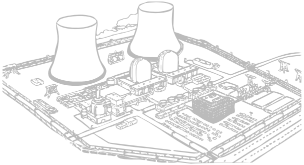

Overview of reliability engineering

Eric Marsden

<eric.marsden@risk-engineering.org>

Context

- I have a fleet of airline engines and want to anticipate when they may fail
- I am purchasing pumps for my refinery and want to understand the MTBF, lambda etc. provided by the manufacturers
- I want to compare different system designs to determine the impact of architecture on availability

Reliability engineering

- ▷ Reliability engineering is the discipline of ensuring that a system will function as required over a specified time period when operated and maintained in a specified manner.

- ▷ Reliability engineers address 3 basic questions:
  - When does something fail?
  - Why does it fail?
  - How can the likelihood of failure be reduced?

Failure

Failure

Loss of ability to perform as required.

- A failure is always related to a required **function**. The function is often specified together with a performance requirement (eg. “must handle up to 3 tonnes per minute”, “must respond within 0.1 seconds”).

- A failure occurs when the function cannot be performed or has a performance that falls outside the performance requirement.

Source: International Electrotechnical Vocabulary (IEV) part 192 on dependability

Fault

Fault

Inability to perform as required, due to an internal state [IEV 192-04-01]

- ▷ While a failure is an event that occurs at a specific point in time, a fault is a state that will last for a shorter or longer period.
- ▷ When a failure occurs, the item enters the failed state. A failure may occur:
  - while running
  - while in standby
  - due to demand

Source: International Electrotechnical Vocabulary (IEV) part 192 on dependability

Error

Error

Discrepancy between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition.

- An error is present when the performance of a function deviates from the target performance, but still satisfies the performance requirement
- An error will often, but not always, develop into a failure

Source: International Electrotechnical Vocabulary (IEV) part 192 on dependability, item 192-03-02

Failure mode

Failure mode

The way a failure is observed on a failed item.

- An item can fail in many different ways: a failure mode is a description of a possible state of the item after it has failed

Source: International Electrotechnical Vocabulary (IEV) part 192 on dependability, item 192-03-17

Failure classification

IEC 61508 classifies failures according to their:

- ▷ Causes:
  - random (hardware) faults
  - systematic faults (including software faults)
- ▷ Effects:
  - safe failures
  - dangerous failures
- ▷ Detectability:
  - detected: revealed by online diagnostics
  - undetected: revealed by functional tests or upon a real demand for activation

Markovian models

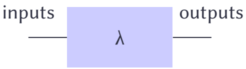

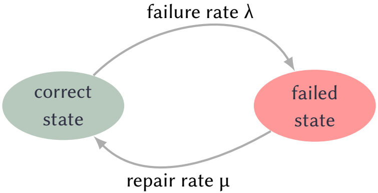

Models the transitions between correct state and failed state.

Assumption: nothing in the past determines future events except for current state.

Failure and repair are stochastic processes.

Availability = proportion of time spent in correct state.

The “safe failure fraction”

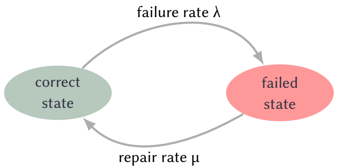

The “safe failure fraction”

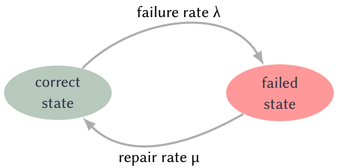

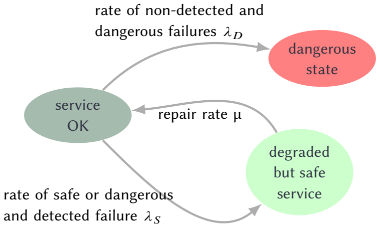

rate of non-detected and dangerous failures λ_D
dangerous state
service OK
repair rate μ
degraded but safe service
rate of safe or dangerous and detected failure λ_S

Not all failures are dangerous: the system may have been designed to tolerate them.

The “safe failure fraction”

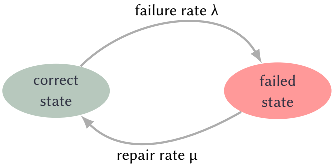

failure rate λ

correct state

failed state

repair rate μ

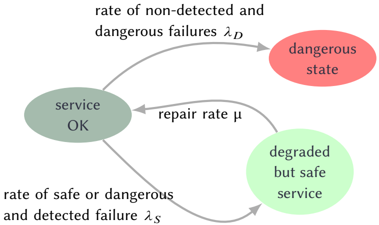

rate of non-detected and dangerous failures λ_D

dangerous state

service OK

repair rate μ

degraded but safe service

rate of safe or dangerous and detected failure λ_S

Not all failures are dangerous: the system may have been designed to tolerate them.

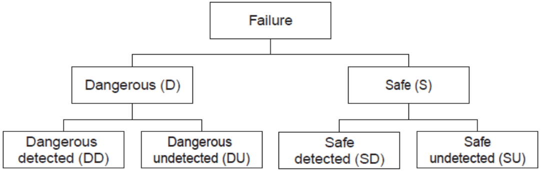

Failure

Dangerous (D)

Safe (S)

Dangerous detected (DD)

Dangerous undetected (DU)

Safe detected (SD)

Safe undetected (SU)

Importance of the **coverage** of the error detection mechanisms, measured by the “safe failure fraction”: conditional probability that a failure will be safe, or dangerous-but-detected.

Failure classification

- Safe undetected (SU): A spurious (untimely) activation of a component when not demanded
- Safe detected (SD): A non-critical alarm raised by the component
- Dangerous detected (DD): A critical diagnostic alarm reported by the component, which will, as long as it is not corrected prevent the safety function from being executed
- Dangerous undetected (DU): A critical dangerous failure which is not reported and remains hidden until the next test or demanded activation of the safety function

Common cause failures

Common cause failure

A failure that is the result of one or more events, causing concurrent failures of two or more separate channels in a multiple channel system, leading to system failure [IEC 61508]

- ▷ Typical examples: loss of electricity supply, massive physical destruction
- ▷ More subtle example: loss of clock function (electronics), common maintenance procedure

Reliability: definitions

Reliability [ISO 8402]

The ability of an item to perform a required function, under given environmental and operational conditions for a stated period of time.

- The *reliability* $R(t)$ of an item at time t is the probability that the item performs the required function in the interval $[0–t]$ given the stress and environmental conditions in which it operates

Reliability: definitions

- If X is a random variable representing time to failure of an item, the *survival function* (or *reliability function*) $R(t)$ is

$$R(t) = \Pr(X > t)$$

- $R(t)$ represents the probability that the item is working correctly at time t
- Properties:
  - $R(t)$ is non-increasing (no rising from the dead)
  - $R(0) = 1$ (no immediate death/failure)
  - $\lim_{t \to \infty} R(t) = 0$ (no eternal life)

RISK ENGINEERING

Interpreting the reliability function

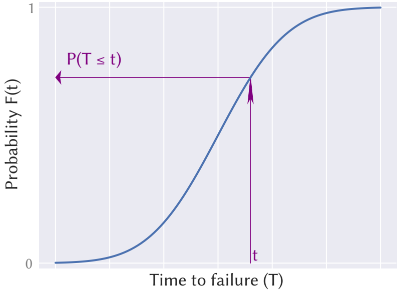

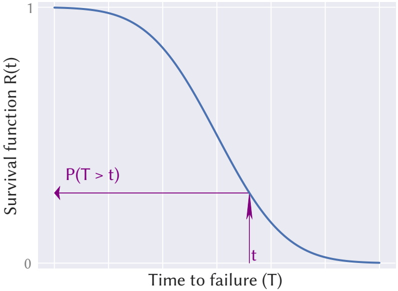

Cumulative distribution function

Tells you the probability that lifetime is ≤ *t*

$$F(t) = P(T \leq t)$$

Reliability function

Tells you the probability that lifetime is > *t*

$$R(t) = P(T > t) = 1 - F(t)$$

Exercise

Problem

The lifetime of a modern low-wattage electronic light bulb is known to be exponentially distributed with a mean of 8000 hours.

Q1 Find the proportion of bulbs that may be expected to fail before 7000 hours use.

Q2 What is the lifetime that we have 95% confidence will be exceeded?

For more on the reliability of solid-state lamps, see `energy.gov`

Exercise

Solution

The time to failure of our light bulbs can be modelled by the distribution

```
dist = scipy.stats.expon(scale=8000)
```

Q1: The CDF gives us the probability that the lifetime is $\leq t$. We want `dist.cdf(7000)` which is 0.583137. So about 58% of light bulbs will fail before they reach 7000 hours of operation.

Q2: We need the 0.05 quantile of the lifetime distribution, `dist.ppf(0.05)` which is around 410 hours.

Exercise

Problem

A particular electronic device will only function correctly if two essential components both function correctly. The lifetime of the first component is known to be exponentially distributed with a mean of 5000 hours and the lifetime of the second component (whose failures can be assumed to be independent of those of the first component) is known to be exponentially distributed with a mean of 7000 hours. Find the proportion of devices that may be expected to fail before 6000 hours use.

Exercise

Solution

The device will only be working after 6000 hours if both components are operating. The probability of the first component still working is

```
> pa = 1 - scipy.stats.expon(scale=5000).cdf(6000)
> pa
0.3011942119122022
```

and likewise for the second component

```
> pb = 1 - scipy.stats.expon(scale=7000).cdf(6000)
> pb
0.42437284567695
```

The probability of both working is pa × pb = 0.127818, so the proportion of devices that can be expected to fail before 6000 hours use is around 87%.

Hazard function

Hazard function

The hazard function or failure rate function $h(t)$ gives the conditional probability of failure in the interval t to $t + dt$, given that no failure has occurred by t.

$$h(t) = \frac{f(t)}{R(t)}$$

where $f(t)$ is the probability density function (failure density function) and $R(t)$ is the reliability function.

It's the probability of quitting a given state after having spent a given time in that state.

Bathtub curve

- Early failure (“burn-in”, “infant mortality” period): high hazard rate due to manufacturing and design problems

- Useful life period: probability of failure is roughly constant

- Wearout period: hazard rate starts to increase due to aging (corrosion, wear, fatigue)

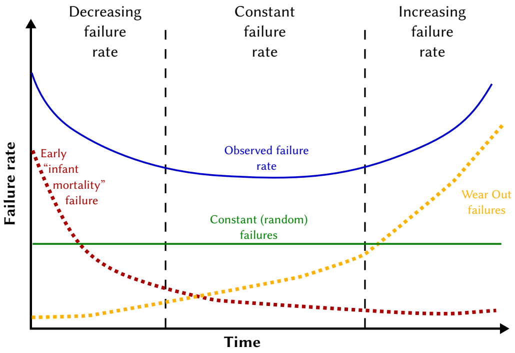

Reliability measures

- Mean time to failure (MTTF) = $\mathbb{E}(T) = \int_0^\infty R(t)dt$
- Often calculated by dividing the total operating time of the units tested by the total number of failures encountered
- Often modelled by a Weibull distribution (systems affected by wear) or an exponential distribution (systems not affected by wear, such as electronics)

Availability

Availability

The ability of an item (under combined aspects of its reliability, maintainability and maintenance support) to perform its required function at a stated instant of time or over a stated period of time [BS 4778]

- The availability $A(t)$ of an item at time t is the probability that the item is correctly working at time t
- Mean availability = $\dfrac{MTTF}{MTTF + MTTR}$

Reliability ≠ availability

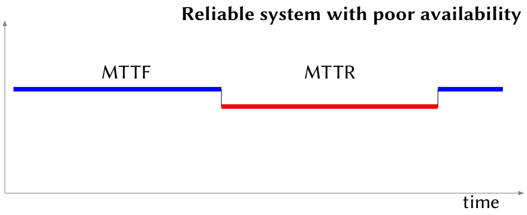

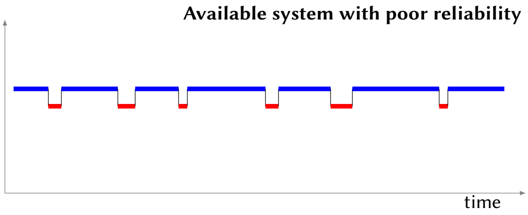

Note the important difference between:

- ▷ reliability (failure-free operation during an interval), measured by the MTTF
- ▷ availability (instantaneous failure-free operation on demand, independently of number of failure/repair cycles), measured by

$$A = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}}$$

Also note that reliability ≠ safety

Maintainability

Maintainability

The ability of an item, under stated conditions of use, to be retained in, or restored to, a state in which it can perform its required functions, when maintenance is performed under stated conditions and using prescribed procedures and resources [BS4778]

- Measured by MTTR: mean time to repair
- Commonly modelled by a lognormal distribution

Reliability measures

$$MTBF = MTTR + MTTF$$

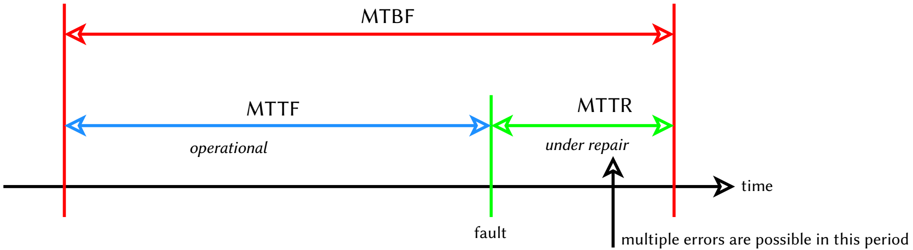

MTBF: mean time between failures

RISK ENGINEERING

Exercise

Problem

For a large computer installation, the maintenance logbook shows that over a period of a month there were 15 unscheduled maintenance actions or downtimes, and a total of 1200 minutes in emergency maintenance status. Based upon prior data on this equipment, the reliability engineer expects repair times to be exponentially distributed. A warranty contract between the computer company and the customer calls for a penalty payment for any downtime exceeding 100 minutes. Find the following:

1. The MTTR and repair rate
2. The probability that the warranty requirement is being met
3. The median time to repair
4. The time within which 95% of the maintenance actions can be completed

Exercise

Solution

1. MTTR = 1200/15 = 80 minutes and the repair rate μ is 1/80 = 0.0125. Our probability distribution for repair times is `dist = scipy.stats.expon(scale=80)`.

2. The probability of time to repair not exceeding 100 minutes is `dist.cdf(100)` = 71%.

3. The median time to repair is `dist.ppf(0.5)` = 55 minutes.

4. The time within which 95% of the maintenance actions can be completed is `dist.ppf(0.95)` = 240 minutes.

Exercise

Problem

From field data in an oil field, the time to failure of a pump, X, is known to be normally distributed. The mean and standard deviation of the time to failure are estimated from historical data as 3200 and 600 hours, respectively.

1. What is the probability that a pump will fail after it has worked for 2000 hours?

2. If two pumps work in parallel (the system can meet performance requirements with a single operating pump), what is probability that the system will fail after it has worked for 2000 hours? Assume that pump failures are independent events.

Exercise

Solution

1. Random variable X can be represented by the model `scipy.stats.norm(3200, 600)`.

   We want to assess $\Pr(X > 2000)$, which is $1 - \Pr(X \leq 2000)$, or `1 - scipy.stats.norm(3200, 600).cdf(2000)`, or 0.977.

2. Let's call Y the random variable representing time to failure of the redundant pump system, and X_1 and X_2 the time to failure of pumps 1 and 2 respectively. We want to determine $\Pr(Y > 2000)$, which is $1 - \Pr(Y \leq 2000)$.

   This is $1 - \Pr(X_1 \leq 2000 \wedge X_2 \leq 2000)$ (given the parallel configuration of the pumps, the system fails when both of the pumps fail).

   Given that pump failure is independent, that's $1 - \Pr(X_1 \leq 2000) \times \Pr(X_2 \leq 2000)$. It's `1 - scipy.stats.norm(3200, 600).cdf(2000)**2`, which is 0.9994.

Further reading

- Wired.com article *Why Things Fail: From Tires to Helicopter Blades, Everything Breaks Eventually* from 2010

For more free content on risk engineering, visit `risk-engineering.org`

Feedback welcome!

Was some of the content unclear? Which parts were most useful to you? Your comments to `feedback@risk-engineering.org` (email) or `@LearnRiskEng` (Twitter) will help us to improve these materials. Thanks!

`@LearnRiskEng`

`fb.me/RiskEngineering`

For more free content on risk engineering, visit `risk-engineering.org`