8 Availability

8.1 Introduction

In an earlier chapter we demonstrated that the availability of a protective system is not fixed by its design; in most cases it can be increased by testing the device to check that it is working, and fixing it if it is not. In general, more frequent checks lead to higher device availability, and less frequent checks deliver lower availability. This chapter develops the relationship between failure-finding task interval and the availability of a protective system.

⇒ *Remember that availability alone should not be used to set failure-finding intervals unless the target level has been derived from a robust, quantitative risk model.*

8.2 Availability and Failure Rate

The availability of a protective device is key to setting up a maintenance policy, but how do we calculate the availability of a real device?

First, what factors contribute to the a protective device such as a low oil pressure sensor being unable to do its job? Among others, we may need to consider the following.

- Failure of the component during its working life
- Installation of a non-functioning component
- Disablement of the switch during a planned test, after which the switch was not reconnected
- Failure of external services such as power, networks and data buses

We have already seen that checking the device at fixed intervals enables us to influence the level of availability achieved, although these checks will have no direct effect on the other causes of unavailability listed above.

Common sense tells us that the availability of a device is directly related to its failure rate. If System A uses a switch whose mean time between failures is 100 years, and System B uses a switch whose mean time between failures is 10 years, then *if they are subject to the same maintenance policy*, we should expect that the switch in System A would demonstrate a far higher availability than that in System B. The less reliable a device is, the lower its availability is expected to be.

In this section we will deal with devices that fail at random. Random failure means that the chance of a failure occurring does not depend on the previous history of the device: not on the age of the device, the time of year, the phases of the moon, the number of starts or stops, or anything else. The chance of a working device failing on any chosen day is exactly the same as that of it failing on any other day.

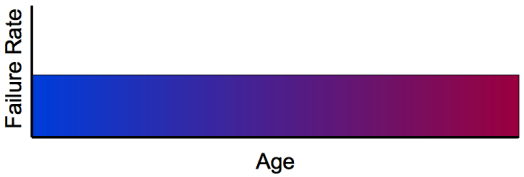

Age-independent random failure: the probability of failure is independent of age

To make the calculation more concrete, the following example determines the availability of a high temperature trip system which has a 10% chance of failure in any one year[1].

The calculation begins when the device is newly installed or has just been checked. In this section we make the assumption that the trip is fully functional at time zero.

If the trip is fully functional at time zero, what is the chance that it is still functional at the end of the first year? If the chance of failure is 10% in any year, the probability that it is still functional is

100%–10% =90%

What is the probability that the trip is still functional at the end of the second year, assuming that it is not checked or replaced during that period?

The chance that the device is still functional at the end of year 2 is given by:

(Chance that the device is working at the end of year 1)
x
(chance that the device does not fail during year 2)

In this case, the chance that the device is still functional after two years is

90% x 90% = 81%

Similarly, the chance that it is functional after three years is

(Probability working at end of year 2)
x
(chance of non-failure in year 3),

or

90% x 90% x 90% = 72.9%

The table below summarises the availability at the end of each of the first ten years.

End of year	Probability that trip functions
Start	100%
1	90%
2	81%
3	72.9%
4	65.6%
5	59.0%
6	53.1%
7	47.8%
8	43.0%
9	38.7%
10	34.9%

Note that the probability that the device is functional at the end of the third year is almost 73%, not 70%. Although there is a 10% chance of failure per year, this is a *conditional probability*: there is a 10% chance of failure of a device that is working at the start of the year. Since there is a 90% chance that the trip is working at the start of the year, the chance of failure during the second year is 90% x 10% = 9%.

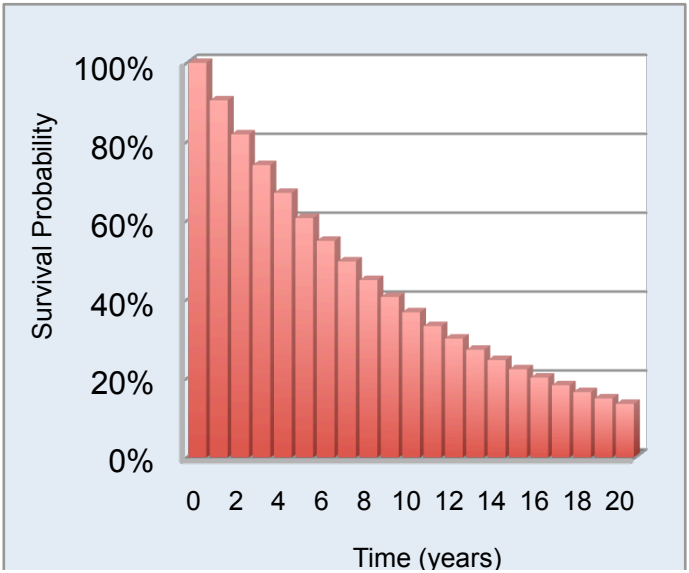

Survival probability to the start of each year for a trip device that has a random failure pattern with a mean time between failures of 10 years

Sometimes it can be helpful to look at this in a different way. Imagine that there are 100 trips at time zero. If every trip is replaced when it fails, there are always 100 trips operational, and about 10 fail every year. However, if trips are not replaced when they fail, then about 90 remain after year 1; therefore the number of failures during year 2 is lower than during the first year because there are fewer working trips which can fail. As the number of working trips diminishes over time, the number of failures decreases as well. The number of failures goes down although the rate of failure per trip stays the same.

Since failure of the trip system is hidden, its availability is given by the probability that the trip is functional at the time of a demand. The availability is 90% after one year; after two years, 81%; after three, 73% and so on. Therefore the graph above shows the relationship between failure rate of the protective system and its availability.

The graph shown above is an approximation. The true relationship between the failure rate of the trip and its availability is

$$A = e^{-t/M_{dev}}$$

Where

A	is the availability of the protective system
t	is the time since the device was installed or tested
M_{dev}	is the mean time between failures of the protective device
e	is the number 2.7182818, the base of natural logarithms

The exact survival curve is shown below.

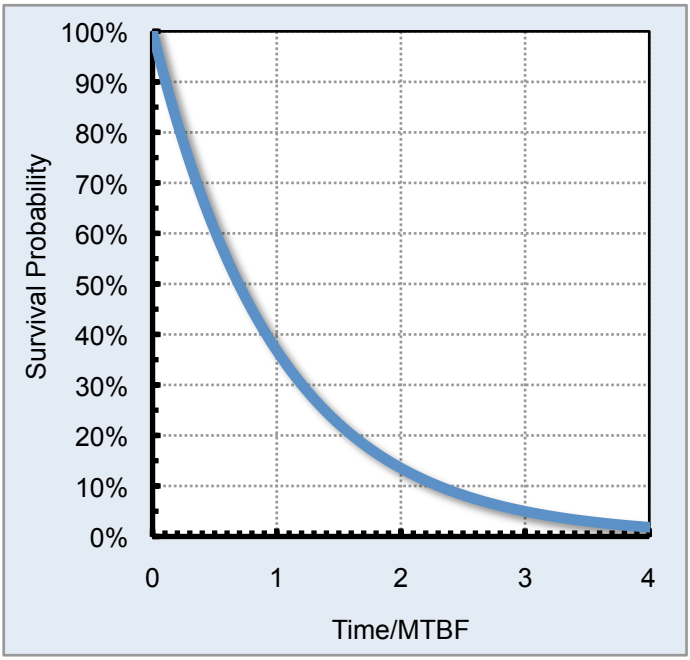

Curve showing the probability of survival to a given time divided by the device's mean time between failures

The survival curve shows the probability that a device functions after a specified time, expressed as a proportion of the device's mean time between failures. Its value is 1 (100%) initially and it decays towards zero, although in theory it never actually reaches it.

One feature of the survival curve is important to the derivation of most of the formulae used in this book. The first part of the curve, up to a time that is around 5% of the mean time between failures, is very nearly a straight line. At 5% of the mean time between failures, the difference between the straight line and the curve is just over 0.1%; at 10%, the difference is under 0.5%. Most of the formulae assume that the relationship between availability and time is a straight line. In order for the formulae to be valid, the following condition must apply.

The failure-finding interval must be less than about 5% of the mean time between failures of the protective device.

The equation of the first part of the survival curve is derived in section 25.3. For times up to about 5% of the mean time between failures of the protective device, its availability is given by the following formula.

$$A = 1 - t / M_{dev}$$

8.3 Minimum Availability Calculations

Availability is the simplest criterion used to set failure-finding task intervals. A target availability of the protective system is chosen, then the failure-finding interval is calculated to achieve at least that level of availability.

If we want to achieve a given minimum availability, we already have all the tools needed to calculate the failure-finding interval for a real device. We choose the availability required and rearrange the formula

$$A = 1 - t / M_{dev}$$

so that we can calculate the failure-finding interval from the device's mean time between failures and the required minimum availability.

$$T_{ff} = (1 - A)M_{dev}$$

Don't start celebrating just yet, though: this is not the calculation that is normally used. For reasons that will become more obvious later when we calculate the rate of multiple failures using availability and demand information, the target is usually the average rather then the minimum availability.

8.4 Availability-based Calculations: Average Availability

The formulae introduced in the previous section calculate the availability of a device at an instant in time. Availability starts at 100% and declines until we test the device, and if necessary repair it. The availability target used in this section is not the instantaneous availability, but the *average* availability of the device assessed over a period of time. The graph below shows how the device's availability changes over the course of time between failure-finding tasks. The dotted line marks the device's average availability.

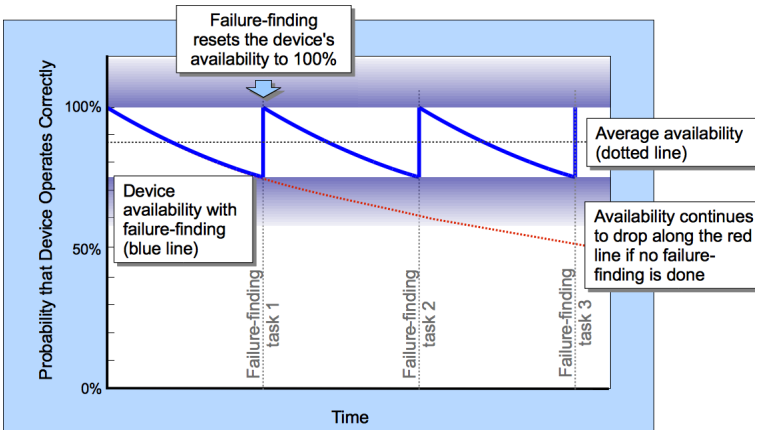

Availability of an ideal protective device with failure-finding (blue) and without failure-finding (red)

A mathematical derivation of the average availability is shown in section 25.4, but if we assume that the graph is a straight line, it is easy to derive the average value visually simply by looking at the graph of availability between two failure-finding tasks.

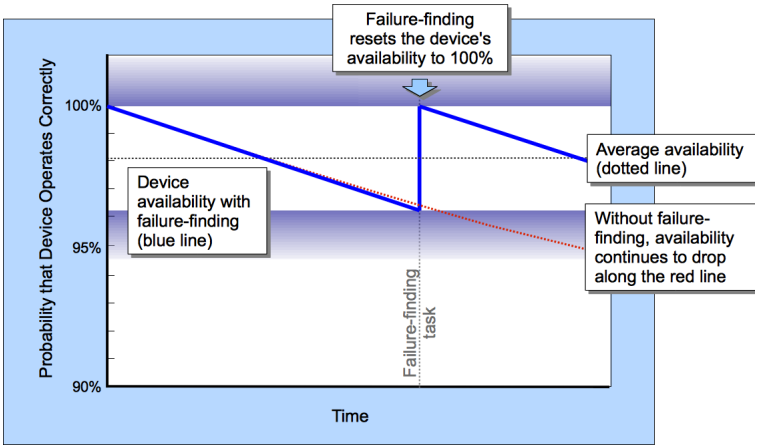

Availability is approximated by a straight line. Note that the left hand axis has been expanded.

The availability of the device immediately after failure-finding (and, if necessary, a repair) is assumed to be 100%. Provided that its availability stays above about 95%, its availability drops approximately according to the equation

$$A = 1 - t / M_{dev}$$

so that when it reaches the failure-finding task interval T_{ff} its availability is

$$A = 1 - T_{ff} / M_{dev}$$

The average availability over the period (shown by the dotted line) is

$$\overline{A} = 1 - T_{ff} / 2M_{dev}$$

Therefore if the target average availability of the protective device is A, then the failure-finding task interval needed to achieve it is

$$T_{ff} = 2M_{dev}(1 - \overline{A})$$

8.5 General conditions

The following conditions apply to the calculations in this chapter.

1 The calculation applies only to one failure mode of a single, simple protective device. It does not apply to multiple failure modes, or to protective systems that consist of several simple devices, such as a pair of pressure relief valves, redundant backup generators, or a 2-of-3 voting system.

2 Failures of the protective device must be random; there must be no relationship between the chance that a device has failed and its age or the time since it was last maintained.

3 The calculation assumes that the failure-finding interval is small compared with the device's mean time between failures. Typically the failure-finding interval should be less than about 5% of the mean time between failures of the protective device.

4 Unavailability of the protective device due to scheduled or unscheduled maintenance is not included in the calculations, and may need to be taken into account in calculating the overall unavailability of the device.

Availability is the simplest criterion that can be used to calculate failure-finding intervals; but this leads immediately to the question of how to determine the availability required. That issue is dealt with in the following chapter.

8.6 Examples

Fan vane switch

A fan is used to dilute boiler flue gases by mixing them with air before they are dispersed at low level. Local regulations state that the CO_2 content of the discharged gas must be below 1%.

If the fan fails for some reason, or if the ducting is blocked, a vane switch shuts down the boiler to prevent discharges with a CO_2 content above the allowed limit.

The mean time between failures of the vane switch is estimated to be 10 years. The required average availability of the shutdown switch is 99.7%.

How often should the vane switch be tested?

The table below summarises the relevant numbers.

Term	Description	Value
M_{dev}	Mean time between failures of the vane switch	10 years
$\bar{A}$	Required average availability of the vane switch	99.7%
T_{ff}	Vane switch failure-finding interval	To be calculated

The failure-finding interval needed to achieve 99.7% availability is

$$T_{ff} = 2M_{dev}(1 - \bar{A})$$

or

$$T_{ff} = 2 \times 10 \times (1 - 0.997) = 0.06 \quad years$$

The required testing interval is 0.06 years, or about every three weeks.

Oil Pressure Switch

A diesel generator's engine contains a low oil pressure warning switch. If oil pressure drops below 1.5 bar, a red light illuminates on the control panel and the operator is expected to take action to prevent damage to the engine.

The manufacturer's data show that the switch's mean time between failures is 200 years. The required average availability is 99%.

How often should the switch be tested?

The table below summarises the relevant numbers.

Term	Description	Value
M_{dev}	Mean time between failures of the low oil pressure switch	200 years
$\bar{A}$	Required average availability of the low oil pressure switch	99%
T_{ff}	Oil pressure switch failure-finding interval	To be calculated

The failure-finding interval needed to achieve 99% availability is

$$T_{ff} = 2M_{dev}(1 - \overline{A})$$

or

$$T_{ff} = 2 \times 200 \times (1 - 0.99) = 4 \quad years$$

So the low oil pressure switch should be checked every four years.

⇒ *This calculation does not include the other components of the alarm system that might fail. Later chapters deal with more complex systems that include more than one component.*

Gas detector

A compartment in an offshore production facility contains a combustible gas detector that should raise an alarm if the gas concentration rises above 10% of the lower explosive limit (LEL).

Records show that the mean time between failures of the detector is about 20000 hours. A quantitative risk assessment implies that the required availability is 99.99%.

How often should the detector be tested?

The table below summarises the relevant numbers.

Term	Description	Value
M_{dev}	Mean time between failures of the gas detector	20000 hours
$\overline{A}$	Required average availability of the gas detector	99.99%
T_{ff}	Gas detector failure-finding interval	To be calculated

The failure-finding interval needed to achieve 99.99% availability is

$$T_{ff} = 2 \times \frac{20000}{8760} \times (1 - 0.9999) = 0.00046 \quad years$$

The detector should be checked every four hours to achieve 99.99% availability. This testing interval is unlikely to be acceptable: this is an indication that the device is incapable of delivering the availability required, and the system should be redesigned in some way to provide better reliability.

8.7 Key Points and Review

Because failure of a protective device is hidden, we cannot be certain whether it will function correctly when a demand occurs.

The availability of a protective device is the probability that it will work at a specific time.

The average availability of a protective device depends on its reliability, as measured by its mean time between failures, and how frequently it is tested.

Under a number of assumptions, there is a direct relationship between the availability of a protective device, its mean time between failures and how frequently it is tested. Therefore it is possible to calculate how often a protective system needs to be tested to achieve the desired level of availability.

Although availability is the simplest criterion for determining failure-finding task intervals, its use is only justified if the availability chosen can be robustly defended.

---

1 It is a slightly odd fact that a failure rate of "10% per year" does not mean that exactly 90% of the devices are working at the end of the first year, or for that matter that 81% (0.9 x 0.9) are working at the end of year two. This is because the technical failure rate is applied continuously, not just at the end of the first year. The formula actually used is $R(t) = \exp(-\lambda t)$, where λ is the failure rate (0.1 or 10%) and t is measured in years. The actual proportion surviving to the end of the first year is 90.48%, dropping to 81.87% at the end of year 2. The difference is similar to that of a bank quoting 10% interest on your account; if interest is added once per year, you have $110 for every $100 in the account. If the bank adds interest every month, you get interest on the interest that has already been added, and at the end of the year you have $110.47 for every $100 invested. The more often interest is added, the more you get, until if the bank adds interest continuously, you have $110.52 for every $100 after one year.