9 Risk

9.1 Introduction

In the previous chapter, failure-finding intervals were set by determining the test frequency that results in the required average availability of the protective device .

It was emphasized that the availability requirement should be derived from a rigorous, robust model. There is no problem if the system has been subjected to a quantitative method such as Fault Tree Analysis (FTA), but for most industrial systems the required availability level simply does not exist. This chapter demonstrates how to derive the availability needed from two numbers:

- The mean time between demands on the protective system
- The minimum tolerable mean time between multiple failures

9.2 Getting to Availability

Although the ultimate objective of managing hidden failures is to reduce or eliminate the risk of multiple failures, the previous chapter's availability calculation does nothing to connect the failure-finding interval to the risk of a multiple failure. Since the whole point of maintaining the protective system is to reduce the chance of a multiple failure happening, it makes sense that the task interval should be based on how often we are willing to allow the multiple failure to occur.

The relationship between availability and multiple failure rate is very simple, but it involves one more parameter, as demonstrated by this example.

> A small 230 volt electrical installation is protected by a residual current detector (RCD) which is intended to cut off the power if current flows to earth, perhaps because of a fault or because someone has accidentally touched the live wire. The device works by comparing the current in the live wire with the return current, and tripping if the imbalance is more than a few milliamps.
>
> Records show that the RCD is tripped in normal use (not during testing) about once per year.

The multiple failure that the device protects against is that there is a fault and the power is not cut, leading to equipment damage, injury, or even death.

What is the relationship between the availability of the RCD and the rate of multiple failures?

We know that the availability of the RCD depends on its reliability and on how often it is tested. Since we already know how to work out its average availability, we are going to treat the availability as a variable and work out the multiple failure rate.

First suppose that the RCD is never tested. Assuming that it worked when it was installed, its availability decays away over time and eventually should be close to zero. Ignoring (just for convenience) the very early part of its life, the RCD will always be in a failed state. A drawing of a typical history might look like this; demands on the RCD are assumed to occur at random.

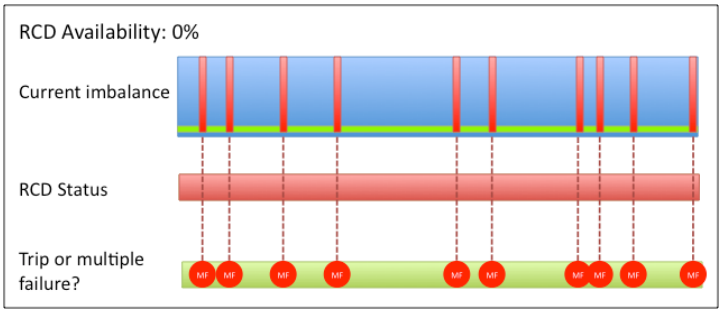

How often does the multiple failure occur? The simple answer is: every time that there is a demand on the RCD, because the device availability is zero. If a demand occurs on average once a year, then the multiple failure also occurs once a year.

Now suppose that the RCD availability is improved a little, so that on average it works 10% of the time. How could the same history look?

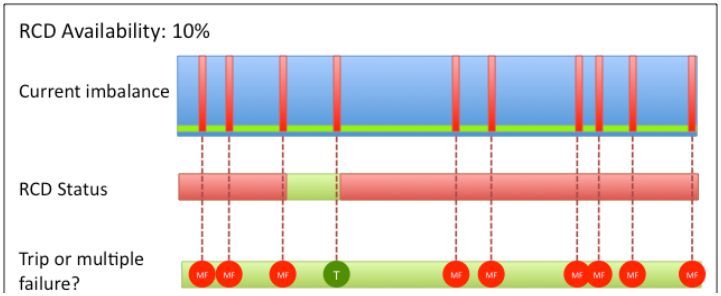

Because failures of the RCD and demands occur at random, it is possible that there could be 10 multiple failures, 9 or any other number. However, on average one in ten demands would lead to a trip, and nine out of ten demands would end in a multiple failure.

If we now increase the RCD's availability to 90%, perhaps by implementing some form of regular failure-finding task, then only one out of 10 demands (on average) would result in a multiple failure.

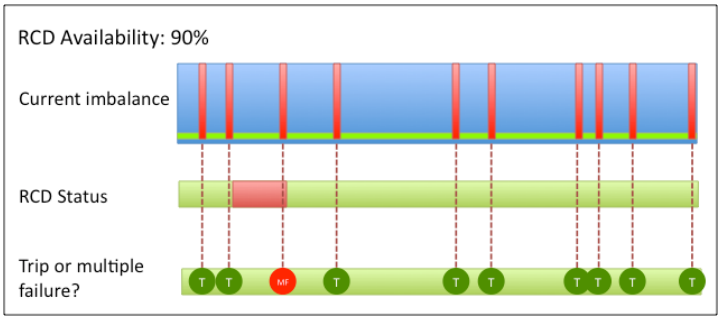

There is a simple relationship between the mean time between demands on the system, the average protective device availability, and the mean time between multiple failures:

$$M_{mf} = \frac{M_{dem}}{(1 - \overline{A})}$$

If the protective device availability is 100%, the mean time between multiple failures is infinite; in other words, the multiple failure never happens.

9.3 Risk-based Calculations

Now that we have a link between demand rate, the mean time between multiple failures and average device availability, we are in a position to work out the mean time between multiple failures that would be achieved if we checked a device at a specific failure-finding interval.

In the previous chapter we found that the average availability of a simple, single protective device that fails at random is

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

From the equation above, the mean time between multiple failures is

$$M_{mf} = \frac{M_{dem}}{(1 - \overline{A})}$$

Putting these two formulae together, the multiple failure rate for a given demand rate and failure-finding interval is

$$M_{mf} = \frac{2M_{dem}M_{dev}}{T_{ff}}$$

In chapter 5 we found that the objective of failure-finding for multiple failures that have safety or environmental consequences is to reduce the rate of multiple failures to a tolerable level. By rearranging the formula, the failure-finding interval needed to achieve a given mean time between multiple failures is

$$T_{ff} = \frac{2M_{dem}M_{dev}}{M_{mf}}$$

To summarise, the failure-finding interval for a risk-based system is determined by the following factors.

Symbol	Description
M_{dev}	Mean time between failures of each protective device
M_{dem}	Mean time between demands on the protective system
M_{mf}	The lowest tolerable mean time between multiple failures

Remember that use of this formula is subject to a number of caveats including the following.

- There is one protective device
- The protective device fails at random
- The protective device is guaranteed to be working immediately after installation
- The failure-finding task is always effective: the task discovers 100% of non-working devices and any devices that have to be repaired are fully functional immediately after the task has been carried out
- Demands on the protective device occur at random
- The failure-finding interval is less than about 5% of the protective device's mean time between failures
- The failure-finding interval is much less than the mean time between demands on the device
- The time taken to repair the protective device is insignificant, or other measures will be taken to prevent multiple failures if a scheduled task discovers a failed protective device

Most of these assumptions will be revisited and relaxed in later chapters, but bear in mind that you need to check your system, data and failure-finding interval carefully to make certain that these conditions are not broken.

9.4 Demand Rate

The failure-finding calculation is now based on the minimum tolerable mean time between multiple failures, but at the expense of adding a new item to data: the mean time between demands on the protective system.

The demand rate is how often the protective system has to operate because of abnormal conditions. If a pressure relief system on a boiler that has been installed for twenty years has been called on to relieve overpressure four times in that period (and if the demands occur at random), then the mean time between demands is five years.

Remember these points when working out the mean time between demands.

- Count any occasion on which the device has had to operate because of genuinely abnormal conditions that could have caused a multiple failure. The key question is, "How many times (or how often) has this device operated because of abnormal conditions." The question is *not*, "How many times has it failed to prevent a multiple failure?"
- Do not count deliberate demands on the system due to routine tests and maintenance
- The calculations assume that demands occur at random. If they are non-random (perhaps they tend to occur just after or just before major maintenance), then the formulae may not give correct answers.

You will probably need to talk to maintainers and operators to find the data that you need. To maximise the chances of obtaining the correct information, it is a good idea to phrase the question in terms that relate to the system under analysis. So, rather than asking

"How many times has a demand occurred on the fire alarm system?"

ask

"How many times has there been a fire in this building?"

If the protective system has been operational over a long period of time, try to be aware of any changes in operating context, such as increases in production rates or the introduction of different technologies, which might influence the demand rate. The demand rate used to set failure-finding intervals should be the anticipated future rate, which may be different from earlier experience.

Some system demand rates are easy to estimate. Domestic and industrial residual current detectors (RCDs), which protect users of mains power supplies, trip sufficiently frequently in normal use that most organisations can estimate the demand rate accurately. Most experienced drivers have at some time encountered a situation where their anti-lock brake system operated, and so they would be able to estimate how often they make a demand on the system.

Even where demands on one device are infrequent, there may be enough of them operating at any time to enable a realistic demand rate to be calculated: although an individual office building's fire alarm may detect be presented with a real fire only every few decades or so, there are plenty of aggregated statistics for different industry sectors, cities, regions and whole countries.

Even so, some demands are infrequent and just about unique to a specific organisation or process. The issue of finding demand rates in a variety of circumstances is discussed in more detail in chapter 12.

9.5 Multiple Failure Rate

The objective of the failure-finding task is to reduce the rate of multiple failures to a tolerable level, or equivalently to deliver a minimum tolerable mean time between multiple failures.

The failure-finding interval depends directly on this number, so it is vitally important that the rate of multiple failures is acceptable for all those who are likely to be affected by the hazard, including duty holders, senior management, company staff and members of the general public. Bear in mind that a single failure mode may represent only a small part of the organisation's risk, and that all significant failure modes need to be included in a complete risk management plan.

There is further discussion of tolerable failure rates, who should be involved in setting risk targets, and methods for determining individual multiple failure rates in chapters 6 and 12.

9.6 Examples

Oil Pipeline Low Pressure

A small lubricant pipeline runs close to an environmentally-sensitive area. A pressure switch is intended to shut down the oil pump if a significant leak occurs.

The mean time between failures of the pressure switch in this application is about two million hours. The low pressure switch has never been activated except during tests, but estimates suggest that it could be called on to operate about once every 10 years.

The minimum tolerated mean time between multiple failures (an undetected pipeline leak) is 10 000 years.

The required failure-finding interval is

$$T_{ff} = \frac{2M_{dem}M_{dev}}{M_{mf}}$$

or

$$T_{ff} = 2 \times 10 \times \frac{2000000}{8760} \times \frac{1}{10000} \quad years = 0.46 \quad years$$

The failure-finding task would have to be carried out every six months.

Standby Generator

A remote medical facility is subject to infrequent power outages that can last for at least several hours. When mains power is not available it relies on a single diesel generator which starts automatically when mains power is lost.

The group reviewing the standby power maintenance policy has decided to treat the generator, its engine, and the cut-in system as a single entity. The overall mean time between failures of similar systems at other installations is about 2 years. The mean time between demands on the system is about one year.

Although higher reliability would be desirable, the review group decided reluctantly that the chance of the generator being unable to produce power when required should be less than 1 in 1000 years.

The required failure-finding interval is

$$T_{ff} = \frac{2 \times 1 \times 2}{1000} \quad years = 1.5 \quad days$$

The failure-finding task would have to be carried out every day.

This interval is might not be acceptable for several reasons:

- It would probably be impractical to carry out the task at this interval
- Performing the task so frequently would place significant stress on the engine and so would contribute to wear and could result in lower reliability
- It indicates a gap between the desired reliability and what the equipment is capable of delivering

The most likely outcome is that the system would be redesigned to make it more reliable, perhaps by providing a second standby generator.

9.7 Time to Repair

The calculation method used in this chapter does not take into account device unavailability that arises during repair of the protective system. In most circumstances this is a reasonable assumption because maintainers and operators take care to reduce or eliminate the risk of a multiple failure during device repair. In most cases the affected system would be shut down, but sometimes the system operators could use alternative protection or closer system monitoring.

Time to repair should be included in the protective system downtime if no additional precautions are taken during the repair period. The modified formulae are derived later chapters.

9.8 Key Points and Review

Availability is not a useful criterion for determining failure-finding intervals unless it is supported by a robust model.

The required device availability can be calculated from two numbers:

- The demand rate on the protective device
- The minimum tolerated mean time between multiple failures

These lead to a simple relationship between the device reliability, demand rate and tolerated mean time between multiple failures.