

We Can't Count on Repairing All Failures Going to Mars

Harry W. Jones¹

NASA Ames Research Center, Moffett Field, CA, 94035-0001

Reliability analysis often assumes that a complex system can be kept operating indefinitely with scheduled maintenance and emergency repair using a stock of spare parts, as long as the spare parts are not depleted. This assumption seems justified for well-tested, widely used, long operational systems with a multigenerational history of failure, redesign, and reliability growth. It seems doubtful that newer, relatively untried, high technology space systems can always be repaired. We cannot assume space systems will have a low rate of random failures that can all be repaired with a few identical spares. New untried systems usually have a high initial failure rate, called infant mortality, due to errors in requirements, design, parts, materials, and operations planning. These problems can cause groups of related failures called Common Cause Failures (CCFs). The practical definition of a CCF is any failure mode that cannot be cured using identical redundant systems or spare parts. Systems with CCFs may fail repeatedly for the same reason. Can a life support system be kept operating on the way to Mars using only redundant systems and spare parts? The failure history of International Space Station (ISS) life support systems suggests that CCFs are likely to occur and will probably require design changes rather than being repairable with spare parts.

Nomenclature

<i>CCF</i>	=	Common Cause Failures
<i>cdf</i>	=	cumulative distribution function
<i>H₂</i>	=	Hydrogen
<i>ISS</i>	=	International Space Station
<i>MTBF</i>	=	Mean Time Before Failure
<i>OGA</i>	=	Oxygen Generation Assembly
<i>ORU</i>	=	Orbital Replacement Unit
<i>pdf</i>	=	probability distribution function
<i>TOC</i>	=	Total Organic Carbon

I. Introduction

THIS paper challenges a common optimistic assumption in reliability analysis that all system failures can be repaired using spare parts. If we assume that failures are random and independent, and that we know the component failure rates, a simple calculation tells us how many spares are needed to achieve any high level of reliability. But what if the system is damaged due to accident, fire, explosion, or toxic contamination? What if a design error or deficient material causes all the spare parts to quickly fail? Is it possible that, on the journey to Mars, the life support system will have a serious problem that cannot be repaired? This paper develops logical arguments, describes reliability models, and investigates International Space Station (ISS) life support failure data to show that we can't count on repairing all failures going to Mars.

II. Discussion of failure and repair

Logical considerations show that it is difficult to achieve high reliability simply by repairing systems using spare parts.

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

A. The total reparability assumption is wrong

Assume a mission uses a variety of hardware systems, that it has a maintenance and repair approach that is well developed and implemented, and that this approach includes the capability to repair or replace or tolerate failure in all the systems. Then if a failure occurs and is traced to a system or component, the options would be to repair, replace, or ignore the failed system or component. For example, if one of three redundant on-line sensors failed, it could be repaired, replaced, or ignored. Robust systems can sometimes use operational slack to avoid repairs, but most systems should be restored to their nominal condition. Large critical systems, such as regenerative life support, must be kept operating nearly continually. If a system or subsystem has failed and been replaced, the failed unit should be repaired in case of future need.

The key assumption here is that all seriously failed systems can be replaced or repaired using spares, and that the replacement or repaired systems will work as expected. This is justified by the further assumption that the failures are random and independent, unpredictable and uncorrelated. This is too strong an assumption. Many failures are not random and uncorrelated. Some, often many failures are due to specification errors, design mistakes, manufacturing problems, and operational errors. All these are built into or equally affect all the identical redundant systems and spare parts. They are called common cause failures (CCFs) and they defeat the use of redundant systems and spare parts, since all the identical systems or components can fail in the same way.

Newly designed systems with limited testing often suffer infant mortality, since design and manufacturing errors cause failures during initial operation. Redesign or process change is then needed. Redundancy and spare parts cannot cure engineering mistakes. A system is made reliable either by making it very similar to earlier proven designs or by careful design and extensive testing. The number of CCFs is much higher for newly designed, recently redesigned, little tested, and relatively unused systems. Reliability analysis assuming total reparability is a too simple approach that produces an over-optimistic, best-case reliability estimate. It assumes that many typical causes of failure do not exist.

B. Achieving high reliability using spares is implausible

If it is assumed that systems can be kept operating by repairing failures using spare parts, it is easy to prove that very high reliability can be easily achieved using a reasonable additional mass of spares. Since this is not actually possible, the assumption must be incorrect.

A typical real world hardware system has many components and subsystems with different mass and failure probability. There are often heavy sturdy components, frames and tanks and panels that have high mass and low failure rate. There are often small delicate components, filters and valves, that have low mass and high failure rate. The smaller, higher failure rate components usually have several spares provided to increase overall system reliability. Providing a small mass of spares can give a large increase in reliability. The ordinary design of systems produces components with reliability inversely related to mass and with maintenance based on lower mass spares.

The extreme worst case of hardware design for a low mass of spares would be if each component had exactly the same mass and failure rate. We assume this worst case. Specifically, suppose a system has mass M and failure probability F . We divide it into N identical components each with mass M/N and failure probability F/M . We provide each component with R identical redundant units. Each component fails only if all R units fail, so the failure probability for redundancy of R is $(F/N)^R$. The overall mass is R times the original mass and the overall failure probability is the sum of the failure rate for N sets of R components.

$$\text{Overall mass} = R M$$

$$\text{Overall failure probability} = N (F/N)^R$$

Component level repair using small spare parts would typically have the number of different replaceable parts, N , to be large. For any level of redundancy, R , even a reasonable $R = 3$, the overall failure probability approaches zero as N becomes large. For example, $F = 0.1$, $R = 3$, $N = 100$, overall failure probability $= N (F/N)^R = 10^{-7}$, one in 10 million. Even for $F = 0.1$, $R = 2$ for an operating system and one spare, and only $N = 10$, overall failure probability is 10^{-3} , one in a thousand. This seems very unrealistic for most real world hardware systems.

Those who suggest the use of spares and low level repair for space life support systems are expecting the kind of reliability improvement that can be achieved if spares can repair all failures. The use of spares can repair the independent random failures assumed in the usual simple reliability model. These failures can all be repaired and the failure rate due to them can be driven down to zero. The problem is that many failures are not random and

independent, but are correlated common cause failures (CCFs) such as design errors and so cannot be cured by spares.

Reliability calculations that assume all failures can be repaired using redundant spare parts are too optimistic. In the real world, CCFs strongly limit the gains of redundancy. New, relatively untested designs are especially difficult to make highly reliable. Assuming that space life support can be kept operational using spares and low level repair is clearly dubious and misleading.

III. Mathematical models of reliability

The usual most common mathematical model of reliability assumes that a working system has a constant failure rate due to unpredictable independent random failures. This model does not consider that systems often have a high initial failure rate, the so-called infant mortality, or that some components wear out, causing a high failure rate at end of life. This failure rate over time is sometimes described as the “bathtub curve,” initially high, decreasing to a constant low level, and then increasing to a high level.

The initially decreasing failure rate is due to burn-in, to failure and replacement of defective components, and to detection and repair of design faults. The low constant failures during useful life are random events sometimes caused by external loads or stresses. The finally increasing failure rate can be caused by mechanical wear out or aging related to chemical or thermal stress. More than two-thirds of systems show infant mortality and then a constant failure rate, but no final aging period. (Hansen, 2001) Shuttle maintenance data show a rapid decrease of the failure rate to a then slowly decreasing baseline, but no later increase due to wear-out effects. (Shishko, 1995) If systems are well tested and designed to last several times the required operating life, their failure rates can be expected to be constant.

Another frequently observed departure from the usually assumed low constant failure rate is the occurrence of clusters of or repeated common cause failures, which are not random or independent. A design error or requirement misunderstanding may cause repeated rapid failures of the same component or a sequential cause-and-effect chain of failures. Common cause failures are often a major factor in infant mortality. The gradual process of operating, trouble-shooting, and improving a system may lead to gradual reliability growth, a slow decline in the failure rate, over the operating life so the system.

Fitting reliability data to a model is interesting in itself, but the major reason to establish a plausible model is to suggest reliability behavior beyond the data. We will discuss the constant failure rate model, the beta factor model for common cause failures, and Weibull model that, with different parameters, describes infant mortality, constant random failures, and wear out.

A. The constant failure rate model

We first assume a constant failure rate, λ , which is the number of times a component or system is expected to fail per unit time, given that it is currently still operating. The failure rate is sometimes given in failures per thousand or per million hours.

The reliability, $R(t)$, is the probability that the system does not fail before time t . If the failure rate, λ , is a constant, the reliability, $R(t)$, is an exponential function.

$$R(t) = e^{-\lambda t}$$

The cumulative distribution function (cdf) of the failure probability, $F(t)$, is the probability that a system does fail before time t .

$$F(t) = 1 - R(t) = 1 - e^{-\lambda t}$$

The failure probability density function (pdf), $f(t)$, is the probability of failure over time.

$$f(t) = dF(t)/dt = -dR(t)/dt = \lambda e^{-\lambda t}$$

This is the exponential probability distribution. The mean value of the exponential pdf is $1/\lambda$. The variance is $1/\lambda^2$. The median is $\ln(2)/\lambda = 0.693/\lambda$.

The Mean Time Before Failure (MTBF) is $1/\lambda$, the mean value. The probability that the equipment will not fail before the MTBF is $R(1/\lambda) = e^{-1} = 0.63$. A constant failure rate is a good model for the long flat "intrinsic failure" portion of the bathtub curve where early failures or wear out are not significant.

B. The beta (β) factor Common Cause Failure (CCF) model

A common cause failure (CCF) occurs when several failures have the same origin. Common cause failures are either common event failures, where the cause is a single external event, or common mode failures, where two systems fail in the same way for the same reason but at different times. Common mode failures can occur because of a design defect and they reduce the dependability of systems that rely on repair using spare parts.

The β factor method assumes that the fraction β of the overall failure probability is due to CCFs. That is, $(1 - \beta)F$ of the failures are random and βF are CCF's. F is the total probability of the system failing due to all events, both independent and common cause. The beta factor model is the most frequently used common cause failure model. Large amounts of data have been gathered, especially on nuclear power systems. (Borcsok et al., 2007) (Bukowski and Goble, 2001) (Stotta et al.)

1. Common cause failures and redundancy

It seems possible to achieve high reliability for a system that has a reasonable initial failure probability F (say less than 0.1), by dividing the system into N subsystems and making each subsystem a redundant pair. Each of the N subsystems has a failure probability of F/N . A redundant pair of subsystems has failure probability $(F/N)^2$, and the series of N pairs has failure probability $N * (F/N)^2 = F^2/N$. For $F = 0.1$ and $N = 10$, $F^2/N = 0.001$.

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

2. Common cause failures in shuttle

Rutledge and Mosleh identified the dependent and common cause failures in all the space shuttle in-flight anomalies that occurred during the first forty flights after the Challenger accident. Of 473 anomalies, 54 (11%) were judged to be common cause failures, 6 due to functional interaction, and 4 due to spatial interaction, for a total of 64 (14%) dependent failures. The frequency of dependent and common cause failures is not significantly different from that found in nuclear power plants. (Rutledge and Mosleh, 1995)

C. The Weibull time varying failure rate model

The Weibull probability density function (pdf) of the failure times $f(t)$ is:

$$f(t) = (k/a) (t/a)^{k-1} e^{-(t/a)^k}$$

In the Weibull pdf, $k > 0$ is the shape parameter and $a > 0$ is the time scale parameter of the distribution. For $k = 1$, the Weibull distribution becomes the exponential distribution used above, with $\lambda = 1/a$.

The Weibull distribution has a failure rate that is proportional to the $k - 1$ power of time, t . The value of k defines three reliability regimes.

$k < 1$ corresponds to "infant mortality," a failure rate decreasing over time.

$k = 1$ indicates that the failure rate is constant over time.

$k > 1$ indicates that the failure rate increases with time, possibly due to wear out. (Wikipedia, Weibull distribution)

The parameter "a" is related to system lifetime and the MTBF. For $k = 1$, the constant failure rate case, $a = 1/\lambda = \text{MTBF}$. For general k , the expected system life is proportional to the inverse of the Weibull probability density function. For $k > 1$, the Weibull probability density function increases over time, so failures become more likely and the MTBF decreases over time. $k > 1$ indicates wear out and deteriorating reliability. For $k < 1$, the Weibull probability density function decreases over time, so that the MTBF increases over time. $k > 1$ indicates infant mortality and reliability growth. (Gaudard and Wright, 2016)

The Weibull probability density function (pdf) has an instantaneous failure rate equal to:

$$v(t) = (k/a) (t/a)^{k-1}$$

For $k = 1$ and $\lambda = 1/a$, $v(t) = \lambda$, the constant failure rate of the exponential distribution. (Scholz, 2008)

1. Weibull reliability plots

Plots of the Weibull probability density function for different k are shown in Figure 1.




Figure 1. Weibull probability density functions.

The failure probability density function is the number of failures expected at any particular time. As units fail and are removed from service, the number of failures drops to zero. All the Weibull pdf plots in Figure 1 have $a = 1$ and roughly similar average lifetimes due to the long tails to the right of most of the pdfs. The case $k = 1$ corresponds to a constant failure rate. That is, for $k = 1$, each system that is still operating has a constant probability of failure. Since the number of operating systems declines over time, the pdf of a system failure declines. The case $k = 0.5$ shows decreasing infant mortality or reliability growth. The case $k = 1.5$ shows wear out, reliability decreasing with time. The case $k = 5$ corresponds to an approximately constant limited life, with most failures occurring near time $a = 1$. Higher k narrows the pdf.

Plots of the Weibull instantaneous failure rate for different k are shown in Figure 1.




Figure 2. Weibull failure rates.

The failure rate is the probability that any one operating unit fails. In Figure 2, $k = 0.5$ shows decreasing failure rate, $k = 1$ constant failure rate, and $k = 1.5$ shows failure rate increasing with time. For $k = 5$, The failure rate increases very rapidly as the life limit is approached.

Figure 3 shows a composite Weibull pdf that models the bathtub curve reliability behavior sometimes observed over time.




Figure 3. A composite Weibull pdf with a bathtub curve failure rate.

The composite Weibull pdf first shows a rapidly declining failure rate. Then the failure rate and pdf increase but the pdf then declines because few systems are still operating.

D. Applying the reliability models

The purpose of theoretical models is to help understand reality. How do these models apply to engineering experience? We consider components, assemblies, and complex systems.

1. Component reliability

Components such as integrated circuits are often tested in large groups to determine their failure rates and lifetimes. They often have constant failure rates per unit over time, $k = 1$. They may be tested at high temperature to identify life limiting flaws and then show wear out, $k > 1$. If design or manufacturing errors cause infant mortality, $k < 1$, the components may be subjected to burn-in before being used. Some components such as adsorbent beds or catalytic converters are simply used up over time and have a characteristic life time in hours or cycles, $k \gg 1$. Components have relatively few failure modes so their reliability data often fits one of the simple models in Figure 1. Components are often operated until failure and then replaced. Assemblies and systems may be maintained or repaired as well as replaced.

2. Assembly reliability

Assemblies such as pumps have multiple components such as motors, bearings, valves, filters, and seals, so they have multiple failure modes. Like components, they can be tested in large groups to determine their failure rates and lifetimes. Their failure data can show the full bathtub curve behavior, infant mortality, a long duration low constant failure rate, and increasing final wear out. Some components may have a short limited life or require midlife scheduled maintenance, such as replacing filters, seals, or bearings.

3. Complex system reliability

Operational hardware that provides service to users can be a complex system, consisting of several assemblies and many components. An automobile is a complex, highly reliable system that can have infant mortality, a long low constant failure rate with occasional repairs and periodic scheduled maintenance, and then final wear out. The user of automobiles wants cost effective and reliable transportation and can optimize. Suppose the user is the manager of a company owned fleet of taxis or loaners. The users can buy a group of all identical automobiles and accept the composite bathtub reliability behavior, or they can continually replace the worst performing ones with similar new models to obtain an approximately constant failure (and repair cost) rate. An individual or family might keep a spare car or rely on sharing or public transportation if a car is in the shop. At the highest user level, it is the reliability of the user service that is critical, not the reliability of any particular system. The user may be served by a system of systems, including alternate hardware and procedural work-arounds. The multiple assemblies in composite systems may have different failure modes and their failure rates and pdfs may have complex time plots.

4. International Space Station (ISS) life support system reliability

The International Space Station (ISS) life support system was designed to be maintained using spare assemblies called Orbital Replacement Units (ORUs). The number of spares made available is based on the failure rates, which have been determined using manufacturers' estimates and operational experience. Some ORUs have been identified as having limited life and are replaced by schedule. The reliability model used for ISS thus includes constant failure rates and limited life or wear out, but not infant mortality and common cause failures. However, the usual simple reliability model that assumes a low constant failure rate does not seem appropriate for ISS systems such as life support. The systems designers explicitly relied on on-board repair and so did not implement the intensive design effort and conduct the extensive testing needed to eliminate most infant mortality and common cause failures. In some cases the failure rate and required number of spares have been higher than originally anticipated. The ISS system development approach seems excessively ambitious, with exaggerated hopes of engineering capability, funding, schedule, and good luck. The work is amazingly hard, expensive, long, and complex. Providing life support for the ISS crew has depended on systems of systems effects, such as supplies of material and spare parts from Earth, reducing crew size and support, and relying on the Russian life support system.

5. Mars life support system reliability

The ISS life support system is the only operational recycling US life support system ever flown in space. It provides valuable lessons for Mars transit life support. (Bagdigian et al., 2015-094), (Takada et al, 2015-115) and others have suggested that life support systems similar to those on ISS can be used for Mars transit, although with extensive redesign to correct known problems. Redesigning and upgrading systems flying on ISS seems especially slow and difficult, especially compared to improving the space shuttle, which was landed and refurbished after each flight. It seems possible that ISS redesigns will introduce a second generation of design errors that will require

further time and effort to correct. It seems very difficult to conduct sufficient testing to eliminate infant mortality and common cause failures. If the estimated failure rate is unrealistically low, or if the likelihood of common cause failures is ignored, the number of spares provided will be too low to keep the systems operational throughout the Mars mission.

As difficult as it is to redesign the systems now operating on ISS, it will be nearly impossible to redesign systems in transit to Mars. Repairs will have to be accomplished with spares and tools already on board. Making this possible requires an intensive design effort and extensive testing to eliminate most infant mortality and common cause failures.

Since it is impossible to receive materials, spares, and new assemblies on the way to Mars, and since the crew cannot return if life support fails, Mars transit has much higher risk than ISS. Mars transit life support must be much more reliable than ISS life support. The engineering approach should be much more conservative, safe, constrained, and technically unambitious. A user focused systems of systems reliability approach similar to that used on ISS should be employed. It could include high reliability systems, stored materials, diverse hardware, and perhaps contingency responses. Using the too simple low constant failure rate reliability model could lead to disaster.

IV. International Space Station (ISS) Oxygen Generation Assembly (OGA) reliability

What does the International Space Station (ISS) life support system reliability data show? Can the reliability models provide insight? We consider the Oxygen Generation Assembly (OGA) now on the International Space Station (ISS).

A. Oxygen Generation Assembly (OGA) reliability data

Two papers summarize the failure and maintenance events of the OGA now on the ISS. (Takada et al., 2015-115 (Bagdigian et al., 2015-094). The operational problems reported in these papers were consolidated and categorized as shown in Table 1.

Table 1. Oxygen Generation Assembly (OGA) operational failures and maintenance events.

Year	H2 sensor	Cell stack	ACTEX	TOC	Other	Year total
2009	4	5	0	0	5	14
2010	4	6	0	3	19	32
2011	3	0	2	0	12	17
2012	5	0	2	1	6	14
2013	4	0	0	0	8	12
2014	4	0	1	0	6	11
Totals	24	11	5	4	56	100

TOC is total organic carbon. The data of Table 1 is plotted in Figure 4.




Figure 4. Oxygen Generation Assembly (OGA) operational failures and maintenance events.

Oxygen Generation Assembly (OGA) problems show an initial increase, partly due to the cell stack problem, and then a gradual decrease. Table 1 and Figure 4 show the number of failures per year in a single operating system, so the plotted number of problems corresponds to the OGA system failure rate. A failure rate increase followed by a decrease is the direct up-side-down inverse of the traditional bathtub curve shown in Figure 3. The OGA failure rate can be approximated using two of the Weibull failure rate curves in Figure 2, a constant failure rate, $k = 1$, and an early, years 2009-2010, limited life bell curve, $k > 1$. These are characteristic of the hydrogen (H₂) sensor and cell stack problems respectively. We clearly have an unusual reliability case. The hydrogen (H₂) sensor, cell stack, ACTEX, TOC, and other problems are discussed.

1. Hydrogen sensor degradation

The hydrogen sensors are triple redundant. Hydrogen sensor replacement was scheduled at 150 day intervals because of calibration limits. Some hydrogen sensors had excessive drift and this was accommodated by an operational procedure change. Improper sensor readings were tolerated if only one of the three sensors failed. Because of this decision, hydrogen sensor replacements have not been necessary before the originally scheduled 150 day intervals. The hydrogen sensor replacements occur about twice per year, but they have required other maintenance attention. The hydrogen sensor replacement, maintenance, and calibration included 24 of the 100 failures and maintenance events. A design change is planned to replace the hydrogen sensors by a hydrogen-oxygen recombiner.

The hydrogen sensor replacement occurs at the constant scheduled rate and is not a failure response. The failures are not of great concern. However, they do indicate an unanticipated design deficiency that caused a sequence of common cause failures and requires a significant redesign using a new approach. Common cause failures and redesigns are indications of an immature design, one containing undetected flaws because it was insufficiently tested before operation.

2. The cell stack failure

The most significant problem was the cell stack degradation and failure. Initially, the OGA recirculating loop pressure increased, apparently because the filters were clogged, and the filters and the Water ORU were replaced, ultimately using new design filters. The problem continued and high pH was noted in the OGA recirculating loop.

This seemed to produce corrosion products that clogged the loop filters. Ultimately the cell stack failed. The failure mechanism was as follows: the electrolysis cell membranes typically degrade, they produce acid and low pH, this caused corrosion products that blocked the filters and contaminated the cell membranes, this increased their resistance, driving up the voltage to the shutdown limit.

The electrolysis cell stack was replaced. Filters, the pump ORU, and the water ORU had been replaced during trouble shooting. After these replacements, the cell stack problem was finally cured by adding a new deionization bed into the water recirculating loop to remove the acid and contaminants. This deionizing bed is the ACTEX, for Activated Carbon/Ion Exchange.

The cell stack failure group includes 11 of the 100 failures and maintenance events. They all occurred during the first year and a half of full operation. The ACTEX replacement events in subsequent years represent continuing effects of the cell stack failure.

The cell stack failure was not infant mortality, it occurred over 18 months, and represents an unanticipated form of wear out. This is another classic case of a sequence of common cause failures that was due to a design error and that could only be cured by a significant design change. Takeda et al. noted that, "Published literature for fuel cell technology identified that there is a known chemical degradation of the cell membrane polymer chain end groups during normal operation." More extensive initial research and ground testing might have discovered this problem before flight. (Takada et al, 2015-115)

3. The ACTEX bed

The ACTEX deionizing bed required installing and several replacements. The difficult task of modifying a flight system was made necessary by the cell stack degradation problem. The ACTEX actions include 11 of the 100 replacement and maintenance events. Adding the ACTEX is a redesign required by the cell stack failure, and can be traced to the same design oversight and failure to test.

4. Total organic carbon (TOC)

The total organic carbon (TOC) in the recirculating loop sometimes increases, apparently randomly, and probably due to contamination of the feed water. The TOC can be reduced by replacement of the ACTEX deionizing bed. It is proposed to add a new capability to periodically refresh the recirculation loop water using a bleed procedure to directly reduce TOC. The TOC actions include 4 of the 100 replacement and maintenance events. Not anticipating the need to refresh the OGA recirculating loop is a specification oversight, as coping with contaminated feed water was not included.

5. Other maintenance events

The other maintenance events in Table 1 and Figure 3 include 56 of the 100 events. They follow the same time pattern as the sum of the hydrogen sensor, cell stack, ACTEEX, and TOC problems, an initial increase in the first two years and then decrease. Many of the other problems are related to the hydrogen, cell stack, ACTEEX, and TOC problems. 18 of the 56 events are repair or replacing ORUs, and the other 38 are trouble shooting including water sampling and cleaning.

B. Modeling and discussion of OGA failure data results

The OGA failure results are related to the Weibull model and their implications for Mars life support are discussed.

1. Weibull modeling

Figure 5 plots the OGA failure replacement data without the scheduled hydrogen sensor replacements.




Figure 5. OGA failure replacements with a Weibull rate fit.

The data plotted earlier in Figure 4 contains 50 component and ORU replacements. 21 of these are scheduled hydrogen sensor replacements that occur at the constant rate of three or four per year, and which are not shown in Figure 5. Only the other 29 failed component and ORU replacements are shown in Figure 5. The Weibull rate fit has two segments. The initial increase is fitted to the Weibull failure rate equation with $k = 1.83$, corresponding to wear out, and $a = 2009.5$, mid 2009. The later year data has a roughly constant failure rate, $k = 1$, and less than one replacement per year.

The initial high failure rate was largely due to the slowly developing cell stack degradation problem and looks like an early wear out or a dampened infant mortality. The initial high failure rate era extended for two years, longer than a Mars round trip, and was followed by many years with low constant failure rate, as is usually expected after initial burn in.

2. *We can't count on repairing all failures going to Mars*

Half or more of the Oxygen Generation Assembly problems are due to design errors and common cause failures that more typically appear during initial operation, infant mortality, burn in, or debugging. These include the hydrogen sensor, cell stack, ACTEEX, and TOC problems and led to many of the other events, such as water sampling, loop flushing, and cleaning. Well-designed and tested mature systems typically have less than 10% of failures due to a common cause. It is practically unheard of to have a level of 50% common cause failures, and this suggests that insufficient attention was paid to design and test.

The large percentage of common cause failures seen here definitely contradicts the usual engineering assumption that most failures are random and independent and so can be effectively repaired using spare parts. We can't count on repairing all failures going to Mars. Design errors require redesign, which is not done with identical spares.

3. *Can we use the ISS Oxygen Generation Assembly for Mars?*

An interesting idea that has been suggested is to use the planned redesigned ISS Oxygen Generation Assembly for Mars transit. (Takeda, et al., 2015-115) The active operational duration of a Mars transit mission is about 15 months, possibly interrupted by a long crew surface stay. Suppose that the original ISS OGA had been flown to Mars. Since the cell stack operated for 18 months on ISS, it would not have failed during the Mars transit. And presumably, as on ISS, there would have been a spare cell stack in case of failure. However there would have been the time consuming and worrisome problems caused by cell stack degradation. The ACTEX has fixed this problem. However, the ISS OGA reliability is not high and the planned redesigns may introduce new problems.

The selection and design of an oxygen system for Mars is a complex systems engineering problem. Oxygen electrolysis has a hydrogen safety hazard and similar Russian systems have failed. The ISS OGA has high mass,

probably higher than the mass of the oxygen that it would produce on a Mars transit. (Jones, 2016-103) The ISS OGA is planned to have extensive redesign, which almost always introduces new errors. Alternate approaches have been considered in the past, and probably should be revisited.

4. *Should we use the ISS life support development approach for Mars?*

The ISS reliability and maintenance approach are not suitable for Mars. The ISS resources for design and testing were limited, probably too much so. It was decided to rely on crew repair, the Russian systems, and resupply and abort responses only available in low Earth orbit.

Mars life support needs thorough design and long testing to provide and demonstrate much higher reliability than ISS has had. We cannot assume, as did ISS, that newly designed or redesigned systems with little testing and limited trouble-free operational experience can be successfully operated by relying on repair using spare parts. Traditional NASA systems engineering is needed to reduce risk and provide cost effective safe operation.

V. Conclusions

It is often assumed that system failures during operations can always be repaired. This is the highly optimistic best case, corresponding to a reliable, mature, well-tested design. Analysis assuming all failures can be repaired can be useful, especially if it proves a system is insufficiently reliable because it requires excessive spares, repairs, and maintenance. But this analysis is not sufficient to prove a system will be reliable, since the assumption of complete reparability can be false.

Design errors happen, and using redundant systems or repairing systems using spares cannot fundamentally cure them. Redundancy and repair using spares can be defeated by design errors, external events, manufacturing errors, and other kinds of common cause failures.

The International Space Station (ISS) Oxygen Generation Assembly (OGA) experienced a group of common cause failures that were due to an oversight and required a design change. Cell membrane degradation over time was known in the literature and would have been found by a reasonably long preflight ground test. It was fixed by a design change that provided filtration to remove degradation products.

If a system design error causes a failure on the way to Mars, it can't be fixed by repairing the system using spares. Assuming that all failures can be repaired using spares has high risk for systems that are new or redesigned and lack extensive testing and operational experience

References

- Bagdigian, R. M., Dake, J., Gentry, G., and Gault, M., "International Space Station Environmental Control and Life Support System Mass and Crewtime Utilization In Comparison to a Long Duration Human Space Exploration Mission," ICES-2015-094, 45th International Conference on Environmental Systems 12-16 July 2015, Bellevue, Washington.
- Borcsok J., Schaefer, S., and Ugljesa, E., "Estimation and Evaluation of Common Cause Failures," IEEE, Second International Conference on Systems (ICONS'07) 2007. Ref 04196343.pdf
- Bukowski, J. V., and W. M. Goble, "Verifying common-cause reduction rules for fault tolerant systems via simulation using a stress-strength failure model," ISA Transactions 40, 2001.
- Gaudard, M., and Wright, L. "Explaining Reliability Growth," wp-explaining-reliability-106026, SAS White Paper, SAS Institute Inc., downloaded Dec. 17, 2016
- Hansen, Robert C., Overall Equipment Effectiveness, Industrial Press Inc., New York, New York, 2001.
- Jones, H. W., "The International Space Station (ISS) Oxygen Generation Assembly (OGA) Is Not Feasible for Mars Transit," ICES 2016-103, submitted to 46th International Conference on Environmental Systems, 10-14 July 2016, Vienna, Austria.
- Rutledge, P. J., and A. Mosleh, "Dependent-Failures in Spacecraft: Factors, Defenses, and Design Implications," IEEE, 1995, Proceedings Annual Reliability and Maintainability Symposium.
- Scholz, F. W., "Inference for the Weibull Distribution," WeibullBounds, University of Washington, May 22, 2008, downloaded Jan. 4, 2016.
- Shishko, R., NASA Systems Engineering Handbook, NASA-SP-6105, June 1995.
- Stotta, J. E., Britton, P. T., Ring, R. W., Hark, F., and Hatfield, G. S., "Common Cause Failure Modeling: Aerospace vs. Nuclear," http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20100025991_2010028311.pdf 20100025991_2010028311.pdf
- Takada, K. C., Ghariani, A. E., and Van Keuren, S., "Advancing the Oxygen Generation Assembly Design to Increase Reliability and Reduce Costs for a Future Long Duration Mission," ICES-2015-115, 45th International Conference on Environmental Systems, 12-16 July 2015, Bellevue, Washington.
- Wikipedia, Weibull distribution, https://en.wikipedia.org/wiki/Weibull_distribution, accessed Jan 7, 2016