A New Approach to SIL Verification

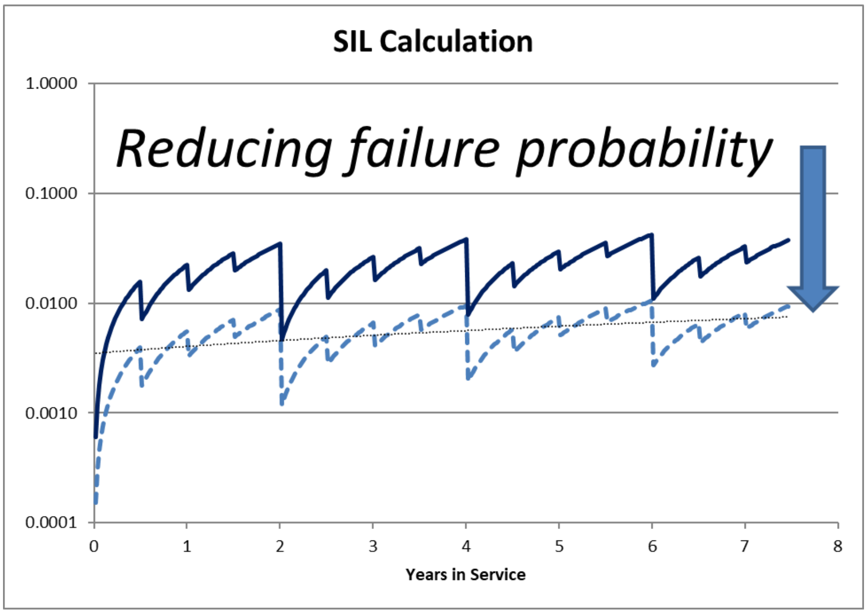

SYSTEMS

I&E Systems Pty Ltd

ACN 069 813 958

Mirek Generowicz

FS Senior Expert (TÜV Rheinland #183/12)

I&E Systems Pty Ltd gratefully acknowledges contributions to this paper from Sam Holden, Harvey Dearden and Bob Weiss.

Contents

Summary

Performance targets

Automated safety functions are designed to achieve performance targets for reliability. The targets are classified in orders of magnitude by safety integrity levels (SIL). Performance is measured either by failure rate or else by probability of failure on demand.

The functional safety standards IEC 61508, IEC 61511 (ANSI/ISA 84.00.01) and IEC 62061 all require calculations to demonstrate that the design of each safety function will meet the performance targets. The objective of the calculations is to verify whether the design of the automated safety system is adequate.

These calculations are often described as 'SIL verification', though that may be a misleading term because verification involves much more than these calculations. The objective of the calculations is to verify whether the design of the automated safety system is adequate to achieve the required failure performance.

This paper provides guidance on the objectives and content of SIL verification. It proposes that SIL verification should always be part of a broader review of SIF compliance. A new model outline is proposed for a SIF compliance report, incorporating SIL verification.

Feasibility of targets

Rules of thumb are given for quick assessment of the risk reduction performance that can be achieved by any safety function. The rules of thumb reveal that performance depends heavily on failure rates. In practice the failure rates vary by more than an order of magnitude between different users. Failure rates are not fixed and constant because most failures are systematic (i.e. preventable) in nature.

Factors affecting performance

Calculations are meaningless if systematic failures are not actively and effectively controlled.

Failure performance depends largely on the effectiveness of the design and maintenance in preventing systematic failures. Two main factors limit the failure performance that can be achieved in practice:

- The suitability of the design for the application, for the environment and for the SIL
- Adequate accessibility and resources to enable effective inspection, testing, maintenance and renewal.

Planning for regular inspection and proof testing must be considered early in the conceptual design - while P&IDs are being developed. Standard architectural design patterns need to be developed for each SIL.

To achieve SIL 3 performance degraded components need to be overhauled or replaced before they fail. If annual planned downtime is insufficient to allow renewal of final elements then a duty/standby design pattern may be required.

Failure measures

The type of failure measure used to characterise a safety function depends on whether it acts continuously or on demand.

Continuous safety functions act continuously to ***maintain*** a safe state. Failure of a continuous function causes a hazardous situation. The failure is treated as an initiating event. The appropriate failure measure is then the failure rate, represented by the Greek letter lambda, λ.

Demand safety functions act to ***put*** a system into a safe state in response to a hazardous situation that has started to develop. They prevent escalation of a hazardous scenario. The appropriate failure measure is the average probability of failure on demand (PFD_{AVG}).

In both cases, the failure measure of the safety function has contributions from three subsystems: the sensors, the logic solver and the final elements:

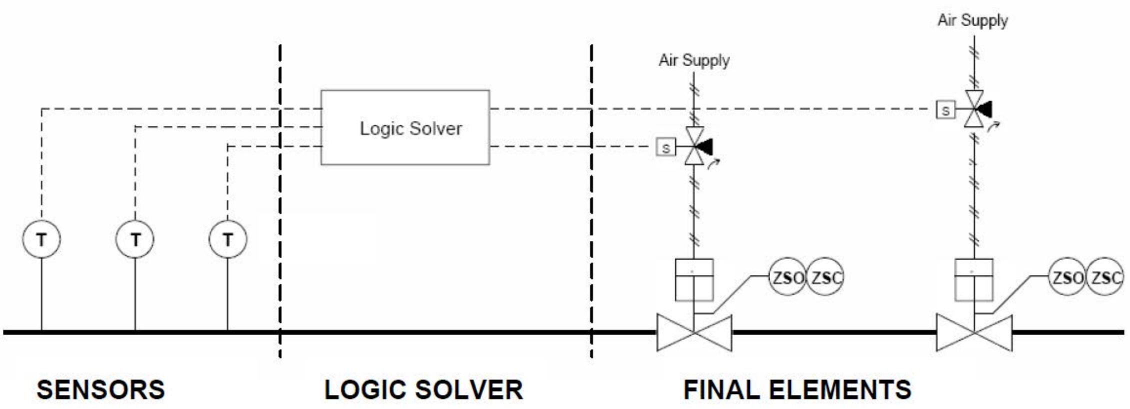

We are concerned primarily with dangerous failures, those failures that prevent successful action.

Safe failures are not considered in detail in the is guideline, though safe failures also need to be analysed because they may lead to spurious trips. Spurious trips can indirectly lead to hazardous situations.

For a continuous function the overall dangerous failure rate is the sum of the dangerous failure rates of the three component sub-systems:

$$\lambda_D^{SIF} = \lambda_D^{S} + \lambda_D^{LS} + \lambda_D^{FE}$$

With continuous functions failures are generally immediately revealed. The calculation of failure rate for continuous functions is relatively simple and is not discussed in further detail in this paper.

With demand functions dangerous failures may not be immediately revealed. Undetected dangerous failures prevent successful performance of functions in response to a demand. The SIL of a demand function is usually characterised by the order of magnitude of the probability of failure on demand. Functions that have a demand rate higher than once year are treated in the same way as continuous functions and are characterised by failure rate.

For a demand function the overall average probability of failure is approximately equal to the sum of the probabilities of failure of the three component sub-systems:

$$PFD_{AVG}^{SIF} \approx PFD_{AVG}^{S} + PFD_{AVG}^{LS} + PFD_{AVG}^{FE}$$

The calculation of probability is based on assumptions about the rate at which hidden failures can accumulate.

A more convenient way of expressing failure measures is in terms of risk reduction factor RRF and the mean time between failures, $MTBF$.

RRF numbers are in the range 10 to 10,000 and these are more convenient to work with than the corresponding PFD_{AVG} numbers which are in the range 0.0001 to 0.1.

SIL 1	RRF range 10 to 100	PFD range 10^{-2} to 10^{-1}
SIL 2	RRF range 100 to 1,000	PFD range 10^{-3} to 10^{-2}
SIL 3	RRF range 1,000 to 10,000	PFD range 10^{-4} to 10^{-3}

Similarly $MTBF$ is usually in the range 10 to 1,000 years. It is easier to work with large whole numbers rather than with λ measured in failures per year or per hour.

Rules of thumb

It is easy to estimate the overall probability of failure for a demand safety function by using some simple rules of thumb. Though these estimates are coarse they are useful for quickly identifying whether a safety function is likely to meet its failure performance target.

The rules of thumb are not presented here as an alternative to detailed calculations. They are presented only to give a quick feel for what performance is achievable.

Final elements (usually) dominate

In process sector applications with clean service the overall failure rate λ and the PFD_{AVG} is usually strongly dominated by the contribution from actuated on-off valves or electrical contactors as the final elements. This is because it is usually not practicable to implement automatic and continuous fault diagnostic functions on mechanical and electromechanical final elements. Without diagnostics dangerous faults remain undetected until the function is inspected and tested or until it fails to respond to a demand.

Sensors and logic solvers tend to have much lower dangerous failure rates than final elements because they are based on electronic devices. Electronic components can usually be equipped with automatic and continuous fault diagnostic functions. Most failures can be detected and remedial action can be taken to ensure that safe operation is achieved or maintained.

As a result, the rates of undetected dangerous failures for electronic components are generally at least an order of magnitude lower than for mechanical and electromechanical components.

Where the final elements are tested on an annual basis the final elements can be expected to contribute between 70% and 90% of the failure measure. The sensors typically contribute up to 25% of the failures and the logic solver less than 5% of the failures.

From this we can derive simple rules of thumb. An explanation of these rules of thumb is included in the Appendix below.

The overall failure rate may be dominated by the sensor failure rates in severe services where reliable sensing is difficult. For example, in minerals processing the sensors can have high failure rates and failures can difficult to detect. Similar rules of thumb could then be applied based on the sensor failure rate instead of the final element failure rate.

Single FE (1oo1)

For a first approximation with a single final element we can use:

$$RRF^{SIF} \approx 1.5 \times \frac{MTBF_{DU}^{FE}}{T}$$

$MTBF_{DU}$, the mean time between **d**angerous **u**ndetected failures of the final element and T is the test interval. Both need to be in the same units of measure and are usually measured in years.

For instance, for an actuated on/off shutdown valve a typical feasible value of $MTBF_{DU}^{FE}$ is in the range 40 y to 100 y. (Refer to the section below on '*Finding failure rates*'.)

With a single shutdown valve as a final element and with annual testing (T = 1 y) it is feasible for a safety function to achieve RRF in the range 60 to 150.

The bare minimum $MTBF_{DU}^{FE}$ of a single final element required to achieve a given RRF can be estimated as:

$$MTBF_{DU}^{FE} > 0.6 \times \ RRF^{SIF} \ \times T$$

To achieve SIL 2 (i.e. RRF> 100) with T = 1 y requires $MTBF_{DU}^{FE}$ of at least 60 y.

This means that the failure performance of the valve will need to be optimised; a $MTBF_{DU}^{FE}$ of 40 y would not be sufficient to achieve SIL 2 with annual testing.

Electrical contactors can typically achieve $MTBF_{DU}^{FE}$ in the region of 200 y, so SIL 2 is easily achieved.

Dual FE (1oo2)

For a first approximation with two similar (1oo2) final elements we can assume $\beta \approx$ 10% and use:

$$RRF^{SIF} \approx 15 \times \frac{MTBF_{DU}^{FE}}{T}$$

With 1oo2 shutdown valve as final elements and annual testing (T = 1 y) and $MTBF_{DU}^{FE}$ in the range 40 y to 100 y it is feasible for a safety function to achieve RRF in the range 600 to 1,500.

The bare minimum $MTBF_{DU}^{FE}$ of dual final elements required to achieve a RRF in a 1oo2 arrangement can be estimated as:

$$MTBF_{DU}^{FE} > \frac{0.6}{\beta} \times \ RRF^{SIF} \ \times T$$

and with β around 10%:

$$MTBF_{DU}^{FE} > {}^{6}/_{100} \times \ RRF^{SIF} \times \ T$$

To achieve SIL 3 (i.e. RRF> 1,000) with two valves and T = 1 y requires a $MTBF_{DU}^{FE}$ of at least 60 y.

With lower $MTBF_{DU}^{FE}$ either the failure performance of the valves will need to be improved, or β will need to be reduced below 10%, or else the inspection and test interval will need to be shortened to less than a year.

Diverse FE (1oo2)

Common cause failures can be minimised by using independent and diverse final elements. For instance, it may be possible to either stop a pump or close a valve to achieve a safe state.

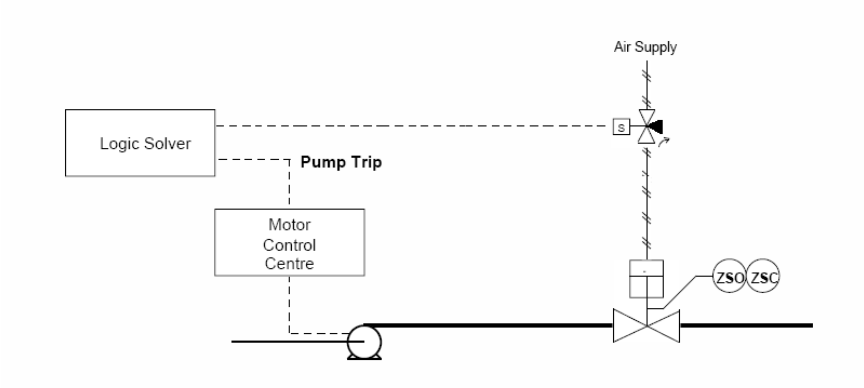

The failure modes and the failure rates for electrical contactors are completely different to failure modes and the failure rates of valves.

The question is sometimes asked: ‘which value of failure rate λ_{DU} should be used in the common cause failure term $\beta.\frac{\lambda_{DU}.T}{2}$, the failure rate of the valve, or of the contactor?’.

A much better question to ask is: ‘what types of failure could possibly affect the valve and the contactor to cause them both to fail dangerously at about the same time?’ and ‘how can we eliminate those common cause failures?’

With diverse elements it is better to identify exactly what the causes of common cause failures might be and what is the expected rate at which those failures will occur. For instance, with example shown above the common cause failures are limited to factors simultaneously affecting multiple output circuits of the logic solver.

It may be possible to eliminate almost all of the common causes of failure by design, so the probability of failure of the final elements would be approximately:

$$PFD_{AVG}^{FE} \approx \frac{(\lambda_{DU}^{valve}.T).(\lambda_{DU}^{contactor}.T)}{3}$$

Safety functions with completely independent and diverse final elements easily achieve SIL 3 levels of risk reduction.

Finding failure rates

There are several useful references for failure rate information:

- OREDA ‘Offshore and Onshore Reliability Handbook’
- *exida* database incorporated into exSILentia software,
  and tabulated in the *exida* SERH ‘Safety Equipment Reliability Handbook’
- SINTEF PDS Data Handbook
  ‘Reliability Data for Safety Instrumented Systems’
- Users’ own failure rate data
- Equipment certificates

The OREDA project provides a useful source of failure rate information gathered over many years by a consortium of oil and gas companies. The OREDA handbooks summarise all failures recorded over the normal useful operating life of equipment (i.e. excluding end-of-life failure). Different editions of the handbook cover different periods of time. The tables can be difficult to interpret but they are useful because they indicate the ranges of failure rates that are achieved in operation.

The failure rate tables published by OREDA show that failure rates recorded by different users typically vary over one or two orders of magnitude. OREDA fits the different reported failure rates into probability distributions to estimate the overall mean failure rate and standard deviation for each type of equipment and type of failure.

It is evident from the OREDA tables that the failure rates are not constant across different users and different applications. Some users consistently achieve failure rates at least 10 times lower than other users. The implication here is that it may be feasible for other users to minimise their failure rates through best practice in design, operation and maintenance.

The SINTEF PDS Data Handbook provides a concise summary of failure rates that are typically achievable. It includes OREDA data as an input source.

The *exida* failure rate database is based on failure mode analysis. It provides estimated failure rates and failure modes for specific makes and models of many commonly used devices. The failure rates are calculated from typical failure rates of the individual components that make up each device. The *exida* failure rates are calibrated with field failure data. The database is regularly updated with current failure measurements from users.

The *exida* failure rates are reasonably consistent with the OREDA data and can be taken as a good indication of failure rates that are achievable in practice. It presents failure rates that typically at least 70% of users are achieving in practice.

The *exida* dataset has some advantages over OREDA: It allows comparison of different designs and different makes and models of devices. It is easier to interpret than OREDA because it presents failure rate data in a more consistent form.

One concern with *exida* failure rates is that they are presented with 3 or 4 significant figures of precision, implying that the rates are fixed and constant. In practice the uncertainty in these rates is as wide as in the OREDA data. The variation in the rates that can be achieved in operation spans at least an order of magnitude.

Failure rates on SIL certificates need to be treated with caution if the quoted failure rates are significantly lower than failure rates in the industry-wide databases. Certifying bodies may specifically exclude systematic failures when evaluating failure rate. Such low failure rates cannot be easily achieved in practice.

Variability in failure rates

One reason for the variability in reported failure rates is that the industry-wide datasets include all failures, systematic failures as well as random failures.

Only a small proportion of failures are purely random. Purely random failures are characterised by a fixed and constant rate that cannot be changed.

Systematic failures are preventable, and the rate of failure depends on the effectiveness of the efforts made to prevent failures. Systematic failure rates vary widely.

Tossing a coin is a good example of a purely random event. The rate at which tosses result in heads is fixed. If a coin is tossed 100 times resulting in 53 heads we can be reasonably certain that the rate of heads is 0.5 heads per toss. The probability of a head on the next toss can be inferred to be approximately 0.5. There is nothing we can do to change the probability.

A football match is a good example of a quasi-random event. If a football team plays 100 matches and wins 53 times we can state that the team's average success rate is 0.53 wins per match. But the probability of the team winning their next match can be anywhere between 0 and 1, depending on many factors.

The probability is influenced by the age, health, condition, skill and motivation of the players compared with those on the opposing team. It also depends on environmental conditions. It may be possible to estimate the team's probability of success and to change that probability if the influencing factors are known, understood and deliberately manipulated.

Purely random failure

A purely random failure of a component or a device is a sudden and complete failure. It occurs without any warning and is virtually impossible to forecast by examining the item. It is the contrary of failures occurring progressively and incompletely. (Ref: Technical Report ISO 12489).

The definition of random failure used in both IEC 61508 and IEC 61511 is much broader:

> *'failure, occurring at a random time, which results from one or more of the possible degradation mechanisms in the hardware*
>
> *Note 1 to entry: There are many degradation mechanisms occurring at different rates in different components and since manufacturing tolerances cause components to fail due to these mechanisms after different times in operation, failures of a total equipment comprising many components occur at predictable rates but at unpredictable (i.e., random) times.'*

Failures resulting from degradation are never purely random, they are partially systematic and can be forecast. The failure rates are measurable and may be predictable but the rates depend on equipment condition. Degradation can be detected and failures can be prevented to some extent.

Purely random failures cannot be prevented. They can be characterised by a fixed and constant rate that depends on the causes and mechanisms of the failure.

In functional safety only electronic components are subject to purely random failure. A good example is the failure of semiconductor elements due to damage from cosmic rays. The rate of failure depends on the cosmic flux and energy levels. The failure rate can be reduced by shielding but the failures can never be prevented entirely. The rate of failures is reasonably constant.

Quasi-random (preventable) failure

By contrast failure of an electronic component due to overheating is only partially random. To some extent failure can be anticipated by measuring the temperature of the device or by inspecting it. Failures can be prevented through design or by controlling the environment.

Examples of quasi-random failures are those caused by:

- Inappropriate specification for the service or environment
- Inappropriate maintenance practices
- Misuse
- Corrosion
- Contamination
- Dust
- Oil
- Vibration
- Temperature
- Humidity

- Power quality
- Electromagnetic interference
- Quality of pneumatic and hydraulic fluid supplies
- Cracking
- Fatigue
- Physical impact damage
- Vermin
- Radiation (heat or UV)
- Aging components
- Worn components
- Deterioration.

The failure rate of a component type may appear to be reasonably constant but the rate depends on the design, the age and on the operation of the components.

Quasi-random failures can be anticipated through inspection, measurement or testing. Failures can be prevented by specifying components to suit the conditions of service and by replacing or renewing components that no longer meet the specification.

Quasi-random failure rates can be minimised by designing the system to enable full access for inspection, measurement, testing, and repair or renewal.

Systematic failure

Purely systematic failures result from pre-existing faults. The failure mechanisms are deterministic and predictable. Purely systematic failures can be eliminated by finding and correcting the faults.

Software coding errors are a good example of purely systematic faults.

The rate at which systematic failures occur indicates the maturity of the quality management. The rate cannot be used to predict future performance with precision. The probability of further systematic faults remaining may be coarsely estimated, but once systematic faults are corrected they should not recur.

Many failures are partially systematic and partially random in nature. If a worn component is not replaced its eventual failure is systematic in nature though the timing of the failure is partially random.

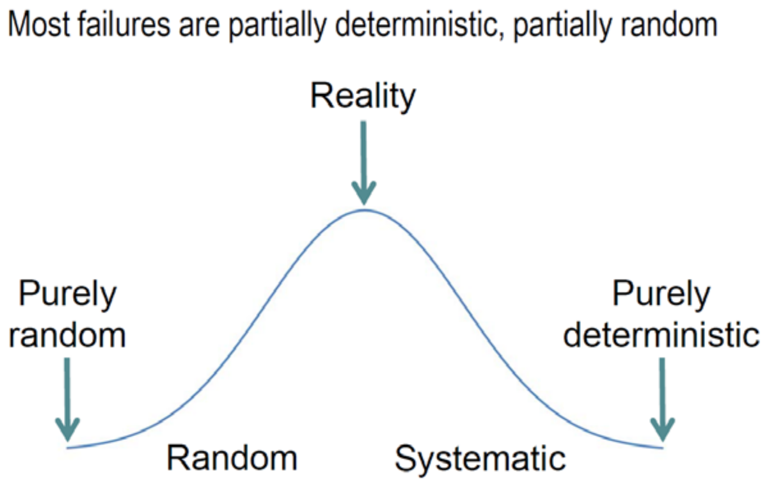

Applicability of probabilistic modelling

Purely random failures can be modelled accurately if the failure rates can be measured with precision. Calculation tools based on Markov models or Petri nets are useful for calculating probability of the purely random failure of electronic components such as logic solver systems and sensor electronics.

These tools can also be applied to estimate probability of quasi-random failure of sensors and final elements but the uncertainty in the result depends on the uncertainty in the failure rates.

The uncertainty in the failure rates of mechanical, electromechanical and pneumatic components is typically more than one order of magnitude. For these types of components sophisticated modelling software cannot provide more accurate results than simplified equations.

Benefits of calculation software

There are benefits in using modelling software even though the results are not as precise as they might seem.

Popular calculation tools such as *exSILentia* include the benefit of inbuilt databases. Failure rates are included for many of the devices commonly used in safety instrumented systems. Comparison with OREDA data shows that the failure rates are realistic and achievable in practice.

The recent V4 release of *exSILentia* now takes into account the wide variability in failure rates. It applies a numeric rating system ('SSI') to gauge the effectiveness of systematic safety integrity management and maintenance competency (http://www.exida.com/SSI). The calculations adjust failure rates with a factor of up to 4 x depending on the SSI score. The SSI is also used to modify the assumed effectiveness of proof testing. Users can easily see the relative importance of the SSI score.

Similarly, the *SISSuite SIL Check* tool allows the user to nominate 'deployment' contributions that modify failure rates. These can include maintenance effectiveness and other factors.

Limitations in precision

Software models allow users to try various techniques to improve failure performance, but the perception of precision in calculation software can lead to disproportionate efforts being made.

For instance, with some effort the RRF of a SIF might be increased from 815 to 1,040 to claim SIL 3. The improvement might be achieved by introducing partial stroke testing or by improving proof test coverage, or by minimising β. It is important to understand the impact of these factors, but it is also important to understand their relative significance. The difference between 815 and 1,040 is a factor of about 1.3. On a logarithmic scale it is about 1/10 of an order of magnitude ($10^{0.1}$).

By contrast the assumptions made in the hazard and risk assessments are usually to within half an order of magnitude at best ($10^{0.5}$, i.e. a factor of 3).

The frequency of initiating events cannot be predicted with precision better than half an order of magnitude. It is pointless to estimate RRF achieved with more precision than the RRF target.

Influencing failure rate

When we understand that most failures are either systematic or quasi-random in nature we can see why failure rates vary so widely and how we can deliberately change failure rates.

Failure rates depend strongly on design and maintenance. Some factors that affect failure rates include:

- Suitability of devices for the process conditions
- Suitability of devices for the ambient environmental conditions
- Suitability of the design for the application
- Quality control in manufacture and installation
- Effectiveness of verification (review, inspection and testing)
- Effectiveness of equipment condition monitoring
- Effectiveness of root cause analysis of failures
- Effectiveness of failure performance measurement
- Accessibility to components for inspection, testing, maintenance and renewal
- Adequate resources to maintain equipment effectively.

Fault exclusion

Some faults can be excluded by design (refer to IEC 61508-2 §7.4.4.1). For instance, an electrical contactor can be designed to minimise the probability of contact welding. The contact current rating could be designed to be several times greater than the load current and the circuit protection can be designed to limit fault let-through current to below the contact rating.

Suitability of devices

Before calculating failure measures we need to establish that the selected devices are suitable for use in safety related service. According to IEC 61511-1 §11.5.2 the suitability of devices can be established by either one of two methods:

- Compliance with IEC 61508-2 and IEC 61508-3, or
- Compliance with requirements for selection of devices based on prior use.

Either way, a significant volume of operational experience with that specific type of device is necessary.

The suitability of devices depends on their systematic safety integrity, which is to do with preventing, avoiding and controlling dangerous systematic failures. This includes ensuring that the specifications of the devices are appropriate for the intended application. It includes ensuring that the quality management is effective in rectifying non-conformance.

Systematic integrity depends on complete and appropriate specifications that are traceable back to the design basis. The design basis needs to cover all environmental conditions as well as process conditions.

Systematic integrity also depends on effective quality control in manufacture, installation and maintenance. Quality control needs to include review, checking, inspection and testing (i.e. verification) against objective acceptance criteria traceable back to the specifications. Complete verification records need to be kept as evidence of quality.

Documentary evidence of suitability is mandatory for all safety systems and devices.

Devices claimed to comply with IEC 61508-2 and IEC 61508-3 must be supplied with this documentation in the form of a safety manual. The requirements for safety manuals are defined in IEC 61508-2 Annex D and IEC 61508-3 Annex D.

Safety manuals include functional specifications for devices. Safety manuals also describe the anticipated failure modes and expected failure rates when the devices are applied in the intended functions. The manuals need to define the constraints on the use of the devices and the assumptions on which the failure behaviour and failure rates are based.

IEC 61511 also requires safety manuals but the requirements are more loosely defined. Documentation equivalent to IEC 61508 compliant safety manuals would be acceptable. Substantial documentary evidence is required by IEC 61511 if devices are selected based on prior use.

Limiting values for MTBF

Measured or calculated failure rates should not be assumed to be precise and fixed constant values. They are simply an indication of the failure rates that are achievable in practice.

If the $MTBF_{DU}$ of an actuated valve assembly is estimated as 40 years the actual $MTBF_{DU}$ achieved could be anywhere between 10 years and 100 years or more.

The upper limit on $MTBF_{DU}$ is set by the proportion of failures that cannot be prevented effectively. This includes:

- Hidden failures that cannot be forecast by examining the item and
- Incipient failures that are not rectified due to a lack of access or a lack of resources.

Failure performance targets

The values of $MTBF_{DU}$ and $MTTR$ (mean time to restoration) that are assumed in failure calculations effectively set performance targets for operations and maintenance.

The operations and maintenance team must be satisfied that these performance targets are feasible given the design and the constraints on accessibility and resources.

Preventive maintenance

Failure of a component due to age, wear or deterioration is an indication that other similar components are likely to fail in a similar manner. Detection of serious deterioration of any device suggests that similar devices of similar age should be inspected promptly.

Corrective and preventive actions should be based on root cause analysis.

The evidence of prior use required by IEC 61511-1 §11.5.3.2 includes evidence of failure performance and evidence of rectification of unsatisfactory performance.

Response to increased failure rates

Experience in operation may reveal failure rates that are worse than the target. A natural response might be to calculate a reduced test interval to meet the failure target. A more appropriate response would be to find the root cause and rectify the problems.

Failure rates that are higher than industry norms may be due to:

- Inadequate maintenance and renewal (worn equipment not repaired or replaced)
- Equipment applied outside its specified performance envelope (e.g. severe corrosion, extreme temperature, extreme vibration)

Maximising RRF

The rules of thumb for estimating RRF reveal that with annual testing:

- Achieving SIL 2 may be difficult with a single valve as a final element
- Achieving SIL 3 may be difficult with two similar valves as the final elements
- Achieving SIL 3 is difficult with extended test intervals for final elements (T >> 1 y)

In conventional SIL verification, strategies for improving RRF have typically been limited to reducing the test frequency or introducing periodic partial testing in between full tests. RRF might also be improved by introducing redundancy or by claiming a lower value of β.

It is now clear that the most important factor in improving RRF is the effectiveness of equipment maintenance (i.e. the systematic safety integrity in maintenance and operation – for instance, as characterised by the *exida* SSI).

The benefit of redundancy in having 1oo2 similar valves is questionable if neither valve can be taken out of service for testing, repair or renewal because of production constraints.

Deterioration must be detected and corrected before failure occurs. If the deterioration is not corrected the rate of failures will increase beyond the target.

A more appropriate strategy is to design the safety function so that each sub-system can be readily accessed for inspection, testing and maintenance. The RRF can be maximised by maximising the $MTBF_{DU}$ as far as is practicable.

Applications such as LNG production may require high availability as well as high reliability. Opportunities for major maintenance may be restricted to intervals as long as 5 or more years.

It may be necessary to configure the valves in a duty/standby arrangement to enable more frequent maintenance:

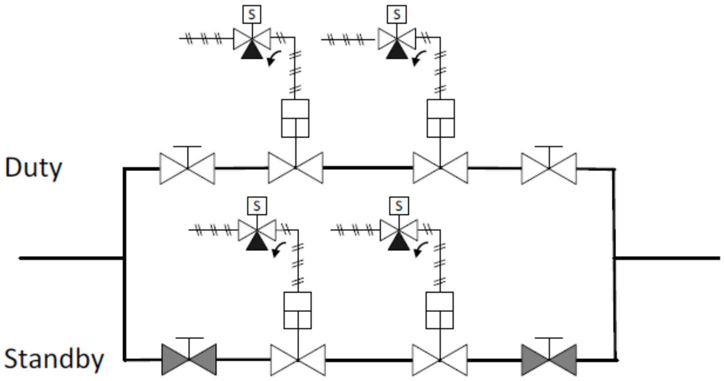

Testability

Proof test coverage

Failure probability calculations must take into account the effectiveness of proof testing and inspection. The calculated PFD_{AVG} includes a contribution from the failures than can never be detected.

Conventional calculation methods assume that '**n**ever **d**etected' failures accumulate at a constant rate λ_{ND} over the mission time T_M (time to planned renewal or replacement). The contribution of these failures to PFD_{AVG} can be estimated as:

$$\frac{\lambda_{ND} . T_M}{2}$$

The fraction of undetected dangerous failures that are never detected can be characterised by a proof test coverage factor, PTC:

$$PTC = 1 - \frac{\lambda_{ND}}{\lambda_{DU}}$$

This approach treats the never detected failures as if they were a random contribution to probability of failure. It may be a useful way of estimating the impact but it is not strictly valid because incomplete testing is a systematic issue rather than a random failure. Limited test coverage is a systematic fault in the design. There is no reason why we should assume that never detected failures would occur at a constant rate.

A good example of a never detected failure is that a level sensor will never trip at a high level because it has been designed for the wrong fluid density, and has been designed so that it can never be fully tested under operational conditions.

FMEDA

Safety functions need to be designed so that they can be fully tested. If complete testing is not possible then it will not be possible to demonstrate that the safety functions meet requirements (i.e. validation will not be possible). It will also not be possible to show that the functions continue to meet requirements. IEC 61511-1 §7.2.1 clearly specifies that testability of the design must be verified. Estimating PTC is one way of evaluating testability.

PTC can be estimated using FMEDA: an analysis of the failure modes, failure effects, failure rates and diagnostic coverage for each component in a sub-system.

One of the benefits in estimating PTC is that the failure modes can be understood and steps can be taken to ensure that all failures can be found, either by diagnostics, proof tests or by inspection.

If the FMEDA reveals that some failures might never be detected then the design of the installation may need to be revised to improve testability.

Incomplete testing and inspection may be acceptable for a SIL 1 function if the overall residual risk is acceptable. The expected frequency of the failures that cannot be found needs to be estimated to show that the risk is acceptable.

SIL 3 safety functions should always be designed so that as far as is practicable they can be fully tested and inspected.

There is no reason to assume that full coverage cannot be achieved. For instance, if the SRS specifies that a valve must be leak tight then in-situ leak testing will be required. This might be achieved by blocking in the valve and measuring the time taken for pressure to rise in a blocked section. Such a test is only possible if the installation is specifically designed to facilitate testing. There is not much point in specifying performance criteria that can never be tested.

As equipment ages the failure rates will obviously increase if deterioration is not detected and remedied. Deterioration will not be remedied if it is never detected by inspection, measurement or testing. The planning for maintenance needs to include planning for overhaul or replacement before the equipment reaches the end of its service life.

What is not so obvious is that many safety functions are designed so that they can never be fully tested, not even during commissioning and validation. A good example is a high level trip in a process separator vessel. Some types of level sensor can only be fully tested by filling the vessel right up to the trip point under normal process operating conditions. The risk associated with such a test might not be acceptable to the operators.

Testability needs to be a prime consideration in the selection and design of safety function sensors.

Common cause failure factor

When voted architectures are applied the failure calculations should always include an estimate of common cause failure factors. The β factor should not simply be assumed to be 10%.

The 2015 SINTEF Report A26922 'Common Cause Failures in Safety Instrumented Systems' suggests that in practice the common cause failure fraction can be expected to be greater than 10%. Typical values achieved are in the range 12% to 15%.

Apply a method such as the one described in IEC 61508-6 Annex D to estimate the value of β factor.

The benefit in carrying out the estimate is in understanding what might lead to common cause failure. The estimate will reveal some obvious strategies for reducing common cause failure.

The quasi-random failures listed above in the section *'Quasi-random (preventable) failure'* are all examples of common cause failures. These failures can all be minimised by design.

One aim of the failure calculations should be to verify that the designers have made an appropriate effort to reduce common cause failures.

Hardware fault tolerance

The hardware fault tolerance requirements in IEC 61511 Edition 2 and IEC 61508 Route 2_H are relatively straightforward. Fault tolerance is required for SIL 3 and for SIL 2 in continuous or high demand mode. Fault tolerance is not required for SIL 1 or for low demand mode SIL 2.

IEC 61508 Route 1_H is useful for assessing complex hybrid architectures where the level of hardware fault tolerance achieved is not immediately obvious. Another reason for using Route 1_H is that it allows SIL 3 without fault tolerance if the safe failure fraction > 90%.

All of the hardware fault tolerance methods require evidence that the assumed failure rates are achievable in practice.

Closely related safety functions

The failure quantification calculations must clearly identify safety functions that share elements and it must clearly identify the relationships between safety functions and hazardous scenarios.

If two safety functions each respond to distinctly different hazardous events with independent causal events then the functions are effectively independent. The risk reduction achieved by the safety system for one hazardous event is independent of the risk reduction for other events.

If two or more safety functions all respond to the same hazardous event then the calculation of overall risk reduction achieved must consider the related safety functions together as a whole.

Take the example of a large gas-fired heater or furnace. The most obvious scenario requiring risk reduction relates to the hazardous consequences of re-ignition of unburnt fuel after a flame-out (i.e. flame failure and loss of combustion).

The following functions are closely related and not independent because they all relate to the same scenario of re-ignition of unburnt fuel following flame failure. They rely on the same final elements:

- Flame failure, tripping the master fuel valves
- Low air flow, tripping the master fuel valves
- Low gas pressure, tripping the master fuel valves
- Fan failure, tripping the master fuel valves.

We cannot take risk reduction credit separately for these safety functions as if they were completely independent. They may be other separate safety functions responding to high tube temperature or high exhaust stack temperature. Those would be completely unrelated and would represent unrelated demands on the master fuel valves.

The risk reduction required for the four flame failure safety functions depends on the consequence of re-ignition of unburnt fuel and the expected frequency of flame failure from all possible causes.

In the calculation of failure probability the four safety functions need to be treated as a single system sharing one common final element subsystem, one shared logic solver subsystem and having four separate sensor subsystems (voted in a '1oo4' arrangement).

The contribution to the overall probability of failure on demand of each sensor subsystem needs to be factored by the proportion of causal events to which that sensor subsystem will respond.

The overall PFD of the system will be dominated by the PFD of the final elements, as they need to function correctly for all flame failure events no matter which initiating event has caused the hazardous scenario.

The final elements may also need to provide risk reduction for scenarios involving failure of tubing carrying the heat transfer fluid. Those scenarios may be completely unrelated to flame failure. The risk reduction required for those scenarios can then ignore the demand on the final elements due to flame failure.

The guiding principles are simple:

- The RRF target of a safety function depends on the consequence and likelihood of the hazardous scenario and on the probability of failure of any other risk controls. The RRF **target** does not depend in any way on the design of the safety functions.
- The risk reduction achieved by a set of safety functions depends only on the design and current condition of the safety functions. The probability of failure of any safety function component on any given day is the same no matter which hazardous scenario has occurred. The RRF **achieved** does not depend in any way on the RRF targets.

SIF Compliance report

Failure quantification calculations are not sufficient on their own because the actual failure rates of SIF components depend heavily on the suitability and maintainability of the components.

The calculations are based on the assumption that preventable failures are effectively prevented. In practice at least 95% of failures are preventable, i.e. systematic. The calculations are meaningless if the design of the safety functions does not achieve sufficient systematic integrity for the SIL.

Rather than producing a 'SIL verification' report consider recording the failure calculations as part of a broader 'SIF compliance' or 'safety analysis' report. The objective of a SIF compliance report is to demonstrate that:

- The safety function design is suitable for the application, for the environment and for the SIL
- The design allows sufficient access to enable inspection, testing, maintenance and renewal
- The planning for regular inspection and proof testing is sufficient to achieve the failure performance targets
- The failure performance targets for each safety function are feasible, given the design of the safety function.

This type of report contributes to the overall 'safety argument', demonstrating due diligence in complying with appropriate standards and established work practices.

A good example of a compliance report is the 'Safety Analysis Report' outlined in the Norsk Olje & Gass Guideline 070.

It may seem to be onerous to produce such a comprehensive report, but the work needs to be done regardless. Evidence of compliance is required in several different contexts:

- System safety manual (IEC 61511-1 §11.2.13 or IEC 61508-2 §7.4.9.3)
- Verification (IEC 61511-1 §7 or IEC 61508-1 §7.18)
- Quantification of failure measures (IEC 61511-1 §11.9 or IEC 61508-2 §7.4.5)
- Audit (IEC 61511-1 §5.2.6.2 or IEC 61508-1 §6.2.7)
- Assessment (IEC 61511-1 §5.2.6.1 or IEC 61508-1 §8)

Producing compliance evidence progressively through the design reduces the workload for audit and assessment. It facilitates compilation of safety manuals at closeout.

Compliance versus verification

'SIL verification' can be a misleading term. Some people confuse SIL verification with verification in general. SIL verification is intended specifically to address the requirements for quantification of failure measures in IEC 61511-1 §11.9 and IEC 61508-2 §7.4.5. It is only one minor aspect of the verification required by IEC 61511-1 §7 and IEC 61508-1 §7.18.

Verification is the process of checking, analysing, reviewing, inspecting or testing that **outputs** are correct and consistent with respect to their corresponding **inputs**.

It needs to be clearly understood that verification is required for ***all*** outputs. This includes specifications, reports, schedules, drawings, data, hardware components, application program code, test plans, test reports, analysis reports.

A SIF compliance report is a precursor to validation. It can be an input to the overall validation process. Validation is the process of showing that the installed and commissioned safety system satisfies all of the safety requirements.

Compliance review stages

Compliance review and failure quantification activities are appropriate at three stages:

- Concept
- Requirements
- Design closeout

Concept

In most applications there should be no surprises about which functions need to be SIL 2 or SIL 3.

In hydrocarbon and hazardous chemical processes we can expect at least SIL 2 for functions that protect against loss of containment. This includes protection against overpressure and against liquid carryover. It includes inventory and feed isolation valves for process units, reactors and burners.

At the concept stage the question needs to be asked: **Can we avoid SIL 3 by applying safety-by-design principles**? Can other (simpler) additional risk reduction layers be applied? SIL 3 may be difficult to achieve and maintain unless the safety function is readily accessible for inspection, testing and maintenance.

Early in the conceptual design develop standard architectural design patterns for SIL 1, SIL 2 and SIL 3 functions. Specify the standard architectural design patterns in a conceptual design specification for the safety system.

Use simple rules of thumb to estimate the RRF that can be achieved by each standard design.

The standard design patterns should be established before the P&IDs are developed, and well before HAZOP studies.

Estimate the proof test coverage (PTC) that can be achieved for the standard design patterns.

Ensure that the design of SIL 2 and SIL 3 functions enables sensors and final elements to be fully tested and readily renewed or replaced when necessary.

Provide redundant sensors and final elements for SIL 2 service if access for repair and renewal is limited. Always provide redundant sensors and final elements for SIL 3 service.

Ensure that sufficient planned downtime and access is available to allow degraded components to be overhauled or replaced.

Consider applying duty/standby architecture for critical applications where accessibility is restricted or where process downtime is limited.

Establishing the feasibility of RRF targets at the conceptual design stage saves work. It avoids having to change the design at a later stage.

Requirements

Prepare the first draft of the SIF compliance report in conjunction with the safety requirements specification.

Estimate the RRF for each chosen architectural design pattern including detailed analysis of each sub-system and considering common cause failures.

Verify that the specifications cover all of the issues commonly leading to common cause failure (refer to the section '*Quasi-random (preventable) failure*' above).

Confirm that the specified SIL targets are feasible with the architecture chosen for each function, including the requirements for hardware fault tolerance.

Confirm that the makes and models of devices nominated on the preferred equipment list are suitable for service at the specified SIL. Confirm that evidence of suitability is available (systematic capability or evidence of prior use).

Confirm that the conceptual design allows ready access to SIL 2 and SIL 3 devices for inspection, full testing and maintenance. Confirm that the PTC targets can be achieved.

Estimate the failure rates that are feasible for the nominated devices. Establish the level of confidence in the estimated failure rates.

Provide a detailed rationale to justify any claimed reduction in failure rates (i.e. increased $MTBF_{DU}$ targets).

Design closeout

Prepare the final version of the SIL compliance report at the completion of the design.

Estimate the RRF for each safety function with the selected devices, including detailed analysis of each sub-system and considering common cause failures.

Confirm that performance targets for each safety function are feasible given the level of access for maintenance and testing as designed.

Verify that safety manuals have been produced and are complete.

Confirm evidence is available to justify claims for systematic capability or prior use.

Three stages of SIF compliance review recommended in the IEC 61511 safety lifecycle:

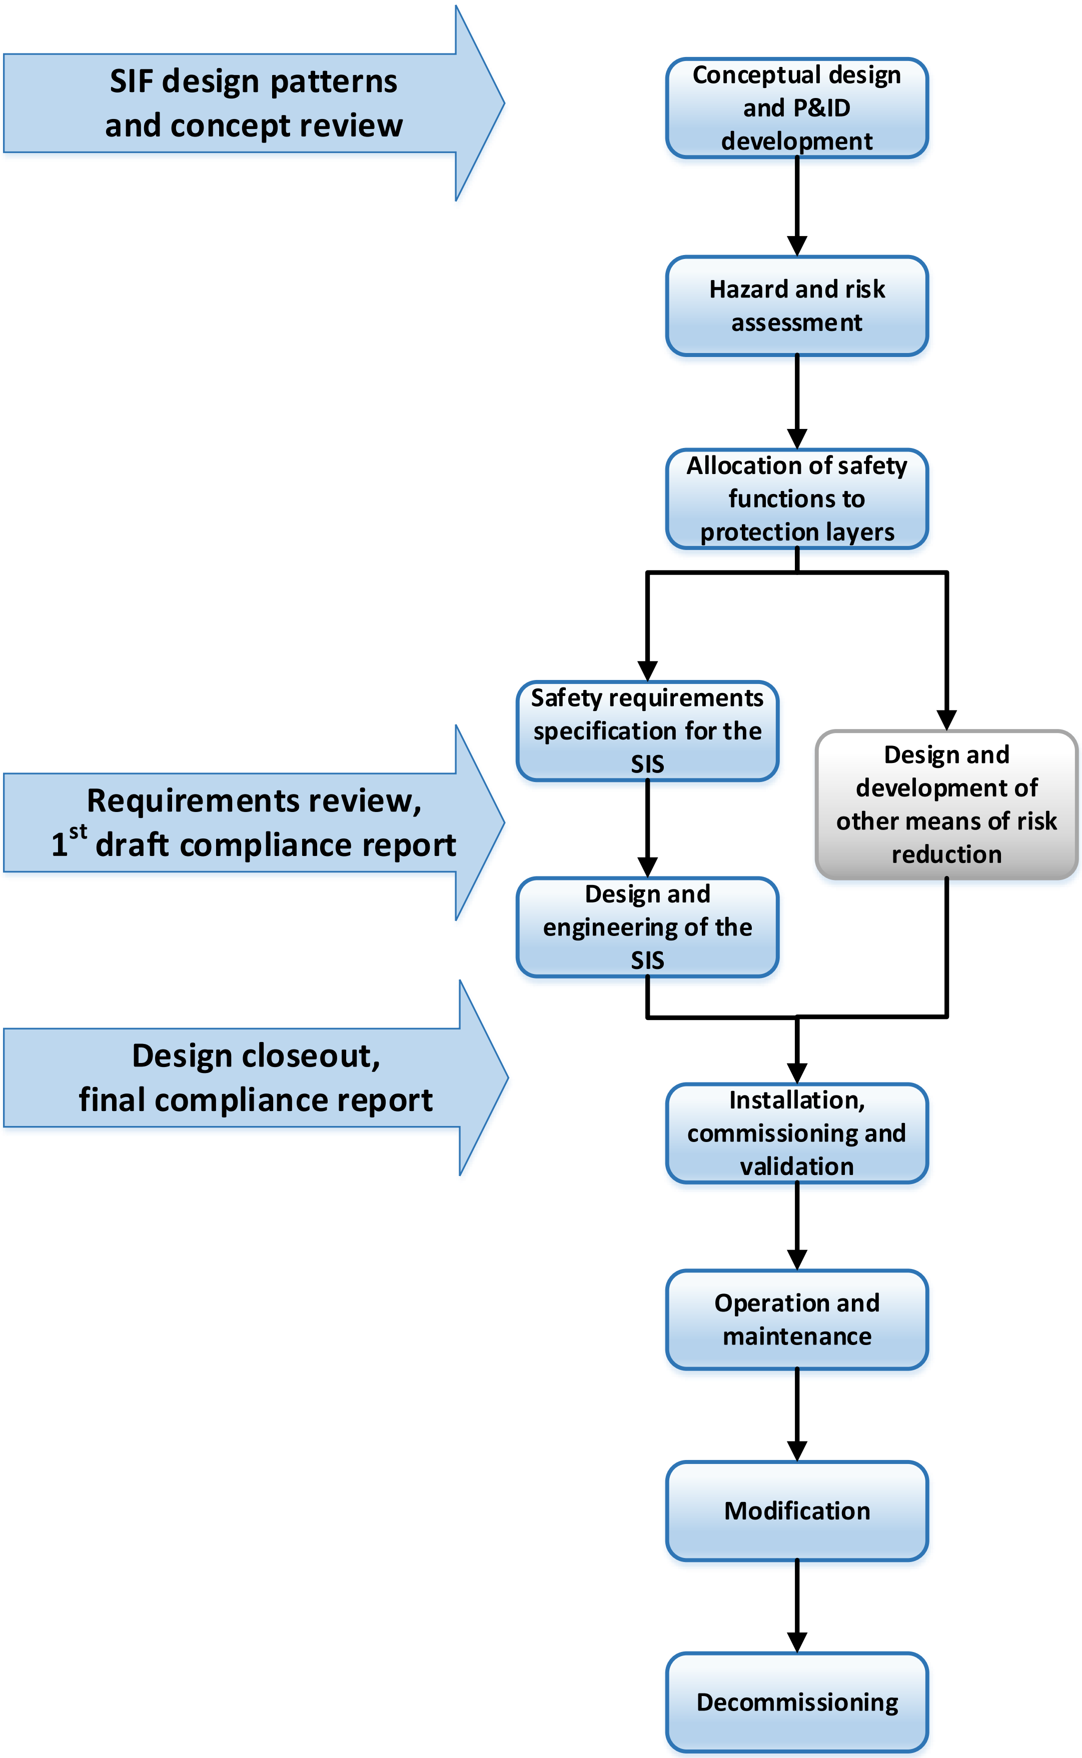

SIF design patterns and concept review

Conceptual design and P&ID development

Hazard and risk assessment

Allocation of safety functions to protection layers

Safety requirements specification for the SIS

Design and development of other means of risk reduction

Requirements review, 1st draft compliance report

Design and engineering of the SIS

Design closeout, final compliance report

Installation, commissioning and validation

Operation and maintenance

Modification

Decommissioning

Outline of a SIF compliance report

1. References
    a. Hazard and risk assessment
    b. Safety requirements specifications
    c. SIF detailed design specifications
    d. Failure rate data sources
2. Tabulation of hazardous scenarios with consequence, causes and frequencies
3. Summary tabulation of safety functions, including:
    a. Unique identifier
    b. Architectural design patterns
    c. SIL and RRF targets
    d. RRF achieved
    e. Systematic capability achieved or prior use claimed
    f. Spurious trip rate
    g. Traceability to hazard and risk analysis
    h. Grouping of safety functions responding to the same hazardous scenario
    i. Grouping of safety functions sharing final elements
4. Suitability of devices
    a. Tabulation of device makes and models
    b. $MTBF_{DU}$ target values and data sources for each device
    c. Confidence level in achievability of $MTBF_{DU}$
    d. Justification for any increases in $MTBF_{DU}$ targets for high SIL service
    e. Evidence of systematic capability or prior use
    f. Verification of safety manuals
    g. Verification of specifications and datasheets for completeness and traceability to safety requirements (including environmental requirements)
5. Common cause failure analysis
    a. Estimation of β factors, including justification of assumptions made
    b. Assessment of environmental design basis
    c. Assessment of dependence on power, air and hydraulic supplies
    d. Assessment of dependence on components shared between safety functions
    e. Assessment of dependence on components shared with other protection layers
    f. Assessment of dependence on interfaces
    g. Verification of freedom from interference in interfaces
6. Verification of diagnostic coverage and proof test coverage
    a. Verification of FMEDA studies
    b. Verification of procedures for response to diagnostic alarms
    c. Verification of planning for periodic inspection and proof testing
    d. Verification of accessibility for maintenance, inspection and testing
7. Assessment of feasibility of restoration time targets
8. Verification of architectural design patterns with regard to SIL capability and hardware fault tolerance
9. Verification of the detailed design of each safety function, including:
    a. Description of the hazardous scenario
    b. Description of the functional requirements
    c. Definition of the safety function components with make and model
    d. Assumptions for λ_{DU} (or $MTBF_{DU}$), β, T, λ_{ND} or PTC, λ_{DD} and $MTTR$
    e. Calculation of RRF
    f. Calculation of spurious trip rate

Appendix: Derivation of Rules of Thumb

Single FE (1oo1)

Calculations of failure probability for demand functions are based on the coarse assumption that **d**angerous **u**ndetected failures occur at a constant rate, λ_{DU}.

In any set of similar devices undetected failures would then accumulate exponentially. Once any individual device fails it remains failed until its failure is revealed by either inspection, a test or a demand on the function.

The probability of failure of any individual device in a given set of devices is proportional to the number of failures that have been allowed to accumulate in the time t since the last full test and inspection:

$$PFD(t) = \int_0^t \lambda_{\mathrm{DU}} \, . \, e^{-\lambda_{DU} \cdot \tau} . \, d\tau = 1 - e^{-\lambda_{DU} \cdot t}$$

The exponential curve is almost linear for $PFD < 0.2$. Safety functions always have $PFD_{AVG} < 0.1$, so it is always valid to use a linear approximation for the curve. There is no benefit at all in making a more precise calculation because the initial assumption of constant failure rate is coarse. Any additional precision would be meaningless because in practice the failure rate is a variable rather than a constant.

The average probability of failure PFD_{AVG} of a single final element can then be approximated by:

$$PFD_{AVG}^{FE} \approx \frac{\lambda_{DU} . T}{2}$$

where T is the test interval.

Failures that are detected make a smaller contribution to the probability of failure because they can be repaired promptly or the system can be taken out of service. The contribution from the rate of detected dangerous failures λ_{DD} is negligible.

A more convenient way of expressing the approximation is in terms of risk reduction factor RRF and $MTBF_{DU}$, the mean time between **d**angerous **u**ndetected failures of the final element.

$$RRF^{FE} \approx 2 \times \frac{MTBF_{DU}^{FE}}{T}$$

RRF is the reciprocal of PFD_{AVG} and $MTBF_{DU}$ is the reciprocal of λ_{DU}.

The overall RRF of the safety function with a single final element is typically in the range:

$$RRF^{SIF} \approx 0.7 \times RRF^{FE} \;\text{ to }\; 0.9 \times RRF^{FE}$$

$$RRF^{SIF} \approx 1.4 \times \frac{MTBF_{DU}^{FE}}{T} \;\text{ to }\; 1.8 \times \frac{MTBF_{DU}^{FE}}{T}$$

For a first approximation use:

$$RRF^{SIF} \approx 1.5 \times \frac{MTBF_{DU}^{FE}}{T}$$

For instance, a typical feasible value of $MTBF_{DU}^{FE}$ for an actuated on/off shutdown valve is in the range 40 y to 100 y.

With a single shutdown valve as a final element and with annual testing (T = 1 y) it is feasible for a safety function to achieve RRF in the range 60 to 150.

The bare minimum $MTBF_{DU}^{FE}$ of a single final element required to achieve a given RRF can be estimated as:

$$MTBF_{DU}^{FE} > 0.6 \times \; RRF^{SIF} \; \times T$$

Dual FE (1oo2)

The PFD_{AVG} of a pair of similar final elements in a 1oo2 arrangement can be approximated by:

$$PFD_{AVG}^{FE} \;\; \approx \; (1-\beta).\frac{(\lambda_{DU}.T)^2}{3} + \beta.\frac{\lambda_{DU}.T}{2}$$

The last term in this equation applies the β factor to represent the proportion of failures that have a common cause. The common cause term will dominate the PFD_{AVG} unless the following is true:

$$\beta \ll \lambda_{DU}.T$$

If the test interval T = 1 y and λ_{DU} = 0.02 pa ($MTBF_{DU}^{FE}$ = 50 y), then the common cause failure term will be greater than the first term unless β < 2%.

A value as low as 2% is difficult to achieve with similar final elements. A recent SINTEF study (Report A26922) showed that β is typically in the range 12% to 15%. With β ≥ 5% we can make the further approximation:

$$PFD_{AVG}^{FE} \;\; \approx 1.1 \; \times \; \beta.\frac{\lambda_{DU}.T}{2}$$

and so

$$RRF^{FE} \approx \; 1.8 \; \times \frac{MTBF_{DU}^{FE}}{\beta.T}$$

For a rule of thumb the common cause factor β can be assumed to be around 10%.

$$RRF^{FE} \approx \; 18 \; \times \frac{MTBF_{DU}^{FE}}{T}$$

The overall RRF of the safety function with 1oo2 final elements is typically in the range:

$$RRF^{SIF} \approx 0.7 \; \times \; RRF^{FE} \quad \text{to} \;\; 0.9 \times \; RRF^{FE}$$

so then

$$RRF^{SIF} \approx 12 \; \times \; \frac{MTBF_{DU}^{FE}}{T} \quad \text{to} \;\; 16 \; \times \; \frac{MTBF_{DU}^{FE}}{T}$$

or roughly,

$$RRF^{SIF} \approx 15 \; \times \; \frac{MTBF_{DU}^{FE}}{T}$$

With 1oo2 shutdown valve as final elements and annual testing (T = 1 y) and $MTBF_{DU}^{FE}$ in the range 40 y to 100 y it is feasible for a safety function to achieve RRF in the range 600 to 1,500.

The bare minimum $MTBF_{DU}^{FE}$ of dual final elements required to achieve a RRF in a 1oo2 arrangement can be estimated as:

$$MTBF_{DU}^{FE} > \frac{0.6}{\beta} \times \; RRF^{SIF} \; \times T$$

and with β around 10%:

$$MTBF_{DU}^{FE} > {}^{6}\!/_{100} \times \; RRF^{SIF} \times \; T$$

To achieve SIL 3 (i.e. RRF> 1,000) with two valves and T = 1 y requires $MTBF_{DU}^{FE}$ of at least 60 y.

Either the failure performance of the valves will need to be optimised, or β will need to be reduced below 10%, or else the inspection and test interval will need to be shortened to less than a year.