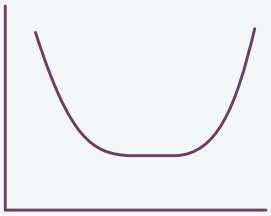
LESSON BOOK — VOLUME I



Reliability Engineering Volume I

Foundations, lifetime models, system logic, risk assessment and repairable systems

Twenty lessons · 509 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Reliability Engineering

Volume I

Foundations, lifetime models, system logic, risk assessment and repairable systems

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. This first volume 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. The 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, and a reference to a chapter of the other volume is written as such.

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. Every symbol is defined in the notation chapter that follows. The vocabulary is the standard one of the discipline, as fixed by Rausand, Barros and Høyland's *System Reliability Theory*, Modarres, Kaminskiy and Krivtsov's *Reliability Engineering and Risk Analysis*, Billinton and Allan's *Reliability Evaluation of Power Systems*, the NRC's *Fault Tree Handbook* 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 **viii**

I Foundations of Reliability Engineering **1**

1 Introduction **2**
1.1 Reliability as an engineering question . . . 2
1.2 Definition of reliability . . . 4
1.3 Component and system reliability . . . 7
1.4 Risk and availability . . . 9
1.5 Reliability of the physical system . . . 12
1.6 Worked example: chiller plant . . . 16
1.7 What to retain . . . 20
1.8 Exercises . . . 21

2 Functions, boundaries and failure **23**
2.1 Required function and performance criterion . . . 23
2.2 System boundaries and hierarchical failure . . . 26
2.3 Multiple requirements and time horizons . . . 29
2.4 Worked examples of boundaries and failure definitions . . . 34
2.5 Functional decomposition and failure-mode analysis . . . 38
2.6 Conditional models and the probabilistic formulation . . . 39
2.7 What to retain . . . 42
2.8 Exercises . . . 43

3 Reliability, availability, maintainability and risk **45**
3.1 Four measures of dependability . . . 45
3.2 Reliability as survival without failure . . . 47
3.3 Availability of a repairable item . . . 49
3.4 Mean times and maintainability . . . 52
3.5 Risk and the frequency of interruptions . . . 56
3.6 Power systems and the RAMS framework . . . 59
3.7 What to retain . . . 62
3.8 Exercises . . . 64

4 Failure mechanisms and FMEA **65**
4.1 Failure cause, mechanism, mode and effect . . . 65
4.2 Physical failure mechanisms . . . 68
4.3 Mathematical failure models . . . 70

4.4 Fatigue as physics-based prediction . . . 75
4.5 Failure mode and effects analysis . . . 79
4.6 Criticality and quantitative risk . . . 83
4.7 What to retain . . . 88
4.8 Exercises . . . 89

II Statistical Reliability Models 91

5 Lifetime distributions and hazard functions 92
5.1 Lifetime as a random variable . . . 92
5.2 Distribution, density and reliability . . . 94
5.3 The hazard function . . . 97
5.4 Conditional failure probability over an interval . . . 100
5.5 Hazard shapes and worked models . . . 102
5.6 Mean life and condition-dependent hazard . . . 107
5.7 What to retain . . . 111
5.8 Exercises . . . 112

6 Exponential, Weibull and lognormal models 114
6.1 The choice of a lifetime distribution . . . 114
6.2 The exponential distribution . . . 116
6.3 The Weibull distribution . . . 119
6.4 The lognormal distribution . . . 124
6.5 Comparing the three models . . . 128
6.6 Estimating and selecting a lifetime model . . . 130
6.7 What to retain . . . 135
6.8 Exercises . . . 136

7 Reliability data and statistical estimation 138
7.1 Reliability from incomplete operating records . . . 138
7.2 Censored lifetime data . . . 141
7.3 The Kaplan–Meier estimator . . . 143
7.4 Likelihood and maximum likelihood . . . 147
7.5 Statistical uncertainty . . . 151
7.6 Diagnostics and engineering credibility . . . 157
7.7 What to retain . . . 161
7.8 Exercises . . . 162

8 Bayesian reliability analysis 164
8.1 Bayesian inference of a failure rate . . . 164
8.2 Gamma–exponential conjugacy . . . 168
8.3 Posterior uncertainty and credible intervals . . . 171
8.4 Posterior predictive reliability . . . 174
8.5 Reliability estimation with zero observed failures . . . 176
8.6 Updating, populations and model credibility . . . 179
8.7 What to retain . . . 183
8.8 Exercises . . . 184

III System Reliability and Fault Trees 186

9 Reliability block diagrams and redundancy 187
9.1 System reliability from component reliability . . . 187
9.2 Series systems . . . 189
9.3 Parallel and k-out-of-n systems . . . 192
9.4 Active and standby redundancy . . . 197
9.5 Independence and common-cause failure . . . 202
9.6 Combined structures, importance and capacity . . . 203
9.7 What to retain . . . 208
9.8 Exercises . . . 209

10 Fault tree construction and logic 210
10.1 Fault tree analysis and the top event . . . 210
10.2 The logic of AND and OR gates . . . 213
10.3 Worked example: cooling-water flow . . . 217
10.4 Boolean algebra and fault-tree reduction . . . 219
10.5 Coherence and other gates . . . 222
10.6 Fault-tree construction in practice . . . 228
10.7 What to retain . . . 230
10.8 Exercises . . . 233

11 Minimal cut sets and quantification 235
11.1 Cut sets and minimal cut sets . . . 235
11.2 Exact quantification of the top event . . . 239
11.3 Rare-event approximation and bounds . . . 242
11.4 Minimal path sets . . . 245
11.5 Computational methods for large fault trees . . . 248
11.6 Application to the data center cooling plant . . . 253
11.7 What to retain . . . 257
11.8 Exercises . . . 259

12 Reliability importance measures 261
12.1 Component importance . . . 261
12.2 Birnbaum importance . . . 265
12.3 Fussell–Vesely, RAW and RRW . . . 267
12.4 Comparison of the four measures . . . 271
12.5 Importance and engineering decisions . . . 274
12.6 Limitations of importance measures . . . 278
12.7 What to retain . . . 280
12.8 Exercises . . . 282

IV Dependencies and Probabilistic Risk Assessment 284

13 Common-cause and dependent failures 285
13.1 Redundancy and the independence assumption . . . 285
13.2 Mathematical dependence . . . 287
13.3 Types of dependence . . . 292
13.4 Shared environmental conditions . . . 297
13.5 Load sharing and conditional hazards . . . 299
13.6 Engineering practice and quantitative risk . . . 304

13.7 What to retain . . . 308
13.8 Exercises . . . 310

14 Beta-factor and alpha-factor models 311
14.1 Common-cause failure in redundant groups . . . 311
14.2 Failure rates and failure probabilities . . . 313
14.3 The beta-factor model . . . 315
14.4 The alpha-factor model . . . 319
14.5 Beta against alpha . . . 326
14.6 Estimating the parameters from data . . . 329
14.7 What to retain . . . 333
14.8 Exercises . . . 335

15 Event trees and accident sequences 336
15.1 Event trees and initiating events . . . 336
15.2 Event-tree branching and sequence probabilities . . . 339
15.3 Worked example: data center power continuity . . . 342
15.4 Worked example: transmission-line outage . . . 346
15.5 Fault-tree links, dependence and recovery timing . . . 348
15.6 Expected consequence, risk and model credibility . . . 353
15.7 What to retain . . . 357
15.8 Exercises . . . 359

16 Integrated probabilistic risk assessment 360
16.1 The integrated assessment . . . 360
16.2 System, functions and dependencies . . . 363
16.3 The integrated event tree . . . 366
16.4 Risk importance and decisions . . . 371
16.5 Uncertainty propagation . . . 375
16.6 Credibility and other systems . . . 379
16.7 What to retain . . . 383
16.8 Exercises . . . 385

V Repairable Systems and Availability 387

17 Repairable systems and counting processes 388
17.1 Repairable systems and recurrent failures . . . 388
17.2 Recurrent failure data and ROCOF . . . 391
17.3 The homogeneous Poisson process . . . 397
17.4 The nonhomogeneous Poisson process . . . 400
17.5 Repair assumptions . . . 403
17.6 Estimation, model selection and applications . . . 406
17.7 What to retain . . . 413
17.8 Exercises . . . 414

18 Renewal processes and maintenance models 416
18.1 The effect of repair on future failures . . . 416
18.2 Perfect repair and renewal processes . . . 419
18.3 Minimal repair . . . 422
18.4 Imperfect repair and virtual age . . . 424
18.5 Renewal-reward theory and age replacement . . . 428

18.6 Optimal age replacement 431
18.7 What to retain 439
18.8 Exercises 440

19 Markov models and state transitions **441**
19.1 From failure counting to system states 441
19.2 Markov property and transition rates 443
19.3 The single component and the Kolmogorov equations 445
19.4 Two repairable pumps 448
19.5 Steady state, repair resources and failure frequency 453
19.6 Extensions, applications and validation 458
19.7 What to retain 463
19.8 Exercises 465

20 Availability and maintenance optimization **467**
20.1 Availability as an engineering objective 467
20.2 The cooling-plant model and repair capacity 471
20.3 Spare parts and repair logistics 475
20.4 Comparison of maintenance strategies 476
20.5 Scheduling and spare-parts inventory 479
20.6 Investment choices under uncertainty 486
20.7 What to retain 489
20.8 Exercises 492

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 I

Foundations of Reliability Engineering

Chapter 1

Introduction to Reliability Engineering

1.1 Reliability as an engineering question

Engineered systems are built to do something: move water, remove heat, carry electric power. They are expected to keep doing it, and they sometimes stop. Whether a system will keep doing what it was built for, over a period that matters to someone, is a question asked every day by plant managers, facility operators and system operators, and the answer is never certain. Reliability engineering is the discipline that states such questions precisely and answers them with probabilities. The question of this chapter is

what it means for an engineered system to be reliable, and how that reliability is quantified.

The chapter answers it in six steps. It first poses three engineering situations and extracts their common structure. It then defines reliability, in words and as a probability, and separates the reliability of a component from that of a system. It distinguishes the probability of failure from risk, which adds consequences, and from availability, which adds repair. It shows that reliability is a property of the physical system, which can fail by missing a performance requirement without breaking, and lists the ways in which failure probabilities are obtained. Finally, it works one example through, a cooling plant of three chillers of which two are required, and states what that calculation leaves out.

Jargon: item, required function, failure, reliability, availability, risk, MTTF, MTTR

An *item* is a component, subsystem or system treated as a unit of analysis. Its *required function* is what it must do for the system to count as operating, stated as a measurable performance criterion, such as "deliver at least $100\,\mathrm{m}^3/\mathrm{h}$ of cooling water." A *failure* is the termination of the item's ability to perform its required function; it is a functional event and need not coincide with physical breakage. *Reliability* is the probability that an item performs its required function under specified conditions for a specified period. *Availability* is the probability that an item is in an operating state at a given time, or the long-run fraction of time it is operating, restoration after failure included. *Risk* is the combination of the probabilities of outcomes and their consequences. *MTTF* (mean time to failure) is the expected operating duration before a failure; *MTTR* (mean time to repair or restore) is the expected duration of the restoration that follows. Both are measured in hours.

Key idea

Reliability is not a number attached to a piece of equipment. It is the probability of a precisely defined event, the item performing a stated function under stated conditions over a stated period, and it changes when any of the four changes; system reliability then follows from the arrangement of the components, and risk and availability are different quantities built on top of it.

1.1.1 Three engineering situations

Three situations, which recur throughout the book with growing complexity, show the kind of question the discipline answers. In the first, a manufacturing facility runs a pump continuously to supply cooling water to a process. The pump has been operating for five years, and the plant manager asks for the probability that it will continue to function for another year. The question is about one item, about a future period, and about an item that has already survived a long time, which matters for the answer (Section 1.2.5).

In the second, a data center is cooled by three chillers, of which only two are required to satisfy the cooling demand. The operator asks whether the facility will continue operating if one chiller fails, and what the probability is of losing sufficient cooling capacity. The question is about a system of identical units with a spare, and about a capacity requirement rather than the survival of any single unit.

In the third, an electric transmission network consists of generators, lines, substations and protective devices. The system operator asks for the probability that equipment failures will prevent the network from supplying its required electrical demand. The question is about a large system whose components interact through the network, and whose ability to carry power depends on the weather as well as on the equipment (Section 1.5.1).

1.1.2 The common structure

The three situations differ in scale and in physics, but they share a structure of four elements. A system is expected to perform a function. Some event, or some combination of events, can prevent it from performing that function. The occurrence of those events is uncertain. Engineering decisions can change their probability or their consequences. Table 1.1 lists each element for the three situations. Its rows are the four elements and its columns the three systems; the entries in the last two rows are typical examples, not an exhaustive list.

Table 1.1: The common structure of three reliability questions. The events, uncertainties and decisions are typical examples.

Element	Cooling-water pump	Chiller plant (two of three)	Transmission network
Function	deliver cooling water to the process	remove the data center's cooling load	supply the required electrical demand
Events that prevent it	pump failure; flow below the requirement	failure of two or more chillers, or of a dependency they share	failures of generators, lines, substations or protective devices; ratings reduced by weather
Uncertainty	when the pump will fail, given five years of service	which chillers fail, when, and the load at the time	which equipment fails; weather; demand
Decisions	maintenance, replacement, a spare pump	redundancy, maintenance schedule, load management	reinforcement, operating limits, redispatch of generation

Reliability engineering provides the methods for analyzing questions of this structure. Each element of the structure becomes a part of the analysis: the function becomes a success criterion, the preventing events become failure events, the uncertainty becomes a probability model, and the decisions are compared by how much they change the probabilities and the consequences. The rest of the chapter defines these parts in turn.

1.2 Definition of reliability

Reliability is defined in words before it is defined as a probability, because the words fix which event the probability is about. This section states the definition and its four elements, writes it as a function of time, and separates survival from a starting condition from survival given that an item has already lasted some time.

1.2.1 The four elements of the definition

A definition that serves engineering work is the following: *reliability is the probability that an item will perform its required function under specified conditions for a specified period of time.* The definition has four elements, and a reliability statement is complete only when each of them is specified. Table 1.2 lists them. Each row is one element; the middle column is the question the element answers, and the right column gives its value for the cooling-water pump.

Table 1.2: The four elements of the definition of reliability, with the values for the cooling-water pump (illustrative).

Element	Engineering question	Pump example
Item	what is being evaluated	the cooling-water pump
Required function	what it must accomplish	deliver at least $100\,\mathrm{m^3/h}$
Operating conditions	under what circumstances	specified temperature, pressure and fluid
Time	over what period	the next 1000 operating hours

It follows that reliability is not an intrinsic numerical property of a piece of equipment, independent of its application. The same pump may be highly reliable under moderate operating conditions and considerably less reliable when it is exposed to high temperatures, to cavitation (the formation and collapse of vapor bubbles where the local pressure falls below the vapor pressure of the liquid) or to abrasive particles in the fluid. Changing the conditions changes the probability, because it changes the rate at which the physical processes that end the pump's function proceed.

1.2.2 Functional failure without physical breakdown

The required function is the element most often left implicit, and it is the one that decides what counts as a failure. A pump may continue to rotate and still fail to deliver the required flow: a worn impeller, a partly blocked suction strainer or a leaking seal can reduce its flow to $85\,\mathrm{m^3/h}$ while the motor runs normally. Against a requirement of $100\,\mathrm{m^3/h}$, that pump has failed, although an inspection that checks only whether it runs would report it as working.

A component therefore does not need to stop operating to fail its required function. Failure is defined relative to the function, and the function is defined by a measurable performance criterion. The distinction is central to the study of physical systems, in which performance degrades continuously before any part breaks, and it returns in Section 1.5.2 in the form of a

performance-based failure event. Chapter 2 treats the definition of functions, system boundaries and failure events in full.

Principle 1.1. *Reliability is function-dependent: successful operation must be defined, as a measurable criterion, before failure can be defined, and an item that runs but misses its criterion has failed.*

1.2.3 The reliability function and the cumulative failure probability

Let T denote the random time to failure of a nonrepairable item, measured in operating hours from a defined starting condition, such as the moment it is put into service in working order; $T \geq 0$. For a specified operating duration t, the reliability of the item is the probability that it survives beyond t,

$$R(t) = \Pr(T > t). \tag{1.1}$$

The function $R(t)$ is called the reliability function, or the survivor function. Its complement is the cumulative failure probability, the probability that the item has failed by time t,

$$F(t) = \Pr(T \leq t) = 1 - R(t). \tag{1.2}$$

The events $\{T > t\}$ and $\{T \leq t\}$ are mutually exclusive and together certain, so their probabilities add to one,

$$R(t) + F(t) = 1. \tag{1.3}$$

In the language of probability, F is the cumulative distribution function of T and R its complement; both are dimensionless and lie between 0 and 1.

As an example, suppose a pump operates successfully for 1000 hours with probability 0.95. Then $R(1000\,\mathrm{h}) = 0.95$ and $F(1000\,\mathrm{h}) = 1 - 0.95 = 0.05$: there is a 5 % probability that the pump fails during those 1000 hours. The number 0.95 is attached to the duration 1000 h; over 2000 h the same pump has a different reliability, which cannot be larger (equation (1.4)). Reliability always refers to a time horizon, and a statement that a pump "has 95 % reliability" without the interval is incomplete.

Reliability cannot increase with the duration. For two durations $t_1 < t_2$, an item that survives beyond t_2 has also survived beyond t_1, so the event $\{T > t_2\}$ is contained in the event $\{T > t_1\}$ and

$$R(t_2) = \Pr(T > t_2) \leq \Pr(T > t_1) = R(t_1). \tag{1.4}$$

The reliability function of a lifetime model is therefore nonincreasing. It starts at $R(0) = 1$ when the item starts in working order, and it tends to zero as t grows if the item is certain to fail eventually.

1.2.4 Reliability as a function of time

Figure 1.1 draws one reliability function. The horizontal axis is the operating time t in hours, from 0 to 10 000; the vertical axis is a probability. The solid curve is the reliability $R(t)$, the dashed curve the cumulative failure probability $F(t) = 1 - R(t)$, and the dots mark $R(t)$ at every 1000 hours. The curve is an illustrative exponential model,

$$R(t) = \mathrm{e}^{-t/10\,000\,\mathrm{h}}, \tag{1.5}$$

which describes an item whose failure rate does not change with age, a constant hazard of $10^{-4}\,\mathrm{h}^{-1}$; its mean time to failure is 10 000 h. The model is the simplest standard lifetime model, and Chapter 5 derives it from the hazard function.

The figure shows the two functions mirror each other about 0.5, as equation (1.3) requires. The reliability falls steadily: 0.905 at 1000 h, 0.607 at 5000 h, and 0.368 at 10 000 h, the mean time to failure. More than a third of such items therefore outlast their mean life, and almost two thirds fail before it: for the exponential model, $F(\text{MTTF}) = 1 - \text{e}^{-1} = 0.632$.

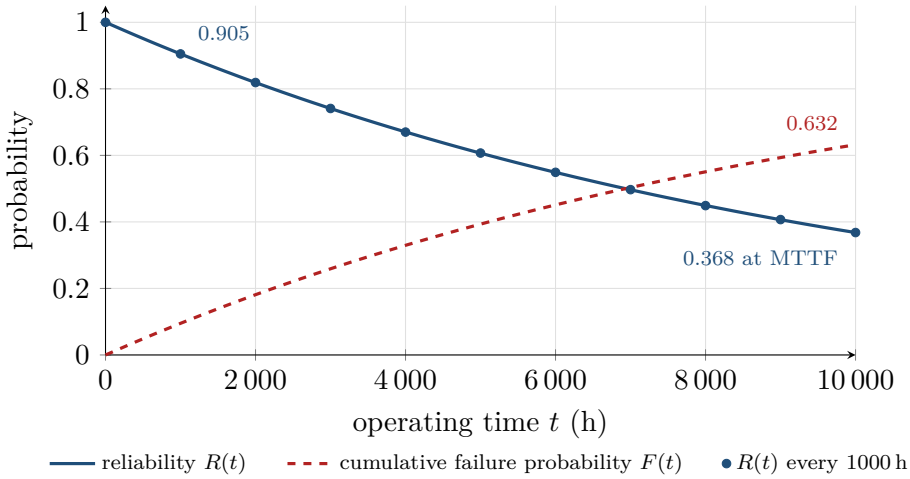


Figure 1.1: Illustrative reliability function $R(t) = \text{e}^{-t/10\,000\,\text{h}}$ (equation (1.5)), an exponential model with constant hazard $10^{-4}\,\text{h}^{-1}$ and MTTF = 10 000 h, and its complement $F(t) = 1 - R(t)$. Reliability decreases with operating time; at $t =$ MTTF it is 0.368.

1.2.5 Conditional reliability

A reliability function cannot increase with the duration, but a reliability prediction can change when new information about the item's condition arrives. The commonest piece of information is that the item has already survived for some time. Two quantities must then be kept apart. The reliability $R(t)$ describes survival over a specified duration measured from the starting condition. The conditional reliability $\Pr(T > t + \Delta t \mid T > t)$ describes survival over a further duration Δt given that the item has already survived to time t. By the definition of conditional probability, and because $\{T > t + \Delta t\}$ is contained in $\{T > t\}$,

$$R(t + \Delta t \mid t) = \Pr(T > t + \Delta t \mid T > t) = \frac{\Pr(T > t + \Delta t)}{\Pr(T > t)} = \frac{R(t + \Delta t)}{R(t)}. \tag{1.6}$$

The two answer different questions. The plant manager of the first situation in Section 1.1.1 asks the conditional one: the pump has survived five years, and the question is about the next year.

Table 1.3 evaluates both quantities for two illustrative lifetime models with the same mean time to failure, 50 000 h, for a pump in continuous operation (1 yr = 8760 h). The first is the exponential model $R(t) = \text{e}^{-t/50\,000\,\text{h}}$; the second is a Weibull model with shape $\beta = 2$ and scale $\eta = 56\,419$ h, $R(t) = \exp[-(t/\eta)^{\beta}]$, a standard model of an item that wears out, treated in a later chapter on lifetime distributions. Each row is an age t. For each model the table gives the unconditional reliability $R(t)$ and the conditional reliability over the next year. As an example of one entry, the Weibull pump at five years has

$$R(5\,\text{yr}) = \exp\left[-\left(\frac{43\,800}{56\,419}\right)^2\right] = 0.5473, \qquad R(6\,\text{yr}) = \exp\left[-\left(\frac{52\,560}{56\,419}\right)^2\right] = 0.4198,$$

and so a conditional reliability for the sixth year of $0.4198/0.5473 = 0.7671$.

Table 1.3: Unconditional reliability $R(t)$ and conditional reliability over the next year, $R(t + 8760\,\text{h} \mid t)$ (equation (1.6)), for two illustrative pump models with MTTF = 50 000 h: exponential, and Weibull with $\beta = 2$, $\eta = 56\,419$ h.

	Exponential		Weibull, $\beta = 2$	
Age t (years)	$R(t)$	next year, given t	$R(t)$	next year, given t
0	1.0000	0.8393	1.0000	0.9762
1	0.8393	0.8393	0.9762	0.9302
3	0.5912	0.8393	0.8050	0.8447
5	0.4164	0.8393	0.5473	0.7671

The table shows two behaviors. Under the exponential model the conditional reliability for the next year is $\mathrm{e}^{-8760/50\,000} = 0.8393$ at every age: from equation (1.6), $\mathrm{e}^{-(t+\Delta t)/\mathrm{MTTF}}/\mathrm{e}^{-t/\mathrm{MTTF}} = \mathrm{e}^{-\Delta t/\mathrm{MTTF}}$, independent of t. A surviving item is as good as new, a property called memorylessness. Under the wear-out model the conditional reliability falls with age, from 0.9762 for a new pump to 0.7671 for a five-year-old one. In both models the conditional value at five years is much larger than the unconditional probability of surviving six years from new (0.3495 and 0.4198), because the first five years are already known to have been survived. The answer to the plant manager's question therefore depends on which lifetime model describes the pump, and the evidence that selects the model is a subject of the statistical chapters that follow.

1.3 Component and system reliability

Systems are built of components, and the reliability of a system follows from the reliabilities of its components and from the way they are arranged. This section derives the two basic arrangements, series and parallel, for a cooling system of two pumps, and states the assumption on which both derivations rest.

1.3.1 Two pumps in series

Consider a cooling system of two pumps, A and B, arranged in series, so that both must function for the water to circulate. Panel (a) of Figure 1.2 draws it as a reliability block diagram: each block is a component, the line runs from the input on the left to the output on the right, and the system functions if a path of functioning blocks connects the two ends. In a series arrangement there is one path, through both blocks. Each pump has reliability $R_A = R_B = 0.95$ over the same period.

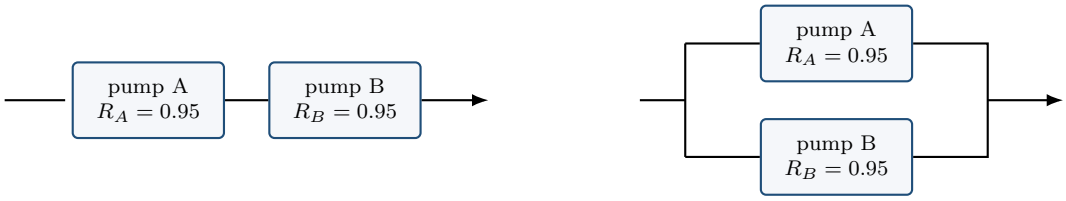
The system functions only if A functions and B functions. If the failures of A and B are independent, the probability of the intersection of the two events is the product of their probabilities,

$$R_{\text{system}} = \Pr(\text{A functions} \cap \text{B functions}) = R_A\, R_B, \tag{1.7}$$

and, for n components in series, $R_{\text{system}} = \prod_{i=1}^{n} R_i$. For the two pumps,

$$R_{\text{system}} = 0.95 \times 0.95 = 0.9025.$$

The system reliability is 90.25 %, although each pump alone has 95 %. The system fails if either pump fails, and its failure probability, $1 - 0.9025 = 0.0975$, is almost the sum of the two component failure probabilities, $0.05 + 0.05 = 0.10$; the difference, $0.05^2 = 0.0025$, is the probability that both fail, which the sum counts twice. A series system is less reliable than its least reliable component, and every component added in series lowers its reliability further.



(a) series: both pumps required
$R_{\text{system}} = 0.95^2 = 0.9025$

(b) parallel: either pump sufficient
$R_{\text{system}} = 1 - 0.05^2 = 0.9975$

Figure 1.2: Reliability block diagrams of a two-pump cooling system, each pump of reliability 0.95 over the same period, with independent failures. (a) Series, equation (1.7): one path through both blocks. (b) Parallel, equation (1.8): two paths, one through each block. The system functions if functioning blocks connect the left end to the right end.

1.3.2 Two pumps in parallel

Suppose now that either pump can provide sufficient cooling on its own, as in panel (b) of Figure 1.2. There are two paths, one through each pump, and the system fails only if both pumps fail. This arrangement is called redundancy: equipment is installed beyond what the function requires, so that the function survives the loss of some of it. The probability that a pump fails is $1 - R_i$; if the failures are independent, the probability that both fail is the product, and the system reliability is its complement,

$$R_{\text{system}} = 1 - \Pr(\text{A fails} \cap \text{B fails}) = 1 - (1 - R_A)(1 - R_B), \tag{1.8}$$

and, for n components in parallel, $R_{\text{system}} = 1 - \prod_{i=1}^{n}(1 - R_i)$. For the two pumps,

$$R_{\text{system}} = 1 - (0.05)(0.05) = 1 - 0.0025 = 0.9975.$$

The same two pumps give 90.25 % in series and 99.75 % in parallel. The failure probability falls from 0.0975 to 0.0025, a factor of 39, and relative to one pump alone from 0.05 to 0.0025, a factor of 20. In series, failure probabilities approximately add; in parallel, they multiply.

1.3.3 The independence assumption

Both results rest on one assumption: the failures of the two pumps are independent, so that the probability of a joint event is the product of the individual probabilities. The assumption fails whenever the pumps share something whose failure affects both: the same power supply, a common cooling-water source, or one controller. Then the failure of the shared element removes both pumps at once, and the multiplication in equation (1.8) no longer describes the system.

An illustrative calculation shows the size of the effect. Let both pumps be fed from one power supply of reliability 0.99 over the same period, and let the pumps otherwise fail independently as before. The supply is in series with the parallel pair, since the system needs the supply and at least one pump, and

$$R_{\text{system}} = 0.99 \times 0.9975 = 0.987525.$$

The system failure probability rises from 0.0025 to 0.012475, five times higher, and four fifths of it comes from the shared supply alone. A component shared by redundant paths sits in series with all of them and caps the reliability of the whole at its own. Failures of several items from one shared cause are called common-cause failures; later chapters on fault trees and on common-cause failure models treat them in detail.

Principle 1.2. *System reliability depends on architecture: series arrangements multiply reliabilities, parallel arrangements multiply failure probabilities, and both rules hold only when the component failures are independent.*

1.4 Failure probability, risk and availability

The probability of failure is one of several quantities used to judge a system, and it is often confused with two others. This section distinguishes failure probability from risk, which weighs outcomes by their consequences, and reliability from availability, which credits restoration after failure.

1.4.1 Failure probability and consequence

Consider two equipment failures in the same facility over one year. The first is the failure of pump A, with an annual probability of 10 %, whose consequence is the temporary loss of a minor process. The second is the failure of the main cooling system, with an annual probability of 1 %, whose consequence is a major data center outage. The first event is ten times more probable; the second may be far more consequential. Ranking the two by probability alone puts the pump first, and that ranking says nothing about which deserves attention, because it ignores what each failure costs.

The two events illustrate why the probability of failure is not the same as risk. Risk combines the probability of an outcome with its consequence, and a rare event with a large consequence can carry more risk than a frequent event with a small one.

1.4.2 Risk as expected loss

A simple quantitative measure of risk for one failure event F is

$$\mathcal{R} = \Pr(F)\, C(F), \tag{1.9}$$

where $\Pr(F)$ is the probability of the specified failure event over a stated period and $C(F)$ is its consequence, conditional on the event occurring, here in Brazilian reais (R$). When the period is one year, $\mathcal{R}$ is an expected annual loss in R$ per year. Table 1.4 applies equation (1.9) to the two events, with illustrative losses. Each row is one event; the columns give its annual probability, the loss if it occurs, and the product. For the pump, $\mathcal{R} = 0.10 \times \text{R\$}\,10\,000 = \text{R\$}\,1000$ per year; for the main cooling system, $\mathcal{R} = 0.01 \times \text{R\$}\,1\,000\,000 = \text{R\$}\,10\,000$ per year. The main cooling system has one tenth of the failure probability and ten times the expected annual loss.

Table 1.4: Failure probability and risk of two events, illustrative losses. The expected annual loss is the product of the annual probability and the loss, equation (1.9).

Event	Consequence	Probability per year	Loss if it occurs	Expected annual loss
Pump A fails	temporary loss of a minor process	0.10	R$ 10 000	R$ 1000
Main cooling fails	major data center outage	0.01	R$ 1 000 000	R$ 10 000

Figure 1.3 places the two events on a plane of probability against consequence. The horizontal axis is the loss if the event occurs, in reais, and the vertical axis the annual probability, both on logarithmic scales. The two dashed lines are lines of constant risk, $\Pr \cdot C = \text{R\$}\,1000$ and R$ 10 000 per year; on logarithmic axes they are straight lines of slope -1. The pump lies on the lower line and the main cooling system on the upper one: it sits lower on the probability axis and higher in risk. Every event on one line carries the same risk, so a hundredfold rarer event with a hundredfold larger loss is equivalent in expectation.

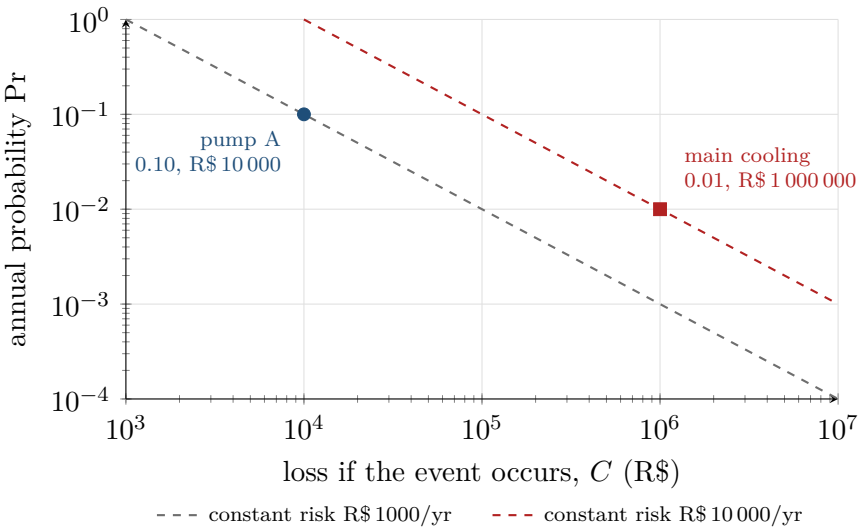


Figure 1.3: Annual probability against consequence for the two events of Table 1.4, logarithmic axes. Dashed lines are lines of constant risk $\Pr \cdot C$. The main cooling system is ten times less probable than the pump failure and carries ten times its risk. Values illustrative.

Several outcomes. When a failure can lead to several outcomes S_i with different consequences, such as a short outage, a long one and equipment damage, the risk is the probability-weighted sum

$$\mathcal{R} = \sum_i \Pr(S_i)\, C(S_i). \tag{1.10}$$

The expression is the expected value of the consequence, and it is correct when the outcomes form a partition, mutually exclusive and together exhaustive (the outcome "no failure" included, with consequence zero), and when the consequences are defined on one scale and over one period. If the outcomes overlapped, a scenario belonging to two of them would be counted twice. Two qualifications sharpen equation (1.9). First, for separate failure events whose losses add, linearity of expectation gives the expected total loss as the sum of the separate products even when the events can occur together: the facility of Table 1.4 has an expected loss of R$ 1000 + R$ 10 000 = R$ 11 000 per year from the two events. Second, an event that can occur more than once in a year contributes the expected number of its occurrences times the loss per occurrence; for a rare event the expected number and the probability of at least one occurrence are nearly equal, which is the case equation (1.9) represents. Risk analysis is therefore concerned not only with the likelihood of failure but with what happens after it, including the cases in which several outcomes are possible.

1.4.3 Reliability and availability

Consider two pumps. Pump P fails every few months but is repaired within a day; pump Q rarely fails, but each repair takes three weeks, the time needed to obtain a spare part. Which pump is more reliable and which is more available are two different questions, and they can have different answers. Reliability concerns successful operation without failure over a specified interval. Availability concerns whether the item is operational at a particular time, or the fraction of time it is operational, and it credits the restoration that follows each failure.

Figure 1.4 draws one illustrative year of operation of each pump. The horizontal axis is operating time over one year, 8760 h; each row is one pump; green bars are operating intervals and red bars repair intervals. Pump P fails four times and is down for 24 h after each failure, 96 h in all; its repair intervals are too short to show at this scale and are drawn at the narrowest

visible width. Pump Q fails once and is down for 504 h. Pump Q runs for long stretches without interruption, so it is the more reliable; pump P is operating for a larger fraction of the year, $8664/8760 = 0.989$ against $8256/8760 = 0.942$ for Q, so in this year it is the more available.

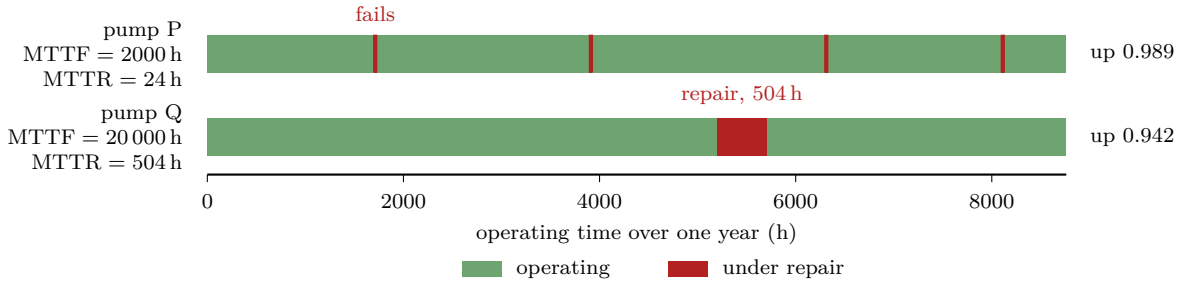


Figure 1.4: One illustrative year of operation of two repairable pumps. Pump P fails four times and is repaired in 24 h each time (red marks drawn at the narrowest visible width); pump Q fails once and is repaired in 504 h. Pump Q is more reliable over any interval; pump P is operating for a larger fraction of the time.

1.4.4 Steady-state availability

For a repairable item that alternates between operating intervals, which end in a failure, and repair intervals, which end in a restoration, the long-run fraction of time it operates follows from a short argument. Let U_k be the length of the k-th operating interval and D_k that of the k-th repair interval, and suppose the cycles are statistically alike and independent, with means $\mathbb{E}[U_k] = \mathrm{MTTF}$ and $\mathbb{E}[D_k] = \mathrm{MTTR}$. Over the first n cycles the fraction of time operating is

$$\frac{\sum_{k=1}^{n} U_k}{\sum_{k=1}^{n} U_k + \sum_{k=1}^{n} D_k} = \frac{\frac{1}{n}\sum_{k} U_k}{\frac{1}{n}\sum_{k} U_k + \frac{1}{n}\sum_{k} D_k},$$

and, by the law of large numbers, each average tends to its mean as n grows. The limit is the steady-state availability,

$$A = \frac{\mathrm{MTTF}}{\mathrm{MTTF} + \mathrm{MTTR}}, \tag{1.11}$$

a standard result for this model, which is called an alternating renewal process. Here MTTF is the mean operating duration before a failure in the alternating model, and MTTR is the mean time to repair or restore, as defined for the application (detection, travel, spare parts and testing included or not). Equation (1.11) is also the long-run probability that the item is operating at a randomly chosen time.

As an example, let $\mathrm{MTTF} = 1000\,\mathrm{h}$ and $\mathrm{MTTR} = 10\,\mathrm{h}$. Then

$$A = \frac{1000\,\mathrm{h}}{1000\,\mathrm{h} + 10\,\mathrm{h}} = \frac{1000}{1010} = 0.990099,$$

a long-run availability of about 99.01 %. The number does not mean that the probability of failure over the next 1000 hours is 0.99 %. With a constant failure rate, $R(t) = \mathrm{e}^{-t/\mathrm{MTTF}}$, and the probability of at least one failure in 1000 h is $1 - \mathrm{e}^{-1000/1000} = 0.632$: the item fails, on average, $8760/1010 = 8.67$ times a year, and it is available 99 % of the time only because each failure is repaired in ten hours.

Table 1.5 compares this item with the pumps P and Q of Figure 1.4. Each row is one item; the columns give its mean times, its steady-state availability from equation (1.11), its reliability over 1000 h with a constant failure rate, and the long-run number of failures and hours of downtime per year, $8760/(\mathrm{MTTF} + \mathrm{MTTR})$ and $(1 - A) \times 8760\,\mathrm{h}$. Pump P has the

higher availability, 0.9881 against 0.9754, and much the lower reliability, 0.607 against 0.951. The ranking of the two pumps reverses between the two measures, so a statement about one says nothing about the other. In the long run pump Q is down about 215 h a year, more than in the single year of Figure 1.4, which is one realization of a random process.

Table 1.5: Steady-state availability (equation (1.11)) and reliability over 1000 h with a constant failure rate, for three illustrative repairable items. Failures and downtime are long-run averages per year of 8760 h.

Item	MTTF (h)	MTTR (h)	A	$R(1000\,\text{h})$	Failures per year	Downtime (h/yr)
Example item	1000	10	0.9901	0.368	8.67	86.7
Pump P	2000	24	0.9881	0.607	4.33	103.9
Pump Q	20 000	504	0.9754	0.951	0.43	215.3

Reliability and availability therefore quantify different properties. Reliability is the right measure when any interruption is costly, as for a mission that cannot be resumed; availability is the right measure when the cost is set by the total time out of service. Chapter 3 treats the two together with maintainability and risk, and later chapters on repairable systems, Markov models and maintenance return to availability as a function of time.

Principle 1.3. *Reliability, availability and risk are different quantities: reliability is the probability of operating without failure over an interval, availability the probability of being in operation, restoration included, and risk weighs the probabilities of outcomes by their consequences.*

1.5 Reliability as a property of the physical system

The calculations so far took the failure probability of each component as a given number. This section asks where such numbers come from, shows with a transmission line that a failure can be the violation of a physical performance limit computed from a model, sets the lifetime and the performance-based formulations of reliability side by side, and lists the approaches by which failure probabilities are obtained.

1.5.1 The conductor heat balance

Consider one overhead conductor of an electric transmission line. The current it may carry is limited by its temperature: a hot conductor expands and sags toward the ground and, held hot for long enough, loses strength. The permissible current, called the ampacity or thermal rating of the line, therefore depends on the conductor's thermal behavior, which is affected by the conductor temperature, the ambient temperature, the wind speed and direction, the solar radiation, the properties of the conductor, and the electrical loading.

Figure 1.5 draws the heat balance of a length of conductor. The box is the conductor, which stores heat; the arrows entering it are heat gains, the arrows leaving it heat losses, and the labels on the outside name the external conditions that set each term. The balance is

$$mc_p \frac{\mathrm{d}T}{\mathrm{d}t} = I^2 R(T) + q_{\text{solar}} - q_{\text{conv}} - q_{\text{rad}}, \tag{1.12}$$

where $T(t)$ is the conductor temperature in °C, mc_p is the heat capacity of the modeled length in J/K, I is the current in amperes, $R(T)$ is the electrical resistance of the length in ohms,

which rises with temperature, q_{solar} is the solar heat gain, and q_{conv} and q_{rad} are the convective and radiative heat losses, all in watts. The first term on the right is the Joule heating; the convective loss grows with the wind speed and with the difference between the conductor and the ambient temperature. The temperature rises when the gains exceed the losses and settles where they balance.

Remark 1.1. In equation (1.12), $R(T)$ is the temperature-dependent electrical resistance of the conductor. It is not the reliability function, which is always written $R(t)$ and has time, not temperature, as its argument. Likewise $T(t)$ here is a temperature, not the time to failure T.

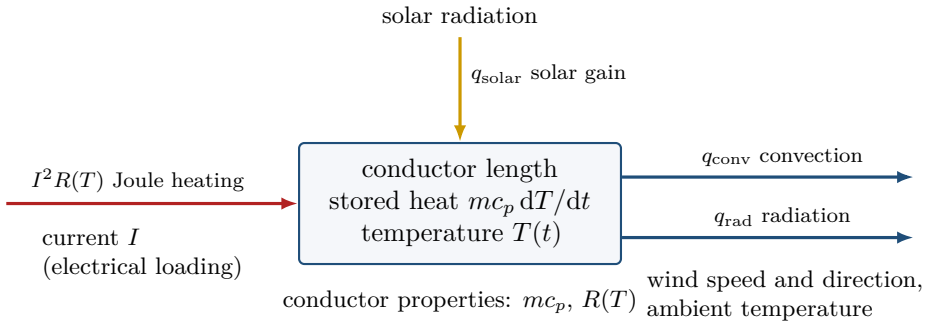


Figure 1.5: Heat balance of a length of overhead conductor, equation (1.12). Joule heating and solar gain enter; convection and radiation leave; the difference changes the stored heat and hence the conductor temperature. The labels outside the arrows name the conditions that set each term.

1.5.2 Failure as a performance event

Suppose the required operating criterion of the line is that the conductor temperature stays at or below a limit, $T(t) \le T_{\max}$, over an operating horizon H. A failure event is then

$$F = \left\{ \max_{0 \le t \le H} T(t) > T_{\max} \right\}, \tag{1.13}$$

the event that the temperature exceeds its limit at some time in the horizon, and its probability is

$$\Pr(F) = \Pr\left(\max_{0 \le t \le H} T(t) > T_{\max} \right). \tag{1.14}$$

The probability arises because the inputs of equation (1.12), the current, the wind, the ambient temperature and the sun, are uncertain over the horizon, so $T(t)$ is a random trajectory.

Figure 1.6 shows three such trajectories. The horizontal axis is the time of day over a horizon $H = 24\,\text{h}$; the vertical axis is the conductor temperature. The dashed red line is an illustrative limit $T_{\max} = 75°\text{C}$. Each curve is the temperature of the same conductor under a different realization of weather and load: all three rise in the afternoon, when load and sun peak. Realizations 1 and 2 stay below the limit; realization 3 crosses it between about 13.3 and 16.7 h (shaded), and in that realization the event F occurs. $\Pr(F)$ is the probability, over all realizations, of a trajectory like the third.

This is a different type of reliability problem from predicting when a mechanical bearing will fail. The conductor may remain structurally intact while its operating temperature exceeds the allowable threshold; the event represents a failure to meet an operating requirement, not necessarily permanent damage. When the temperature falls again, the line is again within its rating, and nothing has to be repaired.

END OF THE PREVIEW

Reliability Engineering
Volume I

These are the opening pages. The complete book runs 509 pages.

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

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com