Design for Reliability (DfR) in Space Life Support

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

The engineering process of Design for Reliability (DfR) is well established in the automotive and aerospace industries. DfR should be useful in the future development of space life support systems. DfR is a sequence of tasks that develop system requirements and plan reliability analysis and testing. First and fundamentally, the reliability requirement is defined. Next the system reliability model is developed, often using a reliability block diagram. The overall system reliability requirement is allocated to the subsystems and an estimate of the attainable reliability is made. This expected reliability can be improved by simplifying the design by removing components or by replacing less reliable components. Improving reliability can require difficult compromises, such as reducing performance requirements, increasing budget, or extending testing. The actual system reliability can be determined only by testing, which should continue long enough to provide the required confidence in the measured value. New systems often have unexpected design errors that cause failures in early testing. The usual reliability improvement process of testing, finding the failure modes, and redesigning to remove them reduces the failure rate and is referred to as "reliability growth." After redesign has been completed, the system should be further tested to determine the actual achieved reliability more accurately. If the final system failure rate is too high, redundant systems can be used to improve overall operational reliability. Adding redundancy simply to increase the one- or two-fault tolerance metric may sometimes reduce reliability. Reliability can be improved in three ways: redesigning the system to include more reliable subsystems and components, reliability growth testing and failure mode removal, and by using parallel redundant systems. DfR should combine these approaches to achieve the required reliability while managing performance, cost, and schedule.

I. Introduction

This paper describes the Design for Reliability (DfR) process and explains how it can be used to improve the reliability of space life support systems. Since the problems caused by poor reliability usually occur during operations and often long after development, DfR can be overlooked during system design, especially if effort is predominantly focused on improving the real time performance demonstrated in an acceptance test. Establishing the DfR process as a series of required steps that are performed during the standard phased system design process should improve reliability.

II. Design for Reliability (DfR) Overview

Many books and journal articles describe the Design for Reliability (DfR) process. [1] [2] [3] [4] [5] DfR should be an integrated part of the system design process and follow similar sequential coordinated steps. Just as system design should begin by understanding the user's system performance needs, reliability design should first identify the user's system reliability requirements.

A reliable system is one that operates as expected for a specified time under the anticipated conditions. "Reliability is the ability of a system to perform as intended (i.e., without failure and within specified performance limits) for a specified time, in its life cycle conditions." [3] Also, reliability is "the probability that an item will perform its intended function for a designated period of time without failure under specified conditions." [4] While defining reliability

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

requirements, the designers must consider the mission duration and the acceptable failure probability of the system. These two numbers define the system reliability requirement, which should be stated formally as, "The system shall have less than a 1 in 1,000 probability of failure over a 1,000-day operating life." The reliability requirement should be verified by a combination of test and analysis, and the system specification should require a documented verification plan. [1] Surprising failures often occur and there is no way to guarantee future reliability. [2] The required reliability should be achievable and verifiable, but the achievable reliability is frequently overestimated and its cost underestimated.

After the requirement is defined, the next step is reliability modeling, which includes reliability analysis and prediction. Understanding reliability is an essential early step in system design. "Reliability practices must begin early in the design process and must be well integrated into the overall product development cycle." [4] The system diagram is used to develop a reliability block diagram that shows the subsystems and their interconnections. The reliability of each subsystem is assessed, based on test data from its components or similar hardware, and the overall system reliability is estimated. If greater reliability is needed, it can be improved by reducing system complexity and removing components, by redesigning the subsystems, or by providing redundant units. Working from the top-down after the bottom-up estimate, the system level reliability requirement can be allocated between the subsystems, defining each subsystem's target reliability requirement. The feasibility of improving system reliability, and the cost of either improving reliability or accepting the current failure probability and adding redundancy to compensate, must be carefully considered. The objective of DfR is "to ensure that customer expectations for reliability are fully met throughout the life of the product with low overall life-cycle costs." [4]

After reliability requirements definition, reliability modeling, and system design, the next step is reliability growth testing. Often initial system testing reveals unexpected early failure modes, sometimes due to errors in requirements, design, manufacturing, or operational procedures. The familiar trouble shooting and redesign process, test-find-fix, is called reliability growth testing. Reliability growth continues as long as testing finds failures and redesigning removes them. It ends when the design is frozen, which can be a decided based on estimates of the achieved reliability and the possibility of further improvement. Reliability growth testing "should continue until the design is considered to be 'acceptable.'" [4] Further testing is often done to more accurately determine the system failure rate, since that determines if and how much redundancy is needed.

The DfR process should continue even after a system is flown. Flight operations should be monitored for anomalies and all failures should be analyzed. Repairs must be planned and tested and design improvements could be considered. The steps of the DfR process are listed in Table 1. [2] (Jamnia and Atua 2020) Similar steps are described in other references. [4] [6] [7] These steps will be considered further.

Table 1. DfR process steps.

1	Reliability requirements
2	Reliability modeling
3	Reliability allocation
4	Reliability estimation
5	Reliability growth testing
6	Reliability performance testing
7	Reliability monitoring

III. Reliability Requirements

The first and most important step in DfR is defining the reliability requirements. The user's reliability needs depend on how long the system will be used and on the consequences of failure. These factors vary greatly for space life support systems. Substantially different system approaches have been used on multi-day shuttle flights and on the multi-decade International Space Station (ISS). Since these missions are in Low Earth Orbit (LEO), the crew can quickly be resupplied or returned to Earth if the life support system fails. A long stay Mars mission would require about 225 days out, a 450-day surface exploration, and a 225-day return, for a roughly 900-day total mission duration. Since the crew cannot be resupplied or return to Earth during a Mars mission, the life support system for Mars must have much greater reliability than for LEO.

Suppose that the top level Mars mission safety requirement is for a total Probability of Loss of Crew, Pr(LOC), of less than 1 in 100, which is roughly the known risk of the later shuttle missions. Most of this accepted risk would be allocated to the unavoidably high risk mission events, including launch, Mars entry, descent, and landing, Mars ascent, and Earth entry, descent, and landing. The Pr(LOC) due to life support would probably be allocated less than one-

tenth of the total, say 1 in 1,000. The Mars mission life support system reliability requirement would be for 0.999 reliability, a 0.001 probability of failure, over the 900-day mission duration. If a 1,000-day mission duration is used to simplify the math, the required life support failure rate for life threatening malfunctions would be less than 1 in 1 million, 10^{-6}, per day.

Achieving and demonstrating such high reliability is a daunting prospect. If the expected system reliability is too low, the Mars mission plan could be modified to reduce risk. Possibilities include developing advanced propulsion to reduce transit time, making only a short 30-day stay on Mars, or using preplaced stored life support materials instead of relying on complex recycling life support systems.

The reliability requirement specifies the expected operating life, here the mission duration, and tolerable failure probability, here the Pr(LOC). The duration and failure probability should always be stated together to avoid confusion. [8] For example, the Mars mission life support system shall have 99.9 percent reliability over the 900-day mission. In addition to the required reliability, the confidence bounds on reliability should be specified. For instance, the 90% lower confidence bound on reliability should be 99.5 percent. [9]

Significant effort should be spent on life support reliability requirements development and analysis. Reliability requirements should be developed using the same process and at the same time as operational performance requirements [1, pp. 31-5] Requirements development is considered the principal challenge in DfR. “It is worth repeating that the sources of most failures are incomplete, ambiguous, and poorly defined requirements.” [5, pp. 5, 130] Fundamental requirements that are often neglected include safety, repairability, logistics, operability, diagnostics, and failure response. Mars life support requirements should include repairability and maintainability, logistics materials and spares, and failure diagnosis procedures.

IV. Reliability Modeling

The two commonly used reliability models are Reliability Block Diagrams (RBDs) and predictive fault trees. They are graphs of the logic that shows how a lower-level subsystem failure can cause an overall system failure. An RBD shows the line connections between a system’s subsystems, which are shown as blocks. Two blocks in series must both operate for the system to operate, so the system’s probability of reliability is the product of the blocks’ probabilities of reliability. R series system = R block 1 * R block 2. If two blocks are in parallel, the system operates as long as either one operates, and the system fails only if both fail. The system failure probability is the product of the blocks’ probabilities of failure. F parallel system = F block 1 * F block 2. Since the probability of reliability is one minus the probability of failure, R = 1 – F, the reliability of a system with two blocks in parallel is, R parallel system = 1 - F block 1 * F block 2. Using parallel subsystems for redundancy adds failure tolerance and improves system reliability. For any connected configuration of subsystems, the RBD shows how to combine the reliabilities of the subsystems to compute the system reliability. [5, pp. 20-3]

The RBD shows the system processing paths that can produce correct operations, while a failure fault tree shows how any specific block’s failure mode can produce a system level failure. A fault tree is a Boolean logic diagram with AND and OR gates. An OR gate result is TRUE if any input condition is TRUE. Two series blocks are shown as connected with an OR gate, so if either block fails the system fails. An AND gate result is TRUE only if all input conditions are TRUE. Two parallel blocks are shown as connected with an AND gate, so the system fails only if both input blocks fail. [5, pp. 20-3] Both these basic reliability models use the subsystem or block reliability to determine the integrated system reliability.

V. Reliability Allocation

The overall system reliability requirement is usually allocated between the different subsystems and then down to lower-level assemblies and components. This provides target reliability specifications for the design engineers at different levels and guides the design toward the needed reliability. [2] [5, p. 23] [4]

Suppose a system has several necessary subsystems. If the subsystems’ reliability is initially unknown, they can all be allocated the same reliability requirement. If the overall system has a required reliability of R, and there are four subsystems, each can be assigned a reliability of $R^{1/4}$. Since all the subsystems must operate for the integrated system to operate, the subsystems can be modeled as a series RBD, and the total system reliability is the product of the four subsystem reliabilities. $(R^{1/4})^4 = R$.

Suppose a Mars life support system has the required failure rate of less than 1 in 1 million, 10^{-6}, per day. The Mars life support has many different functions, including maintaining atmosphere pressure, temperature, and humidity, carbon dioxide removal, trace contaminant removal, oxygen provision, water provision, food storage, human waste management, and fire suppression. Since this list has ten functional subsystems, each could be allocated the required failure rate of less than 1 in 10 million, 10^{-7}, per day. This is clearly unrealistic because the hardware implementations

of the different functions have very different intrinsic reliability. Food storage, if compartmentalized and backed up by sufficient spares, should be very highly reliable. The spacecraft should reliably maintain pressure if provision is made to repair micrometeor punctures and replace lost atmosphere. Carbon dioxide removal, trace contaminant removal, oxygen generation, and water recycling as implemented on the ISS have had unexpected high failure rates, although without harm or danger to the crew. In designing Mars life support, the acceptable failure rate should be allocated to the intrinsically less reliable subsystems. The failure rates for highly reliable functions should be made negligible. Given the numerous life support failures in the International Space Station (ISS), the reliability of current space life support appears inadequate for a Mars mission. [10] [11]

There is often a gap between an optimistic initial reliability allocation and a realistic data-based estimate. [2] The reliability allocation process can be used to plan reliability improvement. A basic reliability allocation process considers only the provision of redundant components in parallel. These may be built-in or hot spares, but space life support usually relies on cold stand-by redundancy. The more general reliability allocation method combines redundancy with improvement of component reliability. [12] Improving component reliability requires redesign and retesting, probably with a significant increase in cost and schedule. Optimizing the planned reliability improvement using both redundancy and component reliability improvement requires estimating their cost and benefit.

A. The Cost-benefit of Improving Component Reliability

The effort to improve reliability has steeply increasing costs and diminishing returns and the required cost increase is difficult to estimate. Several different mathematical models have been suggested. One rule of thumb is that cutting the failure probability in half requires an investment equal to the original development cost. The mathematical relation is Cost = original cost [1 + $\log_2$ (original failure probability/reduced failure probability)]. [13] Other authors suggest a much faster cost increase, exponential rather than logarithmic. Cost = original cost EXP [(original failure probability - reduced failure probability)/reduced failure probability]. [14] [15] (A more flexible cost model has the cost increase proportional to the power, a, of the failure rate improvement ratio. Cost = original cost [original failure probability/reduced failure probability]^a. [16] The $\log_2$ (original failure probability/reduced failure probability) and exp (original failure probability/reduced failure probability) can be approximated by (original failure probability/reduced failure probability)^a for different specific exponents, a. The proportional function cost of reliability seems able to model all the expected cases. [17]

Assessing the opportunity and cost of improving component reliability is crucial in reliability planning. The recycling and regenerative life support systems now on the ISS are unique designs. They were very briefly tested before launch and have no identical copies, no prototypes or test models on Earth. Failures have been unexpectedly frequent and troubleshooting and redesign have been difficult. A strong effort focusing on better reliability could make significant gains. When improving an existing system's reliability is difficult, using a different technology or simplifying the system architecture can improve reliability and even save cost. One example is replacing a complex redundant halon-using fire suppression system proposed for ISS with standard fire extinguishers. [18] A possibility in life support is using stored tanks of water instead of, or as backup to water recycling systems. [19]

B. The Cost-benefit of Adding Redundancy

The cost of adding redundancy is simply that of producing and flying the additional units. The cost of producing each additional unit decreases due to the learning curve, with one well known formula giving cost ~ quantity $^{0.59}$. [20]

It has been suggested that a recycling life support system for a Mars mission be provided with three or four spares of each major subsystem, such as the water processor or oxygen generator [21] (Connelly 1999) This approach is not cost efficient since repairing a single failure uses a full subsystem. In contrast, ISS life support uses multiple Orbital Replacement Units (ORU's) for each subsystem, so that failure replacement units are much smaller and less costly. Most important, a full spare subsystem can repair only a single failure while its set of component ORUs can repair many failures. This produces very much higher reliability using the same mass of spare parts.

General calculations show how system reliability increases when system and component spares are provided. Suppose that each system has a failure probability F over the mission duration. If two parallel redundant components are provided, the probability that both will fail is F^2, which is less than F assuming $F < 1$. If M parallel redundant system are provided, the probability that all M will fail is F^M. Suppose that the system is divided into N components, each with the failure rate F/N. The probability that the system with one single string of components will fail is still F = N (F/N). The probability that a pair of redundant components will both fail is $(F/N)^2$ and the probability that a system with a string of pairs of redundant components will fail is N $(F/N)^2 = F^2/N$. Providing two complete systems reduces the failure probability to F^2, but dividing the system into N components each with 1/N of the total failure probability, and then providing two full sets of N components reduces the failure probability to F^2/N. Similarly, providing M

instead of two redundant full systems reduces the failure probability to F^M, and dividing the system into N components as before, and then providing full sets of N components reduces the failure probability to F^M/N.

This calculation is straightforward because it assumes that the redundant systems are all online and operating, so that they all have an online system's probability of failure. The ISS life support ORU's are in storage and are unlikely to fail before being used. A combinatorial calculation assuming stored spares do not fail gives an even further reduced failure probability. Taking a system, dividing it into N components with equal mass and failure rates, and then providing M copies of each component reduces an initial failure probability of F to $[F^M]/[M!\ N^{M-1}]$. For M = 2, this is half the failure rate of the simplified calculation, reflecting that only half of the components are operating at any one time. [11] [22] [23]

In life support systems design for the space station, the potential points of failure were contained in replaceable ORU's, which can replace filters, motors, sensors, and valves, not but not piping, framing, structure, or panels. This makes improving reliability using spares easier and more effective.

C. Common Cause Failures

Not all failures can be repaired using spares. Sometimes rework or redesign are needed. Common Cause Failures (CCF's) defeat redundancy since they damage both the original parts and the spares. CCF's can be due to deficient parts, but many are caused by design errors, unexpected environmental challenges, or operator errors.

Rutledge and Mosleh investigated all the space shuttle flight anomalies in the first forty flights after the Challenger accident. Of 473 anomalies, 54 (11%) were judged to be CCF's, and 11 more were due to functional or spatial interaction, for a total of 64 (14%) non-independent failures. [24] The ISS Oxygen Generation Assembly (OGA) had 50 component and ORU replacements and as many other maintenance events from 2009 through 2014. About half of these OGA problems were due to design errors and other CCF's. The very high level of CCF's in the ISS OGA probably reflects the great difficulty in testing, trouble shooting, and redesigning a system that has only one operating unit, with that unit located in space. CCF's can be reduced to a few percent of all failures if systems are well designed and thoroughly tested.

CCF's reduce the reliability improvement that would be expected from redundancy. A CCF can be defined as a failure mode that defeats redundancy. If there are two redundant operating units, a common cause that makes one fail will cause both to fail. If a single operating unit fails due to a continuing or repeated common cause, its replacement will fail for the same reason. A random failure, even one with unexpected high probability, is not a CCF since operational reliability can be improved by adding redundant spares. Suppose a component and its spare both have a failure probability of 0.01 per year. If the failures are independent, the probability that both fail is 0.01 * 0.01 = 0.0001 per year, a reduction by two orders of magnitude. But suppose the failure probability of 0.01 per year includes ten percent of CCF's, so the CCF probability is 0.001 per year. If the original part fails due to a CCF, the replacement parts can also be expected to fail. The redundant system failure probability is no less than the CCF probability, no matter how many spares are provided. The probability that two units 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 probabilities. Redundancy cannot reduce the overall failure probability below the common cause failure probability. Reducing the probability of CCFs should be a major objective of space life support design and testing.

D. Problems with Fault Tolerance Redundancy Requirements

NASA uses fault tolerance requirements to improve the safety and reliability of human missions. A two-fault tolerant system is one that can operate satisfactorily after two failures. This requires triple redundancy, one operating system and two spares. Achieving fault tolerance requires redundancy, which is often needed to improve reliability, but redundancy design can create problems. Automated redundancy can add sensors, controllers, logistics, and programmed operations, which add new failure modes including CCF's that may not be sufficiently understood. NASA's fault tolerance requirements can increase risk while mistakenly leading to overoptimistic assessments of reliability. [25]

E. Summary of Reliability Allocation

After reliability requirements and modelling, the next DfR step is reliability allocation. A rough initial reliability allocation can be done theoretically from the top-down, but a bottom-up data-based approach seems preferrable. If reliability data is available for similar past systems or an engineering prototype, it should be used to make the reliability allocation. After reliability allocation, the next step is reliability estimation, which should uncover any forseeable difficulty in achieving the requirements. There may be a significant gap between the requirements and realistic estimates. [2]

VI. Reliability Estimation

"It is highly important to estimate the product's reliability, even with a rough first cut estimate, early in the design phase." [4] Correctly estimating the future reliability of a projected system is very difficult and overestimates are common. There is no certain method for accurate prediction of future reliability. [2] Nevertheless, reliability prediction is useful and necessary. [5] Comparing the expected reliability to the requirement quantifies the anticipated difficulty of DfR.

A. Initial Bottom-Up Component Based Estimate

The usual way to do an initial reliability estimate is bottom up, simply adding the component failure rates. The system design is graphed in a block diagram and the component level RBD is developed. The component failure rates are obtained from vendor data, preliminary tests, physics of failure analysis, and if this is insufficient also using engineering judgment and expert opinion. [4]

Typically, all the components must perform for successful operation, so the estimated overall failure rate is the sum of the component failure rates. This provides an optimistic lower bound on the failure rate. The lowest possible failure rate can be achieved only if the system design does not introduce other failure modes in addition to the component failure sources.

"Not all system failures are caused by parts. Other causes include unexpected interactions between components, tolerance build-up, and software faults. But because of the difficulty of associating these with numerical failure rates the practice has been to use part failure rates as a proxy." [26] Component-based system failure assessments underestimate the true failure rate because they ignore integration and system level failures. Life support reliability has been greatly overestimated using the sum of components failure rate approach. [27]

B. Similar Systems Reliability Estimate

A better and more realistic failure rate estimation should be made based on the actual failure rates of similar systems. Estimation using similar systems' data is called "reference class forecasting," and is based on Nobel prize winning decision analysis described in Daniel Kahneman's book, Thinking, Fast and Slow. [28] Planning is often overconfident and uninformed by relevant experience. Costs, schedules, failure rates, and risks can all easily be underestimated. Reference class forecasting requires: 1) identifying the reference class of similar systems, 2), establishing the class distribution of the estimated parameter, e.g., the reliability, and 3), comparing the target system to the reference class.

The project planning expert Flyvbjerg found that cost under-estimation errors can be caused by technical errors, over-optimism, and deliberate self-serving deception, but he also found widespread ignorance of actual past project costs. He considered reference class forecasting to be "the single most important piece of advice regarding how to increase accuracy in forecasting." [28] Reference class forecasting deliberately avoids the distracting details of system design and project planning and considers only the actual costs of similar projects. Using reference class forecasting is easier and more accurate than component-based bottom-up estimation. [29] [30]

C. Poor Past Life Support Reliability

The ISS life support system has had poor reliability. The actual life support failure rates have been significantly greater than predicted. Almost always, the specific subsystem failure rates were higher than predicted and they were usually a full order of magnitude higher. Water and oxygen storage and resupply were found to be significantly more reliable than physical-chemical recycling processors. [27] [11] Russell and Klaus state the "total ECLSS maintenance for 865 days was found to exceed the design estimate by a factor of 22." A contributing factor was the oxygen generation system's greater than expected failure rate. [10] [11]

D. Reliability Feasibility Analysis

A reliability feasibility analysis should compare the estimated reliability of each subsystem to the allocated reliability requirement. [2] The feasibility analysis is needed to identify gaps, define strategy, and develop the reliability plan. The reliability estimates should be updated based on test results on engineering models and prototypes. [2]

VII. Reliability Growth Testing

"(T)raditional Design-for-Reliability (DFR) principles alone are simply not sufficient when designing highly complex systems," and they should be augmented by a reliability growth program.[31] The purpose of reliability growth testing is to induce failures, identify their causes, and redesign the system to remove them. Initial testing of a

newly developed system often finds an unexpectedly high failure rate, sometimes referred to as "infant mortality." The unanticipated failure causes include poor design, defective materials, an inadequate manufacturing process, and poor operating instructions. [2] As these initial failure modes are identified and removed, the failure rate usually declines exponentially with time. Reliability growth is commonly modeled as continuing indefinitely as long as testing continues, with less probable failure modes continually occurring and being removed, and the failure rate constantly declining toward zero. However, this is frequently a mathematical modeling artefact, due to an early high number of failures being averaged over ever-longer time periods. In actual system operations, it is more typical for the failure rate to fall only to a constant low level. Reliability growth testing should be planned and monitored using a two-phase model, one having a period of exponential decline in failure rate followed by a period of constant failure rate. This corresponds to the first two segments of the traditional 'bathtub curve" of reliability over time. [32] [33]

In the usually assumed exponential model of reliability growth, the cumulative failure rate $n(t)/t = k\ t^{-\alpha}$, where n(t) is the number of failures until time t and k and α are constants. Suppose k failures occur before t = 1 and no further failures occur. Then the cumulative failure rate is $n(t)/t = k/t = k\ t^{-1}$ and cumulative failure rate declines at the fastest possible speed. As t increases, the original failure count is averaged over an increasingly longer period and the cumulative failure rate n(t)/t = k/t approaches 0. It is usually assumed that testing and reliability growth can continue indefinitely. However, mathematically averaging a high initial failure count over an increasingly longer time will produce a declining cumulative failure rate that incorrectly suggests that the actual reliability is improving. Data often show that the initial exponential decline in failures during reliability growth is followed by a constant failure rate. If this is the case, long extended testing will cause the cumulative failure rate to approach the final constant failure rate. This final rate is the operational failure rate that should be used in reliability planning, and it should be more accurately determined by reliability measurement testing discussed in the next section. [32] [33]

One key concern is identifying or determining the expected transition during testing between reliability growth and a constant failure rate. Reliability growth testing is an active process of finding and fixing failure modes. It will encounter diminishing returns and will probably be terminated when its assessed cost begins to exceed its benefit in failure rate reduction. The final failure rate may be due to many random, independent, and low probability failure modes that are considered unnecessary to remove. Reliability growth testing "should continue until the design is considered to be 'acceptable.'" [4] The expected failure rate must be deemed satisfactory.

VIII. Reliability Measurement Testing

After reliability growth testing is completed and the final system design established, further testing should be conducted to establish the expected failure rate of the final design. [2] Sometimes the achieved failure rate is too high to provide the required reliability and confidence over the mission duration and redundant systems are required. Since the measured failure rate has some uncertainty, the reliability of the redundant systems will be overestimated about half the time. Adding more redundant units increases the confidence that the targeted reliability will be achieved. For any chosen number of redundant units, a lower reliability goal will be achieved with higher confidence. [32] [34] [35]

The needed number of redundant units for any required reliability and confidence can be estimated based on the failure history. The record of the failure times determines the failure rate and its distribution so that reliability confidence intervals can be estimated. The needed redundancy can be computed using the cumulative Poisson distribution. The confidence that the expected reliability will be achieved can be computed using the cumulative Poisson distribution or the chi-square distribution. Extended testing should be conducted to more accurately determine the failure distribution and so reduce the variance of the failure rate and the width of the reliability confidence intervals. Narrowing the confidence interval reduces the number of redundant units needed to provide the required reliability and confidence and therefore reduces the total system implementation cost. However, the extended test time increases the cost, so testing cannot continue indefinitely.

As the test time is increased, the test cost increases linearly but the number and cost of the needed spares drops rapidly, at first exponentially. The total cost is the sum of the cost of the redundant units and of the extended test program. There is often an optimum test time that produces the minimum total cost for the system failure rate, mission length, required reliability, and confidence level. If the projected total cost must be cut, either the reliability or confidence level or both must be reduced. The required reliability and confidence directly determine the minimum total cost for the redundant units and extended testing. This cost-based justification should encourage sufficient testing. [32] [34] [35]

The major ISS life support systems, carbon dioxide removal, water recycling, and oxygen recovery, were protoflight systems with little testing before launch and operations. The water recycling system had only thirty hours of integrated testing before starting an operational life that may extend from 2008 to beyond 2025. One of the major

lessons learned from the ISS development experience is the need for extensive ground testing before launching space systems. [36]

IX. Reliability Monitoring in the Field

Most system operations and most failures occur not in testing but when systems are used in the field. And despite all admonitions to "test as you fly," field operations often provide unanticipated causes to fail. Field failures should be analyzed to obtain the information needed to improve operational system use and to prioritize possible future redesigns.

Another important lesson learned from the ISS experience, beyond the requirement for preflight ground testing, is the need to test space systems on orbit. [36] The initial operation of a system in space frequently reveals unexpected failure modes, often caused by launch vibration damage or by operation in microgravity. Fluid flow and heat convection are reduced. Space induced failures of ISS life support include flow clogging and particle migration due to microgravity. The ISS implements operational redundancy by providing onboard spares of ORUs (Orbital Replacement Units). The numbers of spares provided are based on the anticipated failure rates, which were initially estimated and then updated based on failure experience.

X. Conclusion

The established aerospace Design for Reliability process includes reliability requirements, modeling, allocation, estimation, growth testing, performance testing, and monitoring. This approach can obtain verified high reliability at the least cost. Space life support system development should use the DfR process to produce the reliable life support needed to build a sustainable human presence on the Moon and continue human exploration on towards Mars. The challenge of providing sufficient reliability for long missions in deep space is great, especially since reliability must be convincingly demonstrated before the mission. The high reliability requirement will drive the design toward using very reliable components and extensive redundancy. The time before a deep space mission is initiated can be well used in long duration life support reliability testing.

References

[1] Daley, D.T., Design for Reliability, Industrial Press, 2011.

[2] Jamnia, A., and Atua, K., Executing Design for Reliability Within the Product Life Cycle, CRC Press, Boca Raton, FL, 2020.

[3] Kapur, K.C., and Pecht, M., Reliability Engineering (Vol. 86), John Wiley & Sons, 2014.

[4] Mettas, A., "Design for reliability: Overview of the process and applicable techniques," International Journal of Performability Engineering, 6(6), 2010, pp. 577-586.

[5] Raheja, D.G. and Gullo, L.J. eds., Design for Reliability, John Wiley & Sons, 2012.

[6] Silverman, M., and Kleyner, A., "What Is Design for Reliability and What Is Not?" 2012 Proceedings Annual Reliability and Maintainability Symposium, IEEE, 2012, pp. 1-5.

[7] Cota, E.V., Gullo, L. and Mujal, R., "Applying Design for Reliability to increase reliability confidence," 2014 Proceedings Annual Reliability and Maintainability Symposium (RAMS), IEEE, 2014.

[8] Schenkelberg, F., "Establishing product reliability goals," 2013 Proceedings Annual Reliability and Maintainability Symposium (RAMS), IEEE, 2013, pp. 1-6.

[9] Sarakakis, G., Gerokostopoulos, A. and Mettas, A., "Special topics for consideration in a design for reliability process," 2011 Proceedings Annual Reliability and Maintainability Symposium (RAMS), IEEE, 2011, pp. 1-6.

[10] Russell, J. F., and D. M. Klaus, "Maintenance, reliability and policies for orbital space station life support systems," Reliability Engineering and System Safety, Volume 92, Issue 6, June 2007, pp. 808-820.

[11] Jones, H., "Life Support Dependability for Long Space Missions," AIAA 2010-6287, 40th ICES (International Conference on Environmental Systems), 2010.

[12] Kuo, W. and Wan, R., "Recent advances in optimal reliability allocation," IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 37(2), 2007, pp.143-156.

[13] Rechtin, E., Systems Architecting: Creating and Building Complex Systems, Prentice Hall, Englewood Cliffs, NJ, p. 165, 1991.

[14] Mettas, A., "Reliability allocation and optimization for complex systems," 2000 Proceedings Annual reliability and maintainability symposium, pp. 216-221, IEEE, 2000.

[15] Elegbede, A.C., Chu, C., Adjallah, K.H. and Yalaoui, F., "Reliability allocation through cost minimization," IEEE Transactions on reliability, 52(1), pp.106-111, 2003.

[16[Aggarwall, K. K., Reliability Engineering, Springer, 1993.

[17] Jones, H., "High Reliability at Minimum Cost," Submitted to 70th Reliability& Maintainability Symposium (RAMS), Albuquerque, NM, Jan 22, 2024.

[18] Anderson, J.M., Stott, J.E., Ring, R.W., Hatfield, S., and Kaltz, G.M., "Factors which Limit the Value of Additional Redundancy in Human Rated Launch Vehicle Systems," AIAA 2008-3585, SpaceOps 2008 Conference.

[19] Jones, H., "Moon base life support design depends on launch cost, crew size, and mission duration, ICES 2019-17, 49th ICES (International Conference on Environmental Systems), 2019.
[20] Guerra, L., and Shishko, R., "Estimating the Cost of Crewed Space Systems," in W. J. Larson, and L. K. Pranke, eds., Human Spaceflight: Mission Analysis and Design, McGraw-Hill, New York, 1999.
[21] Connelly, J. F., "Mars Design Example," in Larson, W. K., and Pranke, L. K., eds., Human Spaceflight: Mission Analysis and Design, McGraw-Hill, New York, 1999, p. 982.
[22] Jones, H., and Ewert, M. K., "Ultra Reliable Closed Loop Life Support for Long Space Missions," AIAA 2010-6286, 40th ICES (International Conference on Environmental Systems), 2010.
[23] Jones, H., "Ultra Reliable Space Life Support Systems," SAE 2008-01-2160, Society of Automotive Engineers, Warrendale, PA, 38th ICES (International Conference on Environmental Systems), 2008.
[24] Rutledge, P.J., and A. Mosleh, "Dependent-Failures in Spacecraft: Factors, Defenses, and Design Implications," IEEE, 1995, Proceedings Annual Reliability and Maintainability Symposium.
[25] Jones, H., "NASA Should Not Use the Traditional One- or Two-Fault Tolerance Rules to Design for Reliability," 69th Reliability& Maintainability Symposium (RAMS), Orlando, FL, Jan 23, 2023.
[26] Hecht, H., Systems reliability and failure prevention, Artech house, 2004.
[27] Likens, W. C., "A Preliminary Investigation of Life Support Processor Reliabilities," International Conference on Life Support and Biospherics, Huntsville, AL, Feb. 18-20, 1992.
[28] Kahneman, D., Thinking, Fast and Slow, Farrar, Straus, and Giroux, New York, 2011.
[29] Flyvbjerg, B., "From Nobel Prize to Project Management: Getting Risks Right," Project Management Journal, vol. 37, no. 3, August 2006, pp. 5-15. Online at http://flyvbjerg.plan.aau.dk/Publications2006/Nobel-PMJ2006.pdf
[30] Flyvbjerg, B., Garbuio, M., and Lovallo, D., "Delusion and Deception in Large Infrastructure Projects: Two Models for Explaining and Preventing Executive Disaster," California Management Review, vol. 51, no. 2, Winter 2009, pp. 170-193.
[31] Wayne, M., and Modarres, M., "A Bayesian Model for Complex System Reliability Growth Under Arbitrary Corrective Actions," IEEE Transactions on Reliability, vol. 64, no. 1, March 2015.
[32] Jones, H., "Achieving Maximum Reliability Growth in Newly Designed Systems," 64th Reliability& Maintainability Symposium (RAMS), Reno, NV, Jan. 2018.
[33] Jones, H., "The abcd Reliability Growth Model," Submitted to 70th Reliability& Maintainability Symposium (RAMS), Albuquerque, NM, Jan 22, 2024.
[34] Jones, H., "Extended Testing Can Provide Cost-Effective Redundancy With High Reliability and High Confidence," 67th Reliability& Maintainability Symposium (RAMS), Orlando FL, May, 2021.
[35] Jones, H., "Redundancy: How Many Unreliable Spares are Needed for High Reliability and Confidence? Submitted to 70th Reliability& Maintainability Symposium (RAMS), Albuquerque, NM, Jan 22, 2024.
[36] Jones, H., "Lessons Learned in Space Life Support System Testing," AIAA 2021-138, 50th ICES (International Conference on Environmental Systems), 2021.