Common Cause Failures and Ultra Reliability

Harry W. Jones[1]
NASA Ames Research Center, Moffett Field, CA, 94035-0001

A common cause failure occurs when several failures have the same origin. Common cause failures are either common event failures, where the cause is a single external event, or common mode failures, where two systems fail in the same way for the same reason. Common mode failures can occur at different times because of a design defect or a repeated external event. Common event failures reduce the reliability of on-line redundant systems but not of systems using off-line spare parts. Common mode failures reduce the dependability of systems using off-line spare parts and on-line redundancy.

Nomenclature

CCF = Common Cause Failure
FMEA = Failure Modes and Effects Analysis
HA = Hazard Analysis
IAEA = International Atomic Energy Agency
IEC = International Electrotechnical Commission
NRC = Nuclear Regulatory Commission
NUREG = Nuclear Regulatory Commission Regulation
OREDA = Offshore REliability Data
PRA = Probabilistic Risk Assessment

I. Introduction

PEOPLE expect technical systems to work. But they do fail. High expected reliability is sometimes justified by analysis and tests that predict low failure rates. Reliability design attempts to identify and remove failure causes, including the sources of common cause failures. If unanticipated failures occur, they cause concern and require investigation. Errors in design and manufacturing are often found. Since such errors can affect other copies of a system, they are at least potential common causes of multiple failures.

Unidentified common cause failure modes reduce the theoretically achievable reliability of systems using redundancy and spares. Common cause failures have been modeled and measured.

In theory, ultra reliability can be achieved for any system by providing sufficient redundant or spare components. In practice, common cause failures can disable all the back-up components so that the system fails. The common cause may be a design deficiency or an unexpected external stress. Actual multiple failures can have either multiple causes or one common cause. In failure analysis, common cause failures include any that can possibly disable both a component and its backups, even if no backup exists, even if no failure occurs. Common cause failures can defeat redundancy and prevent the achievement of ultra reliability.

About ten per cent of all failures are usually identified as common cause. This means that typically redundancy cannot reduce the individual component failure rate by more than about a factor of ten! Achieving ultra reliable requires both designing very reliable components and using redundancy more effectively by significantly reducing the potential for common cause failures.

II. Reliability, redundancy, and common cause failures

Achieving ultra reliable system performance is difficult. It requires identifying the even very unlikely failure causes and then redesigning the system to remove them. The first step is to select or design highly reliable subsystems and components. But often the best possible subsystems have failure rates that are too high, so that their total sum, the overall system failure rate, is also too high. Then the necessary second step is to provide redundant subsystems. If two redundant subsystems always fail independently of each other, their combined failure rate is the

[1] Systems Engineer, Bioengineering Branch, Mail Stop N239-8, AIAA Senior Member.

product of the two individual subsystem failure rates. For example, suppose a subsystem has a failure probability of 0.01 per year. Then the probability that two identical redundant subsystems both fail is 0.01 * 0.01 = 0.0001 per year. Ultra reliability can be achieved in theory by providing one or more redundant subsystems.

One kind of common cause failure occurs when all the redundant subsystems fail at the same time due to a single cause. Such common cause failures are a serious difficulty in achieving ultra reliability using on-line redundancy.

Common cause failures can occur in surprising ways. Failure of one subsystem may damage its spare, as when an internal failure caused one Apollo 13 oxygen tank to explode and its explosion destroyed the second oxygen tank. This is a cascade type of common cause failure. Simpler common causes of failure are either external or internal. A tsunami disabled all the Fukushima Daiichi reactors' dissimilar cooling pump power sources: the transmission grid, diesel generators, and battery back up. In the most frequent case, an internal design flaw or programming bug can cause all identical redundant subsystems to fail in the same way at the same time.

A. Single points of failure

A common cause failure should be distinguished from a single point of failure. In both cases, one failure cause can disable an entire system. But two or more redundant subsystems must fail in a true common cause failure, while only one must fail at a single point of failure.

Correctly distinguishing single point from common cause failures is not critical during initial safety and reliability analysis. Hazard analysis (HA) and logic models can be used top-down to explain how events such as component failures can cause system failures. Failure modes and effects analysis (FMEA) works bottom-up starting from component failures to identify system level effects. The objective qualitative analysis such as HA and FMEA is to understand what can go wrong and how, to eliminate or mitigate the hazards, and perhaps to achieve fault tolerance for critical systems. Probabilistic risk assessment (PRA) and reliability block diagrams use the quantitative likelihood of failure and to estimate system reliability. Looking for common cause and single point failures helps identify the important failure modes, however they are classified.

B. Reliability estimation allowing for common cause failures

After the obvious and high probability failure modes are eliminated during system design, it is sometimes assumed that only unpredictable random failures will occur during test and operations. The subsystem random failure rate is estimated from historical data and measured in long duration test. The rate of unexpected common cause failures that were missed during design is similarly estimated and measured.

The common cause failure rate must be included to predict the overall failure rate of a system that uses redundancy to achieve high reliability. The previous example used a subsystem with a failure probability of 0.01 per year, which gave the probability that two redundant subsystems both fail of 0.01 * 0.01 = 0.0001 per year, assuming no common cause failures. Common cause failures occur when the redundant subsystems have correlated or dependent failures. Suppose that every time a subsystem fails, it explodes and destroy the redundant subsystem. Then all the failures are common cause failures, similar to the Apollo 13 failure cascade. In this case, the probability that both subsystems fail is 0.01 + 0.01 = 0.02 per year. Two subsystems fail twice as often as one. System design or out of specification operating conditions also could cause two redundant subsystems to fail in the same way at the same time.

Typically common cause failures account for about ten percent of all the failures, so ninety percent of failures are random and independent, as usually assumed. If a subsystem has a stand-alone failure probability of 0.01 per year and a common cause failure probability of 0.001 per year, the probability that two subsystems both fail is 0.01 * 0.01 + 0.001 = 0.0001 + 0.001 = 0.0011 per year, adding the redundant pair and common cause failure rates. Redundant systems reduce the failure rate, but only by one order of magnitude, not two orders. Redundancy does not prevent common cause failures, and it cannot reduce the failure rate for a redundant system to less than the common cause failure rate.

III. Common cause failure effects on ultra reliability

The effects of common cause failures are different for systems using on-line redundant subsystems, the usual concern for common cause failures, than for systems using off-line spares, which are subject to common mode but not common event failures.

A. Common cause failures and redundancy

Previous work has shown that it is possible to achieve ultra reliability for a system that has a reasonable initial failure probability F (reasonable meaning F is less than 0.1 over the mission length), by dividing the system into N

subsystems and making each subsystem a redundant pair. Each of the N subsystems has a failure probability of F/N. A redundant pair of subsystems has failure probability $(F/N)^2$, and the series of N pairs has failure probability $N * (F/N)^2 = F^2/N$. For F = 0.1 and N = 10, $F^2/N = 0.001$.

(This example uses on-line redundancy to provide a simple illustration of how it is theoretically possible to achieve ultra reliability. Operating on-line redundancy is not suitable for spacecraft life support, since the crew can replace any failed subsystems. Operating two subsystems simultaneously would double the power use and number of failures, so it is less effective than using spares.)

Suppose that the system has a common cause failure probability of β F. Dividing the system into N subsystems, each has a common cause failure probability of β F/N. A redundant pair of subsystems has failure probability $(F/N)^2 + \beta F/N$, and the series of N pairs has failure probability $N * [(F/N)^2 + \beta F/N] = F^2/N + \beta F$. The system failure rate cannot be reduced below the original common cause failure rate. For F = 0.1, β = 0.1, and N = 10, $F^2/N + \beta F = 0.001 + 0.01 = 0.011$. Achieving ultra reliability using redundant pairs of subsystems is possible if there are no common cause failures, but is prevented by a high percentage of common cause failures.

B. Common mode failures and spares

The previously suggested approach to ultra reliable life support uses spares, not on line redundancy. The spares are off line, in storage, and ready to be installed if a failure of the operating subsystem occurs. A typical external event, a fire or an operator error, would not affect both the functioning unit and its spare. Any external failure cause would hopefully prevented in the future to avoid damaging the replacement spare, but it could recur nonetheless. An internal problem, defective parts or a design or programming error, can cause the identical subsystems to fail in the same way, but usually not at the same time. Systems that use spares instead of redundancy to achieve high dependability do not experience common event failures, but they can have lower than expected reliability because of common mode failures. Since spare components are not affected by common event failures, their rate of common cause failures is lower than that of redundant components.

IV. Brief history of common cause failure analysis

Reliability analysis was developed to understand the frequent vacuum tube failures in electronic systems before the transistor was invented. Fault tree analysis was developed in the early 1960's by Bell labs and Boeing to analyze missile systems and later commercial aircraft. After the 1967 Apollo 1 fire, fault tree analysis was performed on the Apollo system. Fault tree analysis was used with PRA and included common mode failures. NASA did not use PRA after Apollo or in space shuttle design. The nuclear power industry adopted PRA after the 1979 Three Mile Island accident. Common cause failures were mathematically modeled in the 1970's and 1980's. Several common methods are beta factor, multiple Greek letter, and alpha factor. PRA was readopted by NASA after the 1986 Challenger accident. It was used to analyze the International Space Station, the Orbital Space Plane, and the Constellation Program. The Constellation Program PRA document requires the use of alpha factors in common cause failure modeling. (Pate-Cornell and Dillon) (Seif) (Stotta et al.)

V. Definitions of common cause failures

There are many definitions and distinguishing them is useful in understanding common cause failures. Most work on common cause failures has been done in the nuclear power industry and is largely concerned with on-line redundant systems, but the work has been extended to space systems. (Stotta et al.) (Rutledge and Mosleh)

A. Descriptive definitions

A common cause failure occurs when one event or shared factor causes two or more failures. Because they have one and the same cause, the failures are not statistically independent. This is contrary to the core assumption of usual reliability theory, that all failures are independent and uncorrelated.

An independent failure occurs when the probability of two or more systems failing is equal to the product of their individual failure probabilities. If a system has failure probability F, a redundant pair of subsystems has failure probability F^2. Failures that are not independent are dependent.

A dependent failure occurs when the probability of two or more systems failing is greater than the product of the failure probabilities of each system failing. Dependent failures may have the same cause, multiple unrelated causes, or no known cause.

A common cause failure is a specific type of dependent failure where several failures result from a single shared cause. A common event failure is a specific type of common cause failure where multiple failures result from one single external event. The failures are usually simultaneous or nearly so. Or, common event failures can occur in an

extended cause and effect sequence, called a cascade failure. Common event failures are a concern for on-line redundant systems.

A common mode failure is a specific type of common cause failure where several subsystems fail in the same way for the same reason. The failures may occur at different times and the common cause could be a design defect or a repeated event. This is the common cause failure type applicable to off-line spares.

Similar definitions are given in several European sources. (Borcsok et al.) (R&M)

C. NRC definitions

The Nuclear Regulatory Commission (NRC) defines a dependent failure the same way as above. They define a common cause failure as a "dependent failure in which two or more component fault states exist simultaneously, or within a short time interval, and are a direct result of a shared cause." (NRC, p. 79) This definition emphasizes multiple simultaneous failures and so is more suitable for on-line redundancy than for off-line spares. This is consistent with the NRC's operational safety system concerns. However the requirement that the failures be close in time is relaxed to "components fail within a selected period of time such that success of the probabilistic risk assessment (PRA) mission would be uncertain." (NRC, p. 1) This would include a mission with a limited number of off-line spares and no further supply of spares.

The NRC states "the term "common-mode failure" which was used in the early literature and is still used by some practitioners is more indicative of the most common symptom of CCF (common cause failure) (i.e., failure of multiple components). As such, it is not a precise term for communicating the main character of CCF events." (NRC, p. 24)

D. Concise and elegant definition

Smith and Watson in 1980 studied many different definitions of common cause failures and concluded that the best definition depends on the field of use. They proposed that a common cause failure is the "Inability of multiple, first-in-line items to perform as required in a defined critical time period due to a single underlying defect or physical phenomena such that the end effect is judged to be a loss of one or more systems." (Smith and Watson in Lilleheier) This definition seems clear and correct, but it is difficult to understand without more extensive definitions and discussion. First-in-line items clearly includes on-line redundant systems, but can be stretched to cover spares. The critical time period obviously indicates an acute need period, but can be the duration of a long mission without resupply.

E. Practical definition

A useful practical definition focuses on the essence of the common cause problem. Common cause failures are "dependent failures that defeat the redundancy or diversity employed to improve the reliability of systems." (Mosleh in Bukowski and Goble)

VI. Failure classes and causes

Non-random failures are systematic, understandable, and explainable. They have identifiable causes and familiar sources.

A. Random and systematic failure types

Failures can be classified as random or systematic. Random hardware failures are caused by time and use and occur independently. Random failures are not common cause failures. Systematic failures occur because of a poor specification or design or an unexpected interaction or external stress. Systematic failures can affect all identical components, so systematic failures are the potential common cause failures.

B. Sources of systematic or common cause failures

There are many sources of systematic failures within systems. Specification, design, and manufacturing problems are frequent, and often result from flawed redesign and modification procedures. Other failures occur due to insufficient design for reliability and incomplete testing. Poor monitoring and maintenance are major contributors to systematic failure. Human errors are rare but can have high impact. Software is a major source of common cause failures.

External sources of systematic failures include interfaces, the environment, and major adverse events. The interfaces include power, cooling, material inputs, and external controls. The environment can produce excess

temperature, pressure, vibration, impact, noise, and contamination. Events include earthquake, tsunami, hurricane, tornado, flood, and blizzard.

Most so-called common cause failures are not actual multiple failure incidents. A single failure or even an inspection, test, or analysis can reveal a system problem shared by many or all the similar subsystems, which could potentially lead to a common cause failure during operations.

VII. Explicit common cause failure analysis

Any common cause failure events or modes that are discovered in analysis should be modeled as explicit events in the initial reliability analysis using probabilistic risk analysis or fault trees. The causes and effects should be considered and the probabilities of the failure events estimated. Mitigations should provided as needed.

The explicit modeling of all non-negligible failure causes is the goal of reliability analysis. In reality, many failures modes are not anticipated. The number of unexpected systematic or common cause failures that are not explicitly modeled and not mitigated can be estimated using implicit, probabilistic, common cause failure models and historical data. It is obviously much better to capture all the failure modes in an explicit model and deal with them than to estimate how many have been missed, even with an elegant mathematical model. Since we cannot identify all the failure modes, it is useful to estimate how far we fall short.

VIII. Implicit common cause failure models

Dependent or common cause failures that are not explicitly modeled in the PRA and fault tree can be implicitly modeled using parametric methods. The simplest model is the basic parameter model. The beta factor model, the multiple Greek letter model, and the alpha factor model are adaptations of the basic parameter model.

A. Basic parameter model

The basic parameter model is called a direct model because the probabilities of the different failure events, Q_i, are used directly. Suppose a system has two identical redundant subsystems. The probability that either one of the two subsystems fails independently is Q_1. The probability that both subsystems fail together from a common cause is Q_2. The dual redundant system failure rate is $Q_S = Q_1^2 + Q_2$. Data is needed to estimate Q_1 and Q_2. In a triply redundant or a 2-out-of-3 system, the probability that the three subsystems all fail together due to common cause is Q_3, and additional data is needed on this rare event. The direct basic parameter model requires estimating all possible failure event probabilities. (Borcsok et al.) (Stotta et al.)

B. Beta factor model

The beta-factor model is a simplification of the basic parameter model that requires less data. The subsystem failure rate, Q_{SS}, is assumed to have two components, Q_I due to independent failures and Q_{CC} due to common cause failures. $Q_{SS} = Q_I + Q_{CC}$. The parameter β is the fraction of the total subsystem failure rate due to common cause failures. $\beta = Q_{CC}/Q_{SS} = Q_{CC}/(Q_I + Q_{CC})$. If β is zero, there are no common cause failures. If β is 0.1, ten percent of all failures have common causes. The beta-factor model was used in the above examples computing the system failure rate with common cause failures. (Borcsok et al.) (Bukowski and Goble) (Stotta et al.)

The beta-factor model uses test data or similar system experience to estimate the subsystem failure rate Q_{SS} and the parameter β. Only these two parameters are used in the model. The number of model parameters and the uncertainty due to sparse data on rare events are minimized. The beta factor model is the most frequently used common cause failure model. Large amounts of data have been gathered, especially on nuclear power systems. (Borcsok et al.) (Stotta et al.)

In systems with three or more redundant subsystems, the beta-factor model cannot distinguish between two and three or more failures. The failure rate can be overestimated if it is assumed that all the subsystems fail whenever a common cause failure occurs.

C. Multiple Greek letter model

Systems that use triple or more redundancy or require say 2-out-of-3 (n-out-of-m) subsystems to operate can be modeled more accurately by adding parameters. The multiple Greek letter model includes parameters for the conditional probabilities that the N+1-th subsystem fails given that N identical subsystems have already failed. Thus, β is the conditional probability of the failure of a second subsystem, given that the first has failed, γ is the conditional probability of the failure of a third subsystem, given that two have failed, and δ is the conditional probability of the failure of a fourth subsystem given that three have failed. The method can be extended to four or

more subsystems. (Borcsok et al.) The multiple Greek letter model is sometimes called the extended or multiple beta model. (Bukowskin and Goble)

D. Alpha factor model

The alpha factor model explicitly includes all the possible combinations of multiple failure events. $Q_k^{(m)}$ is the probability that k specific subsystems in a group of m fail. The probability that any k of the m subsystems fail is $Q_k^{(m)}$ multiplied by the number of different possible specific groups of k subsystems, $(m!/(m-k)!\ k!)\ Q_k^{(m)}$. The total probability that one or more of the m subsystems fail is the sum of this expression (for any k of m failure probability) over all k from 1 to m. (Borcsok et al.) (Stotta et al.)

The alpha factor $\alpha_k(m)$ is the fraction of the total subsystem failure rate due to any k out of m failures. It is equal to the probability that any k of the m subsystems fail, $(m!/(m-k)!\ k!)\ Q_k^{(m)}$, divided by the total probability that one or more of the m subsystems fail. (Borcsok et al.) (Stotta et al.)

Suppose there are three identical redundant subsystems labeled A, B, and C. The probability that a particular one (say A) fails is $Q_1^{(3)}$. The probability that any one of the three (A or B or C) fails is $(3!/(2!\ 1!)\ Q_1^{(3)} = 3\ Q_1^{(3)}$. The probability that two specific subsystems (say A and B) fail is $Q_2^{(3)}$. The probability that any two of the three fail (AB, BC, or AC) is $(3!/(1!\ 2!)\ Q_2^{(3)} = 3\ Q_2^{(3)}$. The probability that all three subsystems (ABC) fail is $Q_3^{(3)}$. The total subsystem failure rate, Q_{SS}, is equal to $3\ Q_1^{(3)} + 3\ Q_2^{(3)} + Q_3^{(3)}$. Then $\alpha_1^{(3)} = 3\ Q_1^{(3)}/\ Q_{SS}$, $\alpha_2^{(3)} = 3\ Q_2^{(3)}/\ Q_{SS}$, and $\alpha_3^{(3)} = Q_3^{(3)}/\ Q_{SS}$. (Borcsok et al.) (Stotta et al.)

The $\alpha_1^{(3)}$ corresponds to individual, independent failures, and the $\alpha_2^{(3)}$ and $\alpha_3^{(3)}$ account for two and three dependent common cause failures. For m = 2, the $\alpha_2^{(3)}$ corresponds to the β in the beta factor model. The alpha factor model is sometimes referred to as the multiple beta factor model. (Lilleheier) The Constellation Program PRA Methodology Document specifies using the alpha factor model. (Stotta et al.) It should be used for triple or higher redundancy if adequate data can be obtained.

E. Shock (binomial failure rate) model

The shock model assumes that the system has common cause failures due to external events or shocks that occur at some average rate, μ. When a shock occurs, each of the m subsystems may fail with probability p. The number of failures after a shock varies from 0 to m and the probability of all the different numbers of failures is given by the binomial distribution for probability p. The common cause failure rate is $\mu\ p$, the product of the shock rate, μ, times the probability of a failure after a shock, p.

The usual individual random failures also occur at rate λ, so the shock model has three parameters, λ, μ, and p. The failure rate $Q_{SS} = \lambda + \mu\ p$. The shock model is directly applicable if the frequency and impact of shocks can be identified. If only the common cause failure rate is known μ and p cannot be independently determined, the beta-factor model is more appropriate. Also, if shocks always cause all the subsystems to fail, p = 1, and the beta-factor model applies. (Borcsok et al.) (Lilleheier)

IX. Beta factor model data

The beta factor model is the simplest and most commonly used model. It requires estimating only two parameters, the total subsystem failure rate, Q_{SS}, and β, the common cause fraction of the total failure rate. The independent failures are the remaining non-common cause fraction $(1 - \beta)$ of the total failure rate, $Q_I = (1 - \beta)\ Q_{SS}$.

Obtaining good failure rate data is difficult, since it requires long testing of many units. Failure rates are usually estimated from data on similar components. Getting common cause failure data is even more difficult, because common cause failures are a relatively small portion of the total failures. It is sometimes difficult to decide if a failure is random or due to a common cause that could affect similar units. Collecting and classifying data requires time, effort, and resources. Limited data is the main reason the simple beta factor model is widely used. More elaborate methods require more data. (Rasmuson and Mosleh) (Lilleheier)

Most of the data gathered has been for the beta factor model. Data has been gathered for offshore oil platforms (OREDA in Lilleheier), nuclear power plants (Mosleh et al., NUREG/CR-5485 in Stotta et al.)

The values of beta factors are remarkably similar across totally different systems and environments. (Rasmuson and Mosleh)

A. Typical betas

Various sources estimate the range of β. A survey of electrical equipment suggests that the best possible value of β is 0.01 while the worst is 0.30. (IEC 61508 in Lilleheier) Nuclear power plant data on thirteen types of

components showed β varying from 0.03 to 0.22, with an average of 0.10. (IAEA) Another report on the nuclear power industry indicated β of 0.01 to 0.20. (Rutledge and Mosleh) The value of β is between 0.001 and 0.05 in safety systems when good engineering practices are applied in design, installation, inspection, and maintenance, but β can be substantially higher, up to 0.25, with poor engineering. (Summers and Gentile) (Summers et al.) For hardware failures, β is in the range of 0.001 to 0.10. (Borcsok et al.) The consensus appears to be that β is 0.01 to 0.10 with good common cause failure prevention, and up to 0.25 for inadequate engineering.

B. Estimating beta

If data is unavailable, β can be estimated using various methods to assess engineering practice and common cause failure prevention. These include expert judgment, qualitatively estimating the impact of common cause problems, and the IEC 61508 method of quantitatively estimating the effect of a checklist of measures to reduce common cause failures. The IEC method requires detailed knowledge of design, software, testing, environment, operation, maintenance, and training. (Summers et al.) (Lilleheier)

X. Common cause failures in shuttle

Rutledge and Mosleh identified the dependent and common cause failures in all the space shuttle in-flight anomalies that occurred during the first forty flights after the Challenger accident. Of 473 anomalies, 54 (11%) were judged to be common cause failures, 6 due to functional interaction, and 4 due to spatial interaction, for a total of 64 (14%) dependent failures. The frequency of dependent and common cause failures is not significantly different from that found in nuclear power plants.

A. Shuttle dependant failures, root causes, and coupling factors

The dependent and common cause failures were distributed among the subsystems as shown in Table 1.

Table 1. Space shuttle subsystem dependent failures. (Rutledge and Mosleh)

Subsystem	#	%
Auxiliary power (e.g., hydraulic)	14	22
Propulsion (other than main engines)	13	20
Environmental control	7	11
Instrumentation (e.g., sensors)	7	11
Payload support subsystems (e.g., floodlights)	7	11
Electrical power	4	6
Guidance, navigation, control (e.g., radar altimeter)	3	5
Landing support systems (e.g., brakes)	3	5
Displays and controls	2	3
Life support and crew accommodations	2	2
Command data handling (e.g., computer)	1	2
Communication (e.g., antenna)	1	2
Structures and mechanisms (e.g., actuators)	0	0

Environmental control, life support, and crew accommodations combined have 9 dependent failures, 14% of the 64 total.

The in-flight anomalies included only a few actual multiple dependent or common cause failures. Most of the so-called dependent failures were only potential dependent failures, where the probability of multiple dependent failures appeared high. This widening of the definition is usual, and was necessary to obtain sufficient data for effective analysis.

The dependent failures were analyzed as to root causes and coupling factors. The most important are summarized in Tables 2 and 3.

Table 2. Important dependent failure root causes.

Root cause	#	%
Design error or inadequacy	11	17
Contamination	7	11
Moisture	4	6
Defective calibration or test procedures	6	9

Design error was the most important root cause, but contamination and moisture that affected fluid, mechanical, and electrical systems were together equally important.

Table 3. Important dependent failure coupling factors.

Coupling factor	#	%
Same part design	17	27
Same component location or environment	7	11
Same system design	6	9
Same supporting systems	6	9
Same component calibration characteristics	5	8
Same system interconnections	4	6

Sameness of many different kinds is the coupling mechanism that leads to dependent and common cause failures. Same part design and same system design account for 36% of dependent failures.

B. Defenses against dependant failures

Redundancy must be coupled with technical diversity and physical separation. It is easier to prevent dependent failures by defeating coupling factors rather than by eliminating root causes. The potential defenses were classified based on the coupling factors they prevent. Some specific approaches are:

1. Use diverse components in redundant sets.
2. Use separate locations for redundant components.
3. Connect the components in a redundant set to different supporting systems using diverse interconnection configurations

A set of detailed spacecraft system design guidelines was developed to help prevent dependent and common cause failures. The guidelines cover six different types of physical elements and describe the stressors, failure modes, and defenses.

Rutledge and Mosleh observe that by using the suggested design defenses to reduce common cause failures, "overall spacecraft reliability can be increased without increasing the reliability of the individual components, and that this probably can be done at little, if any, additional cost." With redundancy, diversity, and separation, system reliability can be increased without high reliability components. (Rutledge and Mosleh)

XI. Designing to reduce common cause failures

Rutledge and Mosleh (immediately above) and Summers et al. and Lilleheier (above in the subsection on estimating beta) discuss designing to reduce common cause failures. Bukowski and Goble verified three different design rules to reduce common cause failures by using simulating failures caused by stress. The three rules concern diversity, separation, and reducing total failure rate: design for diversity, reduce the common stress, and the lower the failure rate, the lower is the beta factor.

Ramirez-Marquez and Coit discuss the process of adding redundant components to improve reliability. They support the above rules of Rutledge and Mosleh and use simulations to optimize reliability while considering the system mass and cost. With no common cause failures, the optimum system uses identical copies of the lightest component available. (See also Jones) For even relatively few common cause failures, the best component choice is diversified, notably including heavy and higher reliability components. As the fraction of common cause failures (beta) increases further, the number of components is reduced and their diversity increases. In the absence of common cause failures, the optimum system design uses only the lightest components, with no diversity. When common cause failures are present, diversity is favored.

XII. With common cause failures, diversity produces lower failure rate and mass

If the subsystem technologies have common cause failures, diverse redundant technologies usually produce the lowest failure rate and mass. A simple mathematical model shows this.

Suppose there are very many different technologies that can perform the same function. Each technology has failure rate λ, common cause fraction β, and mass m. Suppose all the technologies have the same β, say 0.10. The technologies can be plotted against increasing failure rate λ and mass m in an X Y plot as in Figure 1. The best technologies are those with the lowest failure rate λ and mass m, which are down and to the left in the plot. The

lower left edge of the group of plotted points defines the "efficient frontier," which is shown as a solid line in Figure 1. Technologies on the efficient frontier are the best selections, since they have either the lowest mass for a given failure rate or the lowest failure rate for a given mass. Technologies above and to the right of the efficient frontier are "dominated," since they are inferior to other selections. A dominated technology can be replaced by one with both lower failure rate and lower mass.

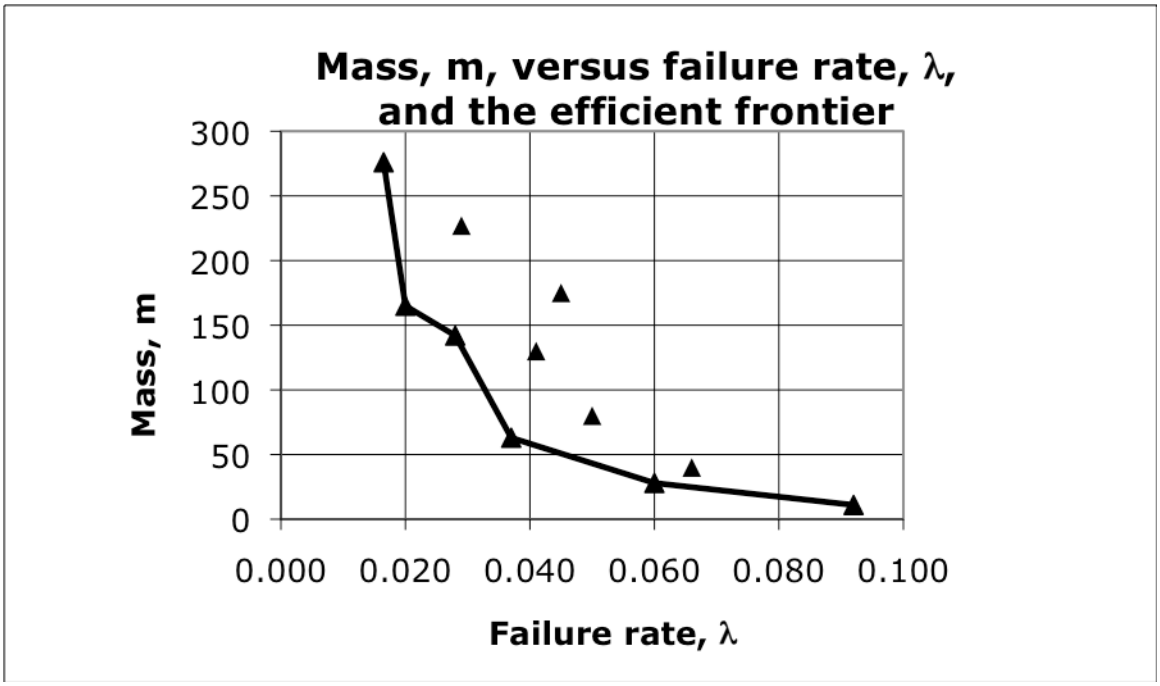

Figure 1. Mass versus failure rate on the efficient frontier.

All this goes to prove that there is no such thing as a free lunch. If lower failure rate is needed, more mass is required. If lower mass is needed, higher failure rate must be accepted. The mass-failure rate tradeoff is defined by the efficient frontier.

For example, suppose technology 1 has $\lambda_1 = 0.1$, $\beta_1 = 0.1$, and $m_1 = 10$, and technology 2 has $\lambda_2 = 0.01$, $\beta_2 = 0.1$, and $m_2 = 100$. Note that lower failure rate requires higher mass. There are three options for a dual redundant system, 11, 12, and 22. For 11, the failure rate $f_{11} = \lambda_1^2 + \beta_1*\lambda_1 = 0.01 + 0.01 = 0.02$, and $m_{11} = 20$. For 12, the failure rate $f_{12} = \lambda_1*\lambda_2 = 0.001$, and $m_{12} = 110$. For 22, the failure rate $f_{22} = \lambda_2^2 + \beta_2*\lambda_2 = 0.0001 + 0.001 = 0.0011$, and $m_{22} = 200$. The 12 dual redundant system has the lowest failure rate by a small margin and average mass.

A general formulation can show diverse redundancy produces the lowest failure rate and mass. Suppose mass is an inverse function of failure rate, m = Inverse Function (λ) = IF(λ). Suppose technology 1 has λ_1, β, and $m_1 = IF(\lambda_1)$, and technology 2 has λ_2, β, and $m_2 = IF(\lambda_2)$.

For 11, the failure rate $f_{11} = \lambda_1^2 + \beta*\lambda_1$, and $m_{11} = 2\ IF(\lambda_1)$. For 12, the failure rate $f_{12} = \lambda_1*\lambda_2$, and $m_{12} = IF(\lambda_1)+IF(\lambda_2)$. The failure fates are set equal and the relative mass is investigated. $f_{11} = \lambda_1^2 + \beta*\lambda_1 = f_{12} = \lambda_1*\lambda_2$, so $\lambda_1 + \beta = \lambda_2$. Then $m_{12} = IF(\lambda_1)+IF(\lambda_2) = IF(\lambda_1)+IF(\lambda_1 + \beta)$. If there are no common cause failures, $\beta = 0$, and $m_{12} = 2\ IF(\lambda_1) = m_{11}$. If the subsystem technologies do not have common cause failures, identical and diverse redundant technologies produce the same failure rate and mass. If there are common cause failures, $\beta > 0$, and $m_{12} = IF(\lambda_1)+IF(\lambda_1 + \beta) < m_{11}$. When the subsystem technologies do have common cause failures, two diverse redundant technologies will produce a given failure rate with lower mass than two identical technologies. It follows that, with common cause failures, two diverse redundant technologies produce a lower failure rate than two identical technologies having the same total mass.

The general rule that diverse redundant technologies produce both lower failure rate and lower mass than identical redundant technologies having common cause failures is based on the assumption that there are many diverse technologies not too far from the efficient frontier. If only one technology exists, diversity is simply not possible. Suppose two technologies exist, but that one dominates the other with both lower failure rate and lower

mass. If the failure rates and masses are similar, the technologies make a good diverse pair. But if one is vastly better than the other, two identical copies of it can be considered. The common cause failure rate, $\lambda\beta$, is the lower bound on the failure rate achievable using dual redundancy. Three or more units could be used to reduce the probability of failure, but the analysis requires a more complex model and data.

XIII. Need to develop several diverse technologies

Because of the need for ultra reliability, research should develop several different technologies for the same function, even if one is clearly better than the others. This is a significant change from the more usual approach of selecting one technology early so as not to waste effort in duplication. Selecting early saves current dollars but it narrows choices, eliminates options, and often costs more in the long run. Picking the winner is usually difficult in a fair and open competition. Developing only one approach because it seems best is a self-fulfilling prophecy. But the future is full of surprises. Developing options is good research, good management, and good politics. Because of common cause failures, diverse technologies are necessary to achieve ultra reliability.

XIV. Conclusion

If common cause failures are present, the failures are not independent as assumed in standard reliability analysis. Common cause failures can prevent achieving high reliability using redundancy. Design flaws or external shocks can make all the redundant units fail due to a common cause. Extensive failure mode analysis, redesign for higher intrinsic reliability, and use of technical diversity and physical separation are all needed to reduce common cause failures.

Common cause failures have been modeled as exceptions to the usual reliability analysis. It is assumed that X percent of all failures have common causes, so that they cannot be reduced by using redundancy. If X is small, corresponding to a rarity of common cause failures, the usual reliability analysis describes most of the failures and its reliability estimate is useful. But if X is large, common cause failures are frequent, and the usual reliability analysis gives too low failure predictions.

The key lesson is that redundancy does not reduce the rate of common cause failures. But common cause failures such as design flaws, manufacturing errors, or external overstresses are not acceptable. The entire program of safety and reliability analysis and design is an attempt to prevent all failures. Reliability analysis and the use of redundancy may fail because of common cause failures, but common cause failures occur because of engineering and design failures. Reliability analysis should both attempt to explicitly identify and remove common cause failure modes and to estimate the remaining level of common cause failures.

References

Borcsok J., Schaefer, S., and Ugljesa, E, "Estimation and Evaluation of Common Cause Failures," IEEE, Second International Conference on Systems (ICONS'07) 2007. Ref 04196343.pdf

Bukowski, J. V., and W. M. Goble, "Verifying common-cause reduction rules for fault tolerant systems via simulation using a stress–strength failure model," ISA Transactions 40, 2001.

IAEA Training Course on Safety Assessment of NPPs to Assist Decision Making.

Jones, H., "The Reliability-Mass Trade-Off in Multi-Criteria Life Support Technology Selection," AIAA-2011-5094, 41st ICES (International Conference on Environmental Systems), 2011.

Lilleheier, T., "Analysis of common cause failures in complex safety instrumented systems," July 28, 2008. www.ntnu.no/ross/reports/stud/lilleheier.pdf

NRC, Common-Cause Failure Database and Analysis System: Event Data Collection, Classification, and Coding, U.S. Nuclear Regulatory Commission, NUREG/CR-6268, Rev. 1 INL/EXT-07-12969, September 2007.

Pate-Cornell, E., and R. Dillon, "Probabilistic risk analysis for the NASA space shuttle: a brief history and current work," Reliability Engineering and System Safety 74, 2001.

R&M, Applied R&M Manual for Defence Systems, Part C – Techniques. (GR-77 Issue 2010), p. 2.

Ramirez-Marquez, J. E., and D. W. Coit, "Optimization of system reliability in the presence of common cause failures," Reliability Engineering and System Safety 92, 2007.

Rasmuson, D. M., and A. Mosleh, "A Brief History of Common-Cause Failure Analysis," IAEA Technical Meeting on CCF in Digital Instrumentation and Control Systems for Nuclear Power Plants, Bethesda, Maryland, June 20, 2007.

Rutledge, P. J., and A. Mosleh, "Dependent-Failures in Spacecraft : Factors, Defenses, and Design Implications," IEEE, 1995 Proceedings Annual Reliability and Maintainability Symposium.

Seife, C., "Columbia Disaster Underscores The Risky Nature of Risk Analysis," Science, vol. 299, 14 February 2003.

Smith, A. M. and Watson, I. A. (1980). Common cause failures - a dilemma in perspective. Reliability Engineering, 1:127–142. Quoted in T. Lilleheier, Analysis of common cause failures in complex safety instrumented systems, July 28, 2008. www.ntnu.no/ross/reports/stud/lilleheier.pdf

Stotta, J. E., P. T. Britton, R. W. Ring, F. Hark, and G. S. Hatfield, “Common Cause Failure Modeling: Aerospace vs. Nuclear,” http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20100025991_2010028311.pdf 20100025991_2010028311.pdf

Summers, A., and M. Gentile, “Random, Systematic, and Common Cause Failure: How do you manage them?” Process Safety Progress, December 2006.

Summers, A., K. A. Ford, and G. Raney, “Estimation and Evaluation of Common Cause Failures in SIS,” Presented at 1999 Loss Prevention Symposium, Houston, TX, March 1999, Published as ““Safeguard Safety Instrumented Systems,” in Chemical Engineering Progress, November 1999.