Developing Reliable Life Support for Mars

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

A human mission to Mars will require highly reliable life support systems. Mars life support systems may recycle water and oxygen using systems similar to those on the International Space Station (ISS). However, achieving sufficient reliability is less difficult for ISS than it will be for Mars. If an ISS system has a serious failure, it is possible to provide spare parts, or directly supply water or oxygen, or if necessary bring the crew back to Earth. Life support for Mars must be designed, tested, and improved as needed to achieve high demonstrated reliability. A quantitative reliability goal should be established and used to guide development. The designers should select reliable components and minimize interface and integration problems. In theory a system can achieve the component-limited reliability, but testing often reveal unexpected failures due to design mistakes or flawed components. Testing should extend long enough to detect any unexpected failure modes and to verify the expected reliability. Iterated redesign and retest may be required to achieve the reliability goal. If the reliability is less than required, it may be improved by providing spare components or redundant systems. The number of spares required to achieve a given reliability goal depends on the component failure rate. If the failure rate is underestimated, the number of spares will be insufficient and the system may fail. If the design is likely to have undiscovered design or component problems, it is advisable to use dissimilar redundancy, even though this multiplies the design and development cost. In the ideal case, a human tended closed system operational test should be conducted to gain confidence in operations, maintenance, and repair. The difficulty in achieving high reliability in unproven complex systems may require the use of simpler, more mature, intrinsically higher reliability systems. The limitations of budget, schedule, and technology may suggest accepting lower and less certain expected reliability. A plan to develop reliable life support is needed to achieve the best possible reliability.

Nomenclature

ISS   = International Space Station
LEO   = Low Earth Orbit
MADS = Maintenance and Analysis Data Set
MTBF = Mean Time Before Failure
ORU   = Orbital Replacement Unit

I.   Introduction

THIS paper suggests an approach to develop reliable human life support systems for Mars. A quantitative reliability goal is assumed. The reliability process begins with designing for reliability, testing to discover failures, and redesigning to remove correctable faults. If the system reliability is still insufficient, spare parts or redundant systems can be used to reach the reliability goal. An example generic recycling water processor is used to illustrate developing reliable life support for Mars.

II.   A process for developing reliable life support

The usual problem in reliability engineering is to improve a well known and long used system. In this usual case, the failure modes are known and the cost and benefit of eliminating them is understood. The reliability improvement is usually incremental, difficult, and expensive, but the process rarely creates major new failure modes.

---

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

Developing life support or other systems for Mars is very different. There are few or only one flight unit. Preflight testing is very limited. The systems often have unexpected problems that require redesign. Some ISS life support systems may not require substantial redesign for Mars, but others will, due to the different mission duration, launch cost, and reliability requirements for Mars.

Developing life support for Mars, especially complex recycling systems, will require a new approach to design substantially new systems for high reliability. The life support requirements to provide oxygen and water and remove carbon dioxide are unchanging. The current systems architecture of partially interconnected water and oxygen loops may have to be simplified. Certainly the hardware developed using the materials and techniques of thirty years ago cannot and should not merely be replicated. The common sense intuition of gradually improving reliability developed over successive generations of many units does not apply here. A fully considered, end-to-end process must be used instead of expecting gradual small step-by-step improvements.

The suggested approach to develop a reliable new system is as follows:
1. Define a quantitative reliability goal.
2. Design the system to be simple, with as few components and interconnections as possible.
3. Select the components for high reliability.
4. Estimate the best possible system failure rate as the sum of the component failure rates.
5. Accept that, due to integration problems, design errors, and flawed components, the actual failure rate may be orders of magnitude higher.
6. Conduct testing to discover any unexpected high failure rates and new failure modes.
7. Redesign to remove all unexpected correctable failure modes. This is heroic failure reduction.
8. Continue testing and redesign either until the actual test failure rate approaches the originally best estimated failure rate or it is necessary to accept the current higher failure rate.
9. Continue testing until the failure rates are closely determined.
10. Use these failure rates to determine the spares or redundancy needed to achieve the reliability goal.

Much of this suggested process is familiar, but there are two new aspects, heroic failure reduction and the needed test duration. Heroic reliability improvement is described in another paper. (Jones, 2017-86) The essential idea is that, if design or component errors occur, they will cause unexpected higher failure rates or new failure modes that will appear during early testing. An idealized process of heroic failure reduction will correctly diagnose and remove all these unexpected failure modes without introducing any new ones. The failure rate then drops rapidly. The heroic failure rate reduction process can be modelled and its expected duration predicted. The expected time needed for the heroic failure reduction process determines the required test duration. If the heroic reliability improvement process is not fully implemented, the system failure rate will be higher than estimated from the sum of the component failure rates. Another reason for extended testing is the need for accurately estimating the failure rates used to compute the number of spares provided. It cannot be assumed that a system will operate much longer than testing has proved it will, even if it has high estimated but unproven reliability.

There are clear development process trade-offs and risks. The less the heroic reliability improvement, the higher the failure rates and the more redundancy needed. The shorter the test, the lower the demonstrated reliability, and again the more redundancy needed to achieve the reliability goal. If a new system is not well tested, it is likely to have unexpected failure modes and surprising high failure rates. If so, the spares provided could be too few and the system may fail. In the absence of resupply or escape options, insufficient spares could result in loss of crew.

The suggested process for developing a reliable new human space system has a complex series of steps. The example of a generic ISS-like recycling water processor is used to illustrate the process.

III. Life support reliability requirements

How reliable must life support be for Mars? The NASA Astronaut Office requested that future launch vehicles have a probability of loss of crew of at most 1 in 1,000 flights, which was a roughly order of magnitude decrease from the then operational space shuttle. (NASA Astronaut Office, 2002) It would then seem reasonable that that the probability of a loss of life support should also be less than 1 in 1,000 per mission. Considering that there are many different life support functions, such as providing oxygen and water, removing carbon dioxide, maintaining the atmosphere, controlling temperature and humidity, and suppressing fire, the probability of losing any particular life support function should then be less than 1 in 10,000.

The assumed reference mission will be a Mars transit requiring about 500 days out and back. For a subsystem failure probability of less than 1 in 10,000 over a 500-day mission, the subsystem failure rate must be less than 2.00E-07 per day or 8.33E-09 per hour. The Mean Time Before Failure (MTBF) is the inverse of the failure rate. For a failure probability of less than 1 in 10,000 over a 500-day mission, the subsystem MTBF must be greater than

5.00E+06 days or 1.20E+08 hours. Five million days is about fourteen thousand years. And to achieve any given subsystem reliability, the failure rates of the components in the subsystems must be roughly an order of magnitude smaller and their MTBFs an order of magnitude longer. Such reliability goals seem impossible to achieve or even measure by test. This suggests that the reliability goal must be reduced.

Rather than top down, consider reliability bottom up. The most reliable life support components have MTBFs of less than one million hours, and more typical MTBFs are about 100,000 hours. A life support subsystem with 10 such components would have an MTBF of 10,000 hours, and a full life support system with 10 subsystems would have an MTBF of 1,000 hours, about 42 days. Such a system would probably fail a dozen times during a 500-day mission. This suggests that achieving sufficient intrinsic system reliability will be difficult and that component or subsystem redundancy will be needed to increase the reliability of Mars transit life support.

To create a plausible quantitative reliability design example, a severely reduced reliability requirement will be used. Defining an acceptable probability of loss of crew is a difficult problem, but not understanding the mathematics would be irresponsible. The space shuttle flew many times after it was accepted that the probability of loss of crew was roughly 1 in 100. Each of these shuttle missions had great value that justified the risk. It might be argued that a manned Mars mission would have a much higher value that would justify a much greater risk, as long as there is an acceptable probability of success.

Assume that, for a 500-day Mars transit mission to be undertaken, the probability of loss of life support should be less than 1 in 10 rather than 1 in 1,000. All the failure probabilities computed here then are larger by a factor of 100 and all the MTBFs cut by a factor of 100. If life support has ten subsystems, the subsystem failure rate must be less 8.33E-07 per hour or 0.000833 per thousand hours. The subsystem MTBF is 1,200,000 hours. This assumed life support system reliability requirement is much less than desired and yet very difficult to achieve, but a quantitative requirement is needed for preliminary reliability planning.

IV. System reliability example

The objective of this analysis is to understand how to develop reliable life support for deep space and Mars by making use of past experience and current knowledge. The general systems architecture of recycling life support has long been established. It was demonstrated in manned closed system ground tests as early as 1970. The International Space Station (ISS) recycling life support system has been operating since about 2009. One suggested approach to developing life support for deep space is to refine the design and improve the reliability of current ISS systems. (Bagdigian et al., 2015-094)

However, simply improving ISS life support would not be sufficient for deep space or Mars transit. The requirements and constraints are significantly different between Low Earth Orbit (LEO) and deep space. Higher launch cost requires lower mass. Greater difficulty of diagnosis, repair and upgrade requires better operability and maintainability. The impossibility of on-demand crew return requires greater reliability and safety. Deep space radiation requires radiation hardening. Waiting during Mars surface operations requires a capability for quiescent operation. (Jones, 2016-103)

Life support must be redesigned for these new requirements of deep space. However, the proven architecture and tested technology of the ISS life support system should be used to the maximum extent possible. It is assumed that the Mars transit life support system will have a recycling architecture similar to that on ISS and will use physical-chemical recycling technologies similar but not necessarily identical to those on ISS.

Developing reliable life support will be investigated using an assumed generic recycling water processor that is similar to the water processors now on ISS. The assumed types of subsystems and their reliabilities are similar to current ISS systems. (Bagdigian et al., 2015-094)

V. Generic water processor reliability

Consider a space habitat water recycling processor, somewhat similar to the US water or urine processors currently on the International Space Station (ISS). These units are designed to be maintained using Orbital Replacement Units (ORUs). Spare ORUs are kept on ISS and used to replaced failed units. The system can be kept operating as long as there are enough spares to replace all the units that fail during the mission. Sufficient spares are provided to ensure that the probability of not having a needed spare is very low, and most of the spares provided will not actually be needed.

Table 1 lists the ORUs of a generic recycling space water processor. The MTBFs, failure rates and masses are given.

Table 1. Generic water processor MTBFs, failure rates, and masses.

#	ORU	MADS MTBF, 1,000 hours	Expected failure rate per 1,000 hours	Unexpected correctable MTBF, 1,000 hours	Unexpected correctable failure rate per 1,000 hours	MADS mass, kg
1	Tank	60	0.017			50
2	Controller	40	0.025			25
3	Separator	60	0.017	4	0.250	30
4	Pump	40	0.025	1	1.000	45
5	Distiller	40	0.025	6	0.167	75
6	Reactor	30	0.033	14	0.071	60
7	Bed	400	0.003			50
8	Valve	180	0.006			4
9	Filter	400	0.003			8
10	Sensor	160	0.006			4
Totals			6.3	0.158	0.67	1.488	351

The generic water processor is assumed to have ten different ORUs. The US water processor on ISS has about fifteen ORUs and the urine processor has five ORUs. When the ISS water processors were designed, the ORU MTBFs were estimated. The MTBFs are recorded in MADS, the ISS Maintenance and Analysis Data Set. (MADS, 2015) MADs also provides the ORU masses, which are used later in minimizing the mass of the additional spares needed to achieve higher reliability. The MADS MTBFs and masses in Table 1 are approximations of the average data for the assumed generic ORUs.

The failure rate is the inverse of the MTBF. The overall generic water system failure rate is the sum of the ORU failure rates, 0.158 per 1,000 hours. The overall system MTBF is the inverse of this failure rate, 6,300 hours.

This is far short of even the lower requirement for subsystem MTBF. The lowered subsystem MTBF of 1,200,000 hours is about 200 times higher than this estimated water processor MTBF of 6,300 hours. This suggests that significant redundancy will be required to meet even the much lower, 100 times reduced reliability requirement.

But there is a more immediate problem in achieving the required reliability. The MADS MTBFs are the estimated and expected MTBFs, based on similar hardware performance and reasonable expectations. However, as is often the case, this reliability is reduced by unexpected problems, such as poor components or design miscalculations. Initial testing often finds that the actual failure rate is 10 or 100 times higher than expected and that redesign is required.

Flight experience with an ISS separator, pump, distiller, and reactor indicates they had much shorter MTBFs and higher failure rates than expected. The flight MTBFs are based on (Bagdigian et al., 2015-094). The overall failure rate is about 10 times the originally expected failure rate and the MTBF is only one-tenth that expected. These unexpected problems are considered to be correctable by redesign or component replacement. These failure modes must be removed if the originally expected failure rate and MTBF are to be achieved.

VI. Heroic reliability improvement

Systems often have excessive failure rates and unexpected failure modes when initially operated. If these failures are eliminated, reliability growth occurs. If all the unexpected failure modes and excessive failure rates are eliminated, the system failure rate will decline to the originally expected failure rate.

Heroic reliability improvement occurs if each unexpected failure mode is eliminated the first time it occurs. Eliminating the failure mode requires that it be correctly diagnosed and corrected without delay or introducing additional failure modes. For the generic water processor of Table 1, the initial failure rate would be 0.158 + 1.488 = 1.646 per 1,000 hours, largely due to the 1.488 per 1,000 hours of unexpected excess failures of the separator, pump, distiller, and reactor. These failures would probably occur not much later than their respective MTBFs of 1,000, 4,000, 6,000, and 14,000 hours. If a heroic reliability improvement effort removed the excess failure modes as they occurred, the system failure rate would drop rapidly over the first 10,000 or 20,000 hours to reach the originally expected rate of 0.158 per thousand hours. Probably none of the other ORUs would have failed, due to their longer MTBFs.

A. Failure rate data

The process of identifying, diagnosing, repairing, and preventing failures can be very complicated. Sometimes failures cause symptoms far from their source. The causes can be misdiagnosed. Failures can cascade. Often there are many failure reports and analyses before the failure cause is discovered and repaired. (Jones, 2016-103)

Reliability growth analysis uses a very simplified and abstract failure data set, a simple list of the failure times of each identified failure mode. After testing is begun, failures occur at times t1, t2, … , tn. The sum of the failures up to time t is N(t). The current measured failure rate at time t is N(t)/t. The failure rate data points that can be plotted are N(t)/t = 1/t1, 2/t2, …, n/tn.

If as often assumed, the failures are acceptable random low rate failures, N(t)/t = c, a constant. For the normal expected failure rate of Table 1, N(t)/t = c = 0.158. All the failure modes continue to occur at the same constant rate. If there are unexpected high rate failures as in Table 1, the initial failure rate is much higher, N(t)/t = 1.646. If no improvements were made, this failure rate would persist. However, if we assume heroic reliability growth, the causes of the unexpected high failure rate are each removed after their first occurrence and the failure rate drops to N(t)/t = c = 0.158. The heroic reliability growth process thus follows the "abcd" reliability growth model. (Jones, 2017-86)

B. The "abcd" reliability growth model

In the "abcd" reliability growth model, the failure rate $f(t) = N(t)/t = a\,t^{-b} + c$. The failure rate at t = 1 is f(1) = a + c. The reliability growth exponent "b" has been defined as the *heroic reliability growth effort metric*. If a heroic effort occurs, b = 1. To see this, suppose that there are N = 4 correctable failure modes, as assumed for the generic water processor of Table 1. The number of occurring unexpected and correctable failures N(t) will increase from 0 to 4 over the first 15,000 or 20,000 hours. If each is corrected when it first occurs, the measured number of correctable failures will not increase above 4. The correctable failure rate is $N(t)/t = 4/t = 4\,t^{-1}$. For heroic reliability improvement, b = 1. For lesser effort, b < 1. For no effort, b = 0 and the failure rate does not decline from the initial value of 4.

The initially expected failure rate is "c," which is also the hoped-for final failure rate after all the unexpected failure modes are corrected. The rate "c" will be achieved in time, as long as the reliability growth effort metric, b > 0. If the process is terminated at t_d, the remaining correctable failure rate will be $d = a\,t_d^{-b}$. The total final failure rate is then c + d. The intended final failure rate is "c," the same as the initially expected failure rate, and it can be achieved if all unexpected failure modes are removed. (Jones, 2017-86)

C. An upper bound on failure rate for heroic failure rate reduction

Analysis shows that, if heroic failure rate reduction is implemented, there is an upper bound on the current failure rate. This has been shown for both a discrete set of correctable failure modes and a continuous distribution of failure rates. It is assumed that each failure mode has a constant failure rate. With heroic failure rate reduction, the probability that the failure is still undetected and uncorrected declines exponentially over time to zero. The maximum failure occurs at the MTBF. The resulting bound on the failure rate is $f(t) \le 1/(e\,t)$, where e is the base of the natural logarithms and t is time. (Bishop and Bloomfield, 1996) (Jones, 2017-86)

For each individual failure mode, the expected failure rate declines exponentially with time. The current expected failure rate is always less than, better than, the bound 1/(e t), regardless of the original failure rate. This is true under the heroic assumption that a failure is immediately corrected, without introducing a new failure mode. The bound directly decreases with increasing test and failure reduction time. The bound is surprising because it proves reliability growth must occur under the heroic assumption and because it allows the maximum future failure rate to be predicted from the current failure rate.

The bound $f(t) \le 1/(e\,t)$ is for one single failure mode. The upper bound on the failure rate for N failures is $f(t) \le N/(e\,t)$. This is an upper bound on the heroic failure rate. If all failure modes are removed when they first occur, the actual failure rate will be less. But if the actual test and redesign process falls short of the ideal heroic failure rate reduction, the bound is more like a lower bound on the failure rate, since continuing uncorrected failures will raise the failure rate over this bound. The bound does not include the continuing constant expected failure rate "c." (Bishop and Bloomfield, 1996) (Jones, 2017-86)

VII. Heroic reliability growth for the generic water processor

During a reliability growth effort, the time of each successive failure is recorded and a cumulative count is kept, N(t). The cumulative failure rate N(t)/t is tracked and compared to the originally expected failure rate c. Since each

excess failure mode is removed when it first appears, the decline in failure rate directly tracks the occurrence of the failures. To illustrate heroic reliability growth for the generic water processor, we assume that failures occur exactly at the MTBF, rather than randomly with an average interval equal to the MTBF. The expected and uncorrected failures reoccur indefinitely. It is assumed that the excess failures are corrected at the first occurrence.

Figure 1 shows constructed water processor failure rate data, the "abcd" model, and the heroic failure rate reduction upper bound on the failure rate.
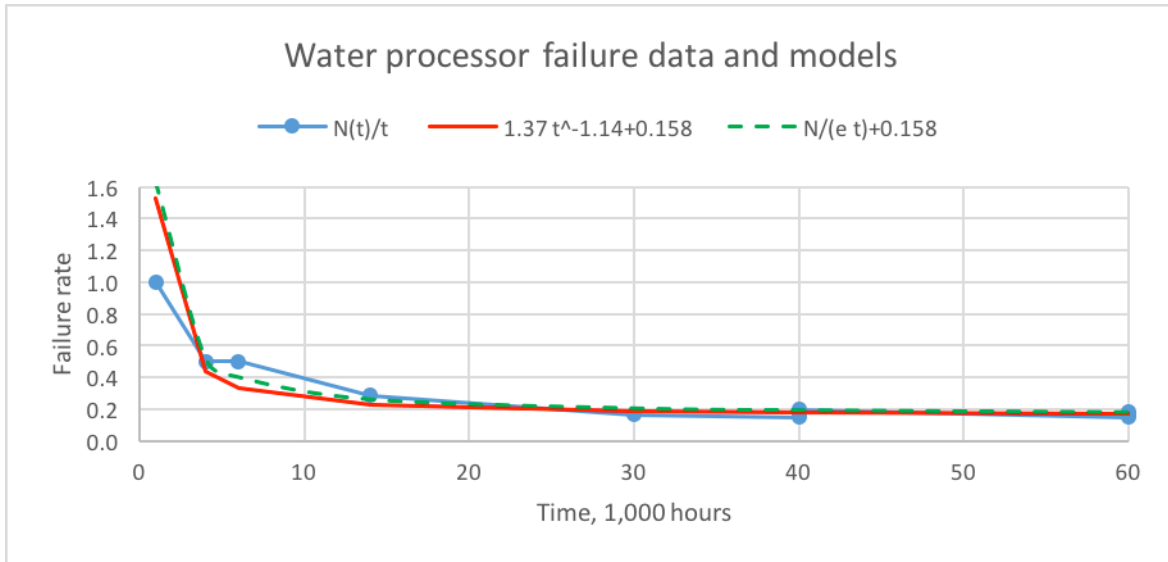

Figure 1. Water processor failure rate data and models.

The water processor failure rate data, N(t)/t, was constructed from the MTBF data. Each ORU in Table 1 was assumed to fail the first time at its MTBF. The unexpected high failure rates were assumed to be corrected at their first occurrence, at 1,000, 4,000, 6,000, and 14,000 hours. Failures then continue for the other ORUs at intervals equal to their expected MTBFs. After 30,000 hours the failure rate is essentially equal to c = 0.158 per thousand hours. The "abcd" model is $f(t) = N(t)/t = 1.37\ t^{-1.14} + 0.158$. The parameters "a" and "b" are computed by curve fitting to N(t)/t. The parameter "b" is 1.14 which is greater than 1 and indicates heroic failure reduction. The failure rate is slightly lower than the upper bound of 4/(e t) + 0.158. The "abcd" model and the N/(e t) upper bound are nearly identical and very close to the data N(t)/t.

VIII.   Spares needed to achieve the required reliability

How many spares are needed to achieve the required life support reliability? The number of spares and their total mass are computed for two different sets of MTBFs. The first set of MTBFs consists of the originally expected failure rates of Table 1. It is assumed that these apply after the heroic failure reduction process removes all the unexpected high failure rates and corresponding low MTBFs. The second set of MTBFs includes the unexpected low MTBFs due to the higher than expected failure rates for four ORUs. It is assumed that none of these high rate failure modes are removed and that heroic failure rate reduction does not occur.

A. Calculating the number of spares

A Mars recycling life support system is expected to replace failed units using stored spares. The operating units are assumed to fail at the predicted rate. The units are assumed to fail only in normal operation and at the expected rate, and not to fail in storage or at turn-on. Life support is lost if a failure occurs and no spare is available. The assumed overall life support system reliability requirement is a probability of failure of less than 10% over the mission length of 500 days or 12,000 hours. Each of the ten life support subsystems, such as a water processor, must have a probability of failure of less than 1% over 12,000 hours.

The probability of some given number of failures over some duration when using non-operating spares is given by the Poisson distribution. The Poisson distribution is used to compute the number of spares needed to have less than 1% probability of not having all needed spares during a 500-day mission.

Each ORU is required to operate over the mission length L = 500 days or 12 thousand hours. The ORU has a failure rate, rate f = 1/MTBF, the number of times it is expected to fail per thousand hours. The number of failures from time 0 to time t is n(t) and has a Poisson distribution.

The Poisson pdf gives the probability, for failure rate f = 1/MTBF, that there will be exactly n(t) = x failures in time t.

$$\text{Poisson pdf } [n(t) = x] = (f\,t)^x\, e^{-f\,t}/x!$$

The Poisson distribution's mean value, which is the expected number of failures during the mission of length L, is $f * L = L/MTBF$. The probability of n(L) or fewer failures over the mission duration L is the summation of the Poisson pdf from x equals 0 to n(L). The probability of n(L) or fewer failures is the probability that the ORU can be kept operating if the ORU is augmented with n(L) spares. The Poisson distribution is used to compute how many spares are required to have less than any particular probability of having too few spares.

B. Spares needed if all correctable faults are removed

Table 2 shows the required numbers of each ORU, the original and spares, needed to reduce the total probability of not having all necessary spares to less than 1% over a 12,000 hour Mars transit mission. The Table 2 number of ORUs, original and spares, is based on the originally expected MTBFs, which are listed.

Table 2. Generic water processor ORUs, expected MTBFs, and number and mass of spares.

#	ORU	Expected MTBF, 1,000 hours	ORU mass, kg	Total ORUs, original and spares	Total ORU mass, kg	Probability of too few spares
1	Tank	60	50	3	150	0.00115
2	Controller	40	25	3	75	0.00360
3	Separator	60	30	4	120	0.00006
4	Pump	40	45	4	180	0.00027
5	Distiller	40	75	3	225	0.00360
6	Reactor	30	60	4	240	0.00078
7	Bed	400	50	2	100	0.00044
8	Valve	180	4	3	12	0.00005
9	Filter	400	8	3	24	0.00000
10	Sensor	160	4	3	12	0.00007
Totals			351	32	1,138	0.01001

The Poisson probability distribution is used to calculate the probability that there will be too few spares for each of the ten ORUs in the generic water processor. The sum of the ten probabilities must be less than 0.01 to meet the assumed requirement. The Poisson probability depends on the ratio of the mission length, L, to the MTBF as well as on the number of spares. A cut and try process is used to add and distribute the spares so as to minimize their total mass, which is also calculated in Table 2.

The ten original operating OURs must be supported with 22 spares, a total of 32 units. The total number of ORUs is 3.2 times the original ten ORUs. The total mass is about 3.2 times the mass of a single set of ORUs. Triple redundancy is reasonable to expect for Mars life support. (Connolly, 2000)

C. Spares needed if all correctable faults remain

The presumed ideal situation is as described in Table 2. The initially expected MTBFs are assumed to be correct and can be relied on to compute the number of spares. However, the heroic failure reduction process may not be completed before the Mars life support system is flown. The ISS life support had very limited integrated testing before flight and unexpected correctable failures similar to those assumed here occurred during flight. The heroic failure reduction process seems very difficult to implement during flight. Many correctable ISS failure modes have not yet been corrected. (Bagdigian et al., 2015-094) (Jones, 2017-86)

Suppose that the testing of the assumed generic water processor does detect all the unexpected correctable failure modes and that their MTBFs can be estimated, but that these correctable failure modes are not removed. To achieve the required reliability, additional spares must be provided for the ORUs with the higher than expected failure rates.

Table 3 shows the required numbers of ORUs, the original and spares, necessary to reduce the total probability of not having all needed spares to less than 1% over a 12,000 hour Mars transit mission. But here the number of

spares is based on the uncorrected MTBFs for the four ORUs with unexpectedly high failure rates. The originally expected MTBFs are used for the other ORUs.

Table 3. Generic water processor ORUs, uncorrected correctable MTBFs, and number and mass of spares.

#	ORU	Uncorrected MTBF, 1,000 hours	Expected failures in 12,000 hours	ORU mass, kg	Total ORUs, original and spares	Total ORU mass, kg	Probability of too few spares
1	Tank	60	0.20	50	3	150	0.00115
2	Controller	40	0.30	25	4	100	0.00027
3	Separator	4	3.00	30	10	300	0.00110
4	Pump	1	12.00	45	23	1,035	0.00305
5	Distiller	6	2.00	75	8	600	0.00110
6	Reactor	14	0.86	60	5	300	0.00190
7	Bed	400	0.03	50	2	100	0.00044
8	Valve	180	0.07	4	3	12	0.00005
9	Filter	400	0.03	8	3	24	0.00000
10	Sensor	160	0.08	4	3	12	0.00007
Totals			18.56	351	64	2,633	0.00912

Again the Poisson probability distribution is used to calculate the probability that there will be too few spares. The sum of the probabilities for the ten ORUs must be less than 0.01. In this case, the total number of ORUs is 6.4 times the original ten ORUs and the total mass is about 7.5 times the mass of a single set of ORUs. The effect of the much lower than expected MTFs occurs largely in the sparing of the affected ORUs. The number of spares for these four ORUs goes from 4, 4, 3, and 4 to 10, 23, 8, and 5. If the MTBFs of some components are much lower than originally expected, the required number and mass of spares will be much higher than originally expected. The required mass of ORUs in Table 3 is 2.3 times larger than in Table 2. This is the mass penalty for not implementing heroic failure rate reduction.

D. What if the too high failure rates are not discovered?
The nominal best case of Table 2 occurs when the initially expected failure rates and MTBFs are achieved by heroic failure rate reduction. Then a reasonable level of redundancy can achieve the needed mission reliability. If higher than expected failure rates occur but are not corrected, the additional spares and a higher level of redundancy of Table 3 are needed.

But the most difficult case occurs if the higher than expected failure rates of Table 3 are not discovered and the number of spares is assigned using the originally expected much lower failure rates and much longer MTBFs of Table 2. The inevitable result would be that an insufficient number of spares would be provided and that the life support system would fail. The probability of this is shown in Table 4, which uses the spares assignment of Table 2, based on the more favorable expected MTBFs, but calculates the probability of having too few spares using the higher than expected MTBFs of Table 3.

Table 4. Probability of too few spares with the low uncorrected MTBFs but spares assuming corrected MTBFs

#	ORU	Expected and correctable MTBF, 1,000 hours	Total ORUs, original and spares, based on unmade corrections to MTBF	Probability of too few spares
1	Tank	60	3	0.00115
2	Controller	40	3	0.00360
3	Separator	4	4	0.35277
4	Pump	1	4	0.99771
5	Distiller	6	3	0.32332
6	Reactor	14	4	0.01145
7	Bed	400	2	0.00044
8	Valve	180	3	0.00005
9	Filter	400	3	0.00000
10	Sensor	160	3	0.00007
Totals			32	0.99901

The probability of having too few spares of ORU #4, Pump, is 0.99771. Only 4 total ORUs were provided but 23 are needed. The overall probability of having too few spares is 0.99901. The probability of success is one in a thousand.

IX. Implications for developing reliable life support

This paper first presented an ideal process for developing reliable life support and then worked through a system reliability example that used a water processor and development process similar to that for ISS. High reliability life support is required for Mars transit but not as much for ISS, since ISS has the possibility of on-demand resupply and emergency crew return. The difficult suggested process for developing reliable life support for Mars was not needed for ISS or even considered. It seems clear that, if Mars life support is developed in the same way as was ISS life support, without extreme emphasis on high reliability, the probability of life support failure may be far too high.

Developing reliable life support for Mars should start with a quantitative reliability requirement and end with the verification that the requirement has been met. The system design process should emphasize achieving high reliability over enhancing other performance parameters or minimizing resources such as cost, schedule, or crew time. An extensive testing and failure rate reduction process is necessary to correct unacceptably high failure rates due to design oversights and other problems and to establish the failure rates and MTBFs that will be used to assign the number of spares.

It is difficult to establish how much testing is needed, especially considering that it is natural to assume no mistakes have been made and the system will work as planned. But testing is needed to provide confidence that the system will work as planned. If the testing and failure rate estimation and reduction process is too limited, the cost-risk trade-off will be less favorable and either cost or risk or both must increase.

Suppose a system has a set of ORUs with their expected MTBFs. The MTBFs range from a few times the mission length to many times. The computed number and mass of spares based on these MTBFs is acceptable. If no testing is done, there is a substantial risk that one or more of the MTBFs will be too low, so the number of spares will be too low, and the system will fail. Considering the history of similar systems and the current development process, this risk may seem unacceptably high so that long duration testing is necessary. The testing should probably extend to one or two times the mission duration to detect any unexpectedly low MTBFs. Ideally, all excessive and correctable failure rates should be removed by a heroic failure reduction process. If some of the unexpected lower MTBFs are acceptable, the number and mass of spares must be increased to regain the required reliability. If little or no failure rate reduction is done, the design may have undiscovered design errors and common cause failure modes. This would suggest using diverse design and dissimilar redundancy, which multiplies the design and development cost. If unexpectedly low MTBFs do not occur or are all corrected, further testing should be conducted to demonstrate the higher, longer MTBFs. For the highest confidence, only the MTBFs verified by test should be used to estimate the needed spares.

X. Conclusion

Developing reliable life support for Mars will be difficult. A well planned, top-down systems engineering process is needed. Some key steps are defining the reliability requirement, designing for high reliability, testing to establish component failure rates, and providing the spares needed to meet the reliability requirement based on the failure rates. Failure rates that are too high or have wide margins of error will increase the spares requirement. Overly limited testing may fail to detect unanticipated or high frequency failure modes that may prevent meeting the reliability requirement.

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.

Bishop, P., and Bloomfield, R., "A Conservative Theory for Long-Term Reliability-Growth Prediction, IEEE Transactions on Reliability, vol. 45, no. 4, 1996.

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.

Jones, H. W., "Heroic Reliability Improvement in Manned Space Systems," submitted as ICES-2017-xxx, 47th International Conference on Environmental Systems, 16-20 July 2017, Charleston, South Carolina.

Jones, H. W., "Using the International Space Station (ISS) Oxygen Generation Assembly (OGA) Is Not Feasible for Mars Transit," ICES-2016-103, 46th International Conference on Environmental Systems, 10-14 July 2016, Vienna, Austria.

MADS, ISS Maintenance & Analysis Data Set (MADS), https://iss-www.jsc.nasa.gov/nwo/apps/mads/web/, accessed Nov 3, 2015.

Memo CB-04-044, From: CB/chief, astronaut o_ce, to: CA/director, Flight Crew Operations," May 2004. In Zwack, M. R., Contrast: A Conceptual Reliability Growth Approach for Comparison of Launch Vehicle Architectures, Ph.D. Thesis, Georgia Institute of Technology, December 2014.