45th International Conference on Environmental Systems
12-16 July 2015, Bellevue, Washington

ICES-2015-047

Diverse Redundant Systems for Reliable Space Life Support

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

Reliable life support systems are required for deep space missions. The probability of a fatal life support failure should be less than one in a thousand in a multi-year mission. It is far too expensive to develop a single system with such high reliability. Using three redundant units would require only that each have a failure probability of one in ten over the mission. Since the system development cost is inverse to the failure probability, this would cut cost by a factor of one hundred. Using replaceable subsystems instead of full systems would further cut cost. Using full sets of replaceable components improves reliability more than using complete systems as spares, since a set of components could repair many different failures instead of just one. Replaceable components would require more tools, space, and planning than full systems or replaceable subsystems. However, identical system redundancy cannot be relied on in practice. Common cause failures can disable all the identical redundant systems. Typical levels of common cause failures will defeat redundancy greater than two. Diverse redundant systems are required for reliable space life support. Three, four, or five diverse redundant systems could be needed for sufficient reliability. One system with lower level repair could be substituted for two diverse systems to save cost.

Nomenclature

A	=	Number of units used in a life test
B	=	Length of a life test in units of mission duration D
C	=	Cost
D	=	Mission duration in years
F	=	Failure probability over the mission duration D
f	=	failure rate, number of failures per unit time, f = F/D
M	=	Number of components in a set of RU's
MTBF	=	Mean Time Before Failure, MTBF = 1/f
N	=	Number of redundant units
ORU	=	Orbital Replacement Units
R	=	Reliability, probability no failure occurs over the mission duration D, R = 1 - F
RC	=	Replacement Components
RU	=	Replacement Units

I. Introduction

HOW can highly reliable life support systems be provided for deep space missions? A future human mission to Mars will depend on life support operating continuously without consumables or spare parts from Earth. Space stations in Low Earth Orbit (LEO) such as Skylab, Mir, and the ISS (International Space Station) and even a future Moon base can operate safely with much less reliable life support. Human space habitats within the Earth-Moon system can receive emergency resupply of materials and equipment. In the worst case, the crew can return to Earth in a few days. Life support failures are much more dangerous in deep space than near Earth. There are no repair facilities or return options on the way to Mars.

Traveling to Mars and back will probably require several years, during which the crew will consume large amounts of water and oxygen. This life support direct supply mass could amount to roughly 20,000 kg for a crew of four. The cost of launching this much mass to LEO, moving it on to Mars, and then either sending it down or Mars or back to Earth could be billions of dollars. To avoid this excessive cost, oxygen and water probably will be recycled on all multiyear human space missions.

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

The ISS now has an operational life support recycling system that provides a strong foundation for all future systems. It has had more failures than expected and is being improved. Clearly, it would have been far too expensive and unnecessary to design ISS life support with the high reliability needed in deep space.

The high cost of directly supplying the oxygen and water for a Mars mission, and the unacceptably large risk in relying on an unimproved recycling system, justify a large research and development effort to develop highly reliable life support systems for deep space. How can this be done?

II. Overview - How to Provide Reliable Life Support

Most of the technical systems used every day are very reliable, but such good performance is typically achieved after decades of design, operation, failure analysis, and redesign. Many units and long operating times are required. It would cost too much and take too long to develop highly reliable life support for space this way. Often only one system is needed on the mission, and the space station life support designers had no need to build test units or run long systems tests.

A practical approach suggested for Mars is simply to use triple redundancy. (Connolly, 2000) Three identical systems are provided for each function. If one system has a reasonably low probability of failure over the mission duration, say one in ten, and if all the failures are independent and not due to a design flaw or common cause, then the probability that all three systems will fail is quite low, only one in a thousand. (The probability that the first system fails is 10%, the probability that the second system then fails is 10%, so the probability that both fail on the mission is 10% of 10% or only 1%. The probability that the third redundant system is needed and then fails is 10% of one percent, 0.1%, or 1 in a thousand.)

The space station life support system does not use full system level redundancy, but a modified version. Each system was designed so that most of the components more likely to fail are contained in an Orbital Replacement Unit (ORU). Since future missions will go beyond LEO, we consider simply a Replacement Unit (RU). An RU can provide most of the reliability improvement obtained by a full system spare and should cost less to build and fly. Designing the RU's so that they are easily replaced is an additional cost.

Providing RU's on a mission does incur significant design, production, and launch costs and it has been suggested that lower level repair or component replacement is better for a deep space mission. (Humphries et al., 1994) Replacing an RU can repair only a single failure while one set of components can replace many different failed ones. If all the failures are independent, a single set of components would give a much greater reliability improvement than an RU. However, the ability to replace components would require not only a stock of components, but tools, materials, test equipment, and workspace, all costly and difficult to provide. If the system has a one in ten probability of failure during the mission, it is 90% likely that no actual repairs would be needed and the expected cost in crew time would be low. Rather than doing detailed training for high proficiency in dozens of different repairs, in might be better to provide instruction if needed.

These approaches provide redundancy using the same system design and identical components. Identical system redundancy works well if all the failures are independent but not if they are due to shared causes or design flaws. To ensure that they do not fail for the same reason, systems must be of diverse design. On the ISS, both the US and Russian segments have fully functioning, always operating life support systems. ISS life support is used as an example of the need for diverse system redundancy. (NASA, 2013) If the individual system reliability is lower than wanted, redundancy is the solution. If common cause failures may be important, dissimilar systems or components are needed to gain the benefit of redundancy.

There are many possible ways to provide highly reliable life support, including improving the basic system reliability, adding redundant units, RU's, or components, and providing diverse redundancy. The best approach depends on the overall reliability achieved and costs incurred. The reliability is computed and costs are compared in detail.

III. Methods and Costs to Provide High Reliability Life Support Systems

Where are we in the effort to develop more reliable recycling systems for LEO, the Moon, and ultimately Mars? What must we do to get the reliability we need for Mars? The problem can be analyzed and solved generally in terms of the cost to provide higher reliability. Suppose the Mars mission will have the duration of D, about two or three years. Suppose that we can build a current design system for a cost of C, some tens of millions of dollars, and that the system is reasonably expected to have an initial failure probability of Fi, for example Fi = 0.1 or 10% over the mission duration D. We would want a much lower probability of failure for the mission, Fm, of say 0.001 on 0.1%. We develop the costs to improve reliability by improving the system or adding redundant units, ORU's or components. Initially it is assumed that all failures are independent. Later non-independent or Common-Cause Failures (CCF's) will be considered.

A. Cost of improving basic system reliability

As a general rule, the cost to develop a system is directly proportional to its Mean Time Between Failures (MTBF). (Rechtin, pp. 127-8) The MTBF is simply the inverse of the failure rate. Cost increases as achieved failure rate decreases, due to the investment in design, test, and failure analysis. If the original system probability of failure is Fi = 0.1 in D years, the original failure rate is fi = 0.1/D per year. The original MTBFi = 1/fi = 1/(0.1/D) = 10 D, ten times the mission duration. This seems good but in fact means that the system has a 10% chance of failure over the mission duration D. We want a much better mission probability of failure of Fm = 0.001 and a failure rate of fm = 0.001/D per year. The needed MTBFm = 1/fm = 1/(0.001/D) = 1,000 D, one thousand times the mission duration. The cost to reduce the system failure rate by 100 is 100 C, one hundred times the original system development cost. In general, if a system with cost C has an initial failure probability of Fi , a system with failure probability Fm over the same duration will have a cost of C Fi/Fm. The cost to build and fly one highly reliable system is 100C, which appears much too expensive.

1. Cost of improved reliability

To improve the failure rate from the initial failure rate fi to the final failure rate ff, the system cost increases by the ratio of the initial failure rate fi to the final failure rate, fi/ff. Since MTBF = 1/f, the system cost increases by MTBFf/MTBFi, the ratio of the increase in the MTBF.

B. Cost of adding redundant systems

Suppose that the original system probability of failure is Fi and we provide N redundant units on the mission. The mission probability of failure is Fm = Fi^{N}. In general, the level of redundancy needed is N = log (Fm/Fi). For Fi = 0.1 = 10^{-1} and Fm = 0.001 = 10^{-3}, N = 3. Triple redundancy reduces the failure probability from Fi = 0.1 to Fm = 0.1 * 0.1 * 0.1 = 0.001. The cost to build a system of three original units is only 3C, much less than the 100C to directly improve system reliability.

1. Required number of redundant systems

The final failure probability Ff is equal to the initial failure probability Fi raised to the Nth power, Ff = Fi^{N}. Therefore, the number of redundant systems N required to decrease the failure probability over the mission duration D from Fi to Ff is given by N = log (Ff)/log (Fi). Since the failure probability F is equal to the failure rate f times the mission duration D, F = f D, and N = log (ff)/log (fi).

C. Trading off system reliability and redundancy

For N = 2 and Fm = 0.001 = 10^{-3}, the two redundant systems should each have Fs2 = $\text{Fm}^{1/N}$ = $10^{-3/2}$ = 0.0316. The cost to improve a system from Fi = 0.1 = 10^{-1} to Fs2= 0.0316 = $10^{-3/2}$ is C (0.1/0.0316) = C ($10^{-1}/10^{-3/2}$) = $10^{1/2}$ C = 3.16 C. The total cost to improve reliability and build two units is 2 * 3.16 C = 6.32 C.

The original system probability of failure of Fi =0.1 was high enough for triple redundancy, N = 3, to provide the required final probability of failure of Ff = 0.001. Higher cost and lower failure probability will allow Fm = 0.001 to be achieved with N = 2. If quadruple redundancy is acceptable, N = 4 and the system reliability could be reduced below the originally assumed Fi. Fs4 = $\text{Ff}^{1/4}$ = $10^{-3/4}$ = 0.178. The required system failure probability can be allowed to increase from Fi = 0.1 to Fs4 = 0.178. The reduced cost for this lower reliability would be C ($10^{-1}/10^{-3/4}$) = $10^{-1/4}$ C = 0.56 C. The cost to build four units with this lower reliability is 4 * 0.56 C = 2.24 C.

D. Cost to develop N redundant full systems

Table 1 shows the cost of providing an Ff = 0.001 reliable system using different levels of redundancy, assuming a system with initial failure probability Fi costs C dollars.

Table 1. Cost of N redundant systems to provide overall probability of failure Ff = 0.001.

Level of redundancy, N	Single system probability of failure, Fs	Cost for one unit with failure probability Fs, in multiples of cost C for Fi =0.1	Cost for N units with failure probability Fs, in multiples of cost for Fi=0.1
N = log (Ff)/log (Fi)	$Fs = Ff^{1/N}$	Cost for one = C Fi/Fs	Cost for N = N C Fi/Fs
1	0.0010	100.00	100.00
2	0.0316	3.16	6.23
3	0.1000	1.00	3.00
4	0.1778	0.56	2.24
5	0.2512	0.40	2.00
10	0.5012	0.20	2.00

The required mission probability of failure, Ff, and the available system's probability of failure, Fs, together determine the required level of redundancy, N = log (Ff)/log (Fi). The required system probability of failure, Fs, and the currently available system's probability of failure, Fi, together determine the cost, Fi/Fs times the cost C for a system with failure probability Fi, to achieve the required Fs. The total cost to achieve Ff using N units each with probability of failure Fs is C N Fi/Fs. Five or ten redundant units meet the failure probability requirement with the minimum cost of 2 C.

E. Using Replacement Units (RU's) instead of redundant full systems.

The cost of providing redundancy can be reduced if smaller Replacement Units (RU's) are used instead of full systems. The purpose of using an RU is to repair system failures using less hardware than the full system The ideal RU would contain the components that have most of the failure modes and are responsible for most of the full system's probability of failure. It would also have a relatively small fraction of the full system mass and development cost. The RU would be designed to include all the more failure prone components, and the reliability of the remaining system would be improved. In the limit, assumed here for ease in analysis, the RU would contain all of the failure modes and failure probability of the full system. If the RU does contain all the system failure modes, proving an RU for backup has the same effect on overall reliability as providing a full redundant system. The above redundancy and reliability analysis applies exactly. The single system probability, Fs, would be the same for the full system and the RU. However the cost to develop and build an RU with a failure probability of Fi = 0.1 would be a fraction of C, the cost to build a full system with the same reliability. For illustration, assume that the cost of the RU is 0.2 C.

Table 2 compares the cost of N redundant full systems with the cost of one full system and N – 1 RU's, assuming that the cost of an RU is 20% of the cost of a full system.

Table 2. Cost of N redundant systems using one full system and N – 1 RU's with Ff = 0.001.

Level of system plus RU redundancy, N	Single system or RU probability of failure, Fs	Cost for one system or RU with failure probability Fs, in multiples of cost C for Fi =0.1	Cost for N units with failure probability Fs, in multiples of cost for Fi=0.1	Cost for one system and N – 1 RU's with failure probability Fs, in multiples of system cost for Fi=0.1
N =log(Ff)/ log (Fi)	$Fs = Ff^{1/N}$	Cost for one = C Fi/Fs	Cost for N = N C Fi/Fs	Cost = (0.8 + 0.2 N) C Fi/Fs
1	0.0010	100.00	100.00	100.00
2	0.0316	3.16	6.23	4.55
3	0.1000	1.00	3.00	1.40
4	0.1778	0.56	2.24	0.90
5	0.2512	0.40	2.00	0.72
10	0.5012	0.20	2.00	0.56

Since the cost for a full system with Fi = 0.1 is C, the cost of an RU with the same Fi is assumed to be 0.2 C. The cost of one system and N – 1 RU's all with Fi = 0.1 is C + 0.2 C (N – 1) = 0.8 C + 0.2 N C. Under the favorable assumptions made here, RU's could save half the cost of using triply redundant systems.

F. Using Replacement Components (RC's) instead of full systems or RU's

In addition to redundant full systems or RU's, Replacement Components (RC's) can be considered. Making the same favorable assumptions as for RU's, suppose that a full set of RC's contains all of the potentially failing components in the system and accounts for all of the failure probability. Then one set of components has the same failure probability as the full system or an all-failure-containing RU. Providing one all-failure-containing set of RC's would reduce the overall probability of failure much more than providing a full system or an all-failure-containing RU. The set of RC's would have much less cost than a full system or an RU. The overall cost impact is difficult to assess, since using RC's would require more tools, materials, test equipment, and workspace. The unlikely event (probability Fs) of actually replacing an RC would require much more crew time than replacing a full system or an RU.

Using sets of RC's for redundancy can be much more effective in improving reliability than using full systems or RU's. Each set of RC's can repair multiple failures, instead of only the one repaired by replacing a full system or RU. Consider a system with failure probability Fs. Suppose that a set of M RC's contains all the failure modes. The system can be considered a failure free platform holding M components, each with some failure probability. If a component fails, it can be replaced as long as spares remain.

The system failure probability Fs is the sum of the M component failure probabilities, FRCj, where j is an index counting from 1 to M. Suppose the best case, that the FRCj are all equal to Fs/M. If there are N redundant sets of RC's, the probability that any single component will fail N times and cause overall system failure is $(Fs/M)^N$. The total probability that any one of the M components will fail N times and cause overall system failure is $M\ (Fs/M)^N$. This is much lower than the probability of overall system failure with N systems or one system plus N - 1 RU's, which is Fs^N. The failure probability for one system plus N - 1 sets of RC's is lower by M^{N-1}. For only 10 components, M = 10, and one system plus two sets of RC's, N = 3, the failure probability improvement of RC's over full systems or RU's is a factor of 100.

For example, suppose that Fs = 0.1. We use one system and one set of RC's. If M = 10 and the FRCj are all equal, FRC = Fs/M = 0.01 and the overall system probability of failure is $10\ (0.01)^2 = 0.001$, which is the originally required Ff. One system and one set of RC's with Fs = 0.1 performs as well as three redundant systems with Fs = 0.1, both achieving Ff = 0.001.

However, assuming all the FRCj are equal is the best case. The worst case occurs when only one component can fail, so it accounts for the entire system's probability of failure. Essentially, there is only one RC and M = 1. Replacing the one fallible component is like replacing a full system or an RU. Only N identical failures can be repaired.

In practice, using RC's should give significant reliability gain. Observed failure rates often approximate the Pareto distribution, which also describes personal incomes. Very many people have a low but non-zero income, some have slightly more income, fewer have moderately higher income, and only very few have very high incomes. Similarly, many components have low failure rates, and fewer and fewer have higher and higher rates. This occurs for failure rates because the more frequent failure modes tend to be observed and eliminated. The design of a system using RC's, including the number of spares of each RC, would be based on the individual RC failure rates.

G. Using one full system, one RU, and one set of RC's

If the initial system design has a probability of failure of Fi = 0.1, the desired overall final probability of failure of Ff = 0.001 can be achieved by using three redundant systems, or one full system and two RU's, or one full system and one or two sets of RC's. The best overall approach would seem to be using one full system, one RU, and one set of RC's. The problem with using one full system and RC's is that the repair of the system might take longer than could be tolerated for a life support system. Having one RU would provide a relatively quick return to service and allow less urgent repair of the failed and removed RU. A set of RC's should provide more reliability improvement than a full system or RU, but require more supporting infrastructure.

H. Practical problems using redundancy

Redundancy works well in theory but is likely to fail in practice because it relies on two key assumptions that are often violated. The first key assumption made but not explicitly stated was that the correct failure rates were known. Usually working failure rates are based on analysis rather than test. The component failure rates are summed up and it is assumed that good design and engineering introduce no degradation or additional failure modes. The estimated failure rates are usually acceptably low but experience often shows that the actual failure rates are orders of magnitude higher than estimated. (Likens, 1992) Measuring the failure rate is considered below.

The second key assumption that was explicitly stated was that the failures were independent and uncorrelated. Failures can be dependent and correlated, in that they depend on each other or the same event or mechanism. An example would be a design flaw, procedure error, or failed component that damages another component. The damaged component may fail, be replaced, and be damaged and fail again until all redundant replacements are used. Such non-independent failures are usually called Common Cause Failures (CCF's). CCF's are considered below.

IV. Why do redundant systems fail unexpectedly?

The two reasons that systems using redundant units can fail more frequently than anticipated are either that the reliability of the redundant units is lower than expected or that several redundant units all fail for the same unanticipated reason, a Common Cause Failure (CCF). In practice, a redundant system failure may involve both unexpected high failure rates and CCF's. An unexpectedly high component failure rate due to a manufacturing problems or an excessively stressful design would probably be called a CCF. Nevertheless, it is important to distinguish between these two causes because their remedies are different. Life testing can help prevent unexpected high failure rates. Careful design and the use of diverse systems can reduce CCF's. But just as the two causes often combine in a redundant system failure, the two cures both reduce both causes. Careful design will reduce the total failure rate, not only that due to CCF's. Extensive life testing will limit the failure rate due to CCF's, as well as the total failure rate.

It can be argued, sometimes it is assumed, that it is simple and easy to develop verified highly reliable redundant systems. Only two things are required, thorough and careful design to reduce the failure rate and long duration multi-unit life testing to demonstrate the low failure rate has been achieved. Reducing the failure rate and measuring a low failure rate can both be impossibly expensive. Since there are exponentially increasing costs for reducing and then measuring a low failure rate, the reliability of redundant systems is limited by their acceptable cost. The objective of this cost-reliability analysis is to understand how to achieve the highest reliability for a given cost, to determine the cost-reliability trade-off. The three fundamental engineering problems are reducing the failure rate, measuring the failure rate, and reducing CCF's.

V. Testing to Measure the Failure Rate

The effect of failure probability errors on performance of systems using redundancy is discussed. The multiunit life testing needed to estimate failure probability is described.

I. Redundancy calculations are highly sensitive to failure rate errors

Any error in the system failure probability has an exaggerated effect on the overall failure rate of a redundant system. Suppose it is thought that Fs = X, but actually Fs = e X, where e is an error factor multiple. For N redundant systems, $Ff = Fs^N$, the expected $Ff = X^N$ but actually $Ff = (e\ X)^N = e^N X^N = e^N$ times the expected Ff. For N = 3, an error factor of e = 2 in Fs becomes an error factor of $e^N = 2^3 = 8$ in Ff.

It was shown that if Fs = 0.1 and Ff= 0.001 is required, N = 3 = log (Ff)/log (Fi) is sufficient. Suppose Fs may be twice as high, Fs = 0.2. Then N = log (0.001)/log (0.2) = 4.3 and N = 5 must be used to meet Ff < 0.001. If we use only N = 3, $Ff = (0.2)^3 = 0.008$, 8 times higher than the requirement.

J. Multi-unit long duration testing is required to measure a low failure rate

The amount of life testing required to measure a failure rate depends on the actual failure rate. Suppose that the estimated system probability of failure over the mission duration D is Fs = 0.1. To measure this failure rate reasonably accurately, roughly ten failures would be needed. Testing A units for B times the mission duration would be expected to produce A*B*Fs failures. For A * B * 0.1 = 10, A*B= 100. Measuring Fs accurately requires testing 100 units over the mission duration D or 50 units over twice the mission duration, etc. This is appears unaffordable.

It could be sufficient to show only that Fs is probably less than 0.1. Suppose we test A systems for B times the mission duration D. In the best case the system no failures will occur. What is the estimated upper bound failure probability F based on the successful test? Since no failures occurred, we can assume that the failure probability is less than ½ for A units tested over B times D years, so the failure probability Fs < (½) [1/(A B D)] per year and the failure probability is Fs < 1/(2 A B) over the mission duration D. Since we want to show Fs $\leq$ 0.1, the test must have 2 A B = 10, A B = 5. Roughly bounding Fs < 0.1 requires testing only 5 units over the mission duration D.

Bounding the failure probability this way is probably the minimum possible testing. Even if Fs = 0.1 exactly, there is a 50% probability that one failure will occur while testing 5 units over one mission duration. If one failure occurs, the measured Fs = 0.2 instead of 0.1, and the performance of redundant systems would be expected to be significantly worse.

Testing to measure or roughly bound the probability of failure requires much more testing for lower failure probabilities. The number of units and the cost of testing for the different system failure rates of Table 1 are shown in Table 3.

Table 3. Number and cost of test units for the different system failure rates.

Level of redundancy, N	Single system probability of failure, Fs	Number of units A for testing over one mission duration	Cost for one unit with failure probability Fs, in multiples of cost C for Fi =0.1	Cost for A units with failure probability Fs, in multiples of cost for Fi=0.1
N = log (Ff)/log (Fi)	$Fs = Ff^{1/N}$	A = 1/(2 Fs)	Cost for one = C Fi/Fs	Cost for N = N C Fi/Fs
1	0.0010	500	100.00	50,000.00
2	0.0316	16	3.16	50.56
3	0.1000	5	1.00	5.00
4	0.1778	3	0.56	1.68
5	0.2512	2	0.40	0.80
10	0.5012	1	0.20	0.20

The cost of higher reliability, lower redundancy test units even for minimal testing greatly exceeds the cost of developing flight units shown in Table 1. The greatly increased cost of testing high reliability systems, much greater than the cost of developing them, makes them very unattractive. Design for high overall liability seems forced to use low reliability systems with high multiples of redundancy.

If RU's are employed, the same number of RU units as in Table 3 for full systems would have to be tested to establish their similar failure rate, since a single RU's would have a failure rate not much lower than a full system. It would be useful to also test the rest of the system not in the RU to further confirm its expected much lower failure rate. The best approach is simply to test full systems containing RU's. The number tested would be the same as in Table 3 for systems not employing RU's.

RC's would have failure rates that are orders of magnitude lower than the full system failure rate, and so would require testing of orders of magnitude more units to confirm the failure rate. If the RC's are off-the-shelf commercial components selected for high reliability, additional life testing is probably not needed. If the RC's are actually not purchased parts but small custom assemblies, something like mini RU's, then extensive testing would be required to establish their failure rates.

K. Combined development and test costs

Table 4 combines the cost of developing N redundant systems with Ff = 0.001 per system from Table 1 with the cost of testing from Table 3. It is assumed that a system with initial failure probability of Fi = 0.1 costs C dollars.

Table 4. Cost to develop and test N redundant systems to provide overall probability of failure Ff = 0.001.

Level of redundancy, N	Single system probability of failure, Fs	Cost for one unit with failure probability Fs, in multiples of cost C for Fi =0.1	Number of units A for testing over one mission duration	Cost for N + A units with failure probability Fs, in multiples of cost for Fi=0.1
N = log (Ff)/log (Fi)	$Fs = Ff^{1/N}$	Cost for one = C Fi/Fs	A = 1/(2 Fs)	Cost for N + A = (N + A) C Fi/Fs
1	0.0010	100.00	500	50,100.00
2	0.0316	3.16	16	56.88
3	0.1000	1.00	5	8.00
4	0.1778	0.56	3	3.92
5	0.2512	0.40	2	2.80
10	0.5012	0.20	1	2.20

A redundancy level of N = 3, 4, or 5 seems reasonable.

VI. Common Cause Failures (CCF's) and their effects

Accident analyses and failure investigations can resemble witch hunts. The justification for assigning blame is the unrealistic expectation of perfection. Good engineering should have guaranteed success. An internal investigation often claims that the failure was unpredictable, unpreventable, a one-in-a-million chance. Or blame is assigned to the lowest level person who might possibly have prevented the accident. External investigations usually identify a sequence of errors leading to the failure and often find that imperfect management is ultimately at fault. All actual failures have specific causes, so usually blame can be assigned and corrective action taken. This does not mean that the system will not fail again for some other reason. Responsible engineering must recognize that failures will occur and must be mitigated to produce the required reliability.

Consider the Fukushima Daiichi nuclear reactor disaster. An earthquake-caused tsunami flooded the reactors and removed all the sources of operating power, the three nuclear generators themselves, off-site power grid connections, backup generators, and batteries. This was thought impossible by most, but the event and the historical record examined later showed otherwise. It was found that this accident "could and should have been foreseen and prevented." (Wikipedia, Fukushima) This was clearly a classic common cause failure that defeated multiple diverse redundant systems and was triggered by an event assigned an unrealistically low probability. Surprising oversights do occur, as also shown by the Apollo 1 fire, which was caused by a dangerous pure oxygen atmosphere.

A. Common cause failure (CCF) analysis

CCF's are well studied and the standard beta factor approach is used here. (Jones, 2012) All technical systems fail sooner or later. All systems have some failure rate over time. In well-designed systems, most of the observed failures are random, unpredictable, and have acceptably low probability. However, often some CCF's occur repeatedly for the same reason. CCF's with high probability and an identifiable cause clearly call for redesign, Anticipated but unidentified CCF's can be mitigated using diverse system redundancy.

The expected fraction of all failures due to CCF's is designated by the Greek letter beta, β. If the overall system failure rate is f, the failure rate due to CCF's is β f. This is the beta factor model for CCF's. The fraction of failures that are CCF's can vary from very small up to a quarter or even a half of all failures. β then ranges from 0 to 0.25 or 0.5. Systems with lower failure rate f tend also to have lower β. Careful design, test, and redesign tend to reduce f and β at the same time. A typical average β is 0.1. (Jones, 2012)

B. What are common cause failures?

The model usually used in computing the effects of redundancy assumes that failures are independent in cause and random in time. This model ignores CCF's, which are not independent, but correlated to each other through some root cause, and are not random in time, but simultaneous or closely sequenced after the same root cause. The use of redundancy mitigates the effect of random failures but not of CCF's. CCF's defeat redundancy. If one system fails due to a CCF, all redundant copies of the identical system may well fail. The practical definition of a CCF is not that the failure can or does repeat, but that it can not be mitigated by the use of redundant systems.

Some types of failures are clearly CCF's. Suppose a system is poorly designed, damaged in manufacturing, or just breaks so that it causes some component to fail. All replacement components will also fail for the same reason, unless the cause is detected and repaired. Suppose a system is exposed to an unspecified, excessively high input voltage and is damaged. All identical replacement systems will be similarly damaged. Suppose the software in a system controller has a failure-causing bug. All copies of the software will fail. Design flaws, workmanship errors, failures in connected systems, improper interfaces, and software bugs are familiar sources of CCF's. All software bugs are design errors and all cause CCF's, since all redundant copies of the software will fail the same way. These examples show how CCF's defeat redundancy.

Some failures seem like CCF's but do not meet the operational definition of defeating redundancy. Suppose that a system has a failure probability of Fs and that extensive long duration testing confirms this failure probability and in the process produces many failures. And further suppose that all the failures are due to the same part failing in the same way. There are many failures, one cause, and good reason to redesign or replace the failing part. How is this not a CCF? N redundant systems will have the overall failure probability of Fs^N, as expected assuming no CCF's. Redundancy is not defeated. Also the failures are random in time, not simultaneous or subsequent to the same triggering event. Even though the failures have the same logical cause, low reliability, they do not have the same precipitating cause. One failure occurring does not change the expected time until a second failure. The failures are random failures with an unusually high probability, not CCF's. Notwithstanding this logic, most failure analysts would still consider this example a CCF. And, if the system design used redundant RC's rather than redundant full systems, redundancy would be defeated! Redundant RC's would produce only the same reduction in failure probability as redundant full systems, not the expected larger reduction. Since CCF's are due to design errors, unexpected stresses, and complex contingencies, they are difficult to prevent or predict. The probability of CCF's, the beta factor, is much more difficult to measure and predict than the overall failure rate. Conservative design would allow for a large fraction of CCF's. Predicting the performance of a redundant design with CCF's is uncertain because the beta factor is difficult to estimate.

C. The redundancy defeating effect of common cause failures

If a CCF occurs, the same cause is expected to produce the same failure in all identical redundant units. The overall failure probability obtained using redundant units does not decrease as much as expected. Suppose we use three redundant systems with Fs = 0.1 and there are no CCF's so βs = 0. Then the overall failure probability Ff = Fs^3 = 0.001. Suppose βs = 0.1. We assume that a CCF, if it occurs, will render all three redundant systems inoperative. The probability that the first unit operates throughout the mission is (1-Fs) = 0.9, and the probability that it fails is Fs = 0.1. The probability that the first unit failure is random is Fs (1 - βs) = 0.09 and the probability that it is a CCF is Fs βs = 0.01. At this point, there is a 0.01 probability that a CCF occurred and all the redundant units fail so that the overall system fails. Given that the first unit failed, the probability that the second unit completes the mission is (1-Fs) = 0.9, and the probability that it fails is Fs = 0.1. Again, the probability the failure is random is 0.09 and the probability that it is a CCF is 0.01. Again, if a CCF occurs the whole redundant system fails. At this point there is a 0.09 * 0.09 = 0.0081 probability that two successive random failures have occurred and the third unit is placed in operation, a 0.01 probability of a CCF in the first unit, and a 0.09 * 0.01 = 0.0009 probability of a CCF in the second unit. The probability that the third unit is placed into service after two previous random failures and itself fails is 0.09 * 0.09 *0.1 = 0.00081.

The total probability of failure is 0.01 for a CCF in the first unit, 0.0009 for a CCF in the second unit, and 0.00081 for any failure in the third unit. The total failure probability is 0.01171 but the 0.01 for a CCF in the first unit dominates the failure probability. Even the lower probability of a second unit CCF is larger than the total probability of random failures disabling the triple redundant system. (Checking, the probability of success is 0.9 that the first unit works, 0.09 * 0.9 = 0.081 that the first unit failure is random and the second unit works, and 0.09 * 0.09 *0.9 = 0.00729 that the first two unit failures are random and the third unit works. The total is 0.98829 = 1 - 0.01171.)

1. Effective redundancy

For N = 3, Fs = 0.1, and no CCF's so βs = 0, Ff = Fs^3 = 0.001. For N = 3, Fs = 0.1, and βs = 0.1, Ff = 0.01171. CCf's have significantly reduced the reliability gain produced by redundancy. For N = 2, Fs = 0.1, and βs = 0, Ff = Fs^2 = 0.01. Having CCF's with βs = 0.1 makes an N=3 system perform slightly worse than an N = 2 system.

The effect of CCF's can be quantified using the concept of effective redundancy. For N level redundancy and no CCF's, Ff = Fs^N. N = log Ff (βs=0)/log Fs. CCF's reduce the Ff achieved for each level of redundancy. We can define the effective redundancy with CCF's as N-effective = log Ff (βs) /log Fs. For N = 3, Fs = 0.1, and βs = 0.1, Ff (βs) = 0.01171. N-effective = log Ff (βs) /log Fs = -1.93/-1 = 1.93. With βs = 0.1, N =3 redundancy is les effective than N= 2 redundancy for βs = 0.

2. General formula for the overall failure probability, Ff, given level of redundancy, N, system failure probability, Fs, and fraction of failures that are CCF's, βs,

The calculation of the total failure probability, Ff, for a system with N level redundancy, a unit failure probability Fs, and fraction of CCF's of βs follows the steps used above. The first unit has failure probability Fs and a CCF failure probability of Fs βs. If the first fails, the second has the same, and so on. The total failure probability for N redundant units is

$$Ff = Fs\ \beta s + Fs\ (1 – \beta s)\ \beta s + [Fs\ (1 – \beta s)]^2\ \beta s + \ldots + [Fs\ (1 – \beta s)]^{N-1}\ \beta s + [Fs\ (1 – \beta s)]^{N-1}\ Fs$$

The term before the last term represents a CCF in the N-1 th unit. The last term represents random failures in the first N-1 units and any failure in the Nth unit. Using a formula for the sum of a series,

$$Ff = Fs\ \beta s\ \{1 - [Fs\ (1 – \beta s)]^{N-1}\}/\ \{1 - [Fs\ (1 – \beta s)]\} + [Fs\ (1 – \beta s)]^{N-1}\ Fs$$

For large N,

$$Ff = Fs\ \beta s/\{1 - [Fs\ (1 – \beta s)]\}$$

The probability of a CCF in the first unit. which is Fs βs, dominates and the additional redundant units have small effect.

3. Table of total failure probability Ff, given N, Fs, and βs,

Table 5 gives the total failure probability, Ff, for a system with N level redundancy, a unit failure probability Fs, and fraction of CCF's of βs.

Table 5. The total failure probability, Ff, for N level redundancy, a unit failure probability Fs, and fraction of CCF's of βs.

Redundancy, N	System failure probability, Fs	CCF beta, βs			
		0.001	0.010	0.100	0.200
N =1	0.001	0.00100	0.00100	0.00100	0.00100
	0.010	0.01000	0.01000	0.01000	0.01000
	0.100	0.10000	0.10000	0.10000	0.10000
	0.200	0.20000	0.20000	0.20000	0.20000
N = 2	0.001	0.00000	0.00001	0.00010	0.00020
	0.010	0.00011	0.00020	0.00109	0.00208
	0.100	0.01009	0.01090	0.01900	0.02800
	0.200	0.04016	0.04160	0.05600	0.07200
N = 3	0.001	0.00000	0.00001	0.00010	0.00020
	0.010	0.00001	0.00010	0.00101	0.00202
	0.100	0.00111	0.00208	0.01171	0.02224
	0.200	0.00822	0.01024	0.03008	0.05152
N = 4	0.001	0.00000	0.00001	0.00010	0.00020
	0.010	0.00001	0.00010	0.00101	0.00202
	0.100	0.00021	0.00121	0.01105	0.02178
	0.200	0.00184	0.00403	0.02541	0.04824
N = 10	0.001	0.00000	0.00001	0.00010	0.00020
	0.010	0.00001	0.00010	0.00101	0.00202
	0.100	0.00011	0.00111	0.01099	0.02174
	0.200	0.00025	0.00249	0.02439	0.04762

For N = 1, Ff = Fs. For N higher, Ff = Fs^N for βs = 0, but Ff increases rapidly for higher βs. The large reductions in failure probability achieved by high redundancy are lost if there is a significant fraction of CCF's.

4. Table of effective redundancy N-effective, given N, Fs, and βs,

The effect of CCF's can be measured by the effective redundancy, N-effective. Table 6 gives the effective redundancy, N-effective, for a system with N level redundancy, a unit failure probability Fs, and fraction of CCF's of βs.

Table 6. The effective redundancy, N-effective, for N level redundancy, a unit failure probability Fs, and fraction of CCF's of βs.

Redundancy, N	System failure probability, Fs	CCF beta, βs			
		0.001	0.010	0.100	0.200
N =1	0.001	1.00	1.00	1.00	1.00
	0.010	1.00	1.00	1.00	1.00
	0.100	1.00	1.00	1.00	1.00
	0.200	1.00	1.00	1.00	1.00
N = 2	0.001	1.90	1.65	1.33	1.23
	0.010	1.98	1.85	1.48	1.34
	0.100	2.00	1.96	1.72	1.55
	0.200	2.00	1.98	1.79	1.63
N = 3	0.001	2.00	1.67	1.33	1.23
	0.010	2.48	2.00	1.50	1.35
	0.100	2.96	2.68	1.93	1.65
	0.200	2.98	2.85	2.18	1.84
N = 4	0.001	2.00	1.67	1.33	1.23
	0.010	2.50	2.00	1.50	1.35
	0.100	3.68	2.92	1.96	1.66
	0.200	3.91	3.43	2.28	1.88
N = 10	0.001	2.00	1.67	1.33	1.23
	0.010	2.50	2.00	1.50	1.35
	0.100	3.95	2.95	1.96	1.66
	0.200	5.15	3.72	2.31	1.89

For N = 1, there is no redundancy and N-effective = 1. For N higher, N-effective = N for βs = 0, but N-effective decreases rapidly for higher βs. The total failure reduction of high redundancy is lost if there is a significant fraction of CCF's.

Tables 5 and 6 show that for N = 3, Fs = 0.010, and βs = 0.010, Ff = 0.0001 and N-effective = 2. The probability that the first unit has a CCF is Fs βs = 0.0001. The possibility of a CCF in the first unit accounts for essentially the entire failure probability. (Checking, the probability that the first unit operates throughout the mission is 0.99. The probability that the first unit fails randomly (no CCF) and the second unit finishes the mission is 0.0099 * 0.99 = 0.0098. The probability that the first two units fail randomly and the third completes the mission is 0.0099 * 0.0099 * 0.99 = 0.0000009. The total probability of no failure is the sum, 0.9999.) Since Ff = Fs^2, N-effective = 2.

The fact that CCF's defeat N large redundancy is highlighted by the observation that, for Fs = 0.010 and βs = 0.010, Ff = 0.0001 and N-effective = 2 for N = 3, 4, and 10, actually for N = 3 and higher. In general, N-effective ≤ 2, if βs ≥ Fs. The requirement βs ≥ Fs ensures that the first unit CCF failure probability of Fs βs ≥ Fs^2, so N-effective ≤ 2. This is important to design because βs is typically 0.1 and thus is often higher than Fs, and it follows that redundancy levels higher than N = 2 are often useless.

VII. Using diverse redundancy to mitigate the effect of common cause failures

If two systems designed for the same function are completely different in concept, architecture, subsystems, and components, the two systems would probably have no internally generated common cause failures. They may or may not fail due to the same unanticipated external condition, an externally initiated common cause failure. If we assume that diverse design prevents CCF's, the failure probability of a redundant set of diverse systems will simply be the product of the individual system failure probabilities, Ff = Fs1 *Fs2 * … * FsN. To minimize the development cost, all the individual system failure probabilities should be made roughly equal, so Ff = Fs1 *Fs2 * … * FsN = Fs^N, just as computed for identical systems with no CCF's.

The cost of developing and testing N technically diverse systems would be somewhat more than the cost of developing N identical systems. However, the cost saving to develop additional copies of the same design was not considered in the cost estimates above. This learning curve cost saving for identical systems is relatively minor compared to the large cost differences that depend on system failure probability.

A. Using N diverse redundant systems

The simplest and most direct design is to use N diverse redundant systems all having the same Fs. The total failure probability is Ff = Fs^N. As before, we set Ff = 0.001. Fs and N can vary in the design. The cost analysis in Tables 1 and 3 applies directly to diverse redundant systems. Table 1 shows the cost of providing an overall failure probability of Ff = 0.001 using N levels of redundancy, assuming that a system with initial failure probability Fi = 0.1 costs C dollars. Using N = 3 or 4 gains most of the cost savings of using redundant lower reliability units. Table 3 shows the number of test units needed to roughly measure the different system failure probabilities, Fs, required at different N, and the cost for these units. Test costs are a significantly lower for N = 3, 4, 5 and 10.

Table 7 combines the costs of development and testing N diverse redundant systems for an overall probability of failure Ff = 0.001. Table 7 shows higher total costs than Table 4 because each of the N diverse designs requires independent testing. The number of units in Table 7 is N (1 + A), not N + A as in Table 4. The costs in Table 7 are roughly double those in Table 4 for N equal to 2 or more.

Table 7. Cost to develop and test N diverse redundant systems for an overall probability of failure Ff = 0.001.

Level of redundancy, N	Single system probability of failure, Fs	Cost for one unit with failure probability Fs, in multiples of cost C for Fi =0.1	Number of units A for testing over one mission duration	Cost for N(1 + A) units with failure probability Fs, in multiples of cost for Fi=0.1
N = log (Ff)/log (Fi)	Fs = $Ff^{1/N}$	Cost for one = C Fi/Fs	A = 1/(2 Fs)	Cost for N(1 + A) = N(1 + A) C Fi/Fs
1	0.0010	100.00	500	50,100.00
2	0.0316	3.16	16	107.59
3	0.1000	1.00	5	18.00
4	0.1778	0.56	3	8.96
5	0.2512	0.40	2	5.97
10	0.5012	0.20	1	3.99

Again, a redundancy level of N = 3, 4, or 5 seems reasonable.

B. Using pairs of identical redundant systems

The analysis of the redundancy defeating effects of CCF's showed that, for a typical βs = 0.1 and Fs = 0.1 to 0.2, with N $\geq$ 3, N-effective can be as high as 2. For βs = 0.1, Fs $\geq$ 0.1, and N = 2, N-effective is 1.72. There is some benefit in using N = 2 identical systems even with the expected level of CCF's. A possible design can use two pairs of systems, with the two in each pair identical to each other but the two pairs pair using diverse technology. This approach could save cost by developing and testing only two different designs and by reducing one unit in each pair to an RU. The cost savings could be 40% if as previously assumed an RU costs only 20% of the full system cost.

However the redundancy reducing effect of CCF's requires all systems to be designed to a lower failure probability, Fs. Reducing the failure probability increases the system cost. However, having only two rather than four different designs reduces the number and cost of test units. These effects roughly cancel each other.

Suppose as before that an overall failure probability of Ff = 0.001 is required. If we use four diverse systems with failure probability Fs, Ff = Fs^4, and Fs = $10^{-3/4}$ = 0.178. If we use two diverse pairs of identical units, all with failure probability Fs and βs = 0.1, the failure probability of each pair is Fs βs + Fs (1 - βs) Fs. The overall failure probability of the two pairs is Ff = $[Fs\ \beta s + Fs\ (1 - \beta s)\ Fs]^2$. Setting Ff = 0.001 and βs = 0.1 and solving by trial and error, Fs = 0.140 for two diverse pairs of identical units. Achieving this is more costly than achieving the higher Fs = 0.178 required when all four systems are diverse. The design cost increase is proportional to failure rate decrease, 0.178/0.140 = 1.27 or 27%. Using the computation of Table 1, the higher cost a system with lower failure probability Fs = 0.140 is (0.1/0.14) C = 0.71C, where C is the cost for a system with Fi = 0.1.

Using the computation of Table 3, the number of units A for testing over one mission duration increases for a lower failure rate. The number of test units per design increases from A = 3 for a system with the higher Fs = 0.178 to A = 1/ (2*0.140) = 3.6, so 4 would be needed for each design, and 8 total for the two designs.

Each set of two identical systems with the higher Fs = 0.178 requires two flight units and four test units, six total. Two diverse sets of two identical systems would require 12 units, for a total cost of 8.57 C.

From Table 7,the cost to design and test four diverse systems with failure probability Fs = 0.178 is 8.96 C.. Designing and testing two diverse pairs of identical units actually costs a little less than designing and testing four diverse systems. The use of replaceable subsystems or sets of components instead of a second identical system could significantly reduce cost.

VIII. Conclusion

This analysis considered different approaches to provide the reliable life support systems needed for long missions in deep space or to Mars. It is prohibitively expensive to develop and test a single very highly reliable system to provide the required overall failure probability. It is much more feasible to achieve the required overall failure probability using three, four, or five identical redundant systems. Each of the redundant systems would have about a hundred times higher failure probability and about one one-hundredth the total development and test cost of a single high reliability system. This is because the system cost and number of test units both decrease at the same rate as the failure probability increases. The use of partial system replacement units, RU's, should further reduce cost. The use of replaceable components, RC's, would improve reliability much more than a single unit or RU. RC's allow one set of components to repair multiple failures but require more tools, materials, workspace, planning, and training. Crew time might be required for repairs on the mission but probably not even one unit would fail. The probably unused redundant units are needed to make the probability of a complete failure very small.

New designs may have failures due to specification, design, and manufacturing problems. Such failures can disable all the identical redundant systems. These common cause failures, CCF's, defeat redundancy. CCF's typically account for one tenth of all failures. If there is the usual significant fraction of CCF's, adding more than two redundant identical units would give very little reliability improvement.

The most feasible way to achieve high reliability systems is to use diverse designs that reduce or eliminate internal system CCF's. External events and unintentional commonalities may still cause CCF's. Probably good designs would use three, four, or five diverse redundant systems. Cost could be reduced if two identically designed systems, one full system with replaceable subsystems or sets of components, are substituted for two diverse design systems.

Developing and testing multiple diverse life support systems for Mars will require a large, long-term program. It would be best to develop and test the diverse systems using different management and engineering organizations. Some should be in different countries. Perhaps twice as many system designs should be begun as we plan to use, since an excessive failure rate can be compensated by increased redundancy. Life testing of multiple units should be conducted on Earth to establish the system failure rates. This may be adequate for Mars surface systems, but Mars transit systems also require microgravity testing. If testing is too limited to confirm the wanted high reliability, the actually confirmedß lower reliability should be used to determine the needed level of redundancy.

References

Connolly, J. F., "Mars Design Example," in Larson, W. K., and Pranke, L. K., eds., Human Spaceflight: Mission Analysis and Design, McGraw-Hill, New York, 2000.

Humphries, W. R., P. K. Seshan, and P. L. Evanich, "Physical-Chemical Life Support Systems," in F. M. Sulzman and A. M. Genin, eds., *American Institute of Aeronautics and Astronautics*, Washington, DC, 1994.

Jones, H. W., "Common Cause Failures and Ultra Reliability," AIAA 2012-3602, 42nd International Conference on Environmental Systems, 15 - 19 July 2012, San Diego, California.

Likens, W. C., "A Preliminary Investigation of Life Support Processor Reliabilities," International Conference on Life Support and Biospherics, Huntsville, AL, Feb. 18-20, 1992.

NASA, "The Importance of Orbital's Recent Successful Demo Flight and Having Multiple Commercial Providers," NASA's Return On Investment Report, Issue 12, May 2013.

Rechtin, E., *Systems Architecting: Creating and Building Complex Systems*, Prentice Hall, Englewood Cliffs, NJ, 1991.

Wikipedia, Fukushima, wikipedia.org/wiki/Fukushima_Daiichi_nuclear_disaster#Tsunami, accessed Oct. 1, 2014.