapter 12 is the same model written in failure probabilities: with $q_i = 1 - p_i$ the probability of the basic event "component i unavailable" and $Q = 1 - R_{\text{system}}$ the system failure probability, $Q(\mathbf{q}) = 1 - h(\mathbf{1} - \mathbf{q})$.

35.1.2 Birnbaum importance

The Birnbaum importance measures the sensitivity of the system reliability to the reliability of one component,

$$I_i^{\text{B}} = \frac{\partial R_{\text{system}}}{\partial p_i}. \tag{35.6}$$

Chapter 12 of Volume I derived it on the failure side, as $I_i^{\text{B}} = \partial Q/\partial q_i = Q(1_i) - Q(0_i)$ (equation (12.12)), where $Q(1_i)$ and $Q(0_i)$ are the system failure probabilities with basic event i certain and impossible. The two definitions give the same number: since $Q = 1 - R_{\text{system}}$ and $q_i = 1 - p_i$, the derivative $\partial Q/\partial q_i$ changes sign twice and equals $\partial R_{\text{system}}/\partial p_i$. On the success side the pivotal decomposition reads

$$R_{\text{system}} = p_i\, h(1_i, \mathbf{p}) + (1 - p_i)\, h(0_i, \mathbf{p}),$$

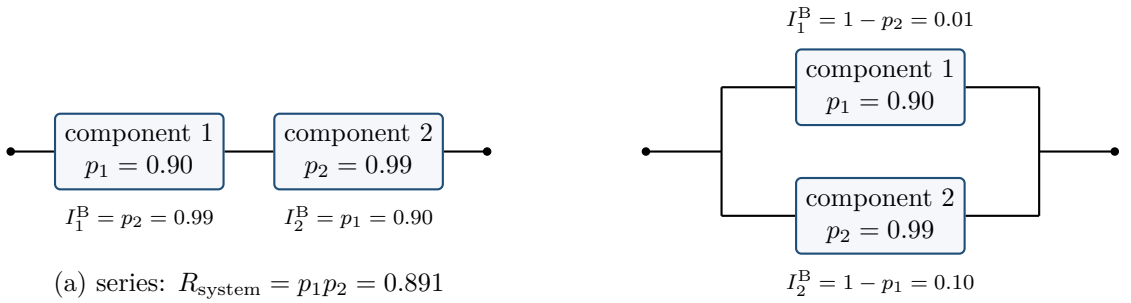
where $h(1_i, \mathbf{p})$ and $h(0_i, \mathbf{p})$ are the system reliability with component i certainly functioning and certainly failed, and every other component keeps its probability. Neither term depends on p_i, so the system reliability is linear in p_i and

$$I_i^{\text{B}} = \Pr(\text{system functions} \mid X_i = 1) - \Pr(\text{system functions} \mid X_i = 0), \tag{35.7}$$

with the other components retaining their specified independent state distributions. The quantity is dimensionless, lies between 0 and 1 in a coherent system, and is the probability that component i is critical, that is, that its state alone decides whether the system functions. A large value means that improving that component's reliability is particularly effective at improving the system reliability, all else equal. Because R_{system} is linear in p_i, the derivative gives the exact effect of any change of p_i alone; it is a local sensitivity only with respect to simultaneous changes of several probabilities.

35.1.3 Series and parallel systems

Figure 35.1 draws the two systems as reliability block diagrams. Each box is a component with its reliability; the system functions when an unbroken path of functioning boxes joins the input node on the left to the output node on the right. Below or beside each box is its Birnbaum importance. Panel (a) is the series system, panel (b) the parallel one; the component reliabilities are the same in both, $p_1 = 0.90$ and $p_2 = 0.99$ (illustrative).



(a) series: $R_{\mathrm{system}} = p_1 p_2 = 0.891$

(b) parallel: $R_{\mathrm{system}} = 1 - (0.10)(0.01) = 0.999$

Figure 35.1: Reliability block diagrams of two independent components in series, $R_{\mathrm{system}} = p_1 p_2$, and in parallel, $R_{\mathrm{system}} = 1 - (1 - p_1)(1 - p_2)$, with $p_1 = 0.90$ and $p_2 = 0.99$ (illustrative). The Birnbaum importance of each component is written beside it; the ranking of the two components reverses between the architectures.

Series system. With $R_{\mathrm{system}} = p_1 p_2$, differentiation gives

$$I_1^{\mathrm{B}} = p_2, \qquad I_2^{\mathrm{B}} = p_1, \tag{35.8}$$

so for $p_1 = 0.90$ and $p_2 = 0.99$, $I_1^{\mathrm{B}} = 0.99$ and $I_2^{\mathrm{B}} = 0.90$. A small increase of p_1 raises the system reliability more than the same absolute increase of p_2; because R_{system} is linear in each p_i, the comparison is exact for any step: raising p_1 by 0.01 to 0.91 gives $R_{\mathrm{system}} = 0.9009$, while raising p_2 by 0.01 to 1 gives 0.9000, against the baseline 0.891. The less reliable component is the better improvement target, which is intuitively reasonable in a series system: component 1 is critical whenever component 2 functions, 99 % of the time.

Parallel system. With $R_{\mathrm{system}} = 1 - (1 - p_1)(1 - p_2)$,

$$I_1^{\mathrm{B}} = 1 - p_2, \qquad I_2^{\mathrm{B}} = 1 - p_1, \tag{35.9}$$

so for the same components $I_1^{\mathrm{B}} = 0.01$ and $I_2^{\mathrm{B}} = 0.10$. The ranking reverses. Improving the already reliable component 2 has the greater marginal value, because a component of a parallel pair is critical only when its partner is down: component 2 decides the system state whenever component 1 is failed, which is 10 % of the time, while component 1 decides it only in the 1 % of the time that component 2 is failed. Importance depends on the system function, not only on the component's own reliability.

Table 35.1 works the two formulas for three pairs of component reliabilities. Each row is one pair; the left block is the series system, the right block the parallel one, each with the system reliability and the two Birnbaum importances. In series the less reliable component always has the larger importance, in parallel the more reliable one, and two equal components tie in both. The table also shows the size of the effect: in the parallel system the two importances are one or two orders of magnitude below those of the series system, because redundancy shields the system from each single component.

35.1.4 Importance is not an investment recommendation

The Birnbaum importance is a sensitivity of the present model, $\partial R_{\mathrm{system}} / \partial p_i$. It states how strongly the system depends on a component at the present parameter values, and it leaves out everything else that an investment decision needs. It does not account for:

- how much p_i can realistically be improved by any available action;
- the cost of achieving that improvement;

Table 35.1: System reliability and Birnbaum importance of two independent components in series and in parallel, equations (35.8) and (35.9) (illustrative values).

	Series			Parallel		
p_1, p_2	R_{system}	I_1^{B}	I_2^{B}	R_{system}	I_1^{B}	I_2^{B}
0.90, 0.99	0.8910	0.99	0.90	0.9990	0.01	0.10
0.95, 0.99	0.9405	0.99	0.95	0.9995	0.01	0.05
0.90, 0.90	0.8100	0.90	0.90	0.9900	0.10	0.10

- failure consequences beyond binary system success, such as the amount of load curtailed or the duration of an outage;
- changes in other components caused by the intervention, for instance a new unit that shares a cause of failure with the old one;
- common-cause dependencies, which the independent model of equation (35.3) excludes by construction;
- uncertainty in the reliability model and its parameters.

Choosing an investment requires a counterfactual comparison: the risk with the action taken against the risk without it, each computed on a model in which the action changes exactly what it changes in the physical system. Section 35.3 carries out such comparisons, after the next section completes the set of importance measures and extends them to the failure and repair rates.

35.2 Risk worth measures and the sensitivity to failure and repair parameters

The Birnbaum importance is a derivative; the two risk worth measures compare the baseline with the two extreme states of a component, and the parameter sensitivities carry the comparison down to the failure and repair rates that engineering actions change. This section states the risk metric, the two worth measures and the comparison of the measures, and then derives the sensitivity of unavailability to the failure and repair rates and its system-level form.

35.2.1 A risk metric for importance

Every risk-importance measure is defined relative to one risk metric, and must say which. For a binary system the natural metric is the system failure probability,

$$Q = \Pr(\text{system failure}) = 1 - R_{\text{system}}. \tag{35.10}$$

When failures differ in their consequences, a consequence-weighted risk is used instead, such as the expected monetary loss $\mathcal{R} = \mathbb{E}[C_{\text{loss}}]$, in reais, or the expected unserved energy $\mathcal{R} = \text{EUE}$ of Chapter 25, in MWh per year. The same component can rank differently under different metrics: a component whose failure is frequent but mild can dominate Q and be minor for EUE. The measures of this section are first written with Q; Section 35.3 replaces it with the expected load curtailment of a transmission network.

35.2.2 Risk reduction worth and risk achievement worth

Let $Q(0_i)$ be the system failure probability after component i's unavailability is set to zero, holding the other modeled mechanisms fixed, and $Q(1_i)$ the system failure probability with

component i forced into its failed state. The risk reduction worth and the risk achievement worth are the ratios

$$\mathrm{RRW}_i = \frac{Q}{Q(0_i)}, \qquad \mathrm{RAW}_i = \frac{Q(1_i)}{Q}, \tag{35.11}$$

as defined and derived in Chapter 12 of Volume I (equations (12.22) and (12.21)), where it is also shown that $\mathrm{RRW}_i = 1/(1 - I_i^{\mathrm{FV}})$ with I_i^{FV} the Fussell–Vesely importance, the fraction of the baseline removed by making the component perfect. A larger RRW indicates a greater potential relative reduction of risk. If the baseline is $Q = 0.01$ and eliminating one component's failures reduces it to $Q(0_i) = 0.004$, then $\mathrm{RRW}_i = 0.01/0.004 = 2.5$, and the component's perfect reliability would reduce the system failure probability by $1 - 0.004/0.01 = 60\,\%$, its Fussell–Vesely importance. An RRW of 2.5 does not mean that any feasible upgrade will achieve that reduction: perfect reliability is a hypothetical reference condition, the upper bound of what working on that component alone can buy.

RAW evaluates how vulnerable the system becomes when the component is unavailable, and identifies the components whose absence leaves the system exposed. A shared cooling-water header may have a low failure probability and a very high RAW: if the header is unavailable, all the redundant chillers lose their ability to deliver cooling. As an illustrative case, add to the cooling plant of Chapter 12, whose failure probability is 0.0184548, a header in series with unavailability $q_h = 0.001$. The plant failure probability becomes $1-(1-0.0184548)(1-0.001) = 0.0194364$, and with the header certainly failed it is 1, so $\mathrm{RAW}_h = 1/0.0194364 = 51.4$, while the header's Fussell–Vesely importance is only $0.001 \times 0.9815452/0.0194364 = 0.0505$ and its RRW $1/(1 - 0.0505) = 1.053$. The header contributes five percent of the present risk, and its loss would multiply the risk fifty-fold.

Table 35.2 applies the two ratios to the two-component systems of Figure 35.1. Each row is one component of one architecture, with the baseline failure probability, the two counterfactual values and the two ratios. In the series system the baseline is $Q = 1 - 0.891 = 0.109$; making component 1 perfect leaves only component 2's failures, $Q(0_1) = 1-0.99 = 0.01$, so $\mathrm{RRW}_1 = 10.9$, and either component certainly failed fails the system, so both have $\mathrm{RAW} = 1/0.109 = 9.17$: RAW does not separate two single points of failure. In the parallel system the baseline is $Q = 0.10 \times 0.01 = 0.001$. Making either component perfect makes the system perfect, $Q(0_i) = 0$, and the RRW is not defined (it grows without bound); the condition $Q(0_i) > 0$ of Chapter 12 fails. RAW, however, separates the two: with component 2 failed the system rests on component 1 and $Q(1_2) = 0.10$, a hundred times the baseline, while with component 1 failed it rests on the reliable component 2 and $Q(1_1) = 0.01$, ten times the baseline. The RAW ranking agrees with the Birnbaum ranking in parallel and is silent in series, where the Birnbaum and RRW rankings agree.

Table 35.2: Risk reduction worth and risk achievement worth of the two components of Figure 35.1, $p_1 = 0.90$, $p_2 = 0.99$, with the system failure probability Q as the risk metric (illustrative values).

Architecture	Component	Q	$Q(0_i)$	$Q(1_i)$	RRW_i	RAW_i
Series	1 ($p = 0.90$)	0.109	0.010	1	10.9	9.17
Series	2 ($p = 0.99$)	0.109	0.100	1	1.09	9.17
Parallel	1 ($p = 0.90$)	0.001	0	0.01	unbounded	10
Parallel	2 ($p = 0.99$)	0.001	0	0.10	unbounded	100

35.2.3 The measures compared

Table 35.3 states the question that each quantity answers. The first three are properties of one component in the present model; the fourth is a property of a real engineering action and is the quantity a decision needs. The three measures are related, since all of them are computed from Q, $Q(0_i)$ and $Q(1_i)$, but they are not interchangeable: the cooling-water header has a modest Fussell–Vesely importance and RRW and a large RAW, and the parallel pair has unbounded RRWs and different RAWs.

Table 35.3: The importance measures and the intervention benefit, with the question each answers.

Quantity	Expression	Question answered
Birnbaum importance	$\partial R_{\text{system}}/\partial p_i$	How sensitive is system reliability to this component's reliability?
Risk reduction worth	$Q/Q(0_i)$	How much risk could be removed by eliminating this component's failure contribution?
Risk achievement worth	$Q(1_i)/Q$	How much would system risk increase if this component were unavailable?
Feasible intervention benefit	$\Delta\mathcal{R}(a) = \mathcal{R}_0 - \mathcal{R}(a)$	How much does a particular real-world upgrade reduce risk?

In models with common-cause failures (Chapters 13 and 14 of Volume I), forcing a component's failure or success requires a carefully defined intervention on the underlying failure mechanisms. In a beta-factor model of three chillers, setting one chiller's unavailability to zero may mean removing only its independent failures, or also its share of the common-cause failures that disable all three; the two definitions give different values of $Q(0_i)$, and the study must state which physical change it represents.

35.2.4 Unavailability as a function of the failure and repair rates

Engineering actions rarely set a component's unavailability directly; they change its failure rate, by better design, derating or preventive maintenance, or its repair rate, by spare parts, crews and procedures. For a repairable component with constant failure rate λ (the failure intensity) and constant repair rate μ (the repair intensity), both in h^{-1} or yr^{-1}, the stationary unavailability is the standard result of the two-state Markov model (Chapter 19 of Volume I),

$$q = \frac{\lambda}{\lambda + \mu}. \tag{35.12}$$

Differentiating with the quotient rule, with respect to the failure rate

$$\frac{\partial q}{\partial \lambda} = \frac{(\lambda + \mu) - \lambda}{(\lambda + \mu)^2} = \frac{\mu}{(\lambda + \mu)^2}, \tag{35.13}$$

and with respect to the repair rate

$$\frac{\partial q}{\partial \mu} = \frac{0 - \lambda}{(\lambda + \mu)^2} = -\frac{\lambda}{(\lambda + \mu)^2}. \tag{35.14}$$

Both derivatives have the units of time, since q is dimensionless and the rates are per unit time. Their signs state that a higher failure rate increases unavailability and a higher repair rate reduces it; their magnitudes depend on the baseline. For the cooling-water pump of

Chapter 20 of Volume I, with $\lambda = 0.001\,\mathrm{h}^{-1}$ and $\mu = 0.05\,\mathrm{h}^{-1}$, $q = 0.001/0.051 = 0.019608$, $\partial q/\partial\lambda = 0.05/0.051^2 = 19.22\,\mathrm{h}$ and $\partial q/\partial\mu = -0.001/0.051^2 = -0.3845\,\mathrm{h}$. The two numbers differ by the factor $\mu/\lambda = 50$ and cannot be compared directly: a change of $0.001\,\mathrm{h}^{-1}$ doubles the failure rate but changes the repair rate by two percent.

35.2.5 Dimensionless elasticities

A dimensionless sensitivity removes the units by comparing relative changes. The *elasticity* of q with respect to λ is the relative change of q per relative change of λ,

$$E_{q,\lambda} = \frac{\lambda}{q}\,\frac{\partial q}{\partial \lambda} = \lambda\,\frac{\lambda+\mu}{\lambda}\,\frac{\mu}{(\lambda+\mu)^2} = \frac{\mu}{\lambda+\mu} = 1 - q, \tag{35.15}$$

and with respect to μ,

$$E_{q,\mu} = \frac{\mu}{q}\,\frac{\partial q}{\partial \mu} = \mu\,\frac{\lambda+\mu}{\lambda}\left(-\frac{\lambda}{(\lambda+\mu)^2}\right) = -\frac{\mu}{\lambda+\mu} = -(1-q). \tag{35.16}$$

For rare failures, $\lambda \ll \mu$, both elasticities are close to one in magnitude, $E_{q,\lambda} \approx 1$ and $E_{q,\mu} \approx -1$: locally, a 1 % reduction of the failure rate or a 1 % increase of the repair rate each reduces the unavailability by about 1 %.

The equivalence is in fact exact in a stronger sense. Equation (35.12) depends on the two rates only through their ratio $\rho = \lambda/\mu$, $q = \rho/(1+\rho)$, so dividing the failure rate by any factor k and multiplying the repair rate by the same k give the same unavailability, for large changes as well as small ones. A 1 % reduction of λ (the factor 0.99) and a 1 % increase of μ (the factor $1/1.01 = 0.990099$ on ρ) differ only at second order. The equivalence concerns proportional changes of the rates and the unavailability only. It does not imply that the two improvements cost the same, and it does not imply that they have the same system-level benefit when the risk metric depends on more than the unavailability: halving λ halves the number of outages, while doubling μ keeps the number of outages and halves their duration, and a metric that counts interruptions, or that charges only outages longer than a ride-through time such as the time to the thermal limit t_{crit} of Chapter 34, distinguishes the two.

Table 35.4 evaluates the derivatives and the elasticities for the cooling-water pump, for the three transmission lines of Chapter 27 with the outage and restoration rates of its Table 27.5 ($\mu = 100\,\mathrm{yr}^{-1}$, a mean outage of 87.6 h), and for two illustrative units with large unavailabilities. Each row gives the two rates, the unavailability, the two derivatives and the two elasticities. The row of L23 shows the arithmetic: $\lambda_{23} = 3.093\,\mathrm{yr}^{-1}$, $q_{23} = 3.093/103.093 = 0.030$, $\partial q/\partial\lambda = 100/103.093^2 = 0.009409\,\mathrm{yr}$ and $E_{q,\lambda} = 1 - 0.03 = 0.97$. The trend is the factor $1-q$: for rare, short outages a proportional change of either rate passes almost entirely to the unavailability, and for a unit that is unavailable half the time it passes only half.

Table 35.4: Sensitivity of the stationary unavailability $q = \lambda/(\lambda+\mu)$ to the failure and repair rates, and the elasticities $E_{q,\lambda} = 1-q$ and $E_{q,\mu} = -(1-q)$. The pump rates are per hour and its derivatives in hours; the line and unit rates are per year and their derivatives in years.

Item	λ	μ	q	$\partial q/\partial\lambda$	$\partial q/\partial\mu$	$E_{q,\lambda}$	$E_{q,\mu}$
Cooling pump	0.001	0.05	0.019608	19.22	−0.3845	0.980	−0.980
Line L13	1.010	100	0.010000	0.009801	−0.0000990	0.990	−0.990
Line L12	2.041	100	0.020000	0.009604	−0.000196	0.980	−0.980
Line L23	3.093	100	0.030000	0.009409	−0.000291	0.970	−0.970
Unit, illustrative	1	4	0.200000	0.160	−0.0400	0.800	−0.800
Unit, illustrative	10	10	0.500000	0.025	−0.0250	0.500	−0.500

35.2.6 System-level sensitivity

Let the system risk be a function of the component unavailabilities, $\mathcal{R} = \mathcal{R}(q_1, \dots, q_n)$, and let each q_i depend on its own rates. The chain rule connects the physical parameters that an intervention changes to the system risk:

$$\frac{\partial \mathcal{R}}{\partial \lambda_i} = \frac{\partial \mathcal{R}}{\partial q_i}\,\frac{\partial q_i}{\partial \lambda_i}, \qquad \frac{\partial \mathcal{R}}{\partial \mu_i} = \frac{\partial \mathcal{R}}{\partial q_i}\,\frac{\partial q_i}{\partial \mu_i}. \tag{35.17}$$

The first factor is the Birnbaum importance of component i on the chosen risk metric, and the second is equation (35.13) or (35.14). Writing the result as an elasticity shows which component's rates matter most:

$$\frac{\lambda_i}{\mathcal{R}}\,\frac{\partial \mathcal{R}}{\partial \lambda_i} = \frac{q_i\, I_i^{\mathrm{B}}}{\mathcal{R}}\,\frac{\lambda_i}{q_i}\,\frac{\partial q_i}{\partial \lambda_i} = I_i^{\mathrm{FV}}\, E_{q_i,\lambda_i}, \tag{35.18}$$

since $q_i I_i^{\mathrm{B}}/\mathcal{R}$ is the Fussell–Vesely importance (Chapter 12 of Volume I, equation (12.20)). The elasticity of the system risk with respect to a component's failure rate is its Fussell–Vesely importance times the component elasticity $1 - q_i$.

For line L23 of the three-bus network, with the expected curtailment $\mathcal{R} = 1.3127\,\mathrm{MW}$ as the metric, the Birnbaum importance is $I_{23}^{\mathrm{B}} = 20.09\,\mathrm{MW}$ per unit of unavailability (Chapter 27, Table 27.7). Then $\partial\mathcal{R}/\partial\lambda_{23} = 20.09 \times 0.009409 = 0.1890\,\mathrm{MW}$ per outage per year, $\partial\mathcal{R}/\partial\mu_{23} = 20.09 \times (-0.000291) = -0.005846\,\mathrm{MW}$ per unit of restoration rate (in yr^{-1}), and the system elasticity is $0.459 \times 0.97 = 0.445$: a 1 % reduction of the outage rate of L23 lowers the expected curtailment by 0.445 %.

Derivatives describe small changes. Consider halving the mean outage duration of L23 from 87.6 to 43.8 h, which doubles μ_{23} to $200\,\mathrm{yr}^{-1}$. Exactly, $q_{23} = 3.093/203.093 = 0.015228$, and since $\mathcal{R}$ is linear in q_{23} (the pivotal decomposition of Section 35.1.2), $\mathcal{R} = 0.71 + 20.09 \times 0.015228 = 1.0159\,\mathrm{MW}$, a reduction of $0.2968\,\mathrm{MW}$. The first-order estimate in the repair rate, $0.005846 \times 100 = 0.5846\,\mathrm{MW}$, nearly doubles it, because q is a strongly curved function of μ (roughly λ/μ). Figure 35.2 shows the same calculation with the mean outage duration $\mathrm{MTTR}_{23} = 1/\mu_{23}$ on the horizontal axis, in hours, and the expected curtailment on the vertical axis, in MW. The solid curve is the exact response with λ_{23} fixed, the dashed line its tangent at the present 87.6 h, the filled dots the present and the halved duration, and the open circle the first-order estimate in μ. In this parameterization the curve is nearly straight, because $q \approx \lambda\,\mathrm{MTTR}$ for short outages, and the tangent predicts a reduction of $0.0066737 \times 43.8 = 0.2923\,\mathrm{MW}$, close to the exact $0.2968\,\mathrm{MW}$. A linearization is only as good as the parameter in which it is taken; for a large change the model is evaluated at the new parameter values.

In more complex systems the rates themselves depend on operating conditions, $\lambda_i = \lambda_i(\mathbf{u}, \mathbf{T}, \mathbf{L})$, where $\mathbf{u}$ are control settings, $\mathbf{T}$ the component temperatures and $\mathbf{L}$ the loadings. A conductor carrying more power runs hotter, and a weather-dependent hazard rises with its temperature (Chapter 28). The sensitivity of the risk to a control setting then includes those physical dependencies,

$$\frac{\mathrm{d}\mathcal{R}}{\mathrm{d}u} = \sum_i \frac{\partial \mathcal{R}}{\partial q_i}\,\frac{\partial q_i}{\partial \lambda_i}\Big(\frac{\partial \lambda_i}{\partial u} + \sum_k \frac{\partial \lambda_i}{\partial T_k}\,\frac{\partial T_k}{\partial u}\Big),$$

and an action that appears to touch one component can change the failure rates of several.

35.3 Counterfactual risk reduction

An investment is evaluated by computing the risk of the system as it would be after the investment and comparing it with the risk of the system as it is. This section states what

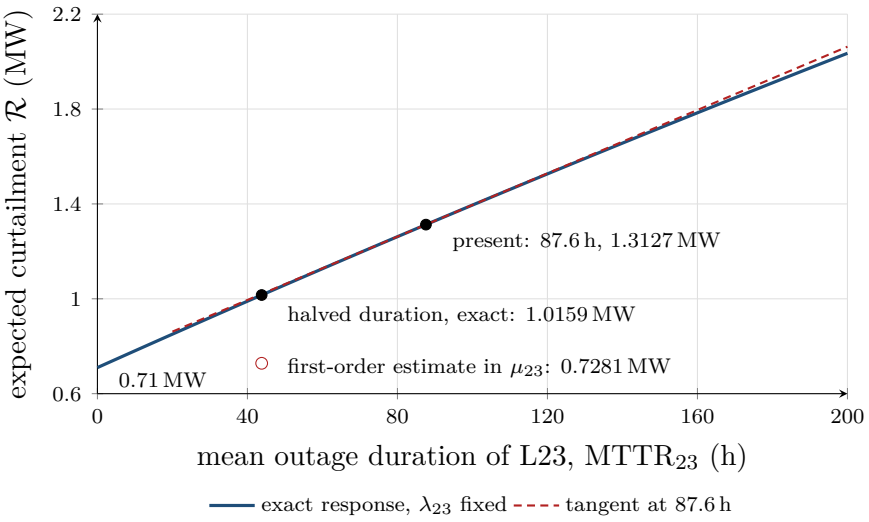


Figure 35.2: Expected curtailment of the three-bus network against the mean outage duration of L23, $\mathcal{R} = 0.71 + 20.09\, q_{23}$ with $q_{23} = \lambda_{23}\mathrm{MTTR}_{23}/(1 + \lambda_{23}\mathrm{MTTR}_{23})$ and $\lambda_{23} = 3.093\,\mathrm{yr}^{-1}$ (the data of Chapter 27; illustrative). The tangent at the present duration has slope 0.0066737 MW per hour. Halving the duration lowers $\mathcal{R}$ to 1.0159 MW; the first-order estimate in the repair rate overstates the reduction by a factor of two.

an intervention changes in the model, distinguishes an intervention from an observation, and evaluates three interventions on the three-bus transmission network: the elimination of one line's unavailability, a finite improvement of it, and a rating increase.

35.3.1 What an intervention changes

The three-bus network of Chapter 27 (Figure 27.2) has three lines with the data used throughout that chapter: L12 with reactance 0.1 pu, rating 70 MW and unavailability 0.02; L23 with 0.1 pu, 70 MW and 0.03; and L13 with 0.2 pu, 80 MW and 0.01. A transfer of 100 MW is required from bus 1, the generator, to bus 3, the load. Under independent line unavailabilities, the DC power-flow constraints and minimum curtailment as the only corrective action, Chapter 27 enumerated the eight network states (Section 27.4.2) and found the expected curtailment

$$\mathcal{R}_0 = \mathbb{E}[\Delta L] = 1.3127\,\mathrm{MW}, \tag{35.19}$$

where ΔL is the part of the transfer that cannot be delivered in a randomly selected network state, the shortfall S of Chapter 25 for this network.

An intervention is represented by modifying exactly the model quantities it changes in the physical system. Better maintenance or faster repair reduces a line's unavailability q_ℓ. Reconductoring raises a thermal rating $P_{ij}^{\max}$. New sensors may improve the estimates of the ratings without changing any physical capacity. A new line changes the topology, the susceptance matrix $\mathbf{B}$ and with it the redistribution of power after every outage. Figure 35.3 draws these paths. The boxes on the left are the actions, the middle column the model quantities each one changes, and the right column the two ingredients of the risk, the state probabilities and the state consequences; the arrows show which quantity feeds which ingredient. The dashed path at the bottom is a measurement: it changes the knowledge of the parameters, which informs the choice of action, and it has no arrow into the physical quantities.

The post-intervention risk is written

$$\mathcal{R}(a) = \mathbb{E}\big[\Delta L \mid \mathrm{do}(a)\big], \tag{35.20}$$

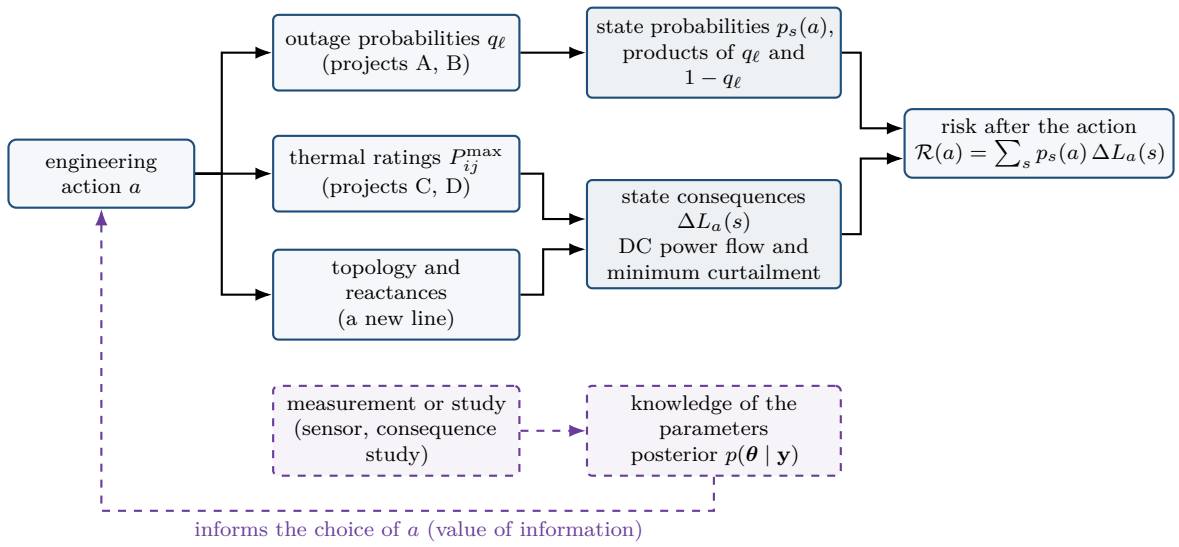


Figure 35.3: What an intervention changes. Physical actions (left) change the outage probabilities, the ratings or the topology; these determine the state probabilities and the state consequences, whose probability-weighted sum is the risk after the action. A measurement (dashed) changes the knowledge of the parameters and informs the choice of action without changing any physical quantity.

and the risk reduction of the action is

$$\Delta\mathcal{R}(a) = \mathcal{R}_0 - \mathcal{R}(a), \tag{35.21}$$

in MW for this network, positive when the action lowers the risk. This is the benefit $\Delta Q(a)$ of Chapter 12 of Volume I (equation (12.30)) with the expected curtailment in the place of the system failure probability.

35.3.2 Intervention and observation

The notation $\mathrm{do}(a)$, from the theory of causal inference, marks that the quantities named in a are *set* by an action on the system, as an intervention, rather than *learned* by observing a favorable state. The difference is that between two questions about line L23. The observational question is: in the hours in which L23 is found in service, what is the expected curtailment? It conditions on an observed state, $\mathbb{E}[\Delta L \mid \text{L23 in service}]$. The interventional question is: in a network in which L23 has been made perfectly reliable, what is the expected curtailment? It changes the model, $\mathbb{E}[\Delta L \mid \mathrm{do}(q_{23} = 0)]$, and leaves everything that the action does not touch as it was.

With independent lines the two numbers coincide. Observing L23 in service tells nothing about the other lines, and both questions give the 0.7100 MW of Chapter 27. They separate as soon as the lines share a cause. Take the storm model of the exercises of Chapter 27: storm weather in 10 % of the hours, with unavailabilities $q_{12} = 0.10$, $q_{23} = 0.15$ and $q_{13} = 0.10$, and the usual values in the other 90 %, the lines independent within each weather state (illustrative). Averaged over the weather, the expected curtailment is $0.9 \times 1.3127 + 0.1 \times 8.875 = 2.0689$ MW. With L23 removed from the outage model, equation (27.18) gives 0.71 MW in fair weather and 5.5 MW in storms. The intervention does not change the weather, so

$$\mathbb{E}[\Delta L \mid \mathrm{do}(q_{23} = 0)] = 0.9 \times 0.71 + 0.1 \times 5.5 = 1.1890\,\text{MW}.$$

The observation does: by Bayes' theorem, finding L23 in service lowers the probability of a

storm from 0.10 to

$$\Pr(\text{storm} \mid \text{L23 in service}) = \frac{0.1 \times 0.85}{0.1 \times 0.85 + 0.9 \times 0.97} = \frac{0.085}{0.958} = 0.08873,$$

and the conditional expected curtailment is $0.91127 \times 0.71 + 0.08873 \times 5.5 = 1.1350$ MW. The observational comparison credits L23 with a reduction of $2.0689 - 1.1350 = 0.934$ MW, of which only $2.0689 - 1.1890 = 0.880$ MW would be delivered by making the line perfectly reliable; the remaining 0.054 MW is the fair weather that an in-service L23 signals. An engineering decision needs the interventional number.

35.3.3 Perfect-reliability attribution

The first counterfactual eliminates one line's unavailability entirely. Table 35.5 reproduces the attribution of Chapter 27 (Table 27.7) and adds the risk achievement worth of each line. Each row is one line, with its unavailability, the expected curtailment with the line made perfect, the reduction, the Birnbaum importance, the RRW, the expected curtailment with the line certainly out and the RAW. The values with the line perfect or certainly out follow from the closed form $\mathbb{E}[\Delta L] = 30\,q_{13}(1-u) + 20\,(1-q_{13})\,u + 100\,q_{13}\,u$ of Chapter 27 (equation (27.18)), with $u = 1 - (1-q_{12})(1-q_{23})$ the probability that the indirect path is broken. For L12, $q_{12} = 0$ gives $u = 0.03$ and $0.291 + 0.594 + 0.030 = 0.9150$ MW, and $q_{12} = 1$ gives $u = 1$ and $20 \times 0.99 + 100 \times 0.01 = 20.80$ MW, so $\mathrm{RAW}_{12} = 20.80/1.3127 = 15.85$; for L13, $q_{13} = 1$ gives $30 \times 0.9506 + 100 \times 0.0494 = 33.458$ MW and $\mathrm{RAW}_{13} = 25.49$.

Table 35.5: Perfect-reliability attribution of the expected curtailment $\mathcal{R}_0 = 1.3127$ MW to the three lines (as in Table 27.7), with the expected curtailment with each line certainly out and its risk achievement worth.

Line	q_ℓ	$\mathbb{E}[\Delta L]$, line perfect	Reduction	I^{B}_ℓ	RRW	$\mathbb{E}[\Delta L]$, line out	RAW
L12	0.02	0.9150 MW	0.3977 MW	19.885 MW	1.435	20.800 MW	15.85
L23	0.03	0.7100 MW	0.6027 MW	20.090 MW	1.849	20.800 MW	15.85
L13	0.01	0.9880 MW	0.3247 MW	32.470 MW	1.329	33.458 MW	25.49

L23 offers the largest perfect-reliability risk reduction, 0.6027 MW, because it has the highest unavailability, while L13 has the largest Birnbaum importance and the largest RAW, because its outage forces the whole transfer onto the two 70 MW lines and its loss together with any other line islands the load. L12 and L23 have equal RAW: with either one certainly out the indirect path is broken and the network depends on L13 alone. None of these numbers makes L23 the best investment. They describe idealized changes; a decision needs the feasible improvements and their costs.

35.3.4 A finite improvement of L23

A practical reliability investment yields a finite improvement, not perfect availability. Suppose the unavailability of L23 is reduced from $q_{23} = 0.03$ to $q'_{23} = 0.01$, every other quantity unchanged. The eight network states keep their transfer capabilities and curtailments, and only their probabilities change. Table 35.6 lists the states; its first column names the lines out, the next two give the state probability in the baseline and with $q_{23} = 0.01$, the next two the curtailment with the present ratings and with the rating of L13 raised to 100 MW (the next subsection), and the last two the products that sum to the expected curtailment of the two interventions. With the factors in the order L12, L23, L13, the probability of the state with

L23 alone out becomes $0.98 \times 0.01 \times 0.99 = 0.009702$, and the sum of the products is

$$\begin{aligned}\mathcal{R}_{\text{new}} &= 0.29106 + 0.19404 + 0.00980 + 0.39204 + 0.01980 + 0.00396 + 0.00020 \\ &= 0.9109\,\text{MW},\end{aligned}$$

a reduction of

$$\Delta\mathcal{R} = 1.3127 - 0.9109 = 0.4018\,\text{MW}. \tag{35.22}$$

The same number follows without enumeration, because the expected curtailment is linear in q_{23} with slope I^{B}_{23}: $\Delta\mathcal{R} = 20.09 \times (0.03 - 0.01) = 0.4018\,\text{MW}$, two thirds of the 0.6027 MW of the perfect-reliability reference, as the step from 0.03 to 0.01 is two thirds of the step to zero.

Table 35.6: The eight network states for a 100 MW transfer: probabilities in the baseline and with $q_{23} = 0.01$; curtailments with the present ratings and with $P^{\max}_{13} = 100\,\text{MW}$; and the products for the two interventions (illustrative data of Chapter 27).

Lines out	p_s	p_s, $q_{23} = 0.01$	$\Delta L(s)$	$\Delta L(s)$, L13 at 100 MW	Product, $q_{23} = 0.01$	Product, L13 at 100 MW
none	0.941094	0.960498	0	0	0	0
L13	0.009506	0.009702	30	30	0.29106	0.28518
L23	0.029106	0.009702	20	0	0.19404	0
L23, L13	0.000294	0.000098	100	100	0.00980	0.02940
L12	0.019206	0.019602	20	0	0.39204	0
L12, L13	0.000194	0.000198	100	100	0.01980	0.01940
L12, L23	0.000594	0.000198	20	0	0.00396	0
all three	0.000006	0.000002	100	100	0.00020	0.00060
Total	1.000000	1.000000			0.91090 MW	0.33458 MW

35.3.5 A rating upgrade of L13

Consider instead raising the rating of L13 from 80 to 100 MW. The outage probabilities are unchanged; the curtailment of some states changes, and it follows from the DC power flow of each topology. Table 35.7 gives, for every topology, the DC flows of the 100 MW transfer, computed as in Chapter 27 with bus 3 as the angle reference, and the transfer capability under three sets of ratings: the present ones, L13 raised to 100 MW, and L12 and L23 raised to 90 MW (project D of Section 35.4). Because the flows are linear in the transfer, the transfer capability of a topology is the smallest ratio of a loaded line's rating to the fraction of the transfer it carries, and the curtailment is $(100\,\text{MW} - \text{TC})^+$.

The rows read as follows. In the intact network the two paths, the direct line of reactance 0.2 pu and the series path of $0.1 + 0.1$ pu, share the transfer equally, so the series lines bind first with present ratings and $\text{TC} = 2 \times 70 = 140\,\text{MW}$. With L13 out the whole transfer crosses L12 and L23 in series, $\text{TC} = \min(P^{\max}_{12}, P^{\max}_{23})$. With L12 or L23 out, bus 2 becomes a terminal bus with no injection, its remaining line carries nothing, and L13 carries the whole transfer, $\text{TC} = P^{\max}_{13}$. With L13 out together with another line, no path joins bus 1 to bus 3. Under the upgrade of L13 the three states in which L13 alone carries the transfer, which curtailed 20 MW each, now curtail nothing; the state with L13 out still curtails 30 MW, and the islanding states still lose the whole 100 MW. The last column of Table 35.6 gives

$$\begin{aligned}\mathcal{R}_{\text{new}} &= 0.28518 + 0.02940 + 0.01940 + 0.00060 = 0.33458\,\text{MW}, \\ \Delta\mathcal{R} &= 1.3127 - 0.33458 = 0.97812\,\text{MW},\end{aligned} \tag{35.23}$$

which reproduces Table 27.8 of Chapter 27.

Table 35.7: DC flows of a 100 MW transfer from bus 1 to bus 3 in each topology (MW, positive in the line orientation), and the transfer capability TC under three sets of ratings $(P_{12}^{\max}, P_{23}^{\max}, P_{13}^{\max})$ in MW.

Lines out	P_{12}	P_{23}	P_{13}	TC, $(70, 70, 80)$	TC, $(70, 70, 100)$	TC, $(90, 90, 80)$
none	50	50	50	140	140	160
L13	100	100	out	70	70	90
L23	0	out	100	80	100	80
L12	out	0	100	80	100	80
L12, L23	out	out	100	80	100	80
L13 with L12 or L23	islanded, no path			0	0	0

This improvement is larger than eliminating any one line's failures. The rating upgrade mitigates the consequences of several different contingencies at once, the outages of L12, of L23 and of both, whose contributions together were 0.97812 MW. It does not reduce the probability of any outage; it changes what the network can deliver after the outages occur. The comparison illustrates the difference between improving component reliability, which lowers the probability of component failures, and reducing their consequences, which makes the system tolerant of failures it still suffers, a property also called resilience to those failures. The transmission loss-of-load probability shows the same split: the upgrade lowers it from 0.058906 to 0.010000, the probability of the states in which L13 itself is out, while the elimination of L23's unavailability lowers it only to 0.029800.

35.4 An investment portfolio

The counterfactual calculations of the previous section become a decision when they are attached to costs, a budget and a valuation of the consequences. This section defines four projects for the three-bus network, recomputes the expected curtailment under each, converts the reductions into annual benefits, formulates the selection under a budget as a 0–1 optimization, and shows why the benefits of the projects cannot be added.

35.4.1 Four projects and the valuation of consequences

Assume that the utility can carry out four interventions. Table 35.8 lists them, with the model quantity each changes and its annualized cost c_i, which includes capital recovery and recurring incremental expenses. The costs are hypothetical. The investment budget is

$$B = \text{R\$ } 100\,000 \text{ per year}, \tag{35.24}$$

and the physical network is the same as in Chapter 27 except for the stated changes. Figure 35.4 draws the network with the four projects attached to the lines they change. The circles are the buses, the thick lines the transmission lines, each labeled with its unavailability and rating; the colored lines below each label name the projects acting on that line and what they change.

Remark 35.2. The budget B (italic, no subscript) is not project B (an upright letter in the text), and neither is the gross or net benefit B_{gross}, B_{net} of Chapter 12 of Volume I, which this section also uses.

To compare consequences in money, assume that the analyzed 100 MW transfer is relevant for $H_{\text{rel}} = 300$ h per year, and that the snapshot distribution of the network states applies in

Table 35.8: The four candidate projects for the three-bus network, the model quantity each changes and its annualized cost (hypothetical).

Project	Description	Model change	Cost c_i per year
A	Reduce the unavailability of L23 from 3 % to 1 %	q_{23}: 0.03 to 0.01	R$ 45 000
B	Reduce the unavailability of L12 from 2 % to 0.5 %	q_{12}: 0.02 to 0.005	R$ 35 000
C	Increase the rating of L13 from 80 to 100 MW	$P_{13}^{\max}$: 80 to 100 MW	R$ 60 000
D	Increase the ratings of L12 and L23 from 70 to 90 MW	$P_{12}^{\max}$, $P_{23}^{\max}$: 70 to 90 MW	R$ 80 000

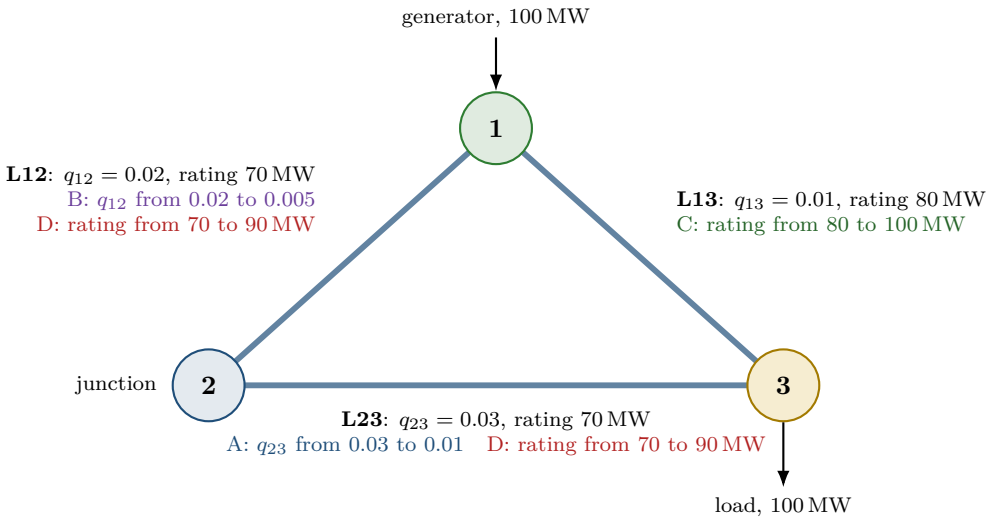


Figure 35.4: The three-bus network of Chapter 27 with the four candidate projects of Table 35.8 attached to the lines they change: A and B reduce outage probabilities, C and D raise thermal ratings. Reactances $x_{12} = x_{23} = 0.1$ pu and $x_{13} = 0.2$ pu are unchanged by every project (illustrative values).

each of those hours; the other hours are outside this illustrative assessment. Let the consequence value be

$$v = \text{R\$ } 10\,000 \text{ per MWh}, \tag{35.25}$$

an illustrative economic penalty for unavailable energy, not a universal value of lost load. With a constant curtailment through each hour, the annual expected unserved energy and the annual expected consequence of a project portfolio a are

$$\mathrm{EUE}(a) = H_{\mathrm{rel}}\,\mathcal{R}(a), \qquad C_{\mathrm{risk}}(a) = v\,H_{\mathrm{rel}}\,\mathcal{R}(a). \tag{35.26}$$

For the baseline,

$$\mathrm{EUE}_0 = 300\,\mathrm{h} \times 1.3127\,\mathrm{MW} = 393.81\,\mathrm{MWh} \text{ per year},$$
$$C_{\mathrm{risk},0} = 10\,000 \times 393.81 = \text{R\$ } 3\,938\,100 \text{ per year}.$$

These are risk-accounting figures for a restricted set of 300 scenario hours, not a complete annual adequacy study of the power system, which would represent the variation of demand, dispatch, ratings and outages hour by hour (Section 27.4.7).

35.4.2 Expected curtailment under each project

Each project is evaluated by enumerating the eight line-availability states under its changes, solving for the largest feasible transfer in each, and recomputing the expected curtailment. Projects A and B change probabilities only; projects C and D change curtailments only. Table 35.6 has done project A (0.9109 MW) and project C (0.33458 MW). Table 35.9 does project B, project D and the two combinations that Section 35.4.4 needs. Its columns are the lines out, the state probabilities with $q_{12} = 0.005$ (project B) and with $q_{12} = 0.005$, $q_{23} = 0.01$ (projects A and B), the curtailment with L12 and L23 at 90 MW (project D, from the last column of Table 35.7), and the four products whose sums are the expected curtailments.

Table 35.9: State probabilities, curtailments and products for project B, project D, projects A and B together, and projects B and C together. The curtailment of the other configurations is that of Table 35.6: present ratings for B and A+B, L13 at 100 MW for B+C.

Lines out	p_s for B	p_s for A+B	$\Delta L(s)$ for D	Product B	Product D	Product A+B	Product B+C
none	0.9554985	0.9751995	0	0	0	0	0
L13	0.0096515	0.0098505	10	0.289545	0.095060	0.295515	0.289545
L23	0.0295515	0.0098505	20	0.591030	0.582120	0.197010	0
L23, L13	0.0002985	0.0000995	100	0.029850	0.029400	0.009950	0.029850
L12	0.0048015	0.0049005	20	0.096030	0.384120	0.098010	0
L12, L13	0.0000485	0.0000495	100	0.004850	0.019400	0.004950	0.004850
L12, L23	0.0001485	0.0000495	20	0.002970	0.011880	0.000990	0
all three	0.0000015	0.0000005	100	0.000150	0.000600	0.000050	0.000150
Total	1.0000000	1.0000000		1.014425	1.122580	0.606475	0.324395

For project B the closed form of Chapter 27 checks the sum: $u = 1 - 0.995 \times 0.97 = 0.03485$ and $30 \times 0.01 \times 0.96515 + 20 \times 0.99 \times 0.03485 + 100 \times 0.01 \times 0.03485 = 0.289545 + 0.690030 + 0.034850 = 1.014425$ MW. Project D changes a single state: with L12 and L23 at 90 MW the outage of L13 leaves a series path of 90 MW and curtails 10 MW instead of 30 MW, while the outages of L12 or L23 still leave L13 alone at its 80 MW; the reduction is $0.009506 \times (30 - 10) = 0.19012$ MW. The higher series ratings also raise the intact transfer capability to 160 MW, which is of no value at a 100 MW transfer.

Table 35.10 collects the four projects; its first columns are the expected curtailment after the project and the reduction $\Delta\mathcal{R}(a)$, and Figure 35.5 draws the reductions as bars. The horizontal axis lists the projects, the vertical axis is the reduction of the expected curtailment in MW, and each bar carries its value. Project C has the largest standalone risk reduction, more than twice that of project A, and project D the smallest, because it improves only the one state with L13 out, of probability 0.0095.

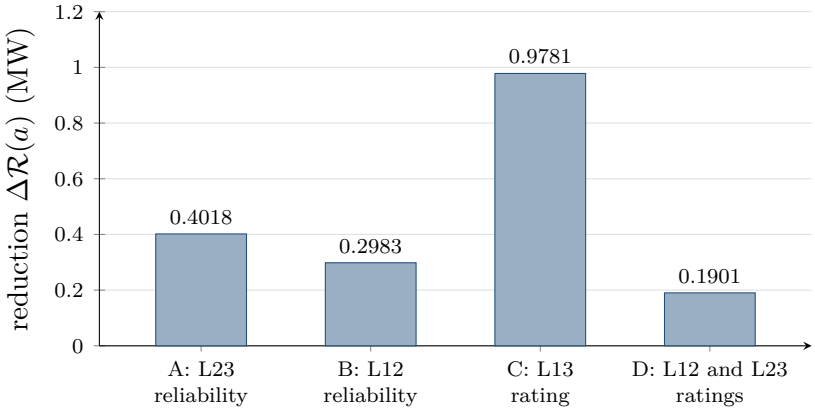


Figure 35.5: Reduction of the expected curtailment of the three-bus network, $\Delta\mathcal{R}(a) = \mathcal{R}_0 - \mathcal{R}(a)$ with $\mathcal{R}_0 = 1.3127\,\text{MW}$, for each project of Table 35.8 carried out alone (Tables 35.6 and 35.9; illustrative values).

35.4.3 Financial benefits

Chapter 12 of Volume I wrote the net benefit of an action as the expected consequence avoided minus the cost of the action, over a common horizon (equation (12.31)). Over one year, with the consequences of equation (35.26) and the annualized cost, the gross and the net annual expected benefits of project a are

$$B_{\text{gross}}(a) = v\,H_{\text{rel}}\,[\mathcal{R}_0 - \mathcal{R}(a)], \qquad B_{\text{net}}(a) = B_{\text{gross}}(a) - c_a. \tag{35.27}$$

For project A, $B_{\text{gross}} = 10\,000 \times 300 \times 0.4018 = \text{R\$}\,1\,205\,400$ per year and $B_{\text{net}} = 1\,205\,400 - 45\,000 = \text{R\$}\,1\,160\,400$ per year. Table 35.10 gives all four, with the annual reduction of the expected unserved energy and the ratio of gross benefit to cost. Project C has the largest net expected benefit among the individual projects, R\$ 2 874 360 per year, and the largest benefit per real spent; project D has the smallest of both. Every project is worth far more than it costs at this consequence value, which is why the remaining question is which combination to choose.

Table 35.10: The four projects evaluated alone: expected curtailment after the project, its reduction, the reduction of the annual expected unserved energy over $H_{\text{rel}} = 300\,\text{h}$, and the gross and net annual benefits at $v = \text{R\$}\,10\,000$ per MWh (illustrative).

Project	$\mathcal{R}(a)$ (MW)	$\Delta\mathcal{R}(a)$ (MW)	ΔEUE (MWh/yr)	B_{gross} (R\$/yr)	B_{net} (R\$/yr)	B_{gross}/c_a
none	1.312700	–	–	–	–	–
A	0.910900	0.401800	120.540	1 205 400	1 160 400	26.8
B	1.014425	0.298275	89.483	894 825	859 825	25.6
C	0.334580	0.978120	293.436	2 934 360	2 874 360	48.9
D	1.122580	0.190120	57.036	570 360	490 360	7.1

35.4.4 Budget-constrained selection

Let a selection variable mark each project,

$$x_i = \begin{cases} 1, & \text{project } i \text{ selected,} \\ 0, & \text{otherwise,} \end{cases} \qquad i \in \{\mathrm{A}, \mathrm{B}, \mathrm{C}, \mathrm{D}\}, \tag{35.28}$$

and collect them in $\mathbf{x} \in \{0,1\}^4$. The budget constraint is

$$\sum_i c_i \, x_i \le B, \tag{35.29}$$

and the selection that minimizes the expected annual consequence plus the annualized project cost is

$$\min_{\mathbf{x} \in \{0,1\}^4} \Big[v \, H_{\mathrm{rel}} \, \mathcal{R}(\mathbf{x}) + \sum_i c_i \, x_i \Big] \quad \text{subject to} \quad \sum_i c_i \, x_i \le \mathrm{R\$}\,100\,000. \tag{35.30}$$

Since $v H_{\mathrm{rel}} \mathcal{R}_0$ is a constant, the same selection maximizes the net benefit $v H_{\mathrm{rel}} [\mathcal{R}_0 - \mathcal{R}(\mathbf{x})] - \sum_i c_i x_i$. The problem is a 0–1 (binary) program of the knapsack type, with one essential difference: $\mathcal{R}(\mathbf{x})$ is not assumed to be a sum of independent project benefits. The network risk is recomputed for every selected combination, by enumerating the network states with all the selected changes applied together.

With four projects there are $2^4 = 16$ portfolios. With a budget of R\$ 100 000 per year, seven are feasible: no project, each project alone, A with B (R\$ 80 000) and B with C (R\$ 95 000). Every other pair exceeds the budget (A and C cost R\$ 105 000, B and D R\$ 115 000, A and D R\$ 125 000, C and D R\$ 140 000), and so does every larger combination. Table 35.11 lists the feasible portfolios with their cost, expected curtailment, gross and net annual benefit. The two combinations are computed in Table 35.9: for A and B, 0.606475 MW and a net benefit of $10\,000 \times 300 \times (1.3127 - 0.606475) - 80\,000 = \mathrm{R\$}\,2\,038\,675$ per year; for B and C, 0.324395 MW and R\$ 2 869 915 per year.

Table 35.11: The feasible portfolios under the budget $B = \mathrm{R\$}\,100\,000$ per year, with the network risk recomputed for each combination (illustrative).

Portfolio	Annualized cost (R$/yr)	Expected curtailment (MW)	Gross annual benefit (R$/yr)	Net annual benefit (R$/yr)
none	0	1.312700	0	0
A	45 000	0.910900	1 205 400	1 160 400
B	35 000	1.014425	894 825	859 825
C	60 000	0.334580	2 934 360	2 874 360
D	80 000	1.122580	570 360	490 360
A + B	80 000	0.606475	2 118 675	2 038 675
B + C	95 000	0.324395	2 964 915	2 869 915

Figure 35.6 plots the portfolios. The horizontal axis is the annualized cost in thousands of reais, the vertical axis the expected curtailment in MW; filled dots are feasible portfolios, open circles the pairs that exceed the budget, shown for reference, and the dashed vertical line is the budget. The lowest feasible point is B with C, at 0.324395 MW; C alone lies just above it at 60 000 instead of 95 000 reais.

The answer depends on the objective. If the objective is strictly to minimize the modeled expected unserved energy within the budget, the choice is B with C, at 0.324395 MW. If the objective is equation (35.30), the expected loss plus the annualized project cost, the choice is C

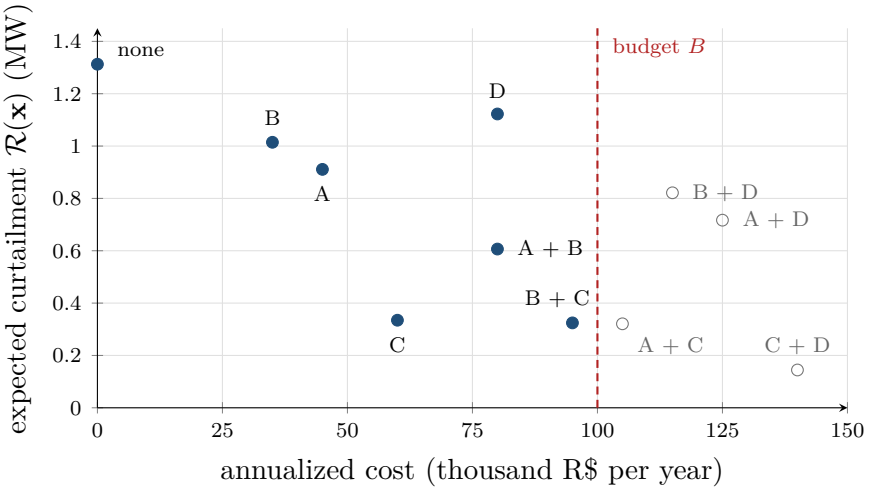


Figure 35.6: Annualized cost and expected curtailment of the portfolios of the four projects. Filled dots: the seven portfolios within the budget of R$ 100 000 per year (Table 35.11); open circles: pairs above the budget, each recomputed on the network. The lowest feasible risk is B + C; the largest net benefit is C alone (illustrative values).

alone, whose net benefit exceeds that of B with C by $2\,874\,360 - 2\,869\,915 =$ R$ 4 445 per year. Adding B after C reduces the expected curtailment by only $0.334580 - 0.324395 = 0.010185$ MW, worth $10\,000 \times 300 \times 0.010185 =$ R$ 30 555 per year, which does not quite cover B's annualized cost of R$ 35 000 under the chosen consequence valuation.

35.4.5 Interaction between projects

Project B alone reduces the expected curtailment by 0.298275 MW. After project C has been installed, it reduces it by

$$\Delta\mathcal{R}(\mathrm{B} \mid \mathrm{C}) = 0.334580 - 0.324395 = 0.010185\,\mathrm{MW}, \tag{35.31}$$

a thirtieth of its standalone value, so that

$$\Delta\mathcal{R}(\mathrm{B} \mid \mathrm{C}) \neq \Delta\mathcal{R}(\mathrm{B}). \tag{35.32}$$

The reason is physical. Project C has already eliminated the curtailment of every state in which L12 is out and L13 is in service, since L13 can then carry the whole transfer. After C, every state that still curtails has L13 out, and the expected curtailment is $q_{13}\,[30\,(1-u) + 100\,u] = q_{13}\,(30 + 70\,u)$. Improving L12 then only converts some islanding states (100 MW) into states with L13 alone out (30 MW): u falls from 0.0494 to 0.03485, and the reduction is $0.01 \times 70 \times 0.01455 = 0.010185$ MW. The two projects are partial substitutes: they protect against the same contingencies. Adding their standalone reductions would predict $0.298275 + 0.978120 = 1.276395$ MW, against the true 0.988305 MW, an overstatement of 0.28809 MW or R$ 864 270 per year.

Interactions can also run the other way. Projects A and B improve the two lines of the same indirect path, which is intact only when both are in service, so $u = 1 - (1 - q_{12})(1 - q_{23})$ falls by more when both probabilities fall together: $\Delta u = 0.0494 - 0.01495 = 0.03445$ against $0.0196 + 0.01455 = 0.03415$ for the two separately. Since the expected curtailment with present ratings is $30\,q_{13} + (20 + 50\,q_{13})\,u$, linear in u, the joint reduction exceeds the sum by $20.5 \times 0.0003 = 0.00615$ MW: A and B are mild complements. In general the benefit of a project depends on which others are in place, and portfolio optimization must account for this interaction.

Ranking standalone benefit-to-cost ratios and selecting the highest-ranked projects until the budget is spent can therefore produce a suboptimal portfolio. With the ratios of Table 35.10, that rule takes C (48.9), cannot afford A next (26.8; the total would be R$ 105 000), and adds B (25.6), reaching B with C, whose net benefit is R$ 4 445 per year below that of C alone. Had the standalone net benefits been added, the rule would have promised R$ 3 734 185 per year for a portfolio that delivers R$ 2 869 915.

Principle 35.1. *The benefit of a project depends on the projects already in place. Portfolio decisions recompute the system risk for every combination considered; standalone benefits, and ratios built from them, are not additive.*

35.5 The value of an upgrade against the value of information

The portfolio analysis treated the consequence value v as known. When it is not, a further option appears: to learn more before investing. This section distinguishes an upgrade from a measurement and works a two-scenario decision on project C through the decision without information, the expected value of perfect information and the value of an imperfect study.

35.5.1 Upgrade and measurement

A reliability upgrade changes the physical system. Raising the rating of L13, $P_{13}^{\max}$ from 80 to 100 MW, changes the largest feasible transfer of several network states, and the risk falls whether or not anyone learns anything. A measurement changes what is known about the system. A conductor-temperature sensor on L13 may reduce the uncertainty about the line's present or future thermal capability without adding a single megawatt to it. The two are different actions. An upgrade may reduce risk directly. A measurement may improve the choice of whether the upgrade is necessary, or allow a less costly operating decision. Its economic value depends on whether its result can change a decision: as Chapter 24 showed for an inspection (equation (24.18)), the expected value of sample information is the reduction in the expected loss of the best decision obtained by observing before deciding, and it is zero when every possible result leads to the same action. Chapter 32 treats the value of information in general.

35.5.2 Two consequence scenarios

Suppose project C is being considered, but the economic consequence of unavailable transfer is uncertain. The project costs $c_{\mathrm{C}} = \mathrm{R\$}\,60\,000$ per year, and its reduction of the annual expected unserved energy is

$$\Delta\mathrm{EUE} = 0.97812\,\mathrm{MW} \times 300\,\mathrm{h} = 293.436\,\mathrm{MWh\ per\ year}.$$

Two hypotheses are considered about the value of that energy, listed in Table 35.12: H_1, a high consequence of R$ 10 000 per MWh, with prior probability $\Pr(H_1) = 0.02$; and H_2, a low consequence of R$ 100 per MWh, with $\Pr(H_2) = 0.98$. These are hypothetical, epistemic uncertainties about the project's applicable economic value, for instance whether the load behind bus 3 has a costly process that a curtailment interrupts; they are not random line failures, which the expected curtailment already averages.

The gross annual benefit of the project under each hypothesis is

$$B_{\mathrm{gross}}(H_1) = 293.436 \times 10\,000 = \mathrm{R\$}\,2\,934\,360, \qquad B_{\mathrm{gross}}(H_2) = 293.436 \times 100 = \mathrm{R\$}\,29\,343.60.$$

Under the high-consequence hypothesis the project is clearly worthwhile; under the low-consequence hypothesis its gross benefit is less than its annualized cost of R$ 60 000.

Table 35.12: The two consequence scenarios for project C and the gross annual benefit of the project under each (hypothetical values).

Scenario	Value of unavailable energy	Probability	Gross benefit of C (R$/yr)
H_1, high consequence	R$ 10 000/MWh	0.02	2 934 360.00
H_2, low consequence	R$ 100/MWh	0.98	29 343.60

35.5.3 The decision without further information

Without further information, the expected gross benefit is

$$\mathbb{E}[B_{\text{gross}}] = 0.02 \times 2\,934\,360 + 0.98 \times 29\,343.60 = 58\,687.20 + 28\,756.73 = \text{R\$}\,87\,443.93,$$

and the expected net benefit

$$\mathbb{E}[B_{\text{net}}] = 87\,443.93 - 60\,000 = \text{R\$}\,27\,443.93 \text{ per year.} \tag{35.33}$$

The optimal present choice under expected monetary value is to implement project C. Because the benefit is linear in v, the same conclusion follows from the expected consequence value $\mathbb{E}[v] = 0.02 \times 10\,000 + 0.98 \times 100 = \text{R\$}\,298$ per MWh, against the break-even value $60\,000/293.436 = \text{R\$}\,204.47$ per MWh. Equivalently, the project pays when $\Pr(H_1)$ exceeds

$$p^* = \frac{60\,000 - 29\,343.60}{2\,934\,360 - 29\,343.60} = 0.01055,$$

and the prior 0.02 is less than twice that threshold: the decision depends on a small probability assigned to a high-impact outcome.

In the loss notation of Chapter 24, the actions are a_1, implement C, and a_0, do not, with losses $L(a, H)$ measured as cost minus benefit relative to doing nothing: $L(a_1, H_1) = 60\,000 - 2\,934\,360 = -\text{R\$}\,2\,874\,360$, $L(a_1, H_2) = 60\,000 - 29\,343.60 = \text{R\$}\,30\,656.40$, and $L(a_0, H) = 0$. The expected losses are $\rho(a_1) = -\text{R\$}\,27\,443.93$ and $\rho(a_0) = 0$, and the expected loss of the best present decision is $C_{\text{now}} = \min_a \rho(a) = -\text{R\$}\,27\,443.93$. (The loss $L(a, H)$ is not the demand L.)

35.5.4 Expected value of perfect information

If the true scenario were known before deciding, project C would be implemented under H_1, for a net benefit of R$ 2 874 360, and not implemented under H_2, for zero. The expected net benefit with perfect information is $0.02 \times 2\,874\,360 = \text{R\$}\,57\,487.20$, so $C_{\text{perfect}} = -\text{R\$}\,57\,487.20$, and

$$\text{EVPI} = C_{\text{now}} - C_{\text{perfect}} = 57\,487.20 - 27\,443.93 = \text{R\$}\,30\,043.27 \text{ per year.} \tag{35.34}$$

The same number reads directly as the loss that perfect information avoids: the present decision implements C in every case, and only under H_2 would perfect information change it, saving the shortfall of benefit against cost there,

$$\text{EVPI} = 0.98 \times (60\,000 - 29\,343.60) = 0.98 \times 30\,656.40 = \text{R\$}\,30\,043.27.$$

This is the largest expected improvement of the decision obtainable from knowing which of the two economic scenarios applies, before the cost of the information (Chapter 24, equation (24.19)).

35.5.5 An imperfect study

Suppose a targeted study can distinguish the two scenarios imperfectly. A positive result Y^+ is evidence favoring the high-consequence hypothesis. The study has sensitivity $\Pr(Y^+ \mid H_1) = 0.90$ and specificity $\Pr(Y^- \mid H_2) = 0.80$, so its false-positive probability is $\Pr(Y^+ \mid H_2) = 0.20$. By the law of total probability,

$$\Pr(Y^+) = 0.02 \times 0.90 + 0.98 \times 0.20 = 0.018 + 0.196 = 0.214, \qquad \Pr(Y^-) = 0.786.$$

By Bayes' theorem the posterior probability of the high-consequence hypothesis is, after a positive result,

$$\Pr(H_1 \mid Y^+) = \frac{0.02 \times 0.90}{0.214} = \frac{0.018}{0.214} = 0.0841,$$

or 8.41 %, and after a negative result,

$$\Pr(H_1 \mid Y^-) = \frac{0.02 \times 0.10}{0.786} = \frac{0.002}{0.786} = 0.00254,$$

or 0.254 %. Although the positive posterior remains small, it is eight times the threshold $p^* = 0.01055$, and the expected net benefit of C after a positive result is $0.0841 \times 2\,934\,360 + 0.9159 \times 29\,343.60 - 60\,000 =$ R\$ 213 690.77: implementing C is favorable. After a negative result the posterior is below the threshold, and the expected net benefit is $0.00254 \times 2\,934\,360 + 0.99746 \times 29\,343.60 - 60\,000 =$ −R\$ 23 264.50: C does not justify its cost. The optimal policy is

$$\text{implement C if } Y^+, \qquad \text{do not implement C if } Y^-. \tag{35.35}$$

Its expected net benefit is obtained by summing over the joint outcomes in which C is implemented, the positive results,

$$\begin{aligned} \mathbb{E}[B_{\text{policy}}] &= \Pr(H_1, Y^+)\, B_{\text{gross}}(H_1) + \Pr(H_2, Y^+)\, B_{\text{gross}}(H_2) - \Pr(Y^+)\, c_{\text{C}} \\ &= 0.018 \times 2\,934\,360 + 0.196 \times 29\,343.60 - 0.214 \times 60\,000 \\ &= 52\,818.48 + 5\,751.35 - 12\,840.00 = \text{R\$ } 45\,729.83, \end{aligned} \tag{35.36}$$

so that $C_{\text{sample}} = -$R\$ 45 729.83, and the expected value of the study's information is

$$\text{EVSI} = C_{\text{now}} - C_{\text{sample}} = 45\,729.83 - 27\,443.93 = \text{R\$ } 18\,285.90 \text{ per year.} \tag{35.37}$$

The study captures 18 285.90/30 043.27 = 61 % of the perfect-information bound. A study costing less than R\$ 18 285.90, on the same annualized decision basis, has a positive expected net decision value under this model.

Figure 35.7 draws the decision as a tree. Squares are decision nodes, where a branch is chosen; circles are chance nodes, whose branches carry probabilities; the numbers at the right are the net annual benefits at the end of each path, in reais, those after the study being expected values over the posterior. The boxed numbers are rolled back from the right, taking the probability-weighted average at a chance node and the maximum at a decision node. Thick branches are the ones chosen: study, then implement after a positive result and not after a negative one.

Interpretation. The study is valuable not because it makes the system physically more reliable, but because it helps to avoid an economically unfavorable investment when the low-consequence scenario is true: it moves the expected loss of the H_2 branch, without touching the network. That distinction matters in instrumentation and model-development decisions. A sensor may be valuable even though installing it does not reduce any failure probability, and an upgrade that does reduce one may be worth less than a study that decides whether it is needed.

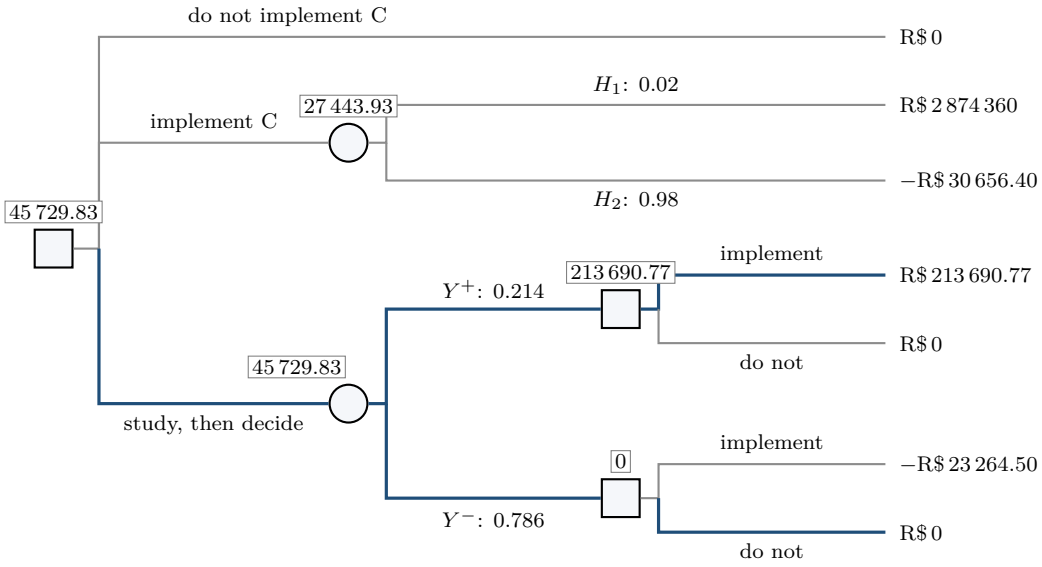


Figure 35.7: Decision tree of project C under uncertain consequence value. Squares are decisions, circles chance outcomes with their probabilities; net annual benefits in reais at the right, after the study as expected values over the posterior ($\Pr(H_1 \mid Y^+) = 0.0841$, $\Pr(H_1 \mid Y^-) = 0.00254$); boxed values are rolled back. The study cost is excluded; the study branch exceeds the best branch without it by the EVSI = R$ 18 285.90 per year (hypothetical values).

35.6 Uncertainty in risk importance and the engineering decision framework

The investment calculations so far assumed that the line unavailabilities, the physical consequences and their valuation were known. This section treats the uncertainty of the rankings, the case in which the risk reduction is not identifiable, and the choice between expected-value and risk-constrained optimization, and then assembles the chapter's steps into a workflow and separates attribution from optimization.

35.6.1 Uncertain importance rankings

In practice the inputs of the risk model are estimates. If the unavailabilities q_{12}, q_{23} and q_{13} have posterior uncertainty, the risk reduction of an intervention a is itself uncertain,

$$\Delta\mathcal{R}(a, \boldsymbol{\theta}) = \mathcal{R}_0(\boldsymbol{\theta}) - \mathcal{R}_a(\boldsymbol{\theta}), \tag{35.38}$$

where $\boldsymbol{\theta}$ collects the uncertain failure, repair, weather, rating and consequence parameters (the parameter vector of Chapter 31, not the bus angles). Instead of the single number $\Delta\mathcal{R}(\mathrm{A}) = 0.4018\,\mathrm{MW}$ a study reports the posterior mean of the reduction, a credible interval, the probability that the intervention meets a target, and the probability that one intervention is better than another,

$$\Pr(\Delta\mathcal{R}(\mathrm{A}) > \Delta\mathcal{R}(\mathrm{B}) \mid D), \tag{35.39}$$

a probability about relative intervention performance under the model and the observed data D.

As an illustration, let the three unavailabilities be independent with beta posteriors whose means are the values used so far and whose standard deviations are half of them, as in the propagation of Chapter 27 (Section 27.6.4): $q_{12} \sim \mathrm{Beta}(3.90, 191.10)$, $q_{23} \sim \mathrm{Beta}(3.85, 124.48)$ and $q_{13} \sim \mathrm{Beta}(3.95, 391.05)$, whose parameters follow from the mean m and standard deviation

s as $m\,k$ and $(1-m)\,k$ with $k = m(1-m)/s^2 - 1$. Represent projects A and B as proportional improvements, A dividing q_{23} by three and B dividing q_{12} by four, which reproduce the stated values at the point estimates. The closed form then gives

$$\Delta\mathcal{R}(\mathrm{A}) = \tfrac{2}{3}\,q_{23}(1-q_{12})(20+50\,q_{13}), \qquad \Delta\mathcal{R}(\mathrm{B}) = \tfrac{3}{4}\,q_{12}(1-q_{23})(20+50\,q_{13}),$$
$$\Delta\mathcal{R}(\mathrm{C}) = 20\,(1-q_{13})\,u,$$

in MW. Table 35.13 summarizes 4×10^6 samples from the posteriors. Each row is one project, with its point estimate, posterior mean and 90 % credible interval (the 5 % and 95 % quantiles). Because each reduction is multilinear in independent unavailabilities, its posterior mean equals its value at the mean parameters, and the first two columns agree. The intervals are wide, roughly a factor of six from end to end, and those of A and B overlap over most of their length.

Table 35.13: Posterior summaries of the risk reductions of projects A, B and C, with beta posteriors of the unavailabilities whose standard deviations are half their means (4×10^6 samples; illustrative).

Project	Point value (MW)	Posterior mean (MW)	90 % credible interval (MW)
A, L23 reliability	0.4018	0.4018	0.136 to 0.779
B, L12 reliability	0.2983	0.2983	0.101 to 0.578
C, L13 rating	0.9781	0.9782	0.482 to 1.614

The ranking probabilities say more than the intervals. The posterior probability that A outperforms B is $\Pr(\Delta\mathcal{R}(\mathrm{A}) > \Delta\mathcal{R}(\mathrm{B}) \mid D) = 0.654$ (a one-dimensional numerical integration over q_{12} gives 0.6540, confirming the sampling): the point estimates prefer A by a third, but the data support that preference only at about two to one. The probability that C outperforms A is one to within the sampling, since in every sample C's reduction exceeds A's by at least a factor 1.36: C mitigates the curtailment of both L12 and L23 outages, and its advantage does not depend on which of them is more frequent. The probability that adding B after C pays its cost, $\Pr(v H_{\mathrm{rel}}\,\Delta\mathcal{R}(\mathrm{B} \mid \mathrm{C}) > 35\,000)$, is 0.31: the conclusion of Section 35.4.4 that C alone is preferred holds in about two thirds of the posterior.

35.6.2 Non-identifiable risk reduction

Suppose the benefit of upgrading a transmission line depends on its uncertain weather-dependent thermal rating. If the available measurements cannot distinguish stronger conductor cooling from a bias of the conductor-temperature sensor, the parameters that set the rating are not identifiable (Chapter 31), and the predicted benefit of the upgrade depends strongly on which explanation the model assumes. An optimizer still returns a preferred investment when it is given numerical inputs, but its answer may not be robust to plausible alternative explanations of the same data. This is why the measurement analyses of the preceding chapters belong to the decision chain:

$$\text{identifiability (Chapter 31)} \rightarrow \text{credible posterior (Chapters 30 and 32)} \rightarrow \text{risk prediction} \rightarrow \text{investment decision},$$

with observability (Chapter 29) as the first condition of the first link. If a key parameter is weakly identified, the corresponding risk ranking may remain uncertain however the optimizer is configured. A useful engineering recommendation therefore places each option in one of four categories:

- a robustly preferred intervention, such as C against A above;

- an intervention preferred only under particular assumptions, such as A against B;
- a decision requiring more measurements, when the value of information exceeds its cost;
- a decision in which no available option satisfies the risk constraints.

35.6.3 Expected-value and risk-constrained optimization

Minimizing the expected loss may not satisfy every engineering requirement. A portfolio with a small expected unserved energy can keep a non-negligible probability of an extremely severe outage, and an organization may impose a limit on that probability, a *chance constraint*,

$$\Pr(E_{\text{loss}} > E_{\text{crit}}) \le \varepsilon, \tag{35.40}$$

where E_{loss} is the random unserved energy over the period assessed, E_{crit} a critical level and ε a tolerated probability, or a limit on the tail of the monetary loss,

$$\text{CVaR}_{\alpha}(C_{\text{loss}}) \le C_{\max}. \tag{35.41}$$

The *value at risk* of a loss X at level α is its α-quantile, $\text{VaR}_{\alpha}(X) = \min\{x : \Pr(X \le x) \ge \alpha\}$. The *conditional value at risk* measures the expected loss in the upper tail beyond that quantile. For a continuous loss it is $\mathbb{E}[X \mid X \ge \text{VaR}_{\alpha}(X)]$; for a discrete loss, whose distribution has atoms at the quantile, the form that handles the atoms correctly is

$$\text{CVaR}_{\alpha}(X) = \text{VaR}_{\alpha}(X) + \frac{\mathbb{E}\big[(X - \text{VaR}_{\alpha}(X))^{+}\big]}{1 - \alpha}, \tag{35.42}$$

the mean of the worst $100(1-\alpha)\,\%$ of the probability mass, with a split atom counted in part. In finance the same quantity is also called the expected shortfall, which is not the expected shortfall $\mathbb{E}[S]$ of the adequacy chapters, the mean of the whole curtailment distribution. The general constrained formulation is

$$\min_{\mathbf{x}} \Big\{ \mathbb{E}[C_{\text{loss}}(\mathbf{x})] + C_{\text{investment}}(\mathbf{x}) \Big\} \quad \text{subject to} \quad \sum_{i} c_i x_i \le B, \quad \Pr(E_{\text{loss}} > E_{\text{crit}} \mid \mathbf{x}) \le \varepsilon, \tag{35.43}$$

together with the applicable reliability, operational and safety requirements.

Table 35.14 applies tail measures to the feasible portfolios of Table 35.11, using the distribution of the curtailment in one assessed hour. Each row is a portfolio, with its net annual benefit, its probability of complete interruption $\Pr(\Delta L = 100\,\text{MW})$ (the islanding probability $q_{13}\,u$), and its value at risk and conditional value at risk at $\alpha = 0.999$. For the baseline, $\Pr(\Delta L \le 20\,\text{MW}) = 0.99 < 0.999$ and $\Pr(\Delta L \le 30\,\text{MW}) = 0.999506 \ge 0.999$, so $\text{VaR}_{0.999} = 30\,\text{MW}$ and $\text{CVaR}_{0.999} = 30 + 0.000494 \times 70/0.001 = 64.58\,\text{MW}$. Project C, the economic optimum, leaves the tail exactly as it is: it mitigates the moderate 20 MW states and does nothing for the islanding states. A chance constraint $\Pr(\Delta L = 100\,\text{MW}) \le 4 \times 10^{-4}$ per assessed hour, or equivalently here $\text{CVaR}_{0.999} \le 55\,\text{MW}$, excludes C alone, D and the baseline; the best feasible portfolio is then B with C, which buys the tail improvement for R\$ 4 445 per year of expected net benefit. With the tighter limit $\text{CVaR}_{0.999} \le 50\,\text{MW}$ only A with B remains, at a cost of R\$ 835 685 per year of expected net benefit against C alone.

The economically optimal portfolio and the risk-constrained optimal portfolio therefore differ, and which is right depends on criteria that the model does not supply. Neither answer is meaningful until the decision criteria, the objective and the constraints, are defined.

35.6.4 A risk-informed investment workflow

The steps of this chapter form a repeatable workflow, drawn in Figure 35.8. Each box is one step, numbered in the order of execution; the arrows run along the top row from left to right

Table 35.14: Tail measures of the hourly curtailment for the feasible portfolios: probability of complete interruption, value at risk and conditional value at risk at $\alpha = 0.999$ (illustrative).

Portfolio	Net benefit (R$/yr)	$\Pr(\Delta L = 100\,\text{MW})$	$\text{VaR}_{0.999}$ (MW)	$\text{CVaR}_{0.999}$ (MW)
none	0	0.0004940	30	64.58
A	1 160 400	0.0002980	30	50.86
B	859 825	0.0003485	30	54.40
C	2 874 360	0.0004940	30	64.58
D	490 360	0.0004940	20	59.52
A + B	2 038 675	0.0001495	30	40.47
B + C	2 869 915	0.0003485	30	54.40

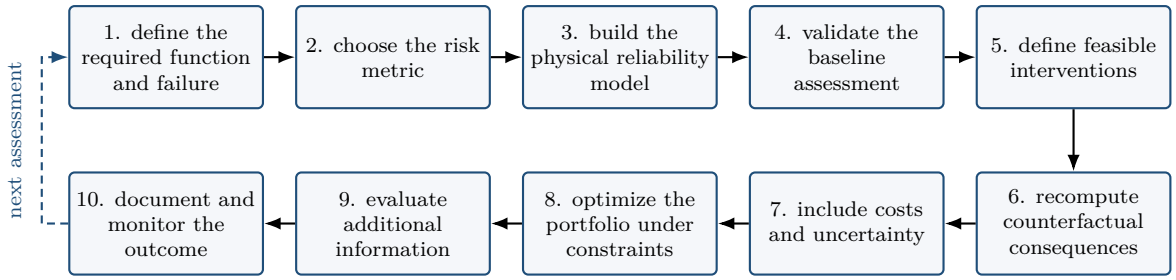


Figure 35.8: The risk-informed investment workflow: ten steps from the definition of the required function to the monitoring of the outcome, which feeds the next assessment.

and along the bottom row from right to left, and the return arrow from the last step to the first marks that the monitored outcome feeds the next assessment.

1. *Define the required system function*: specify what constitutes successful operation, failure, curtailment or interruption (Chapter 2 of Volume I).
2. *Choose the risk metric*: availability, outage frequency, EUE, expected loss, tail risk, or a stated combination.
3. *Construct the physical reliability model*: component states, dependencies, constraints, restoration and operating conditions.
4. *Validate the baseline assessment*: data quality, identifiability, model adequacy and the numerical implementation.
5. *Define feasible interventions*: which parameters, capacities, controls or dependencies each investment changes.
6. *Recompute the counterfactual system consequences*: the risk after each intervention and after each relevant combination.
7. *Include costs and uncertainty*: life-cycle cost, parameter uncertainty, model uncertainty and the confidence in each benefit.
8. *Optimize the intervention portfolio*: apply the budget, safety, availability and operational constraints.
9. *Evaluate additional information*: determine whether a further measurement could change the preferred decision, and whether its value exceeds its cost.
10. *Document and monitor the outcome*: retain the evidence, the assumptions and the acceptance criteria, and compare the post-implementation performance with the prediction.

35.6.5 Attribution and optimization

The chapter has answered three different questions, and Table 35.15 pairs each with the quantity appropriate to it. Which component has the greatest marginal effect on system reliability is answered by the Birnbaum importance. Which underlying failure mechanism contributes most to the chosen risk is answered by a risk importance measure, such as the Fussell–Vesely importance or the state contributions of Chapter 27, or by a causal attribution. Which feasible investment to implement is answered only by counterfactual risk reduction and constrained decision optimization.

Table 35.15: Three questions of a reliability investment study and the quantity appropriate to each.

Question	Appropriate quantity
Which component has the greatest marginal effect on system reliability?	Birnbaum importance
Which underlying failure mechanism contributes most to the chosen risk?	Risk importance (Fussell–Vesely, RRW, state contributions) or causal attribution
Which feasible investment should be implemented?	Counterfactual risk reduction and constrained decision optimization

Three cautions follow. A high-importance component does not automatically justify an expensive upgrade: L13 has the largest Birnbaum importance of the three lines, and the most valuable project is not an improvement of its availability but an increase of its rating. A high-risk state is not necessarily preventable by the proposed intervention: the islanding states, which carry the tail of the curtailment distribution, are untouched by the economically optimal project. And a component improvement that appears beneficial in isolation may provide little additional value after other projects are installed, as project B after project C. A defensible reliability investment study connects the component analysis to realistic counterfactual engineering actions. The last chapter, Chapter 36, assembles the methods of the book into one complete assessment, from a physical system and its observations to an evidence-supported intervention decision.

35.7 What to retain

Definition 35.1 (Risk importance and decision optimization)**.** Importance measures describe a component's role in the present model: the Birnbaum importance $\partial R_{\text{system}}/\partial p_i$ its sensitivity, $\text{RRW}_i = Q/Q(0_i)$ the potential of its perfect reliability, $\text{RAW}_i = Q(1_i)/Q$ the vulnerability to its loss. Interventions are chosen by credible system-level benefit, not by component-level importance alone: by the counterfactual reduction $\Delta\mathcal{R}(a) = \mathcal{R}_0 - \mathcal{R}(a)$, recomputed for every candidate and every combination, set against cost under the stated budget and risk constraints.

Reliability importance depends on the system architecture, and the most failure-prone component is not always the most important. Feasible interventions are evaluated through their actual effects on the physical system, and increasing capacity, reducing failure frequency, shortening restoration and reducing consequence severity are different reliability strategies. Intervention benefits overlap and interact, so portfolio decisions recompute the risk of each combination. Optimizing reliability and optimizing financial value are different objectives, and budget and risk constraints set the tradeoff. A measurement can have economic value by improving a decision without improving any equipment, and uncertain parameters make the rankings themselves uncertain, so the preferred intervention is assessed for robustness and credibility. A reliability model becomes a decision tool only when its physical assumptions, risk metrics, interventions and costs are explicitly connected. Table 35.16 lists the chapter's terms.

Table 35.16: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Reliability function	system reliability from component probabilities	$R_{\text{system}} = h(\mathbf{p})$	–	35.1.1
Birnbaum importance	sensitivity; probability component is critical	$\partial R_{\text{system}}/\partial p_i = \partial Q/\partial q_i$	–	35.1.2
Risk metric	quantity the measures refer to	Q, $\mathbb{E}[C_{\text{loss}}]$, EUE	varies	35.2.1
Risk reduction worth	factor of reduction, component perfect	$Q/Q(0_i)$	–	35.2.2
Risk achievement worth	factor of increase, component failed	$Q(1_i)/Q$	–	35.2.2
Rate sensitivities	unavailability per unit rate	$\mu/(\lambda+\mu)^2$, $-\lambda/(\lambda+\mu)^2$	h, yr	35.2.4
Elasticity	relative change per relative change	$E_{q,\lambda} = 1-q$, $E_{q,\mu} = -(1-q)$	–	35.2.5
System-level sensitivity	chain rule to the rates	$I_i^{\text{B}}\,\partial q_i/\partial\lambda_i$	varies	35.2.6
Intervention	quantities set by an action	$\mathcal{R}(a) = \mathbb{E}[\Delta L \mid \text{do}(a)]$	MW	35.3.1
Counterfactual risk reduction	benefit of an action	$\Delta\mathcal{R}(a) = \mathcal{R}_0 - \mathcal{R}(a)$	MW	35.3.1
Consequence mitigation	lower consequence, same outages	e.g. rating increase	–	35.3.5
Relevant hours, consequence value	valuation of the curtailment	H_{rel}, v	h/yr, R\$/MWh	35.4.1
Gross and net benefit	annual value of a project	$vH_{\text{rel}}\,\Delta\mathcal{R}(a)$, minus c_a	R\$/yr	35.4.3
Budget-constrained selection	0–1 portfolio choice	problem (35.30)	R\$/yr	35.4.4
Interaction	benefit depends on other projects	$\Delta\mathcal{R}(\text{B} \mid \text{C}) \neq \Delta\mathcal{R}(\text{B})$	MW	35.4.5
EVPI	value of resolving the uncertainty	$C_{\text{now}} - C_{\text{perfect}}$	R\$/yr	35.5.4
EVSI	value of an imperfect study	$C_{\text{now}} - C_{\text{sample}}$	R\$/yr	35.5.5
Ranking probability	posterior support for an ordering	$\Pr(\Delta\mathcal{R}(\text{A}) > \Delta\mathcal{R}(\text{B}) \mid D)$	–	35.6.1
Chance constraint	limit on a tail probability	$\Pr(E_{\text{loss}} > E_{\text{crit}}) \le \varepsilon$	–	35.6.3
VaR, CVaR	tail quantile; mean beyond it	equation (35.42)	loss units	35.6.3

Further reading. *Risk Assessment and Risk Management for the Chemical Process Industry*, for risk contribution, consequences and mitigation decisions; *Reliability Evaluation of Engineering Systems*, for reliability importance and system analysis; *Risk Analysis: A Quantitative Guide*, for uncertainty propagation and risk-based decisions; *Decision Analysis for Management Judgment*, for decision criteria, the value of information and alternatives under uncertainty.

35.8 Exercises

Exercise 35.1. (a) State what the Birnbaum importance measures. (b) Define the risk reduction worth. (c) Define the risk achievement worth. (d) Explain why increasing a transmission line's

rating can outperform improving its failure probability. (e) Explain why the benefits of a portfolio cannot always be added. (f) State the difference between an upgrade and a measurement. (g) State what the EVPI provides. (h) Explain why the economically optimal portfolio can differ from the portfolio with the lowest EUE. (Answers: (a) the sensitivity of the system reliability to a component's reliability, $\partial R_{\text{system}}/\partial p_i$; (b) the baseline system risk divided by the risk with the specified component made perfectly reliable; (c) the risk with the component forced unavailable divided by the baseline risk; (d) a rating increase can mitigate the consequences of several outage configurations at once, as project C does for the outages of L12, of L23 and of both; (e) projects may overlap, substitute for one another or interact physically, as $\Delta\mathcal{R}(\text{B} \mid \text{C}) = 0.010185\,\text{MW}$ against $\Delta\mathcal{R}(\text{B}) = 0.298275\,\text{MW}$; (f) an upgrade changes the physical system, a measurement primarily changes knowledge; (g) an upper bound on the decision improvement obtainable by perfectly resolving the specified uncertainty; (h) the costs and the consequence valuation enter the economic objective, so a cheaper portfolio with slightly higher EUE can have the larger net benefit, as C alone against B + C.)

Exercise 35.2. A system has component 1 in series with a parallel pair of components 2 and 3, all independent, with $p_1 = 0.95$ and $p_2 = p_3 = 0.90$. (a) Compute R_{system} and Q. (b) Compute the Birnbaum importances of components 1 and 2. (c) Compute the RRW and the RAW of components 1 and 2, with Q as the risk metric. (d) In the three-bus network, line L12 is replaced by a line with failure rate $\lambda = 2\,\text{yr}^{-1}$ and mean outage duration 50 h. Using $I^{\text{B}}_{12} = 19.885\,\text{MW}$ and the expected curtailment 0.915 MW with L12 perfect, compute its unavailability, its two elasticities and the expected curtailment. Compare halving the failure rate with halving the outage duration, for the expected curtailment and for the outage frequency $f = \lambda(1 - q)$. (Answers: (a) $R_{\text{system}} = 0.95 \times (1 - 0.1^2) = 0.9405$, $Q = 0.0595$; (b) $I^{\text{B}}_1 = 1 - (0.1)(0.1) = 0.99$, $I^{\text{B}}_2 = p_1(1 - p_3) = 0.095$; (c) component 1: $Q(0_1) = 0.01$, $Q(1_1) = 1$, $\text{RRW}_1 = 5.95$, $\text{RAW}_1 = 16.81$; component 2: $Q(0_2) = 0.05$, $Q(1_2) = 1 - 0.95 \times 0.90 = 0.145$, $\text{RRW}_2 = 1.19$, $\text{RAW}_2 = 2.44$; (d) $\mu = 8760/50 = 175.2\,\text{yr}^{-1}$, $q = 2/177.2 = 0.011287$, $E_{q,\lambda} = -E_{q,\mu} = 0.9887$, $\mathbb{E}[\Delta L] = 0.915 + 19.885 \times 0.011287 = 1.1394\,\text{MW}$; both halvings give exactly $q = 1/176.2 = 0.0056754$ and 1.0279 MW, because q depends only on λ/μ, but halving the failure rate lowers the outage frequency from 1.977 to $0.994\,\text{yr}^{-1}$, while halving the duration leaves it at $1.989\,\text{yr}^{-1}$ with outages of 25 h.)

Exercise 35.3. A fifth project, E, reduces the unavailability of L13 from 0.01 to 0.005 at an annualized cost of R$ 25 000, with $H_{\text{rel}} = 300$ h per year and $v =$ R$ 10 000 per MWh. (a) Compute the expected curtailment with E alone, its reduction and its gross and net annual benefits. (b) Compute the expected curtailment with C and E together, and the reduction that E adds after C; compare it with E's standalone reduction and explain the difference. (c) List the feasible portfolios that contain E under the budget of R$ 100 000 per year, and find the portfolio with the largest net benefit among all feasible ones. (Answers: (a) the expected curtailment is linear in q_{13} with slope $I^{\text{B}}_{13} = 32.47\,\text{MW}$, so $\mathcal{R} = 1.3127 - 32.47 \times 0.005 = 1.15035\,\text{MW}$, a reduction of 0.16235 MW, $B_{\text{gross}} =$ R$ 487 050 and $B_{\text{net}} =$ R$ 462 050; (b) after C the expected curtailment is $q_{13}(30 + 70u) = 0.005 \times 33.458 = 0.16729\,\text{MW}$, so E adds $0.33458 - 0.16729 = 0.16729\,\text{MW}$, more than its standalone 0.16235 MW: after C every remaining curtailment state has L13 out, the risk is proportional to q_{13}, and halving q_{13} halves it, so C and E are complements; (c) E (R$ 25 000, net R$ 462 050), A + E (R$ 70 000, 0.753450 MW, net R$ 1 607 750), B + E (R$ 60 000, 0.8557125 MW, net R$ 1 310 962.50) and C + E (R$ 85 000, 0.16729 MW, net R$ 3 351 230); C + E has the largest net benefit and the lowest expected curtailment of all feasible portfolios.)

Exercise 35.4. In the consequence problem of Section 35.5, let the prior probability of the high-consequence hypothesis be $\Pr(H_1) = 0.01$ instead of 0.02, everything else unchanged. (a) Compute the expected net benefit of project C and the best present decision. (b) Compute the EVPI. (c) For the study with sensitivity 0.90 and specificity 0.80, compute $\Pr(Y^+)$, the two

posteriors, the optimal policy and the EVSI. (d) Compare the EVSI with the value at a prior of 0.02 and explain the difference. (Answers: (a) $\mathbb{E}[B_{\text{gross}}] = 0.01 \times 2\,934\,360 + 0.99 \times 29\,343.60 =$ R\$ 58 393.76, a net benefit of −R\$ 1 606.24, so C is not implemented and $C_{\text{now}} = 0$; (b) EVPI $= 0.01 \times 2\,874\,360 =$ R\$ 28 743.60; (c) $\Pr(Y^{+}) = 0.009 + 0.198 = 0.207$, $\Pr(H_1 \mid Y^{+}) = 0.0435$ (net benefit of C R\$ 95 648.66, implement), $\Pr(H_1 \mid Y^{-}) = 0.00126$ (net −R\$ 26 993.08, do not); the policy is worth $0.009 \times 2\,934\,360 + 0.198 \times 29\,343.60 - 0.207 \times 60\,000 =$ R\$ 19 799.27, which is the EVSI; (d) the EVSI rises from R\$ 18 285.90 to R\$ 19 799.27 although the prior halves, because the prior 0.01 lies just below the break-even 0.01055, where a positive result reverses the present decision; information is worth most near a decision boundary.)

Chapter 36

Integrated Reliability Assessment

36.1 An integrated reliability assessment

The two volumes of this book have developed reliability engineering from several directions. Volume I treated component failure statistics and lifetime models (Chapters 5 to 8 of Volume I); reliability block diagrams, fault trees, redundancy and common-cause failure (Chapters 9 to 14); event trees and integrated probabilistic risk assessment (Chapters 15 and 16); and repairable systems, Markov models and availability (Chapters 17 to 20). Volume II added degradation, fault diagnosis and remaining useful life (Chapters 21 to 24); power-system adequacy and contingency risk (Chapters 25 to 28); observability, state estimation and parameter identifiability (Chapters 29 to 31); Bayesian updating, model adequacy and the value of information (Chapter 32); and reliability standards, investment optimization and engineering assurance (Chapters 33 to 35). Each method answers part of an engineering question. The methods become most useful when they are applied together, and this last chapter applies them together to one problem. Its question is

whether a defensible, end-to-end reliability assessment can be constructed that states what is likely to fail, why it matters, what remains uncertain, and which action is justified.

The problem is an apparent cooling-performance anomaly in a data center. The objective is not merely to calculate a probability but to reach a decision that can be explained, tested, challenged and revised when new evidence arrives. Everything that follows is an illustrative engineering case with hypothetical data and explicitly assumed probabilities. The calculations are reproducible, but their results are not validated performance estimates for an actual facility.

Jargon: integrated reliability assessment, interruption event, competing hypotheses, predictive probability, counterfactual, acceptance criterion, verification, validation, model credibility, decision statement

An *integrated reliability assessment* carries one defined system from its required function and failure event, through its physical and reliability models, the audit of its measurements, the diagnosis of what is wrong, the risk and its consequences, and the comparison of interventions, to a documented decision. The *interruption event* is the loss of the required function itself, here the loss of IT service because a thermal limit is exceeded, as distinct from the equipment and sensor events that may lead to it. *Competing hypotheses* are mutually exclusive physical explanations of the evidence, each with a probability; a *predictive probability* averages a conditional probability over them. A *counterfactual* is the model's prediction of what would happen under an action that has not been taken. An *acceptance criterion* (risk constraint) is an approved limit that a risk measure must satisfy, independently of cost. *Verification* checks that the model is computed correctly; *validation* checks, against independent evidence, that it represents the physical system; *model credibility* is the degree to which both have been established for the intended use. A *decision statement* is the written conclusion of the assessment: the recommended action, the evidence and assumptions it rests on, and the conditions under which it would change.

Key idea

A reliability decision is defensible when a traceable chain connects the physical model, the failure model and the observations to quantified uncertainty and consequences, and those to the preferred action; an assessment that cannot resolve the diagnosis can still decide, provided it carries the competing explanations through to the risk instead of silently selecting one.

This section lays out the eight steps of the assessment, describes the facility and the reliability question, defines the failure event, and lists the assumptions on which every later number rests.

36.1.1 The eight steps of the assessment

The assessment proceeds in eight steps, drawn in Figure 36.1. Each box is one step; its first line names the step, and its second line gives the section of this chapter in which it is carried out. The solid arrows give the order, along the upper row from left to right and then along the lower row from right to left. The dashed arrow from step 7 back to step 3 marks the loop that a recommended measurement opens: once the measurement is taken, its result enters the audit of the evidence and the diagnosis is repeated.

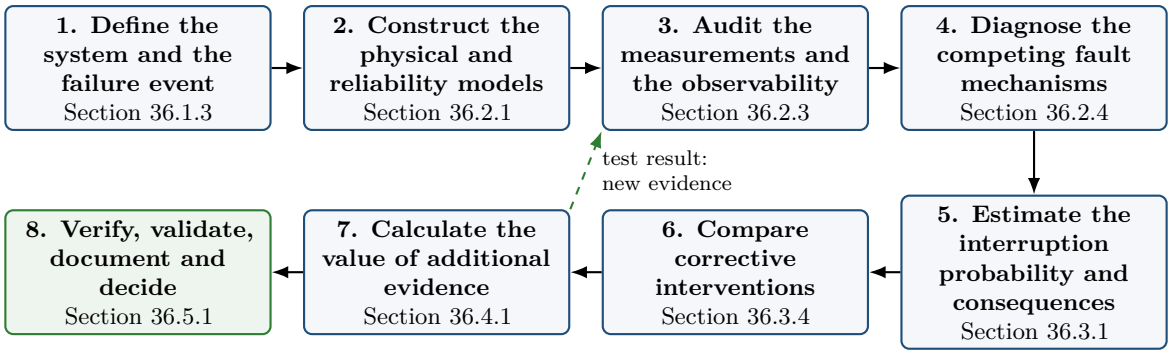


Figure 36.1: The eight steps of the integrated reliability assessment and the sections of this chapter that carry them out. The dashed arrow is the return of a recommended measurement's result to the audit of the evidence.

Steps 1 and 2 fix what is being assessed and how it behaves physically; steps 3 and 4 establish what the evidence can and cannot say about its present condition; steps 5 and 6 turn the remaining uncertainty into risk and compare the actions that would reduce it; step 7 asks whether it pays to reduce the uncertainty first; and step 8 establishes how far the whole chain can be trusted and writes it down. The order matters: each step consumes the output of the one before it, and an error in the definition of the failure event in step 1 propagates unchanged to the decision in step 8.

36.1.2 The physical system

The facility is the data center of Chapter 34: an IT workload of 800 kW served by a chilled-water cooling plant. Its main infrastructure comprises the electrical supply with its uninterruptible power supply (UPS) and standby generation; chilled-water production; two circulating pumps, either of which can meet the normal hydraulic demand; the cooling-water distribution and the computer-room air handlers (CRAH units), which transfer the heat of the room air to the chilled water; temperature, flow, pressure and electrical-load sensors; and the supervisory controls and the maintenance personnel who operate them.

The assessment concentrates on one apparent cooling-performance anomaly while preserving its dependencies on electrical power, on pumping and on the heat generated by the IT equipment. The system boundary is therefore the thermal environment of the IT equipment and the plant that removes heat from it; the electrical system enters only as the source of the IT heat and as a support that the cooling plant needs, and its own interruption modes are outside the scope.

36.1.3 The interruption event

The required IT operating condition is that the representative inlet temperature of the IT equipment, $T_{\mathrm{in}}(t)$, stay at or below a limit throughout a specified operating mission,

$$T_{\mathrm{in}}(t) \le T_{\max} = 33^{\circ}\mathrm{C}, \tag{36.1}$$

the limit of Chapter 34. The value 33°C is hypothetical and is not a universal safe limit for data centers. The failure event of the assessment is the loss of the required function,

$$F = \{\text{required IT service is lost because the thermal limit is exceeded}\}. \tag{36.2}$$

This event is distinct from a failed temperature sensor, a reduced pump capacity, a chiller fault, a brief cooling-equipment outage, and an elevated but still permissible temperature. Each of these may contribute to F, but none is equivalent to it: a chiller fault that is restored before the room reaches $T_{\max}$ is not an interruption, and a sensor that reads high while the room is cool is not one either. Defining the failure event as the loss of the required function, and not as a component state, is the first decision of the assessment, as Chapter 2 of Volume I established.

The assessment horizon is one operating year, $H = 1\,\mathrm{yr}$. The assessment seeks the probability $\Pr(F \text{ occurs during the year})$ and, more specifically, the expected consequences of defined cooling disturbances. It must also support a decision: whether to repair equipment, add cooling capacity, reduce the IT heat load, or first collect additional measurements.

Remark 36.1. The symbols of this chapter keep their meanings from earlier chapters: T_{in} is a temperature, not the random time to failure T; the heat rates $\dot{Q}$ are unrelated to the top-event probability Q_{top}; the thermal capacitance C_{eff} is not a consequence C; the hypotheses H_{S}, H_{C} and H_{H} are not the horizon H; and the test outcomes Y^{+} and Y^{-} are not the service indicator $Y(t)$ of Chapter 34.

36.1.4 Assessment assumptions

Table 36.1 lists the principal assumptions. Each row gives a quantity, its symbol and its hypothetical value. The IT electrical load is also the IT heat production, since the electrical power drawn by IT equipment is dissipated as heat in the room; the other internal heat gains (lighting, fans, the distribution losses inside the room, people) add 100 kW. The thermal capacitance lumps the room air, the equipment mass and the thermally active part of the chilled-water loop into one effective value.

Table 36.1: Principal assumptions of the assessment (hypothetical values).

Quantity	Symbol	Hypothetical value
IT electrical load (IT heat production)	$\dot{Q}_{\mathrm{IT}}$	800 kW
Other internal heat gains	$\dot{Q}_{\mathrm{other}}$	100 kW
Total baseline heat gain	$\dot{Q}_{\mathrm{heat}}$	900 kW
Nominal cooling capacity	$\dot{Q}_{\mathrm{cool}}$	950 kW
Suspected degraded cooling capacity	$\dot{Q}_{\mathrm{cool}}$	650 kW
Initial representative temperature	T_0	25°C
Hypothetical thermal limit	$T_{\max}$	33°C
Effective thermal capacitance	C_{eff}	120 MJ/K
Assessment horizon	H	one year
Critical cooling-disturbance frequency	f_{IE}	$1\,\mathrm{yr}^{-1}$

The disturbance frequency of one per year is an assumed rate for the defined risk scenario, the initiating-event frequency of Chapter 15 applied to critical cooling disturbances (high load, a pump trip, a hot day). It is not inferred from the temperature observations. The detailed event model of Section 36.3 distinguishes the probability of interruption conditional on each physical hypothesis.

36.2 The physical model and the measurement audit

This section writes the energy balance of the IT thermal environment, computes the time to the thermal limit under the healthy condition and under two explanations of the anomaly, and then audits the evidence: what is measured, which hypotheses it leaves open, and which quantities it can and cannot separate.

36.2.1 The energy balance and the baseline

The physical model is an energy balance on one lumped thermal state. The heat stored in the IT thermal environment changes at the rate at which heat enters it minus the rate at which the cooling plant removes it:

$$C_{\mathrm{eff}}\,\frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} = \dot{Q}_{\mathrm{IT}} + \dot{Q}_{\mathrm{other}} - \dot{Q}_{\mathrm{cool}} \equiv \dot{Q}_{\mathrm{net}}, \tag{36.3}$$

where C_{eff} is the effective thermal capacitance in joules per kelvin, T_{in} the representative IT thermal state, $\dot{Q}_{\mathrm{IT}}$ the IT heat production, $\dot{Q}_{\mathrm{other}}$ the other applicable heat gains, $\dot{Q}_{\mathrm{cool}}$ the heat removed by the cooling plant, all three in kilowatts, and $\dot{Q}_{\mathrm{net}}$ the net heat accumulation rate. It is the balance of Chapter 34 with the heat sources written separately, because the diagnosis turns on which of them is in error.

Under the baseline conditions of Table 36.1, $\dot{Q}_{\mathrm{IT}} + \dot{Q}_{\mathrm{other}} = 900\,\mathrm{kW}$ and $\dot{Q}_{\mathrm{cool}} = 950\,\mathrm{kW}$, so

$$C_{\mathrm{eff}} \frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} = 900 - 950 = -50\,\mathrm{kW},$$
$$\frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} = \frac{-0.050\,\mathrm{MJ/s}}{120\,\mathrm{MJ/K}} = -4.17 \times 10^{-4}\,\mathrm{K/s} = -0.025\,\mathrm{K/min}.$$

The healthy plant has 50 kW of capacity in reserve, and the model predicts a tendency toward cooling rather than a temperature increase at that instant. In operation the plant's controls throttle the cooling to hold the set point, so the negative rate expresses a margin, not a sustained fall of temperature.

36.2.2 Thermal ride-through under two explanations

Suppose the cooling capacity has degraded to $\dot{Q}_{\mathrm{cool}} = 650\,\mathrm{kW}$. The net heat accumulation becomes $\dot{Q}_{\mathrm{net}} = 900 - 650 = 250\,\mathrm{kW}$, and with $250\,\mathrm{kW} = 0.250\,\mathrm{MJ/s}$,

$$\frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} = \frac{0.250\,\mathrm{MJ/s}}{120\,\mathrm{MJ/K}} = 2.083 \times 10^{-3}\,\mathrm{K/s} = 0.125\,\mathrm{K/min},$$

or 0.125°C per minute. This rate is a measurable physical implication of the reduced-capacity hypothesis.

When the heat rates are constant, equation (36.3) integrates to a straight line, $T_{\mathrm{in}}(t) = T_0 + \dot{Q}_{\mathrm{net}}\, t/C_{\mathrm{eff}}$, from the initial temperature T_0. Setting $T_{\mathrm{in}} = T_{\max}$ gives the time to the thermal limit, the thermal ride-through,

$$t_{\mathrm{crit}} = \frac{C_{\mathrm{eff}}\,(T_{\max} - T_0)}{\dot{Q}_{\mathrm{net}}}, \qquad \dot{Q}_{\mathrm{net}} > 0, \tag{36.4}$$

in seconds when C_{eff} is in megajoules per kelvin and $\dot{Q}_{\mathrm{net}}$ in megawatts. The available temperature rise is $T_{\max} - T_0 = 33 - 25 = 8\,\mathrm{K}$, so under the degraded-capacity hypothesis

$$t_{\mathrm{crit}} = \frac{120 \times 8}{0.250} = 3840\,\mathrm{s} = 64\,\mathrm{min}:$$

if these conditions persist, the representative thermal state reaches the limit about 64 minutes after the deficit begins.

A different explanation is that the cooling plant is functioning normally but 200 kW of previously unaccounted heat is being generated. Then $\dot{Q}_{\mathrm{heat}} = 1100\,\mathrm{kW}$, $\dot{Q}_{\mathrm{net}} = 1100 - 950 = 150\,\mathrm{kW}$, the temperature rises at $0.150/120 = 1.25 \times 10^{-3}\,\mathrm{K/s} = 0.075\,\mathrm{K/min}$, and

$$t_{\mathrm{crit}} = \frac{120 \times 8}{0.150} = 6400\,\mathrm{s} = 106.7\,\mathrm{min}.$$

Both mechanisms produce an elevated temperature, but they create different thermal trajectories and different times to the limit. Table 36.2 compares the three conditions: for each, the heat generated, the heat removed, the net accumulation and the idealized time to the limit.

Figure 36.2 draws the three trajectories. The horizontal axis is the time since the condition began, in minutes; the vertical axis is the representative temperature; the red horizontal line is the limit $T_{\max} = 33$°C, and the dotted vertical lines mark the two crossing times. The healthy trajectory slopes gently downward and never reaches the limit. The two fault trajectories rise at different rates and cross the limit 43 minutes apart. The physical model therefore gives a way to distinguish the explanations through their predicted consequences, provided the measurements can tell which trajectory the facility is on, the question of the rest of this section.

Table 36.2: The three physical conditions under the energy balance (36.3) with $C_{\text{eff}} = 120\,\text{MJ/K}$, $T_0 = 25°\text{C}$ and $T_{\max} = 33°\text{C}$ (hypothetical values).

Explanation	Heat generation	Cooling removal	Net heat accumulation	Rate (K/min)	Idealized time to limit
Baseline healthy	900 kW	950 kW	−50 kW	−0.025	no upward crossing under fixed conditions
Reduced cooling capacity	900 kW	650 kW	+250 kW	0.125	64 min
Additional internal heat	1100 kW	950 kW	+150 kW	0.075	106.7 min

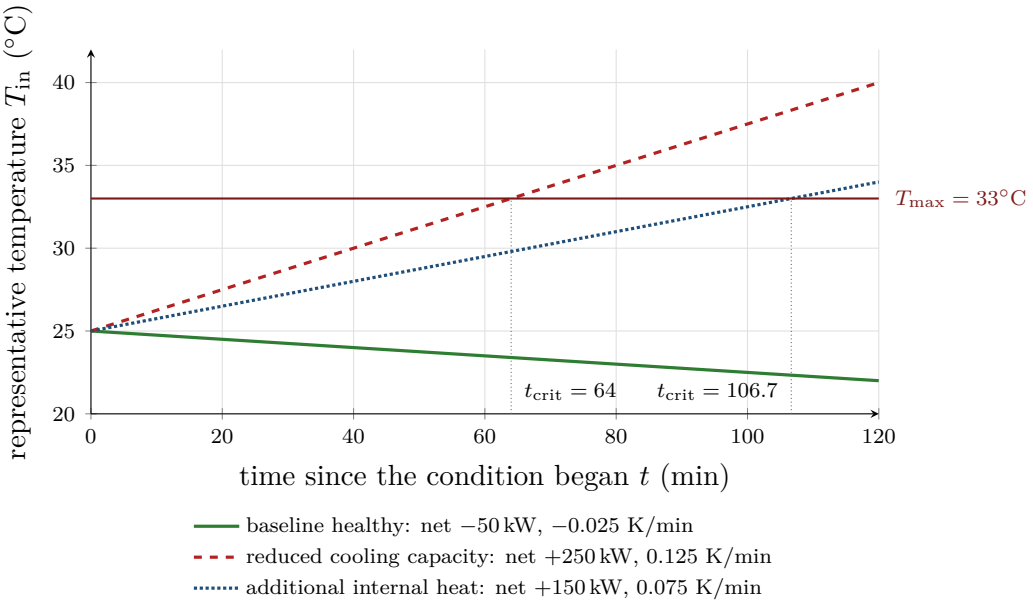


Figure 36.2: Representative temperature under the three conditions of Table 36.2, $T_{\text{in}}(t) = T_0 + \dot{Q}_{\text{net}}\, t / C_{\text{eff}}$ with $T_0 = 25°\text{C}$ and $C_{\text{eff}} = 120\,\text{MJ/K}$ (hypothetical values). The two fault explanations cross the limit at 64 and 106.7 minutes; the healthy plant does not cross it.

Table 36.3: The available evidence and the credibility concern attached to each measurement (hypothetical case).

Measurement	Observed information	Credibility concern
Return-air temperature	higher than expected	possible sensor bias
IT electrical demand	approximately 800 kW	some loads may be unmetered
Chilled-water supply temperature	within the normal range	not sufficient to determine cooling capacity
Pump command	both pumps commanded to run	a command does not prove actual flow
Pump electrical current	normal-looking values	current alone does not establish hydraulic performance
Chilled-water flow	incomplete measurement coverage	insufficient independent evidence of capacity

36.2.3 The available evidence

The facility reports an apparent thermal anomaly. The measurements available to the assessment are listed in Table 36.3: each row is one measurement, the second column says what it shows, and the third column states the concern about its credibility as evidence for the diagnosis.

The table shows the gap. The only measurement that departs from normal is one temperature, and none of the others measures the heat removed by the plant: a normal supply temperature says that the chillers produce cold water, not how much of it reaches the coils; a run command and a normal motor current are consistent with a pump that turns but delivers reduced flow, through a worn impeller or a partly closed valve; and the flow measurement does not cover the plant. The elevated temperature should therefore not be interpreted automatically as reduced cooling capacity. At least three mechanisms are plausible.

36.2.4 Competing hypotheses

Define three diagnostic hypotheses, in the notation of Chapter 22:

- H_{S}: the return-air temperature sensor is biased, and the room is not in fact hotter than expected;
- H_{C}: the cooling capacity is actually deficient;
- H_{H}: additional, unmetered heat is being generated.

For this case they are assumed to be mutually exclusive dominant explanations during the assessment, so that their probabilities sum to one. After the existing evidence is considered, they are assigned the probabilities

$$\Pr(H_{\mathrm{S}}) = 0.55, \qquad \Pr(H_{\mathrm{C}}) = 0.25, \qquad \Pr(H_{\mathrm{H}}) = 0.20. \tag{36.5}$$

These probabilities summarize hypothetical initial evidence; they are not derived from an actual data audit. A complete study would justify them from fault histories, the existing measurements, expert elicitation or other defensible information, by the methods of Chapter 8 of Volume I and Chapter 22. They are called priors in what follows because they precede the test of Section 36.4, although they are already posterior to the evidence of Table 36.3.

36.2.5 Identifiability of the heat balance

Suppose the rate of temperature rise $\mathrm{d}T_{\mathrm{in}}/\mathrm{d}t$ can be estimated accurately, C_{eff} and the metered IT load $\dot{Q}_{\mathrm{IT}}$ are known, but neither $\dot{Q}_{\mathrm{other}}$ nor $\dot{Q}_{\mathrm{cool}}$ can be measured independently. The only relation the data supply is the balance itself,

$$C_{\mathrm{eff}} \frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} - \dot{Q}_{\mathrm{IT}} = \dot{Q}_{\mathrm{other}} - \dot{Q}_{\mathrm{cool}}.$$

An increase of 200 kW in unmetered heat and a decrease of 200 kW in cooling removal have the same effect on the aggregate temperature derivative, an additional $0.200/120 \times 60 = 0.1\,\mathrm{K/min}$. The measurement constrains a difference of two physical quantities; it does not identify either one separately.

In the language of Chapter 31, take the parameter vector $\boldsymbol{\theta} = (\dot{Q}_{\mathrm{other}}, \dot{Q}_{\mathrm{cool}})^{\mathsf{T}}$ and the measured quantity $\dot{Q}_{\mathrm{net}}$. The sensitivity matrix of the measurement with respect to the parameters is the row

$$\mathbf{S}_\theta = \frac{\partial \dot{Q}_{\mathrm{net}}}{\partial \boldsymbol{\theta}} = \begin{bmatrix} 1 & -1 \end{bmatrix}, \tag{36.6}$$

whose rank is one. There are two unknown quantities and only one independent measurement relation, so the problem is locally underdetermined: the null space of $\mathbf{S}_\theta$ is spanned by $\Delta\boldsymbol{\theta} \propto (1, 1)^{\mathsf{T}}$, and moving both heat rates up or down by the same amount leaves every prediction unchanged. The Fisher information $\mathbf{S}_\theta^{\mathsf{T}} \mathbf{W} \mathbf{S}_\theta = \sigma^{-2} \left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$, for a measurement

of standard deviation σ, has the singular values $2/\sigma^2$ and 0; no estimator can reduce the variance along the null direction. Repeated measurements of the same aggregate balance at unchanged conditions add identical rows to $\mathbf{S}_\theta$ and leave its rank at one, so they do not resolve the ambiguity, however many are taken. An independent measurement of the cooling plant's heat removal (the chilled-water flow times the water temperature difference across the coils), or an electrical-load audit that closes the heat-generation balance, adds a row with a different direction and raises the rank to two.

Figure 36.3 draws the ambiguity. The horizontal axis is the unmetered heat $\dot{Q}_{\text{other}}$ and the vertical axis the cooling removal $\dot{Q}_{\text{cool}}$, both in kilowatts. Suppose the measured temperature rise corresponds to $\dot{Q}_{\text{net}} = 150\,\text{kW}$ with $\dot{Q}_{\text{IT}} = 800\,\text{kW}$; then every pair on the solid line $\dot{Q}_{\text{cool}} = \dot{Q}_{\text{other}} + 650\,\text{kW}$ explains it exactly. The red point is a cooling deficiency of 200 kW (cooling 750 kW, other heat 100 kW), the blue point an additional heat of 200 kW (cooling 950 kW, other heat 300 kW), and the open circle the healthy design point, which lies off the line. The arrow along the line is the null direction of $\mathbf{S}_\theta$. The dashed horizontal line is what an independent heat-removal measurement reading 750 kW would add: it crosses the solid line at one point and identifies both quantities.

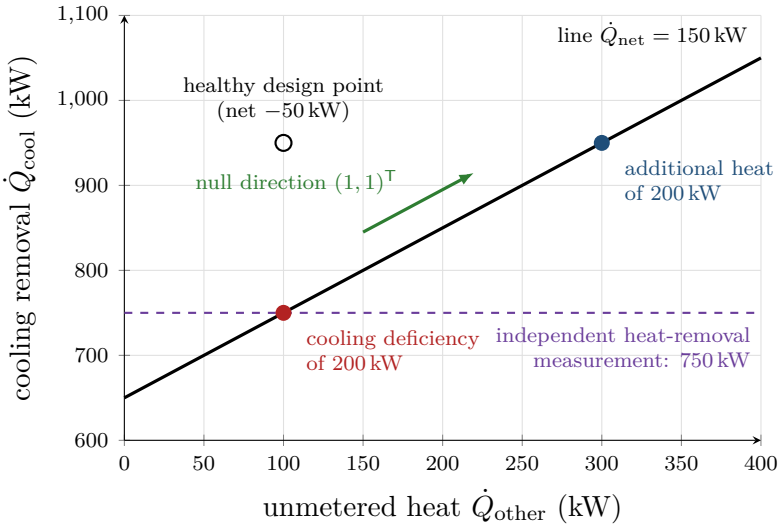


Figure 36.3: The parameter pairs consistent with one measured net heat accumulation, $\dot{Q}_{\text{IT}} + \dot{Q}_{\text{other}} - \dot{Q}_{\text{cool}} = 150\,\text{kW}$ with $\dot{Q}_{\text{IT}} = 800\,\text{kW}$ (illustrative). The sensitivity matrix $[1 \quad -1]$ has rank one, and the two explanations lie on the same line; an independent heat-removal measurement (dashed) selects one point.

36.2.6 Sensor bias as a second ambiguity

The temperature itself carries a second ambiguity. Suppose the return-temperature sensor reports

$$T_{\text{rep}} = T_{\text{act}} + b_T + \varepsilon_T, \tag{36.7}$$

where T_{act} is the actual temperature, b_T an unknown bias and ε_T a random error with zero mean. With only this sensor, a higher reported temperature may reflect a higher actual temperature, a positive bias, or both. With the unknowns $\boldsymbol{\theta}_T = (T_{\text{act}}, b_T)^\mathsf{T}$, the measurement model is

$$y = \begin{bmatrix} 1 & 1 \end{bmatrix} \boldsymbol{\theta}_T + \varepsilon_T,$$

and the sensitivity matrix $[1 \quad 1]$ again has rank one, with null direction $(1, -1)^\mathsf{T}$: a degree more in the room and a degree less of bias produce the same reading. The two unknowns cannot

be separated from one measurement. This is a structural limitation of the information, not a failure of the estimation algorithm, and it is the sensor-fault ambiguity that bad-data detection faces when a measurement is critical (Chapter 30). A second, independent temperature sensor at the same location, $y_2 = T_{\text{act}} + \varepsilon_2$, adds the row $[1 \;\; 0]$ and makes the matrix $\left[\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right]$ of rank two; so does a calibration of the first sensor against a reference.

36.2.7 The measurement sufficiency map

The audit ends with a map of what the existing evidence can and cannot establish, Table 36.4. Each row is a question the assessment must answer; the second column says whether the present measurements suffice. The first rows concern the state of the facility, the middle rows its diagnosis, and the last rows what the assessment can still do.

Table 36.4: The measurement sufficiency map: what the existing evidence establishes.

Question	Existing evidence sufficient?
Is the reported temperature elevated?	yes
Is the actual representative temperature elevated?	not securely, without a bias assessment
Is the cooling capacity reduced?	not uniquely
Is additional heat generation present?	not uniquely
Can the thermal response be predicted under a specified hypothesis?	yes, conditionally
Can the physical cause be diagnosed definitively?	no
Can actions be compared under the remaining uncertainty?	yes, with explicit probabilistic assumptions

The last two rows carry the method of the rest of the chapter. A useful decision can sometimes be made without completely resolving the diagnosis, but only if the remaining uncertainty is preserved: the three hypotheses are carried, with their probabilities, through the risk and the comparison of actions, instead of one explanation being selected silently and the others discarded.

> **Principle 36.1.** *When the measurements constrain only a combination of the unknown quantities, the competing explanations are carried with their probabilities into the risk and the decision. Repeating the same measurement does not separate them; a measurement with a different sensitivity does.*

36.3 From diagnosis to risk, and the comparison of interventions

This section connects the diagnosis to the reliability risk: it assigns a conditional interruption probability to each hypothesis, combines them into a predictive probability and an annual risk, attaches the financial consequence, and then evaluates three interventions by recomputing the risk under each.

36.3.1 Conditional interruption probabilities

The thermal calculations of Section 36.2.2 show how each mechanism shapes the temperature trajectory. An annual interruption risk needs more: the distributions of the disturbance duration, the actual cooling capacity, the operating load, the thermal capacitance, the success of restoration, and the behavior of sensors and controls. For this assessment, suppose that a

separately developed and validated event model provides the probability of IT interruption during one critical cooling disturbance under each hypothesis,

$$p_F(H_i) = \Pr(F \mid H_i, \text{critical disturbance}), \tag{36.8}$$

with the values in the third column of Table 36.5: 0.01 under sensor bias, 0.45 under cooling deficiency and 0.20 under additional heat. They are hypothetical inputs to the decision example. They are not computed from the deterministic crossing times of Table 36.2 alone; Section 36.5.3 states the integral over disturbances, restoration times and operating conditions that would compute them.

Because the hypotheses are mutually exclusive and exhaustive, the law of total probability gives the predictive interruption probability per disturbance,

$$\begin{aligned} p_F &= \sum_i \Pr(H_i)\, p_F(H_i) = 0.55 \times 0.01 + 0.25 \times 0.45 + 0.20 \times 0.20 \\ &= 0.0055 + 0.1125 + 0.0400 = 0.158, \end{aligned} \tag{36.9}$$

so that $\Pr(\text{IT interruption} \mid \text{critical disturbance}) = 15.8\,\%$. This probability combines the uncertain diagnosis with the conditional behavior of the system. Table 36.5 lists each hypothesis with its probability, its conditional interruption probability, their product and the share of p_F that it contributes.

Table 36.5: Conditional interruption probabilities per critical disturbance and their contributions to the predictive probability p_F (hypothetical values).

Hypothesis	$\Pr(H_i)$	$p_F(H_i)$	$\Pr(H_i)\, p_F(H_i)$	Share of p_F
H_{S}: sensor bias	0.55	0.01	0.0055	3.5 %
H_{C}: cooling deficiency	0.25	0.45	0.1125	71.2 %
H_{H}: additional heat	0.20	0.20	0.0400	25.3 %
Total	1.00		0.1580	100 %

The most probable hypothesis, sensor bias, contributes 3.5 % of the interruption probability; the cooling deficiency, with a quarter of the probability, contributes 71 %. Ranking the explanations by their probability and ranking them by their contribution to the risk give opposite orders, which is why the less probable explanations cannot be dropped.

36.3.2 Expected annual count and annual probability

Assume that critical disturbances occur as a homogeneous Poisson process with rate $f_{\mathrm{IE}} = 1\,\mathrm{yr}^{-1}$ and that, conditional on a fixed dominant hypothesis H_i, each disturbance independently results in an interruption with probability $p_F(H_i)$. By the thinning property of the Poisson process (Chapter 17 of Volume I), the interruptions given H_i form a Poisson process with rate $f_{\mathrm{IE}}\, p_F(H_i)$, so the number of interruptions in the year, N_F, satisfies $N_F \mid H_i \sim \mathrm{Poisson}(f_{\mathrm{IE}}\, p_F(H_i))$. By the law of total expectation, the expected annual number of interruptions is

$$\mathbb{E}[N_F] = \sum_i \Pr(H_i)\, f_{\mathrm{IE}}\, p_F(H_i) = f_{\mathrm{IE}}\, p_F = 1 \times 0.158 = 0.158 \text{ interruptions per year.} \tag{36.10}$$

This is an expected count, not the probability of one or more interruptions in the year.

If p_F were a fixed, known value, the interruptions would form one Poisson process with rate $0.158\,\mathrm{yr}^{-1}$ and

$$\Pr(N_F \geq 1) = 1 - \mathrm{e}^{-f_{\mathrm{IE}}\, p_F} = 1 - \mathrm{e}^{-0.158} = 0.1462,$$

about 14.6 %. In this model, however, the dominant hypothesis is uncertain and stays fixed for the year: the sensor is either biased or not, the plant either deficient or not. The prior-predictive annual probability then averages the conditional probabilities, not the rates:

$$\Pr(N_F \geq 1) = \sum_i \Pr(H_i)\big[1 - \mathrm{e}^{-f_{\mathrm{IE}}\, p_F(H_i)}\big], \tag{36.11}$$

which gives

$$\begin{aligned}\Pr(N_F \geq 1) &= 0.55\,(1 - \mathrm{e}^{-0.01}) + 0.25\,(1 - \mathrm{e}^{-0.45}) + 0.20\,(1 - \mathrm{e}^{-0.20}) \\ &= 0.55 \times 0.009950 + 0.25 \times 0.362372 + 0.20 \times 0.181269 = 0.1323,\end{aligned}$$

about 13.2 %. The mixture is smaller than the fixed-rate value because $1 - \mathrm{e}^{-x}$ is concave, so its average over the hypotheses lies below its value at the average rate (Jensen's inequality).

The difference is a consequence of dependence. Under a persistent but unknown condition, the disturbances of one year are not independent trials: they share the condition, as redundant components share a common cause in Chapter 13 of Volume I. An interruption is evidence about the condition. After one interruption, Bayes' rule gives $\Pr(H_{\mathrm{C}} \mid \text{interruption}) = 0.1125/0.158 = 0.712$, and the probability that the next disturbance also interrupts service rises from 0.158 to $0.0348 \times 0.01 + 0.712 \times 0.45 + 0.253 \times 0.20 = 0.371$. The interruptions are therefore clustered: $\Pr(N_F \geq 2)$ is 0.0224 under the mixture against 0.0112 for a fixed rate, and the variance of N_F is $\mathbb{E}[\mathrm{Var}(N_F \mid H)] + \mathrm{Var}(\mathbb{E}[N_F \mid H]) = 0.158 + 0.0337 = 0.192$, larger than the Poisson variance 0.158. Fewer years see an interruption, and more of those that do see two. The expected count is the same, 0.158 per year, under both models, because expectation is linear and is unaffected by the dependence.

36.3.3 Financial consequences

Suppose each IT interruption has an average modeled consequence of

$$C_F = \text{R\$ } 500\,000 \text{ per event},$$

covering contractual penalties, recovery work and business interruption. It is a hypothetical consequence, not a universal data center outage cost. The expected annual interruption loss is the risk of Chapter 16 of Volume I, frequency times conditional probability times consequence summed over the scenarios, $\mathcal{R} = \sum_k f_{\mathrm{IE}}\, \Pr(S_k \mid I)\, C_k$ (equation (16.1)), here with one initiating-event family and one consequence,

$$\mathcal{R}_0 = f_{\mathrm{IE}}\, p_F\, C_F = 1 \times 0.158 \times 500\,000 = \text{R\$ } 79\,000 \text{ per year}. \tag{36.12}$$

The same value follows hypothesis by hypothesis: the annual risk is R$ 5000 under sensor bias, R$ 225 000 under cooling deficiency and R$ 100 000 under additional heat, and $0.55 \times 5000 + 0.25 \times 225\,000 + 0.20 \times 100\,000 = 2750 + 56\,250 + 20\,000 = \text{R\$ } 79\,000$.

The assessment now has a clear decision metric: the expected annual financial loss attributable to the modeled thermal-interruption events. It is not the total reliability cost of the facility, and it excludes electrical-service interruptions unrelated to these cooling disturbances.

36.3.4 Candidate interventions and the counterfactual model

Three interventions are feasible, each with an annualized cost c_a that spreads its investment and running costs over its service life:

- action A, *repair the suspected cooling fault*: corrective work on the affected cooling equipment and its recommissioning, $c_{\mathrm{A}} = \text{R\$ } 25\,000$ per year;

- action B, *add independent cooling capacity*, which provides margin and some protection against high thermal load, $c_{\mathrm{B}} = \mathrm{R\$}\,45\,000$ per year;
- action C, *implement IT load and thermal-demand management*, which reduces heat generation during stressed conditions through controlled load flexibility, $c_{\mathrm{C}} = \mathrm{R\$}\,15\,000$ per year.

Doing nothing has $c = 0$. The interventions do not act on every hypothesis in the same way, so their effect must be stated per hypothesis. Assume that the physical reliability model predicts the conditional interruption probabilities $p_F(H_i, a)$ after each action listed in Table 36.6. Its rows are the hypotheses with their probabilities; its columns are the four actions; its last row is the predictive probability $p_F(a)$ of each action. These are hypothetical counterfactual predictions, which would have to be supported by modeling and validation before they were used for a real investment.

Table 36.6: Counterfactual model: conditional interruption probability per critical disturbance under each hypothesis after each action, and the predictive probability $p_F(a)$ (hypothetical values).

Hypothesis	$\Pr(H_i)$	No action	A: repair	B: capacity	C: load
H_{S}: sensor bias	0.55	0.01	0.01	0.01	0.01
H_{C}: cooling deficiency	0.25	0.45	0.05	0.02	0.25
H_{H}: additional heat	0.20	0.20	0.20	0.10	0.04
$p_F(a) = \sum_i \Pr(H_i)\, p_F(H_i, a)$		0.158	0.058	0.0305	0.076

The entries have physical interpretations. Repair A is highly effective if the cooling equipment is deficient, but does nothing for excess heat. Added capacity B helps under both real thermal mechanisms, because extra cooling offsets a deficit from either side of the balance. Load management C has a moderate effect on a cooling deficiency and a strong effect on excess heat, since it removes heat at its source. None of the actions changes the sensor bias, and under that hypothesis the small interruption probability, 0.01, is that of a healthy facility, which no action improves in this simplified table.

The risk after action a averages the counterfactual probabilities over the hypotheses and multiplies by the frequency and the consequence:

$$\mathcal{R}(a) = f_{\mathrm{IE}}\, C_F \sum_i \Pr(H_i)\, p_F(H_i, a) = f_{\mathrm{IE}}\, C_F\, p_F(a). \tag{36.13}$$

The quantity to be minimized is the total expected annual cost, the expected loss of action a,

$$\rho(a) = c_a + \mathcal{R}(a) = \sum_i \Pr(H_i)\, L(a, H_i), \qquad L(a, H_i) = c_a + f_{\mathrm{IE}}\, C_F\, p_F(H_i, a), \tag{36.14}$$

where $L(a, H_i)$ is the annual loss of action a when hypothesis H_i is true, in reais per year.

36.3.5 Risk and cost of each action

For action A, only the cooling-deficiency term changes:

$$p_F(\mathrm{A}) = 0.55 \times 0.01 + 0.25 \times 0.05 + 0.20 \times 0.20 = 0.0055 + 0.0125 + 0.0400 = 0.058,$$

so $\mathcal{R}(\mathrm{A}) = 0.058 \times 500\,000 = \mathrm{R\$}\,29\,000$ per year and, with its cost, $\rho(\mathrm{A}) = 29\,000 + 25\,000 = \mathrm{R\$}\,54\,000$ per year. For action B,

$$p_F(\mathrm{B}) = 0.55 \times 0.01 + 0.25 \times 0.02 + 0.20 \times 0.10 = 0.0055 + 0.0050 + 0.0200 = 0.0305,$$

so $\mathcal{R}(\mathrm{B}) = 0.0305 \times 500\,000 = \mathrm{R\$}\,15\,250$ and $\rho(\mathrm{B}) = 15\,250 + 45\,000 = \mathrm{R\$}\,60\,250$ per year. For action C,

$$p_F(\mathrm{C}) = 0.55 \times 0.01 + 0.25 \times 0.25 + 0.20 \times 0.04 = 0.0055 + 0.0625 + 0.0080 = 0.076,$$

so $\mathcal{R}(\mathrm{C}) = 0.076 \times 500\,000 = \mathrm{R\$}\,38\,000$ and $\rho(\mathrm{C}) = 38\,000 + 15\,000 = \mathrm{R\$}\,53\,000$ per year.

The risk reductions $\Delta\mathcal{R}(a) = \mathcal{R}_0 - \mathcal{R}(a)$ come from identifiable hypotheses. Action A removes $0.25 \times (0.45 - 0.05) = 0.100$ of interruption probability, all under H_{C}, worth R\$ 50 000 per year. Action B removes $0.25 \times 0.43 + 0.20 \times 0.10 = 0.1275$, worth R\$ 63 750. Action C removes $0.25 \times 0.20 + 0.20 \times 0.16 = 0.082$, worth R\$ 41 000. Each action is worth taking on expected cost alone, since each risk reduction exceeds its annualized cost; the question is which one.

36.3.6 Comparison of the actions and the risk constraint

Table 36.7 collects the results. Its columns are the predictive interruption probability per disturbance, the annual risk, the risk reduction, the annualized cost of the intervention, the total expected annual cost, and whether the action meets the risk constraint introduced below. Figure 36.4 draws the total expected annual cost of each action as a stacked bar: the lower part is the expected interruption loss $\mathcal{R}(a)$, the upper part the annualized intervention cost c_a, and the label above each bar gives the total and the interruption probability.

Table 36.7: Comparison of the four actions (hypothetical values, $f_{\mathrm{IE}} = 1\,\mathrm{yr}^{-1}$, $C_F = \mathrm{R\$}\,500\,000$).

Action	$p_F(a)$	$\mathcal{R}(a)$ (R\$/yr)	$\Delta\mathcal{R}(a)$ (R\$/yr)	c_a (R\$/yr)	$\rho(a)$ (R\$/yr)	$p_F(a) \le 0.05$
No action	15.8 %	79 000	–	0	79 000	no
A: repair	5.8 %	29 000	50 000	25 000	54 000	no
B: add capacity	3.05 %	15 250	63 750	45 000	60 250	yes
C: manage heat load	7.6 %	38 000	41 000	15 000	53 000	no

Action C has the lowest total expected cost, R\$ 53 000 per year, and action A follows within R\$ 1000. Action B has the lowest probability of interruption, 3.05 %, and the highest cost of the three interventions. These are different objectives, and they select different actions. Suppose the organization has an approved acceptance criterion on the per-disturbance interruption probability,

$$p_F(a) \le 0.05. \tag{36.15}$$

Action C, at 0.076, does not meet it, and neither does A, at 0.058. Action B does, with $p_F(\mathrm{B}) = 0.0305$. Under the constraint, B is the preferred feasible action among these candidates, at an expected cost R\$ 7250 per year above the unconstrained optimum; that difference is the expected price of the acceptance criterion. The conclusion assumes that the risk limit applies to the defined per-disturbance event; a limit stated on the annual probability of an interruption would be checked against equation (36.11) instead. The engineering decision cannot be separated from its approved acceptance criteria, which is the role of the standards and assurance requirements of Chapter 33.

Principle 36.2. *The action with the lowest expected cost and the action with the lowest risk are generally different, because intervention costs and the value assigned to consequences change the objective. An approved acceptance criterion is a constraint on the optimization, not one more cost term.*

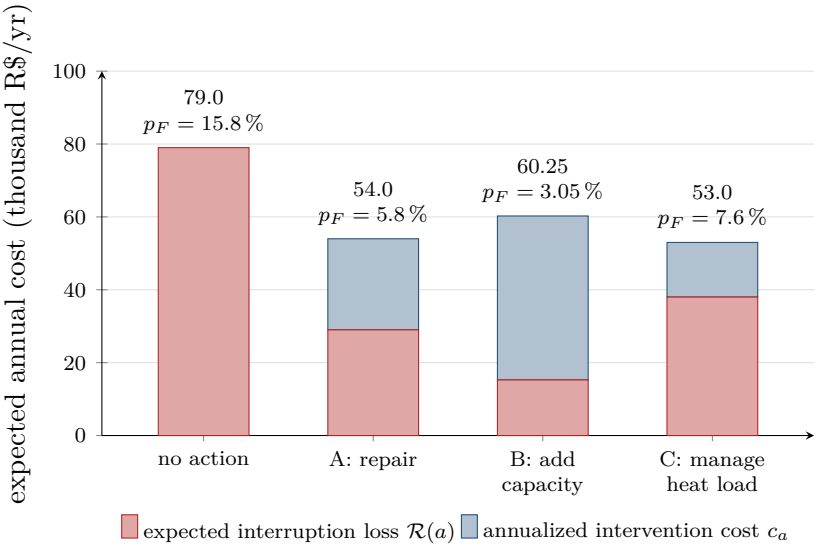


Figure 36.4: Total expected annual cost of each action, $\rho(a) = \mathcal{R}(a) + c_a$, in thousands of reais per year, with the predictive interruption probability per critical disturbance above each bar (hypothetical values).

36.4 The next measurement and the value of information

The comparison of Section 36.3.6 was made under the present uncertainty about the diagnosis. This section asks whether that uncertainty should be reduced before acting: it defines a diagnostic test, updates the hypotheses after each result, recomputes the best action for each, and evaluates the test by the expected value of sample information and its bound, the expected value of perfect information, as Chapters 24 and 32 defined them.

36.4.1 The diagnostic test and the probability of each outcome

The decision is sensitive to the diagnosis. The cooling deficiency has probability 0.25, yet the expected costs of actions A and C differ by only

$$\rho(\mathrm{A}) - \rho(\mathrm{C}) = 54\,000 - 53\,000 = \mathrm{R\$}\,1000 \text{ per year},$$

and A is far better than C if the deficiency is present while C is far better if it is not. A diagnostic measurement that bears on H_{C} is therefore likely to be valuable. Suppose an independent cooling-flow and capacity test can be performed: it measures the chilled-water flow and the temperature difference across the coils, which is the independent heat-removal measurement that Section 36.2.5 found missing. Its outcome is Y^{+}, evidence of reduced cooling capacity, or Y^{-}. Assume the likelihoods of a positive result

$$\Pr(Y^{+} \mid H_{\mathrm{S}}) = 0.05, \qquad \Pr(Y^{+} \mid H_{\mathrm{C}}) = 0.85, \qquad \Pr(Y^{+} \mid H_{\mathrm{H}}) = 0.10,$$

hypothetical test-performance values: the sensitivity of the test to a cooling deficiency is 0.85, and its false-positive probability is 0.05 or 0.10 according to the alternative. The probability of a negative result under each hypothesis is the complement, $\Pr(Y^{-} \mid H_i) = 1 - \Pr(Y^{+} \mid H_i)$.

The probability of a positive result follows from the law of total probability over the hypotheses:

$$\begin{aligned} \Pr(Y^{+}) &= \sum_i \Pr(Y^{+} \mid H_i)\,\Pr(H_i) = 0.55 \times 0.05 + 0.25 \times 0.85 + 0.20 \times 0.10 \\ &= 0.0275 + 0.2125 + 0.0200 = 0.26, \end{aligned} \tag{36.16}$$

and $\Pr(Y^-) = 0.74$. Of the positive results, $0.2125/0.26 = 82\,\%$ come from a real deficiency.

36.4.2 Bayesian diagnosis after each result

Bayes' rule (Chapter 8 of Volume I) gives the posterior probability of each hypothesis after a result $y \in \{+, -\}$:

$$\Pr(H_i \mid Y^y) = \frac{\Pr(Y^y \mid H_i)\,\Pr(H_i)}{\Pr(Y^y)}. \tag{36.17}$$

Table 36.8 carries out the computation. For each hypothesis it lists the prior, the likelihood of each result, the joint probability (likelihood times prior, whose column sum is the probability of the result) and the posterior (the joint probability divided by that sum). After a positive result, for instance, $\Pr(H_{\mathrm{C}} \mid Y^+) = 0.2125/0.26 = 0.8173$; after a negative result, $\Pr(H_{\mathrm{S}} \mid Y^-) = 0.5225/0.74 = 0.7061$.

Table 36.8: Bayesian update of the three hypotheses after a positive and after a negative cooling-capacity test (hypothetical test performance).

Hypothesis	$\Pr(H_i)$	$\Pr(Y^+ \mid H_i)$	joint, Y^+	$\Pr(H_i \mid Y^+)$	$\Pr(Y^- \mid H_i)$	joint, Y^-	$\Pr(H_i \mid Y^-)$
H_{S}: sensor bias	0.55	0.05	0.0275	0.1058	0.95	0.5225	0.7061
H_{C}: cooling deficiency	0.25	0.85	0.2125	0.8173	0.15	0.0375	0.0507
H_{H}: additional heat	0.20	0.10	0.0200	0.0769	0.90	0.1800	0.2432
Total	1		0.26	1		0.74	1

A positive result strongly favors the cooling deficiency, raising its probability from 25 % to 81.7 %, and under that posterior the repair becomes attractive. A negative result lowers it to 5.1 %, which makes a substantial capacity deficiency unlikely relative to the alternatives and removes most of the economic case for repairing the cooling equipment. Neither result resolves the diagnosis completely: after a negative result the sensor bias and the additional heat remain in the proportion $0.7061 : 0.2432$, close to their prior ratio $55 : 20$, because the test barely distinguishes them.

36.4.3 Expected costs after each result and the optimal policy

For each test result the four actions are evaluated again with equation (36.14), the posterior probabilities replacing the priors:

$$\rho(a \mid Y^y) = c_a + f_{\mathrm{IE}}\, C_F \sum_i \Pr(H_i \mid Y^y)\, p_F(H_i, a).$$

After a positive result, for action A, $p_F(\mathrm{A} \mid Y^+) = 0.1058 \times 0.01 + 0.8173 \times 0.05 + 0.0769 \times 0.20 = 0.0573$, and $\rho(\mathrm{A} \mid Y^+) = 25\,000 + 500\,000 \times 0.0573077 = \mathrm{R\$}\,53\,654$ per year. Table 36.9 lists the interruption probability and the total expected cost of every action after each result.

The least-cost action differs between the two results, shown in bold:

$$a^*(Y) = \begin{cases} \text{A, repair the cooling equipment,} & Y = Y^+, \\ \text{C, manage the heat load,} & Y = Y^-. \end{cases} \tag{36.18}$$

The test has decision value because it changes which intervention is preferred. This is the connection between fault diagnosis and decision optimization that Chapter 24 established for one pump, here with three hypotheses and four actions.

Table 36.9: Interruption probability per disturbance and total expected annual cost of each action after each test result (hypothetical values).

Action	$p_F(a \mid Y^+)$	$\rho(a \mid Y^+)$ (R\$/yr)	$p_F(a \mid Y^-)$	$\rho(a \mid Y^-)$ (R\$/yr)
No action	0.3842	192 115	0.0785	39 257
A: repair	0.0573	**53 654**	0.0582	54 122
B: add capacity	0.0251	57 548	0.0324	61 199
C: manage heat load	0.2085	119 231	0.0295	**29 730**

Figure 36.5 draws the decision as a tree in the convention of Chapter 24. Squares are decision nodes, where the operator chooses a branch; the circle is the chance node of the test result, whose branches carry probabilities; the numbers at the right are the total expected annual costs of each action, in reais per year; and the boxed numbers beside the nodes are the expected costs obtained by rolling the tree back from the right, taking the minimum at a decision node and the probability-weighted average at the chance node. Thick branches are the ones chosen. The upper branch decides now, without the test; the lower branch tests first and then decides.

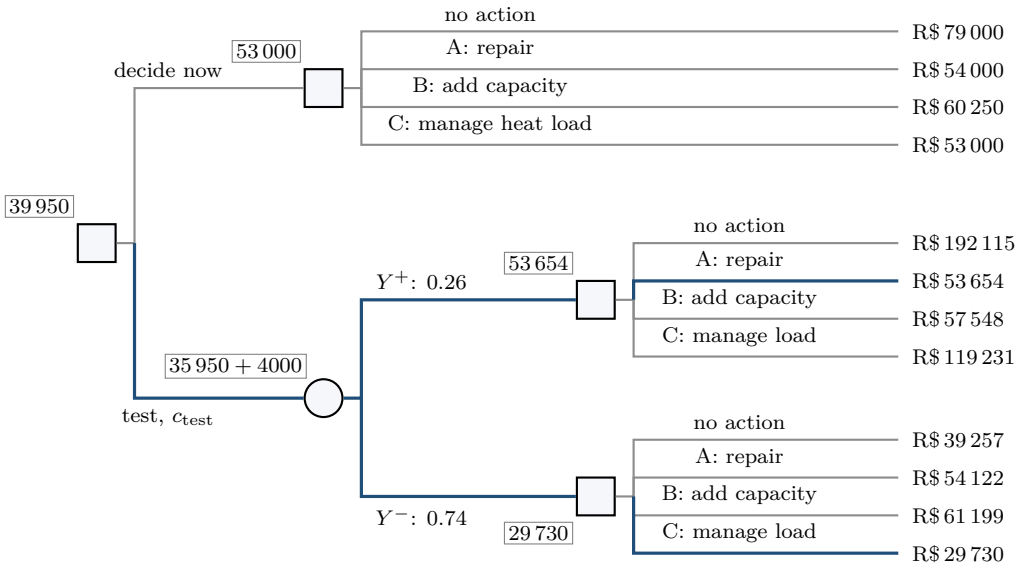


Figure 36.5: Decision tree of the cooling-anomaly decision. Squares are decisions and the circle is the test result with its probabilities; total expected annual costs in reais per year at the right; boxed values are the rolled-back expected costs. Thick branches are chosen: test, then repair after Y^+ and manage the heat load after Y^- (hypothetical values).

36.4.4 Expected value of sample information

Without the test the preferred action is C, with expected cost

$$C_{\text{now}} = \min_a \rho(a) = \rho(\text{C}) = \text{R\$ } 53\,000 \text{ per year.}$$

With the test, the action is chosen after the result is known, so the expected cost, before the cost of the test, is the average over the results of the least cost given each result:

$$\begin{aligned} C_{\text{sample}} &= \Pr(Y^+) \min_a \rho(a \mid Y^+) + \Pr(Y^-) \min_a \rho(a \mid Y^-) \\ &= 0.26 \times 53\,653.85 + 0.74 \times 29\,729.73. \end{aligned} \tag{36.19}$$

The two terms are exact: since $\Pr(Y^y)\,\Pr(H_i \mid Y^y) = \Pr(Y^y \mid H_i)\,\Pr(H_i)$, the first is $25\,000 \times 0.26 + 500\,000\,(0.0275\times0.01 + 0.2125\times0.05 + 0.0200\times0.20) = 6500 + 500\,000\times0.0149 = \text{R\$}\,13\,950$ and the second is $15\,000 \times 0.74 + 500\,000\,(0.5225 \times 0.01 + 0.0375 \times 0.25 + 0.1800 \times 0.04) = 11\,100 + 500\,000 \times 0.0218 = \text{R\$}\,22\,000$. Hence $C_{\text{sample}} = \text{R\$}\,35\,950$ per year, and the expected value of sample information is

$$\text{EVSI} = C_{\text{now}} - C_{\text{sample}} = 53\,000 - 35\,950 = \text{R\$}\,17\,050 \text{ per year.} \tag{36.20}$$

If the test costs $c_{\text{test}} = \text{R\$}\,4000$ on the same annualized decision basis, its expected net benefit is $17\,050 - 4000 = \text{R\$}\,13\,050$ per year, and the expected cost of the test-informed strategy, including the test, is $35\,950 + 4000 = \text{R\$}\,39\,950$ per year, lower than the R\$ 53 000 of choosing C immediately. Under the model, the preferred immediate decision is therefore to perform the diagnostic test, then repair if it is positive or implement load management if it is negative. The policy presumes that the test is available before the intervention, that its cost is as stated, and that conducting it introduces no material additional interruption risk or decision delay; a test that took months, during which the facility ran with an undiagnosed deficiency, would carry a cost that this calculation omits.

36.4.5 Expected value of perfect information

Suppose the dominant physical hypothesis could be learned with certainty. The least-cost action would then be chosen under each hypothesis separately, and the expected cost would be the average of those minima. Table 36.10 lists the annual loss $L(a, H_i)$ of equation (36.14) for every action and hypothesis; the last column is its average over the priors, $\rho(a)$, and the bold entry of each hypothesis column is its minimum. Under sensor bias, no action is cheapest, R\$ 5000; under cooling deficiency, the repair A, $25\,000 + 0.05 \times 500\,000 = \text{R\$}\,50\,000$, which is below the R\$ 55 000 of added capacity B; under additional heat, load management C, R\$ 35 000.

Table 36.10: Annual loss $L(a, H_i) = c_a + f_{\text{IE}}\, C_F\, p_F(H_i, a)$ of each action under each hypothesis, in reais per year; the bold entries are the least-cost actions under perfect information (hypothetical values).

Action	$L(a, H_{\text{S}})$	$L(a, H_{\text{C}})$	$L(a, H_{\text{H}})$	$\rho(a)$
No action	**5000**	225 000	100 000	79 000
A: repair	30 000	**50 000**	125 000	54 000
B: add capacity	50 000	55 000	95 000	60 250
C: manage heat load	20 000	140 000	**35 000**	53 000
$\Pr(H_i)$	0.55	0.25	0.20	

The expected cost with perfect information is

$$\begin{aligned} C_{\text{perfect}} &= \sum_i \Pr(H_i)\, \min_a L(a, H_i) = 0.55 \times 5000 + 0.25 \times 50\,000 + 0.20 \times 35\,000 \\ &= 2750 + 12\,500 + 7000 = \text{R\$}\,22\,250, \end{aligned} \tag{36.21}$$

and the expected value of perfect information is

$$\text{EVPI} = C_{\text{now}} - C_{\text{perfect}} = 53\,000 - 22\,250 = \text{R\$}\,30\,750 \text{ per year.} \tag{36.22}$$

Table 36.11 collects the values. As required of a coherent decision model,

$$0 \le \text{EVSI} \le \text{EVPI}, \qquad 0 \le 17\,050 \le 30\,750,$$

and the proposed test captures $17\,050/30\,750 = 55\,\%$ of the bound. The difference, R\$ 13 700 per year, is the additional decision benefit that a perfect diagnosis could provide beyond this particular imperfect test. Most of it lies in the separation of sensor bias from additional heat, which the test hardly resolves: a test that identified the cooling deficiency perfectly, but left the other two hypotheses in their prior ratio, would lead to A under H_{C} and to C otherwise, at the expected cost $0.25 \times 50\,000 + 0.75 \times 24\,000 = \text{R\$ } 30\,500$, and would be worth R\$ 22 500; the remaining R\$ 8250 of the EVPI can be earned only by distinguishing sensor bias from additional heat.

Table 36.11: Value of the information for the cooling-anomaly decision, in reais per year (hypothetical values).

Measure	Value (R$/yr)
Expected cost of the best action now, C_{now}	53 000
Expected cost after the test, before its cost, C_{sample}	35 950
Expected cost with perfect information, C_{perfect}	22 250
EVPI	30 750
EVSI of the proposed test	17 050
Test cost c_{test}	4000
Net expected value of the test, EVSI $- c_{\mathrm{test}}$	13 050
Expected cost of the test-informed policy, $C_{\mathrm{sample}} + c_{\mathrm{test}}$	39 950

36.4.6 Interpretation of the information values

The EVPI does not imply that spending up to R\$ 30 750 per year on instrumentation is automatically justified. It is an upper bound on the expected decision improvement from perfectly resolving the specified hypothesis uncertainty, for this decision and this horizon. Sensor costs, operating constraints, test delays and other objectives still matter, and a real instrument is imperfect, so its value is an EVSI, computed for its own likelihoods as above. The bound serves as a screen: a proposed measurement that costs more than the EVPI cannot pay for itself through this decision, whatever its accuracy.

The acceptance criterion (36.15) changes the values. After a positive result the repair leaves $p_F(\mathrm{A} \mid Y^+) = 0.0573$, above the limit 0.05, and the only compliant action is B, at R\$ 57 548; after a negative result C remains compliant, with 0.0295. The constrained test-informed policy (B after Y^+, C after Y^-) has expected cost $0.26 \times 57\,548 + 0.74 \times 29\,730 = \text{R\$ } 36\,963$ per year before the test, against R\$ 60 250 for the best compliant action without the test, B. Under the constraint the test is worth R\$ 23 288 per year, more than without it, because it allows the costly capacity addition to be confined to the cases in which a deficiency is likely.

Principle 36.3. *A diagnostic test is valuable when its possible outcomes lead to different actions and so improve the expected decision. Its value is the EVSI, bounded by the EVPI, and it is worth performing when the EVSI exceeds its cost, including the cost of the delay it imposes.*

36.5 Verification, validation, credibility and the assessment report

A coherent calculation is not yet a credible assessment. This section states what the calculation has and has not demonstrated, lists the verification checks and the validation evidence, writes the propagation of uncertainty that would replace the assumed probabilities, keeps model inadequacy visible, tabulates how the preferred action depends on the diagnosis, and gives the structure of the assessment report with an example decision statement.

36.5.1 What has been demonstrated, and verification

The calculations so far form a coherent decision model. They have defined the system function; constructed a simplified physical energy balance; identified an underdetermined diagnosis problem; assigned probabilities over the plausible mechanisms; propagated them into the interruption risk; compared corrective actions; and computed the value of an additional measurement. They have not established that the hypothetical model is adequate for a real facility. That distinction, between a model computed correctly and a model that represents the system, is the distinction between verification and validation.

Verification checks the implementation. The model should be checked for probability normalization, consistent units, correct Bayesian updates, correct application of the heat balance, correct calculation of the expected risk, proper comparison of the conditional actions, consistency between the baseline and the counterfactual models, and numerical stability and reproducibility. Applied to this chapter:

1. the priors sum to one, $0.55 + 0.25 + 0.20 = 1$, and so does each posterior: $0.1058 + 0.8173 + 0.0769 = 1$ and $0.7061 + 0.0507 + 0.2432 = 1$ (Table 36.8);
2. the posteriors average back to the priors over the results, the law of total probability: $0.26 \times 0.8173 + 0.74 \times 0.0507 = 0.2125 + 0.0375 = 0.25 = \Pr(H_{\mathrm{C}})$;
3. units are consistent: $\mathrm{kW/(MJ/K)} = 10^{-3}\,\mathrm{K/s}$ in the balance, frequencies in yr^{-1}, probabilities per disturbance dimensionless, consequences in reais per event and risks in reais per year;
4. the no-action column of Table 36.6 reproduces the baseline probabilities of Table 36.5, and the risk computed through the predictive probability equals the risk averaged hypothesis by hypothesis, R$ 79 000 both ways;
5. every action is compared on the same basis, annualized cost plus expected annual loss, before and after the test;
6. the value of information satisfies $0 \le \mathrm{EVSI} \le \mathrm{EVPI}$.

If any of these conditions failed, the numerical implementation or the probability definitions would be wrong. The last check also catches a specific error: had the least-cost action under H_{C} been taken as B (R$ 55 000) instead of A (R$ 50 000), the computed EVPI would have been R$ 29 500 instead of R$ 30 750; the inequality would still hold, which shows that the check is necessary but not sufficient, and that each minimum must be taken over the full table.

36.5.2 Validation evidence

For deployment, every element of the model would need supporting evidence from the physical system. Table 36.12 lists the elements in the first column and the evidence that would validate each in the second.

The most consequential missing evidence is the conditional interruption model. The values 0.01, 0.45 and 0.20 of Table 36.5 are assumed, and the cooling-deficiency value alone carries 71 % of the risk. Without supporting evidence, the final numerical decision is a demonstration of method, not an approved engineering conclusion.

Table 36.12: Validation evidence required for each model element.

Model element	Validation evidence
IT heat generation	electrical-load audit and heat balance
Cooling removal	flow, temperature-difference and heat-transfer measurements
Thermal capacitance	transient temperature-response tests
Sensor-bias hypothesis	independent calibration or a redundant temperature measurement
Equipment degradation	inspection and condition-monitoring records
Conditional interruption probabilities	event simulation, historical disturbances, or suitably validated engineering models
Intervention effectiveness	commissioning tests and physical performance models
Annual disturbance frequency	facility or relevant fleet reliability records
Financial consequences	contractual, operational and business-impact analysis

36.5.3 Uncertainty propagation

The conditional probabilities can be derived rather than assumed. Suppose the probabilities over the hypotheses are themselves uncertain, because the likelihoods and priors were estimated from limited data, and suppose the thermal capacitance, the cooling performance and the restoration time are uncertain as well:

$$C_{\text{eff}} \sim p(C_{\text{eff}}), \qquad \dot{Q}_{\text{cool}} \sim p(\dot{Q}_{\text{cool}} \mid H_i), \qquad T_{\text{R}} \sim p(T_{\text{R}} \mid H_i, a),$$

where T_{R} is the time to restore normal cooling after the disturbance. Collect the uncertain operating conditions, thermal parameters and restoration trajectories in a vector $\boldsymbol{\xi}$. For each value of $\boldsymbol{\xi}$ the energy balance (36.3) gives a temperature trajectory over the disturbance window $[0, \tau]$, and the interruption occurs when that trajectory crosses the limit. The probability of the failure event under hypothesis H_i and action a is then the integral of the indicator of a crossing:

$$p_F(H_i, a) = \int \mathbf{1}\Big\{ \max_{0 \le t \le \tau} T_{\text{in}}(t; \boldsymbol{\xi}) > T_{\max} \Big\}\, p(\boldsymbol{\xi} \mid H_i, a)\, \mathrm{d}\boldsymbol{\xi}. \tag{36.23}$$

This converts the assumed interruption probabilities into quantities derived from the physical and stochastic model; the simplest instance is the thermal ride-through probability of Chapter 34, $\Pr(T_{\text{R}} > t_{\text{crit}})$, in which $\boldsymbol{\xi}$ is the restoration time alone.

The predictive risk of action a is then the expectation of the loss over both the conditions and the hypotheses,

$$\rho(a) = \mathbb{E}_{\boldsymbol{\xi}, H}\big[L(a, H, \boldsymbol{\xi})\big], \tag{36.24}$$

which reduces to equation (36.14) when the loss depends on $\boldsymbol{\xi}$ only through the interruption indicator. Such a model can be evaluated by analytical approximation, by dynamic programming over the disturbance and restoration sequence, or by Monte Carlo simulation: draw a hypothesis from its probabilities and a vector $\boldsymbol{\xi}$ from its conditional distribution, integrate the energy balance, record the loss, and average. Any correlations among failure modes, weather, electrical supply, cooling performance and repair times must be retained in the joint distribution of $\boldsymbol{\xi}$; a hot day raises the IT and cooling loads together and may also lengthen restoration, and sampling these independently would understate the probability of the severe combinations, as Chapter 16 of Volume I showed for a shared cause.

36.5.4 Model inadequacy

Suppose the heat-balance model consistently underpredicts the observed temperature rise even after plausible parameter adjustments, so that

$$C_{\text{eff}} \frac{\mathrm{d}T_{\text{in}}}{\mathrm{d}t} \neq \dot{Q}_{\text{heat}} - \dot{Q}_{\text{cool}}$$

beyond the expected measurement and approximation error. The model is then inadequate in the sense of Chapter 31, and the residual is a model discrepancy $\delta(t)$, not measurement noise. Possible explanations include an omitted heat source, incorrect system boundaries, an inaccurate thermal storage, nonuniform airflow (hot spots and recirculation that a single representative temperature cannot represent), missing dynamic transport delays, incorrect sensor readings, and unmodeled behavior of the cooling controls.

The discrepancy should not simply be forced into one parameter. Absorbing it into C_{eff} or $\dot{Q}_{\text{other}}$ would produce a precise fit, a precise and wrong parameter, and a precise and wrong interruption probability. A finding of model inadequacy may invalidate otherwise precise reliability predictions, and it must remain visible in the assessment, as a stated limitation and as a reason to collect the measurements that would identify the missing mechanism.

36.5.5 Sensitivity of the preferred action to the diagnosis: a worked table

The point-estimate optimum of Section 36.3.6 rests on $\Pr(H_{\text{C}}) = 0.25$. To see how the preferred action depends on the diagnosis, let $q = \Pr(H_{\text{C}})$ vary and redistribute the remaining probability between sensor bias and additional heat in their original ratio $55 : 20$:

$$\Pr(H_{\text{S}}) = \tfrac{0.55}{0.75}\,(1-q) = 0.7333\,(1-q), \qquad \Pr(H_{\text{H}}) = \tfrac{0.20}{0.75}\,(1-q) = 0.2667\,(1-q).$$

With the counterfactual probabilities of Table 36.6, one disturbance per year and $C_F =$ R\$ 500 000, each action's interruption probability is linear in q,

$$p_F(a;q) = (1-q)\,\bar{p}_a + q\,p_F(H_{\text{C}},a), \qquad \bar{p}_a = \frac{0.55\,p_F(H_{\text{S}},a) + 0.20\,p_F(H_{\text{H}},a)}{0.75},$$

with $\bar{p}_a = 0.06067$ for no action and for A, 0.034 for B and 0.018 for C; and $\rho(a;q) = c_a + 500\,000\,p_F(a;q)$. One row of the arithmetic, $q = 0.50$: $\Pr(H_{\text{S}}) = 0.3667$ and $\Pr(H_{\text{H}}) = 0.1333$, so for action A $p_F = 0.5 \times 0.06067 + 0.5 \times 0.05 = 0.05533$ and $\rho = 25\,000 + 27\,667 =$ R\$ 52 667, against $p_F = 0.5 \times 0.018 + 0.5 \times 0.25 = 0.134$ and $\rho = 15\,000 + 67\,000 =$ R\$ 82 000 for action C. Table 36.13 gives the results for six values of q; its columns are the interruption probability and the total expected cost of each action, the action with the least expected cost, and the least-cost action among those that meet the acceptance criterion (36.15).

Figure 36.6 draws the total expected cost of each action as a function of q. The horizontal axis is the probability of a cooling deficiency, the vertical axis the total expected annual cost in thousands of reais; each action is a straight line. The shaded band on the left is the range of q in which action C meets the acceptance criterion, and the dashed vertical line is the assessed value $q = 0.25$.

Four trends follow. First, the least-cost action changes from load management to repair where $\rho(\text{A};q) = \rho(\text{C};q)$, that is where $-10\,000 + 500\,000\,[(0.018 - 0.06067)(1-q) + 0.20\,q] = 0$, which gives $q^* = 31\,333/121\,333 = 0.258$. The assessed value 0.25 lies just below the switch, which is why the two actions differ by only R\$ 1000 and why a diagnostic test is valuable. When the probability of a deficiency is raised to 0.80, the repair becomes substantially more attractive, R\$ 51 067 against R\$ 116 800 for load management. Second, the repair always beats added capacity on expected cost: their difference, $\rho(\text{B};q) - \rho(\text{A};q) = 6667 - 1667\,q$ reais per year, is positive for every q. Third, as q falls toward zero, load management remains preferable to doing

Table 36.13: Interruption probability per disturbance and total expected annual cost of the four actions as functions of the probability q of a cooling deficiency, with the remaining probability split 55 : 20 between sensor bias and additional heat (hypothetical values).

	$p_F(a;q)$ (%)				$\rho(a;q)$ (thousand R$/yr)				Preferred action	
q	none	A	B	C	none	A	B	C	least cost	with $p_F \le 0.05$
0	6.07	6.07	3.40	1.80	30.3	55.3	62.0	**24.0**	C	C
0.10	9.96	5.96	3.26	4.12	49.8	54.8	61.3	**35.6**	C	C
0.25	15.80	5.80	3.05	7.60	79.0	54.0	60.3	**53.0**	C	B
0.50	25.53	5.53	2.70	13.40	127.7	**52.7**	58.5	82.0	A	B
0.80	37.21	5.21	2.28	20.36	186.1	**51.1**	56.4	116.8	A	B
1	45.00	5.00	2.00	25.00	225.0	**50.0**	55.0	140.0	A	A

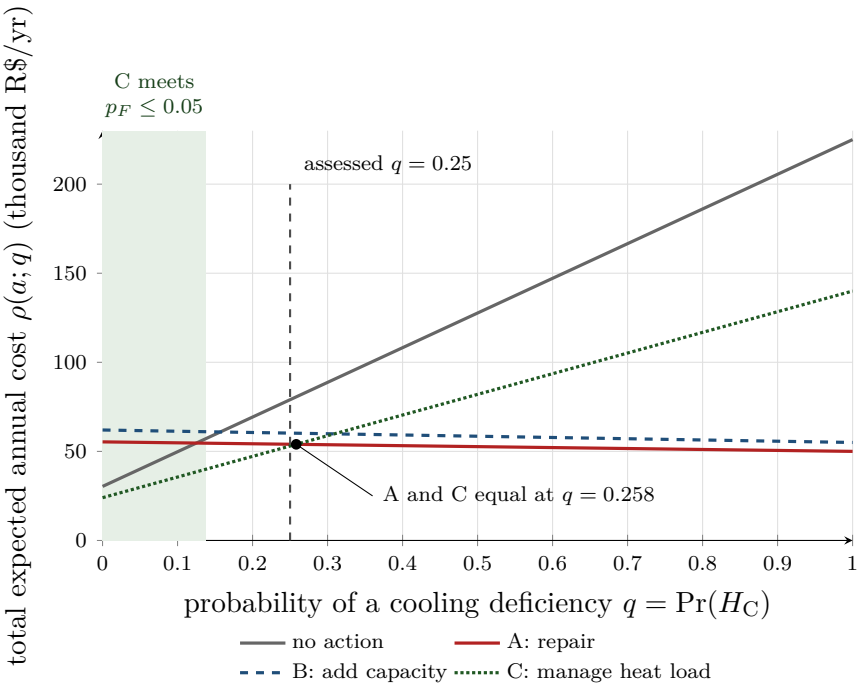


Figure 36.6: Total expected annual cost of the four actions as linear functions of the probability of a cooling deficiency, $\rho(a;q) = c_a + C_F[(1-q)\,\bar{p}_a + q\,p_F(H_C, a)]$, with the counterfactual probabilities of Table 36.6 and $C_F = \text{R\$}\,500\,000$ (hypothetical values). The least-cost action changes from C to A at $q = 0.258$.

nothing at a consequence of R$ 500 000 per event, R$ 24 000 against R$ 30 333 at $q = 0$; doing nothing becomes the least-cost action at $q = 0$ only when the consequence per interruption falls below $15\,000/(0.06067 - 0.018) = \text{R\$ } 351\,563$. The preferred action thus depends on both the diagnosis and the assumed economic consequence. Fourth, under the acceptance criterion the preference is simpler: C is compliant only while $0.018 + 0.232\,q \leq 0.05$, that is for $q \leq 0.138$, A only at $q = 1$, where it reaches the limit exactly, and B everywhere; the constrained choice is C below $q = 0.138$ and B above, except at certainty of a deficiency.

The table ranks immediate actions only; the value of first testing is evaluated separately, as in Section 36.4.4. It shows why the optimum of a point-estimate calculation should not automatically be treated as a robust decision: near a switching point, a small change of one uncertain probability changes the preferred action, and the assessment should report the switching point together with the optimum.

36.5.6 The assessment report

A practical integrated reliability assessment makes its conclusions auditable. Table 36.14 gives an appropriate structure for the final deliverable: each row is one section of the report and the second column states the content it must hold. The order follows the steps of Figure 36.1, with the decision summary placed first for the reader who acts on it and the limitations and traceability placed last for the reviewer who checks it.

Table 36.14: Structure of an integrated reliability assessment report.

Report section	Required content
Executive decision summary	preferred action, expected benefit, principal uncertainties
System definition	physical boundaries, required function, operating envelope
Data audit	available measurements, their quality, gaps and provenance
Physical model	energy and mass balance equations, topology, parameters
Reliability model	failure events, rates, dependencies, restoration
Diagnostic inference	competing hypotheses and their posterior probabilities
Model credibility	identifiability, verification, adequacy, validation
Risk assessment	conditional failure probabilities, expected consequences, tail risks
Decision analysis	feasible interventions, costs, counterfactual outcomes
Information value	candidate tests, EVSI, EVPI, test recommendations
Limitations	assumptions, missing mechanisms, sensitivity
Traceability	requirements, evidence, model versions, responsible reviewers

The report preserves the distinction between observed evidence, modeling assumptions, calculated results and engineering judgment. In this assessment, the elevated return temperature is evidence; the priors, the conditional interruption probabilities and the counterfactual table are assumptions; $p_F = 0.158$, R$ 79 000 per year and the EVSI are calculated results; and the choice of the acceptance criterion is a judgment for the organization. Each statement in the report belongs to one of the four classes and says which.

36.5.7 An example decision statement

An appropriate conclusion for this illustrative case would read as follows.

> The current evidence is consistent with several competing explanations for the reported cooling anomaly. The assumed probabilities are 55 % for sensor bias, 25 % for a cooling-capacity deficiency and 20 % for additional heat generation.

> Under the hypothetical conditional interruption model, the expected thermal-interruption consequence without intervention is R$ 79 000 per year. Among the immediate actions considered, IT load management produces the lowest expected annual cost, R$ 53 000, while adding cooling capacity yields the lowest modeled interruption probability, 3.05 % per critical disturbance, and is the only immediate action that meets the acceptance criterion of 5 %.
>
> An independent cooling-performance test is expected to improve the intervention decision by R$ 17 050 per year before the test cost. With an assumed test cost of R$ 4000, the preferred decision is to perform the test, repair the cooling equipment after a positive result, and implement load management after a negative result; if the acceptance criterion applies, add cooling capacity instead of repairing after a positive result.
>
> This recommendation is conditional on validating the diagnostic likelihoods, the thermal dynamics, the disturbance frequency, the intervention-effectiveness models and the interruption consequences. Until those inputs are substantiated, the numerical result is an illustrative decision-analysis finding, not a field-validated reliability claim.

The point of the statement is not the particular choice. It is the chain of evidence that connects the measured anomaly to the preferred action, every link of which the reader can inspect and challenge: the assessment reached a recommendation from a temperature reading that, on its own, could not even establish that the room was hot.

36.6 Synthesis of the two volumes

This last section places the assessment of the chapter in the whole of the two volumes. It maps the engineering questions of reliability to the methods that answer them, writes the decision problem that all of them serve as one formulation, and states the principle that joins them.

36.6.1 Engineering questions and the methods that answer them

The discipline of reliability engineering can be organized by the questions it answers about an engineered system. Table 36.15 lists thirteen of them in the order in which an assessment meets them. For each question it names the methods developed in this book and the chapters that develop them; Chapters 1 to 20 form Volume I and Chapters 21 to 36 Volume II.

The order of the rows is the order of the assessment of this chapter. Its first step answered the first question by defining the required function and the interruption event; its physical model answered the fifth, how long the room can ride through a deficit; its audit answered the sixth and seventh, finding the condition detectable but its cause unidentifiable from the existing measurements; its risk and intervention steps answered the tenth and eleventh; its value-of-information step the twelfth; and its verification, validation and report the ninth and the thirteenth. No single row produces a decision. The decision is the product of the whole table.

36.6.2 The integrated decision formulation

The questions culminate in one formulation. Let $\mathbf{x}_t$ be the physical operating and degradation states of the system at time t; $\boldsymbol{\theta}$ the uncertain model parameters; $\mathbf{y}_{1:t}$ the observations available up to t; a a proposed engineering intervention from the feasible set $\mathcal{A}$; and $L(a, \mathbf{x}, \boldsymbol{\theta})$ the loss of the intervention, its costs and the monetary value of its physical consequences. The decision problem is

$$\begin{aligned} a^{*} &= \arg\min_{a \in \mathcal{A}} \; \mathbb{E}\big[L(a, \mathbf{x}, \boldsymbol{\theta}) \mid \mathbf{y}_{1:t}\big] \\ &= \arg\min_{a \in \mathcal{A}} \int L(a, \mathbf{x}, \boldsymbol{\theta})\, p(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{y}_{1:t})\, \mathrm{d}\mathbf{x}\, \mathrm{d}\boldsymbol{\theta}, \end{aligned} \tag{36.25}$$

subject to the required physical, operational, reliability and safety constraints, for example $\Pr(F \mid a, \mathbf{y}_{1:t}) \le p_{\max}$. Every part of the two volumes supplies one ingredient. The lifetime,

Table 36.15: The engineering questions of reliability, the methods that answer them, and the chapters that develop them.

Engineering question	Methods developed	Chapters
What can fail?	functional analysis, FMEA, fault trees	2, 4, 10
How does it fail?	failure mechanisms, degradation models	4, 21
How frequently does it fail?	lifetime distributions, repairable processes, Bayesian reliability	5–8, 17, 18
What happens when it fails?	system reliability, contingency analysis, event trees	9, 11, 15, 27
How long does the system remain affected?	Markov models, restoration processes, physical dynamics	3, 19, 28, 34
Can the condition be detected?	condition monitoring, state estimation, residuals	22, 30
Can the cause be determined?	observability, identifiability, Bayesian diagnosis	22, 29, 31
Can future failure be predicted?	degradation processes, RUL distributions	21, 23
Can the model be trusted?	verification, validation, model adequacy	31, 36
What is the resulting risk?	availability, LOLE, EUE, expected consequences	3, 16, 25, 26
What should be improved?	importance measures, counterfactuals, portfolio optimization	12, 35
What should be measured next?	expected information gain, EVSI, EVPI	24, 32
Can the decision be justified?	requirements, traceability, standards, assurance	2, 33

degradation and counting models supply the stochastic evolution of $\mathbf{x}_t$; the fault trees, event trees and physical balances map $\mathbf{x}$ to the failure event and its consequence; the estimation, state estimation and diagnosis chapters supply the posterior $p(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{y}_{1:t})$; the observability and identifiability chapters say which parts of that posterior the data can sharpen; the value-of-information chapters choose the next observation that would sharpen it; and the standards chapter supplies the constraints.

The assessment of this chapter is the smallest nontrivial instance of equation (36.25). The state is the dominant hypothesis $H \in \{H_{\mathrm{S}}, H_{\mathrm{C}}, H_{\mathrm{H}}\}$ together with the conditions $\boldsymbol{\xi}$; the posterior is the set of probabilities $\Pr(H_i \mid \mathbf{y})$; the loss is $L(a, H_i)$ of equation (36.14); the expectation is $\rho(a)$; and the constraint is $p_F(a) \leq 0.05$. The test of Section 36.4 is the choice of the next element of $\mathbf{y}$ before the minimization.

The expectation alone is not enough. The posterior must be supported by appropriate evidence, the consequences must follow from a credible physical model, and the resulting decision must meet the applicable acceptance criteria. An optimization carried out on an unsupported posterior or an inadequate model returns a precise answer to the wrong problem, which is why the verification, validation and adequacy checks of Section 36.5 are part of the formulation and not an appendix to it.

36.6.3 The closing principle

Figure 36.7 draws the principle that joins the two volumes. The three boxes at the top are the inputs of every assessment: a physical model of the system, a failure model of how its components and functions fail, and the observations of its condition. The arrows merge them into the middle box, credible uncertainty and consequences, and that box leads to the bottom box, defensible engineering decisions. The notes at the right of each level name what the level requires.

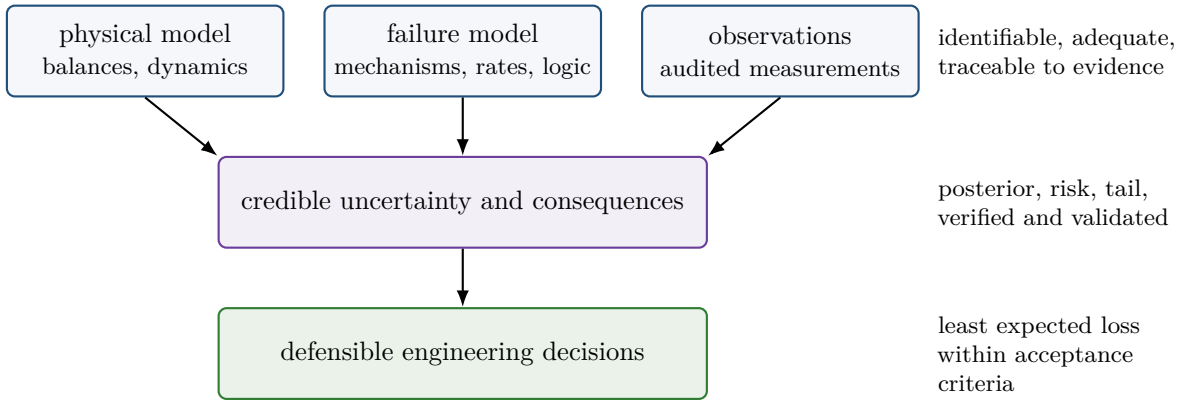


Figure 36.7: The principle of the two volumes: a physical model, a failure model and observations, joined, yield credible uncertainty and consequences, and those yield defensible engineering decisions.

Principle 36.4. *Physical model, failure model and observations together yield credible uncertainty and consequences; credible uncertainty and consequences yield defensible engineering decisions. A reliability result is worth what the weakest link of that chain is worth.*

A natural extension of this assessment is a reproducible computational benchmark: one defined physical system, controlled synthetic observations, competing fault hypotheses and several interventions, with reference answers against which alternative reliability methods can be tested. Such a benchmark would test not only whether a method produces answers but

whether those answers remain correct when sensors are biased, when the model is inadequate, and when the decision depends on the uncertainty, the three conditions this chapter met in one small facility.

The two volumes began with a definition: failure is the termination of the ability to perform a required function, and reliability is the probability that the function is performed under stated conditions for a stated period. Every later chapter refined one part of that sentence: what the function is and how its loss is observed, how long items survive and how their failures depend on one another, how systems are repaired and restored, how degradation and diagnosis reveal the present condition, how networks of components carry load and fail in sequence, what measurements can and cannot establish, and what an engineer should do with what remains uncertain. This chapter put the parts back together. The outcome is a way of working: define the required function before computing anything, model the physics that turns component events into the loss of that function, keep competing explanations alive until the evidence separates them, propagate what is uncertain to the decision that depends on it, measure next what would change that decision, and write down the chain so that it can be checked. That is the practice of reliability engineering, and with it the two volumes end.

36.7 What to retain

Definition 36.1 (Integrated reliability assessment)**.** An integrated reliability assessment defines a system's required function and failure event, models the physics that leads from equipment and sensor conditions to that event, carries the competing explanations of the evidence with their probabilities through to the risk, compares interventions by recomputing the risk under each within the approved acceptance criteria, values the next measurement by the decision it would change, and documents the chain from evidence to action so that it can be verified, validated and challenged.

Seven results carry forward. The failure event is the loss of the required function, $T_{\text{in}} > T_{\max}$ with loss of service, not a component or sensor state. A lumped energy balance turns each explanation into a trajectory and a time to the limit, 64 and 106.7 minutes here, but a single temperature cannot separate them: the heat balance and the sensor-bias model both have rank-one sensitivity matrices. The diagnosis is carried, not resolved: the predictive interruption probability 0.158 per disturbance averages three hypotheses, the least probable of which carries most of the risk, and a persistent unknown condition clusters the interruptions within a year, so the annual probability (13.2 %) is not $1 - \mathrm{e}^{-0.158}$ (14.6 %). The least-cost action (C, R$ 53 000 per year) and the least-risk action (B) differ, and an acceptance criterion selects between them. A test is worth its EVSI, R$ 17 050 per year here, because its results lead to different actions, and no test is worth more than the EVPI, R$ 30 750. And a coherent calculation is a demonstration of method until its inputs are validated and its model shown to be adequate. Table 36.16 lists the chapter's terms, their meaning and units, and the subsection that defines each.

Table 36.16: The terms of this chapter and where each is defined.

Term	Meaning	Expression	Units	Defined in
Interruption event	loss of service from the thermal limit	F, equation (36.2)	–	36.1.3
Energy balance	lumped thermal state	$C_{\text{eff}}\dot{T}_{\text{in}} = \dot{Q}_{\text{IT}} + \dot{Q}_{\text{other}} - \dot{Q}_{\text{cool}}$	kW	36.2.1
Time to the thermal limit	thermal ride-through	$C_{\text{eff}}(T_{\max} - T_0)/\dot{Q}_{\text{net}}$	s, min	36.2.2
Competing hypotheses	exclusive explanations	H_{S}, H_{C}, H_{H} with $\Pr(H_i)$	–	36.2.4
Rank-deficient sensitivity	combination identified, not parts	$\mathbf{S}_\theta = [1 \;\; -1]$, rank one	–	36.2.5
Measurement sufficiency map	what the evidence establishes	Table 36.4	–	36.2.7
Conditional interruption probability	interruption given a hypothesis	$p_F(H_i)$, $p_F(H_i, a)$	–	36.3.1
Predictive interruption probability	averaged over hypotheses	$\sum_i \Pr(H_i)\, p_F(H_i)$	–	36.3.1
Prior-predictive annual probability	mixture of Poisson processes	$\sum_i \Pr(H_i)[1 - \mathrm{e}^{-f_{\text{IE}} p_F(H_i)}]$	–	36.3.2
Risk after an action	expected annual loss	$f_{\text{IE}}\, C_F\, p_F(a)$	R$/yr	36.3.4
Expected loss	cost plus risk	$\rho(a) = c_a + \mathcal{R}(a)$	R$/yr	36.3.4
Counterfactual model	outcome under an untaken action	Table 36.6	–	36.3.4
Acceptance criterion	approved risk limit	$p_F(a) \le 0.05$	–	36.3.6
EVSI	value of the test	$C_{\text{now}} - C_{\text{sample}}$	R$/yr	36.4.4
EVPI	bound on any test	$C_{\text{now}} - C_{\text{perfect}}$	R$/yr	36.4.5
Verification, validation	computed correctly; represents the system	checks and evidence tables	–	36.5.1
Derived interruption probability	integral over conditions	equation (36.23)	–	36.5.3
Model inadequacy	discrepancy not explained by parameters	$\delta(t)$	kW	36.5.4
Switching point	where the preferred action changes	$q^* = 0.258$	–	36.5.5
Decision statement	documented conclusion	Section 36.5.7	–	36.5.7
Integrated decision	least expected loss given evidence	equation (36.25)	varies	36.6.2

Further reading. This chapter draws on no sources beyond those of the chapters whose methods it combines; the reading for each step is listed at the end of the chapter that develops it, in particular Chapter 16 of Volume I (integrated probabilistic risk assessment) and Chapters 24 (decision trees and the value of inspection), 31 (identifiability and model adequacy), 32 (the value of information), 33 (reliability standards and assurance) and 34 (data center power and cooling reliability).

36.8 Exercises

Exercise 36.1. (a) State what must be defined before the reliability of a system is calculated. (b) State the difference between the snapshot unavailability of the cooling infrastructure and an interruption of IT service. (c) Suppose both mechanisms of Section 36.2.2 act together:

the cooling capacity has fallen to 650 kW and 200 kW of unmetered heat is present. With $C_{\text{eff}} = 120\,\text{MJ/K}$, $T_0 = 25°\text{C}$ and $T_{\max} = 33°\text{C}$, compute the net heat accumulation, the rate of temperature rise and the time to the limit. (d) Under the degraded-capacity hypothesis alone, compute the time to the limit if the room starts at 27°C. (Answers: (a) the required function, the failure event and the assessment horizon. (b) thermal inertia, restoration and operational response can prevent an equipment loss from interrupting service, so an infrastructure deficiency is not an interruption until the thermal limit is crossed. (c) $900 + 200 - 650 = 450\,\text{kW}$; $0.450/120 = 3.75 \times 10^{-3}\,\text{K/s} = 0.225\,\text{K/min}$; $120 \times 8/0.450 = 2133\,\text{s} = 35.6\,\text{min}$. (d) $120 \times 6/0.250 = 2880\,\text{s} = 48\,\text{min}$.)

Exercise 36.2. (a) State what a rank-deficient sensitivity matrix indicates. (b) State why competing diagnostic hypotheses are retained instead of the most probable one being selected. (c) A second, independent temperature sensor is installed beside the return-air sensor of equation (36.7). Write the sensitivity matrix of the two readings with respect to (T_{act}, b_T) and give its rank. (d) Suppose the audit had assigned $\Pr(H_{\text{S}}) = 0.40$, $\Pr(H_{\text{C}}) = 0.40$ and $\Pr(H_{\text{H}}) = 0.20$, with the conditional probabilities of Table 36.5. Compute p_F, the expected annual number of interruptions, the annual probability of at least one interruption under a fixed rate and under the mixture, and the baseline risk. (Answers: (a) some unknown quantities, or combinations of them, cannot be inferred independently from the stated measurements; it is a property of the measurements, not a defect of the optimizer. (b) different mechanisms lead to different risks and different intervention benefits, and the least probable hypothesis here carries most of the risk. (c) $\left[\begin{smallmatrix}1 & 1\\ 1 & 0\end{smallmatrix}\right]$, rank two. (d) $p_F = 0.004 + 0.180 + 0.040 = 0.224$; 0.224 per year; $1 - \text{e}^{-0.224} = 0.2007$ against $0.40 \times 0.00995 + 0.40 \times 0.36237 + 0.20 \times 0.18127 = 0.1852$; $\mathcal{R}_0 = 0.224 \times 500\,000 = \text{R\$}\,112\,000$ per year.)

Exercise 36.3. (a) State why the intervention with the least expected cost can differ from the intervention with the least risk. (b) With the priors of equation (36.5) and the counterfactual probabilities of Table 36.6, suppose the consequence per interruption is $C_F = \text{R\$}\,300\,000$. Compute $\rho(a)$ for the four actions and name the least-cost action and the least-cost action that meets $p_F(a) \leq 0.05$. (c) With the probabilities of part (d) of the preceding exercise, compute $\rho(a)$ for the four actions at $C_F = \text{R\$}\,500\,000$ and compare the preferred action with that of Table 36.13. (Answers: (a) the intervention costs and the value assigned to the consequences change the objective; the least-risk action may cost more than the risk it removes. (b) no action $0.158 \times 300\,000 = \text{R\$}\,47\,400$; A $17\,400 + 25\,000 = \text{R\$}\,42\,400$; B $9150 + 45\,000 = \text{R\$}\,54\,150$; C $22\,800 + 15\,000 = \text{R\$}\,37\,800$; least cost C, compliant B. (c) p_F of 0.224, 0.064, 0.032 and 0.112 give R\$ 112 000, R\$ 57 000, R\$ 61 000 and R\$ 71 000; A is preferred, as in the rows of Table 36.13 with $q = 0.50$ and above, although the remaining probability is split 2 : 1 instead of 55 : 20.)

Exercise 36.4. (a) State what makes a diagnostic test valuable. (b) State why model adequacy is necessary even when the mathematics has been verified. (c) State the final purpose of an integrated reliability assessment. (d) A better test has $\Pr(Y^+ \mid H_{\text{S}}) = 0.02$, $\Pr(Y^+ \mid H_{\text{C}}) = 0.95$ and $\Pr(Y^+ \mid H_{\text{H}}) = 0.05$. With the priors of equation (36.5), compute $\Pr(Y^+)$, the posterior probability of a cooling deficiency after each result, the least-cost action after each result, C_{sample} and the EVSI, and check the EVSI against the EVPI. (Answers: (a) its outcomes can change the subsequent action enough to improve the expected decision; a test that leaves the action unchanged has no decision value, however accurate. (b) a correct computation can still rest on incorrect physical assumptions or omit mechanisms, and verification does not establish physical validity. (c) a traceable, uncertainty-qualified engineering decision. (d) $\Pr(Y^+) = 0.011 + 0.2375 + 0.010 = 0.2585$; $\Pr(H_{\text{C}} \mid Y^+) = 0.2375/0.2585 = 0.9188$ and $\Pr(H_{\text{C}} \mid Y^-) = 0.0125/0.7415 = 0.0169$; A after Y^+ (R\$ 52 050), C after Y^- (R\$ 25 866); $C_{\text{sample}} = 25\,000 \times 0.2585 + 500\,000 \times 0.013985 + 15\,000 \times 0.7415 + 500\,000 \times 0.016115 = 13\,455 + 19\,180 = \text{R\$}\,32\,635$; EVSI $= 53\,000 - 32\,635 = \text{R\$}\,20\,365$ per year, below the EVPI of R\$ 30 750.)