How can we accurately quantify β in fault-tolerant safety functions?

Summary

The hardware safety integrity of a safety function can be maximised by using a fault-tolerant architecture. Fault tolerance is '*the ability of a function to continue to perform a required function or operation in the presence of faults or errors*'. [IEC 61508-4]

Fault-tolerance is commonly achieved by applying redundant elements or paths within a function.

The hardware safety integrity of a fault-tolerant safety function is limited by single fault events or conditions that affect multiple channels in common, causing complete failure of the safety function. These are described as common cause failures. Common cause failure is the main factor that limits hardware safety integrity performance in most fault-tolerant architectures.

This leads to an important question:
How can we accurately estimate the likelihood (frequency or probability) of common cause failure?

Most estimates of hardware safety integrity use a β factor to model the proportion of failures that result from common causes. The alternative is to apply failure mode and effects analysis (FMEA) to identify and model common cause failures explicitly.

It is not possible to model common cause failure accurately with either β factor or FMEA.

β factor models and FMEA can only produce coarse estimates for the likelihood of common cause failure. The uncertainty in the estimates is always more than +/- 50% because the models are subjective and semi-quantitative.

The main benefit of these estimates is in enabling understanding of the reasons for common cause failure and their impact on integrity. Better understanding enables failure to be avoided, prevented or revealed by improving design, operation and maintenance practices.

The practical way to deal with common cause failure is:

Understand potential failure modes and effects	Use FMEA or inspection of fault trees to identify failures that might affect more than one channel at about the same time.
Prevent failures	Design systems to prevent common cause failures as far as practicable, or at least reduce the expected frequency of failure.
Add diagnostics	Apply frequent diagnostics to detect failures that cannot be prevented.
Facilitate periodic testing and inspection	Conduct routine testing and inspection to reveal undetected failures.
Add fault tolerance	Improve performance by applying redundant elements or paths to increase the number of faulty channels that can be tolerated; improve diversity and independence between channels.
Stagger maintenance, inspection and testing	Reduce the average time to reveal faults by staggering maintenance and test intervals for different channels. Review all other channels if a fault is found in any one channel.
Evaluate the effectiveness of controls	Set achievable performance targets for common cause failures; monitor performance against the targets throughout operation.

Published by I&E Systems Pty Ltd August 2025

Introduction

A safety function is described as having a fault-tolerant architecture if it contains redundant elements or paths.

The term 'M out of N' (or MooN) describes an architecture with a total of N channels of which only M channels are needed to function correctly. The safety function will continue to perform its specified action successfully with N-M faulty channels. The function will fail if N-M+1 or more channels are faulty, i.e. MooN architecture has hardware fault tolerance = N-M.

A single fault event or condition will cause complete failure of a safety function if it causes more than N-M channels to fail at about the same time. This is described as 'common cause failure'. The safety function then behaves as if it had a single-channel architecture.

A safety function has three sub-systems: sensor, logic, and final element. Each of the sub-systems within a safety function may have a different architecture. A combination of architectures may be applied within a sub-system.

The standards IEC 61508, ANSI/ISA 84 and ISO/TR 12489 provide detailed calculation models for estimating frequency of failure or probability of failure of safety functions with MooN architectures. The models take common cause failures into account.

Analysis of those models [1] reveals that the performance of any fault-tolerant safety function will usually be limited by common cause failures.

It is essential to understand and model common cause failures when designing safety functions. All common cause failures result from events and/or conditions (such as contamination or deterioration) that can be anticipated and controlled to some extent. The hardware safety integrity achieved by any safety function depends on how effectively causal events and conditions are anticipated and controlled.

β factor models estimate the proportion of all failures within a sub-system that are from a common cause. Different values of β may be applied for different categories of failure (dangerous or safe, and detected or undetected). Different sub-systems may have different values of β. These models are useful for producing coarse estimates of the likelihood of common cause failures in 1oo2 and 2oo3 architectures. The results can be extrapolated for other MooN architectures.

The models are useful for understanding the typical range of β and how it can be reduced.

A β factor of around 10% is feasible for 1oo2 architecture if typical engineering practices are applied. Deliberate effort is needed to reduce β factor to 5% or lower. The impact of common cause failures can be reduced by increasing the independence and diversity between redundant channels.

The uncertainty in the predicted β factor is always more than +/- 50% at best because the models are coarse. The estimates produced for β are usually in discrete steps: 0.5%, 1%, 2%, 5% and 10%.

The actual β factor achieved in operation depends on human factors and on environmental stress factors. It depends on the effectiveness of procedures applied in design, installation, and maintenance to achieve reliability. These dependencies lead to further uncertainty in the β factor that will be achieved because variation is unavoidable. Engineering practices vary. Environmental factors vary.

The uncertainty in β factor models increases further with the number of channels. Several different β factor models have been published. The different models produce significantly different estimates for architectures if there are 3 or more channels. Explicit analysis (such as FMEA) may be more useful for systems with more than 3 channels.

Common cause failure models

There are two commonly used ways of modelling common cause failures. ANSI/ISA TR84.00.02 describes them as the 'explicit model' and the 'implicit model'.

Explicit models

Explicit models depend on explicit identification of the conditions and events that might lead to concurrent failure of more than N-M channels. The likelihood of common cause failure may be estimated by using techniques such as FMEA or inspection of fault trees.

Explicit analysis enables common cause failures to be avoided by design. The likelihood of common cause failure can be minimised or eliminated by improving independence, separation and diversity between redundant channels.

IEC 61508-6 Annex D describes the binomial failure rate model (Shock model) as an example of an explicit model. It is based on fault tree modelling.

If channels A and B are completely independent and have the probability of failure $P(A)$ and $P(B)$, then we can estimate the probability of both failing at the same time as the product of the two individual probabilities: $P(A \cap B) = P(A) \times P(B)$

Explicit models add a separate term to the equation to account for probability of common cause failures. For example, a model with two channels A and B might include a third term $P(C)$ for probability of common cause failures.

The overall probability of failure of the safety function would be $P(A) \times P(B) + \boldsymbol{P(C)}$.

Implicit models

Implicit models are heuristic methods that depend on a qualitative analysis of common cause failures. The expected frequency of common cause failures is estimated as a fraction of the expected total frequency of failures in the individual channels.

These implicit models are based on many years of industry experience. They are limited to providing coarse estimates because they are qualitative models. The models rely on subjective judgements.

Implicit models consider the probability of channel B failing given that channel A has failed, $P(B|A)$ or vice versa, the probability of channel A failing given that channel B has failed, $P(A|B)$.

The overall probability of the safety function failing is the probability of channel B failing given that channel A has already failed, multiplied by the probability of A failing: $P(A \cap B) = P(B|A) \times P(A)$. It can also be expressed as $P(A \cap B) = P(A|B) \times P(B)$.

The $\boldsymbol{\beta}$ factor method is a commonly used implicit model.

The model is useful for safety functions that rely on N *similar channels* which can be expected to have a similar estimated failure frequency or probability.

Explicit models are more useful for systems that have diverse channels.

Explicit models can also be applied in conjunction with β factor models. For example, if a power failure or other systematic failure is known to affect all channels, then the power supply can be modelled separately as a common single-channel element. It can be included explicitly as a separate serial element in the reliability model.

The $\boldsymbol{\beta}$ factor is the fraction of channel failures that are expected to affect similar channels in a similar way and at a similar time.

Conditional probability of failure may be inferred from the β factor. With 2 channels A and B, the β factor represents the fraction of failures that cause concurrent failure of both A and B.

If it is given that channel A has already failed, then the probability of a similar failure occurring at a similar time in channel B can be modelled by β.

In effect: $\beta = P(B|A)$, or similarly $\beta = P(A|B)$

So $P(A \cap B) = \beta \times P(A)$ or $= \beta \times P(B)$

Similarly, if channels A and B can be expected to fail at some similar rate λ, then the expected rate of common cause failures could be estimated as $\beta \times \lambda$.

Note that β may be applied separately to different modes of failure: detected dangerous failures, undetected dangerous failures, detected safe failures and undetected safe failures.

β factor method

IEC 61508-6 Annex D provides a β factor methodology that is commonly used. The method is based on a weighted scoring against a set of question under these headings:

- Separation/segregation
- Diversity/redundancy
- Complexity/design/application/maturity/experience
- Assessment/analysis and feedback of data
- Procedures/human interface
- Competence/training/safety culture
- Environmental control
- Environmental testing

The scoring results in an estimate of β factor in steps of 1%, 2%, 5% and 10% for sensor and final element sub-systems using a 1oo2 architecture. The steps for logic solver sub-systems are 0.5%, 1%, 2%, and 5%. Scaling factors are applied to estimate β for other MooN architectures. The scaling is explained below.

The scoring method in IEC 61508-6 Ed.2 usually results in an estimated β factor of 10% for sensors and final element subsystems. The estimate may be reduced to 5% if a deliberate effort is made to address causes of common cause failure. A revised scoring table has been proposed for IEC 61508-6 Ed. 3. It provides more guidance on how the β factor might be reduced.

IEC 62061 describes a similar method. It uses more optimistic scoring and usually results in an estimate of 5%.

The coarse steps in these two models limit mean that the uncertainty in any estimate of β factor for 1oo2 architecture is always at least +/- 50%.

The uncertainty is worse than +/- 50% because of the assumptions made when applying the models. It results from a lack of knowledge or information rather than from uncertainties or errors in measurement. Uncertainty in assumptions cannot be evaluated using statistical techniques. For example, system designers need to make assumptions about how the equipment will be operated and maintained in the future. The actual effectiveness of future maintenance practices cannot be known with certainty.

SINTEF paper A26922 [2] summarises results from a review of common cause failures recorded in industrial applications. The review concluded that the β factor achieved in operation is typically in the range 12% to 20% for sensor and final element sub-systems.

The initial β factor estimate is specifically for 1oo2 architecture. Scaling factors are applied when β factor is estimated for other MooN architectures. Those scaling factors add further uncertainty.

For example, a 1oo5 voting architecture is much less susceptible to common cause failure. Only 1 out of the 5 channels needs to work correctly for successful action. It will still work correctly with 4 faulty channels.

A 4 out 5 voting architecture might be selected to reduce the rate of spurious trips, but it is much more susceptible to common cause failure. At least 4 devices must work correctly to trip the function. A 1oo2 architecture and a 4oo5 architecture will both fail if 2 or more channels are faulty at the same time. The chance of having 2 faulty channels out of 5 is greater than having 2 faulty channels out of 2.

This is similar to how the probability of the independent failures increases by the factor $\binom{N}{N-M+1}$, i.e. how many different ways can we have N-M+1 unrelated faults in N channels. Unfortunately, there is no exact way of modelling the increased likelihood of common cause failure: what is the likelihood that the same event will affect at least N-M+1 channels in a similar way. Models have been proposed, but there is no clear agreement, and the models have not been validated.

IEC 61508-6 Table D.5 suggests β scaling factors for MooN architectures with N up to 5. For example, if the β factor is estimated at 0.1 for 1oo2 voting, then IEC 61508 suggests it should be increased to 0.15 for 2oo3 voting.

The SINTEF PDS Method Handbook [3] provides an explanation of how these scaling factors are derived. It describes complex mathematical modelling as the basis of 'configuration factors' C_{MooN}. It provides configuration factors to scale the β factor for MooN architectures with N up to 6.

Some of the configuration factors suggested in the PDS Method Handbook are higher than given in IEC 61508-6. For example, if the β factor is estimated at 0.1 for 1oo2 voting, then the PDS Method suggests it should be increased to 0.2 for 2oo3 voting, compared with 0.15 in IEC 61508.

PDS Method suggests a scaling of 2.8 for 3oo4 and 1.1 for 2oo4 voting. The corresponding scaling factors in IEC 61508 are 1.75 for 3oo4 and 0.6 for 2oo4.

The different models give different results because they are subjective. The models are supposedly based on observation. Neither the PDS Method nor IEC 61508 refer to objective empirical evidence to validate the models.

The β factor achieved in operation with a 2oo3 architecture could be anywhere between 10% and 30%. For 3oo4 architectures it could be in the range 20% to 40%.

The actual value of β that we achieve during the operation of safety functions will vary widely. It depends on the design, the installation, on human factors and on environmental factors.

Selecting a value of $\boldsymbol{\beta}$ effectively sets performance targets for the design, installation, operation and maintenance of a safety function.

Multi-channel architectures with N > 3

β factor models are not likely to be practicable with HFT = 1 and N > 3 because of the increasing uncertainty in the models. An explicit approach (such as FMEA) is more useful. It is better to identify the causes of common cause failure so that they can be prevented or detected.

Occasionally safety function architectures may be applied with N as high as 10. These applications generally rely on higher levels of fault tolerance and/or higher diagnostic coverage.

For example, a pressure let down station might need to close at least 9 out of 10 pressure control valves. Each valve would have continuous position control and diagnostics. A shutdown solenoid

bypasses the positioner on each valve to force closure in response to a high-pressure trip. An overall flow rate of 10 % of full flow might be tolerable.

A β factor model could be applied, but the β factor would be somewhere in the range of 50% to 100%.

Common cause failure should be managed

β factor models should not be applied without question. The models do not compensate for poor practice in managing systematic issues. It is more useful to conduct FMEA to prevent, avoid or control common cause failure.

Neither β factor models nor FMEA should be expected to produce precise or accurate estimates of common cause failure rate.

The conditions and events leading to common cause failure can be anticipated. The likelihood of common cause failure can be managed by:

- Specifying devices that are better suited for the application
- Improving diagnostic coverage
- Increasing the level of fault tolerance
- Staggering maintenance intervals
- Staggering inspection and testing intervals.

Are common cause failures random or systematic?

The topics covered by the β factor evaluation questions are obviously mostly systematic rather than random. This seems to contradict the objective of estimating the probability of random hardware failures that have a common cause. However, it is difficult to distinguish between systematic and random failure in practice. There is no clear difference.

The functional safety standards try to distinguish two different types of safety integrity: systematic safety integrity and hardware safety integrity (for random failure). The distinction is not clear. The definitions are ambiguous. The following explanation is based on the definitions given in IEC 61508-4.

Systematic safety integrity

Systematic safety integrity relates to the probability of a safety function performing correctly with respect to systematic failures.

Systematic failures are those that are related in a deterministic way to a certain cause. Systematic failure may be eliminated, prevented or controlled by dealing with the causes of failure or by preventing fault conditions or events developing into failure.

Systematic failures do not occur at predictable rates because systematic failures *should not* recur after they have been found and remedied. They may recur if they have not been resolved effectively.

Software errors provide a good example of purely systematic faults that can lead to failure. They should not recur after they have been corrected, but similar faults may remain hidden in the software.

Methods such as HEART and THERP can be applied to estimate the probability of future systematic failures (refer to Smith, D.J. [6]). The probability depends on the level of effort put into finding and addressing the causes of systematic failure. IEC 61508 provides detailed guidance on the application of techniques and measures to achieve systematic integrity.

Hardware safety integrity

Hardware safety integrity relates to the probability of a safety function performing correctly with respect to random hardware failures.

Failures are classed as random failures if they occur at a constant rate that can be predicted with reasonable accuracy. The times at which these failures occur is random. It cannot be predicted.

ISO/TR 12489 provides more detailed explanation of failure types. It attempts to clarify the distinction between random and systematic. It uses the term 'catalectic' failure to describe purely random failure of individual components. Catalectic failures result from some sort of stochastic process, a series of independent trial events that has no memory. Catalectic failures are sudden failures that occur without warning. They are impossible to forecast by examining an item. These failures cannot be prevented. Catalectic failures are characteristic of simple components with constant failure rates.

IEC 61508-4 introduces some confusion by defining random hardware failures as those that result from one or more of the possible degradation mechanisms in the hardware. This means that they are not purely random in the mathematical sense because they do not result from stochastic processes. These failures are expected to occur at a reasonably constant and predictable rate, but that rate will vary. The rate depends on the level of effort put into managing the causes of degradation and preventing the development of degradation into failure.

This means that random hardware failures are always at least partially systematic. They cannot be eliminated but they can always be *prevented to some extent*. The rate of failure may be reasonably constant, but that rate can be changed through deliberate or inadvertent action.

The OREDA datasets [4, 5] show that equipment failure rates recorded by different users vary over at least a range of 30:1. It is clear from OREDA that there is no consistent way of classifying and counting failures. The point at which degraded performance results in failure of equipment or failure of a function depends on the performance required for each specific application. For example, a valve may be designed to close within 20 seconds. The closing time might degrade to 30 seconds as the valve deteriorates. That would be classed as a dangerous failure if a hazard occurs after the flow continues for 25 seconds. It is not a failure if a closing time longer than 30 seconds is acceptable.

ISO 14224 provides useful guidelines for the collection of reliability data, but it also reveals the difficulties in comparing reliability data between different applications.

Probability of hardware failure

The probability of hardware failures may be estimated from mathematical models.

IEC 61508-6, ANSI/ISA 84 and ISO/TR 12489 describe techniques for evaluating the probability of random hardware failure. The models are based on the fundamental assumption that hardware failure rates remain constant. The models reveal that the probability of failure varies in almost direct proportion to the variation in failure rates.

The random hardware failure probability models all include common cause failures in the calculation. Common cause failures typically account for more than 75% of the overall random hardware failure probability estimate for any fault-tolerant safety function. The overall random hardware failure probability is typically no more than about 30% higher than the probability of common cause failure.

This causes concern for many people, because common cause failures are not random.
They result from causes that may be eliminated, prevented or controlled. They are at least partially systematic. Systematic failures should be treated by improving systematic integrity.

The pragmatic approach

There is no need to be too concerned about a dichotomous classification of failures into systematic and random. The failure models are still useful even though we cannot clearly distinguish between systematic and random hardware failures.

The primary approach in functional safety is to prevent failures as far as is necessary to meet the safety integrity target. It is not practicable to prevent all failures.

Failures may be treated as if they were random to the extent that they cannot be reasonably prevented, and to the extent that a reasonably consistent rate of failure can be achieved.

It does not matter that the rate can vary. What matters is that the rate can be controlled to some extent. Safety integrity may be improved by:

- Preventing preventable failure as far as is practicable through application of effective procedures, techniques and measures
- Adding continuous automatic diagnostic functions to detect failures that cannot be prevented
- Conducting regular periodic inspection and testing to find the failures that cannot be detected by diagnostics
- Applying fault-tolerant designs so that the performance of the function can be maintained in the presence of one or more faults
- Identifying and addressing common cause failures
- Maintaining the condition of the safety function equipment throughout its lifetime
- Monitoring reliability performance to ensure that the integrity targets are achieved.

Precise prediction vs prevention

Physical processes may be measured and modelled with extreme precision.

Parameters such as pressure, temperature, flow, level, voltage and current may be measured with an accuracy of better than +/- 0.1%

Measurements of process or machine parameters are typically presented with 3 or 4 significant figures of precision (i.e. #.## x $10^{\#}$ or #.### x $10^{\#}$).

Some physical systems may be modelled accurately. For example, electrical and electronic systems can usually be modelled to within +/- 0.1%. The models can be validated through physical measurements. The results will be repeatable. The models can account for variations in temperature and impedance.

Chemical, mechanical and structural models usually include many non-linear variables. The chaotic nature of the interactions in these systems results in more uncertainty. These systems may be modelled with a reasonable level of accuracy, typically in the range of +/- 1% to +/- 10%.

It is more difficult to model equipment failures. Failures cannot be easily categorised and measured. The point at which degraded performance results in failure depends on each individual application.

Equipment or device failure rate statistics are published by manufacturers and by some certifying authorities (such as *exida* and the TÜV companies).

The failure rates are either based on failures recorded during operation of a specific make and model, or on FMEA. Device failure rates can be estimated through FMEA using industry-wide failure rate statistics for the individual components used in a device.

The published failure rates are based on long-term averages measured over large populations. The failure rates may be presented with at least 3 significant figures of precision (#.## x $10^{\#}$) because long-term averages are stable. They include thousands of measurements, so any variation is gradual. The

failure rate of any one set of devices should be expected to vary from the long-term averages by at least a factor of 10. The failure rate of devices in a specific application can be up to 10 times higher or lower than average. This variation is clearly evident in OREDA.

The failure of individual electrical and electronic components can be characterised by reasonably constant failure rates, but the failure rates may vary over at least a range of 10:1 (Smith, D.J. [6]).

IEC 61709 describes stress factors that contribute to variation in failure rates and how the change in failure rate might be predicted. It reveals that all failure rates can be influenced if those stress factors can be controlled. It also shows that individual component failure rates will vary over at least a range of 10:1 because of variations in stress and strength.

Several failure probability calculation software packages are available commercially. Some are specifically intended for estimating the failure probability or failure rate in safety functions. Some are general fault tree models that can be applied more widely.

These software packages are all based on the assumption that failure rates remain fixed and constant. They report results with at least 3 significant figures of precision. Unfortunately, this may create the impression that the results are accurate to within +/- 1%. The reality is that failure rates can only modelled to within a range of 10:1 at best. The actual likelihood of failure may be 2 or 3 times worse than estimated.

The main benefit in modelling safety functions mathematically is that the models reveal the parameters that determine safety function performance. Those parameters can be controlled.

Safety integrity of safety functions depends primarily on these factors:

λ_{DU}	The rate of undetected dangerous failures
T	The weighted average time interval for revealing undetected failures (e.g. revealed either by periodic inspection and testing, or revealed at the time of renewal, or revealed by failure on demand)
MTTR	The mean time to restoration after a fault or failure is detected by a diagnostic function (applies if equipment is left in operation while out of repair and without compensating measures)
MooN	The voting architecture (with fault tolerance = N-M)
β	The fraction of failures that have a common cause

All these factors can be controlled.

Conclusion

Failure probability (including probability of common cause failure) depends primarily on the effectiveness of procedures, techniques and measures applied to prevent failures.

The conclusion is that our primary concern should be in preventing safety function failure. It requires a deliberate and concerted effort.

All failure rates vary. Equipment failures are not caused by events that can be characterised by fixed and predictable rates.

Common cause failures dominate the probability of failure in most fault-tolerant safety functions.

The fraction of failures with a common cause β will always vary. It is never fixed or constant.

Failure rates and β factors can be thought of as a performance targets.

We need to select targets that are feasible. We need to design, install, operate and maintain the safety systems so that the targets will be achieved.

Failure probability cannot be calculated with precision, but it can be controlled. Probability calculations are useful in evaluating whether the techniques and measures are likely to be sufficiently effective to achieve the safety integrity target.

The estimated probability of failure is not a precise guarantee of performance. It is essentially an indication of the safety integrity that can be achieved if the performance targets are met.

References

TABLE 1. STANDARDS AND CODES

Number and date	Title
IEC 60812: 2018	Analysis techniques for system reliability – Procedure for failure mode and effects analysis (FMEA)
IEC 61508-4: 2010	Functional safety of electrical/electronic/programmable electronic safety-related systems Part 4: Definitions and abbreviations
IEC 61508-6: 2010	Functional safety of electrical/electronic/programmable electronic safety-related systems Part 6: Guidelines on the application of IEC 61508-2 and IEC 61508-3
IEC 61511: 2016	Functional safety — Safety instrumented systems for the process industry sector
IEC 61709:2017	Electronic components – Reliability – Reference conditions for failure rates and stress models for conversion
IEC 62061: 2021	Safety of machinery – Functional safety of safety-related control systems
ISA TR84.00.02-2022	Safety Integrity Level (SIL) Verification of Safety Instrumented Functions
ISO/TR 12489: 2013	Petroleum, petrochemical and natural gas industries — Reliability modelling and calculation of safety systems
ISO 14224: 2016	Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment

TABLE 2. REFERENCE DOCUMENTS

Ref	Title
1	I&E Systems Pty Ltd '*Are simple approximations for safety function calculations too good to be true?*' White paper published April 2025 https://www.iesystems.com.au/uploads/2025/04/Error-in-safety-function-PF-approximations.pdf
2	SINTEF report A26922, *'Common Cause Failures in Safety Instrumented Systems; Beta-factors and equipment specific checklists based on operational experience. '*, 2015
3	SINTEF Reliability Prediction Method for Safety Instrumented Systems PDS Method Handbook. 2013
4	OREDA '*Offshore Reliability Data Handbook'* Volume 1, 5th Ed. SINTEF. 2009
5	OREDA '*Offshore and Onshore Reliability Data Handbook*' Volume 1, 6th Ed. SINTEF. 2015
6	Smith, Dr David. J. *'Reliability, Maintainability and Risk'*, 9th Ed. Butterworth Heinemann. 2017

Creative Commons Licence

The document was prepared by Mirek Generowicz and Blake Merritt of I&E Systems Pty Ltd. The authors acknowledge comments and contributions from Ray Martin.