High Reliability Requires More Than Providing Spares

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

It is sometimes optimistically hoped that a space life support system can be kept working throughout a long duration mission by repairing failed components, as long as sufficient spares are flown. It is usually assumed that the components have constant known failure rates. Then the needed numbers of spares can be computed to have any particular probability that all failed components can be replaced by available spares. This approach can provide high reliability if its favorable assumptions, including constant known failure rates, are satisfied. Other favorable assumptions are that the failures are statistically independent, repair will be successful without causing further failures, and all failures are due to internal component failures. These assumptions are not usually justified. The failure rates may be estimates that are inadequately verified because of insufficient testing. Failure rates may change due to materials substitutions, manufacturing changes, redesigns to fix failures, and new failures caused by redesigns. Failures that are not statistically independent may result from one common cause, such as a design or manufacturing error or a cascade of cause and effect, possibly caused by an external event such as a power outage. Repair may be unsuccessful or cause damage. Many failures occur at component interfaces or at the overall systems level, not within isolated components. Other failures causes are completely external to the system, due to assembly, maintenance, and operational errors or to unexpected environmental challenges. Replacement with sufficient spares can compensate for expected internal component failures but may not be able to cope with unpredictable design and manufacturing flaws, human errors, and environmental impacts. Reliability estimates based on providing sufficient spares to compensate for expected failures may be far too high. They are essentially upper bounds on the reliability that might be approached if many frequent but often unconsidered failure causes can be eliminated.

Nomenclature

ACTEX	=	Activated Carbon/Ion Exchange
CCF	=	Common Cause Failure
CDRA	=	Carbon Dioxide Removal Assembly
ECLSS	=	Environmental Control and Life Support System
FCPA	=	Fluids Control and Pump Assembly
ISS	=	International Space Station
MTBF	=	Mean Time Before Failure
OGA	=	Oxygen Generation Assembly
OGS	=	Oxygen Generation System
ORU	=	Orbital Replacement Unit
PRA	=	Probabilistic Risk Analysis
Pr(LOC)	=	Probability of Loss of Crew
UPA	=	Urine Processor Assembly
WPA	=	Water Processor Assembly

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

I. Introduction

THE reliability risks that engineers are aware of are usually mitigated, so that many actual failures are surprises. In general, unexpected and unwelcome events show that our assumed model of reality does not let us predict and adequately control actual reality. The common reliability model assumes that all system failures are due to component failures and that they can always be repaired using spares. System behavior can be brought closer to such expectations if careful design and testing eliminates most system level problems and if external operational and environmentally caused failures are set aside as not faults of the system itself. The model falls notably short if there are system level failures or if the system is required to be reliable under actual operational and environmental conditions.

Suppose it is assumed that a deployed system, such as recycling life support system travelling to Mars, can be kept operating throughout the mission by maintenance and repairs, as long as sufficient spares are provided. If a component has a fixed operating life of 365 days and the Mars transit requires 1,000 days, 3 units are needed. Rather than a fixed lifetime, it is usually assumed that components have a constant failure rate per unit time, so that the expected number of failures has a defined probability distribution. Logistics analysis can then estimate the number of spares required to provide any required probability that the number of spares will be sufficient. This number is computed using the estimated failure rates of each component. Considering the error or statistical variance in the failure rate estimates increases the number of spares needed. Reducing this error by testing, like reducing the failure rates themselves by redesign, reduces the needed number of spares. The total number or mass of the spares can be minimized for any particular reliability requirement.

Spares logistics analysis can be useful but it can create problems. The most basic is that, like any mathematical model, spares logistics analysis depends on assumptions which may not always apply and which may not be considered when using the model predictions. Usually the logistics model developers are careful to state the assumptions; constant failure rates, statistically independent failures, and perfect repair. In reality, failure rates may change due to manufacturing changes, redesigns to fix failures, and even errors in redesign. Failures that are not independent may result from a common cause, such as bad material or a design or manufacturing error, and even from a cascade of cause and effect, triggered by an internal or external event. Attempts to repair may be unsuccessful or cause other problems. Externally caused failures can be due to assembly, maintenance, and operational errors. Most failures are not predicted by previous reliability analysis. Many occur at the systems level or at component interfaces, not within components. One approach to account for failure causes not included in the fundamental component failure rate is simply to multiply the component failure rate by some K-factor, and to increase the number of spares accordingly. This seems wrong. There is little reason to expect that the number of system level or externally caused failures is proportional to the number of internal component failures. Providing spares is appropriate to compensate for system component failures, but design and manufacturing errors, environmental impacts, and human errors seem to call for other directly targeted remedies.

The list of serious accidents in space life support includes Apollo 1, Soyuz, and Apollo 13. None was caused by a lack of spares, but Apollo 13 was saved by an on-board built spare. The history of human space flight suggests there is a tendency to ignore risk, to assume unfavorable events will be handled, and to expect higher than reasonable reliability. The risk to safety must be considered and minimized as far as possible at every step of a program, through mission planning, systems design, testing, and operations. An overall Probabilistic Risk Analysis (PRA) of the overall mission risk should be developed and should include life support reliability and spares analysis. The system design should consider reliability at every stage of development.

Achieving high reliability requires more than providing spares. The paper will identify the objections to the idea that spares are sufficient to achieve high reliability, classify the objections in groups, describe some examples, provide data showing how the approach can fail, and suggest a better approach.

II. Objections to the idea that providing spares alone can achieve high reliability

This section reviews some previous objections to the idea that spares are sufficient to achieve high reliability. Many papers repeat the same concerns, but do not suggest methods to cope with them.

Work analyzing the International Space Station (ISS) failure experience found that it had “uncovered failure rate underestimates that would have resulted in an order of magnitude increase in risk had they not been discovered and corrected.” The risk increase is due to the probability of insufficient spares. Also, it was necessary to multiply the expected failure rate by a K factor “to account for external influences on maintenance demands such as

environmental effects, crew error, or false alarms," but the K factor was not estimated. (Owens and de Weck, 2018-198)

Work criticizing traditional approaches to space reliability found that, "Ultimately, the approach to maintenance and repair applied to ISS, focusing on the supply of spare parts, may not be tenable for deep space missions." Beyond random failures, "There are numerous other factors that may contribute to reliability and will also have to be evaluated and considered." "The K-Factor defines the probability of 'induced failures' in spacecraft components and systems." "Component K-Factors are defined based on past experience and may also involve a high degree of uncertainty for deep-space missions. Finally, there is the issue of " 'design and manufacturer errors.' A significant fraction of failures on-board ISS have been attributed to 'other' causes, including some number associated with deficiencies in design or manufacture of components." In light of "the inherent uncertainty in reliability," the authors suggest that, "rather than manifesting a more closed-loop system (e.g., water, air, waste, etc.) and all of the required spares, it will be more efficient (and likely less risky) to simply manifest the consumables themselves." (Stromgren et al., 2016-5307)

Work on the limitations of reliability for long duration spaceflight found that increasing component reliability required greatly increased test time that led to diminishing returns. The use of spare parts would be much more cost-effective. Spares reliability calculations assumed "that all repairs are completed successfully and there are no manufacturing defects, environmental effects, or unanticipated interactions." "In addition, this analysis did not include common cause failures such as manufacturing defects, human error, or environmental effects. Operational experience on the ISS has shown that unanticipated issues continue to appear in human spaceflight, even after years of testing and operations. These unanticipated issues will most likely continue to appear in future missions, especially when new environments are encountered for the first time. Common cause failures can significantly limit the risk coverage provided by spare parts and redundancy, since a common failure mode that causes an unexpectedly low reliability in a particular component is likely to also decrease the reliability of all of its spares." (Owens and de Weck, 2016-5308)

An MIT Ph.D. thesis that computed the sufficient number of spares observed that, "every additional spare contributes progressively less to the overall probability of having sufficient spares. It is important to note here that this calculation assumes that all repairs are successfully performed and that common cause failures do not occur - hence the repeated reference to this result being a "probability of having sufficient spares" rather than a "system reliability". This quantity can therefore be considered to be a lower bound on the total system reliability." (Do, 2016, p. 149) As usual, repairs were assumed to restore good-as-new condition and failures were assumed to be independent with no cascades or common cause failures.

Work on the supportability of long duration human missions emphasized the unknown unknowns. Supportability risk analysis "can only account for known or anticipated factors. Unknown or unanticipated effects can have significant implications for system supportability that must be kept in mind during system development. For example, the ISS ECLSS system has experienced several instances of component failure and/or degraded performance due to unanticipated issues, including significantly reduced component lifetimes due to dust and debris impingement, seal material degradation, component manufacturing and assembly errors, or even changes in the concentration of calcium in crew urine due to physiological changes resulting from the microgravity environment." "These 'unknown unknowns' are, by definition, unknown. However, historically they have had a detrimental impact on system supportability, introducing additional risk factors. Therefore, probabilistic assessments of risk should be seen as a bound, rather than a guarantee; when risk factors beyond the scope of the analysis (known and unknown) are included, the actual risk will almost certainly be higher." (Owens et al., 2017-5124)

A paper directly challenged the idea that all failures can be repaired using spares and emphasized the importance of Common Cause Failures. "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." "Redundancy and repair using spares can be defeated by design errors, external events, manufacturing errors, and other kinds of common cause failures." (Jones, 2016-113)

That paper continued, "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." (Jones, 2016-113)

A work investigating in-space manufacturing noted that "ISS experience has shown that unanticipated issues arise during the operation of complex spacecraft systems, even after years of testing and on-orbit experience. This is

particularly true for the regenerative Environmental Control and Life Support (ECLS) systems." "These unanticipated failures often result from environmental interactions or manufacturing defects." And, "it is unlikely that all eventualities can be identified and mitigated ahead of time for a multi-year mission to a new destination." Furthermore, the availability of only a single system "means that significant uncertainty remains in failure rate estimates." "(S)ystem failure due to insufficient spares is only one of many situations that could lead to loss of crew," but it "sets a limit on the minimum risk that can be achieved." (Owens and de Weck, 2016-5394)

An investigation of space station life support maintenance noted a "need to consider failures in the facility and in the supply chain." (Russell and Klaus, 2007) Facility failures may prevent normal system operations. The space shuttle accidents resulted in long periods where life support spares and materials were difficult to supply.

III. Different types of objections to the sufficient spares assumption

Many objections to the idea that all failures can be repaired with spares have been cited. They can be grouped into general areas or types as shown in Table 1.

Table 1. Different types of objections to the idea that all failures can be repaired with spares.

Type of objection	Specific objection
Incorrect model assumptions	
	Failure rates are not accurately known
	Failure rates are often only estimates
	Estimated failure rates are often too low
	The model ignores most failure causes
Repair process problems	
	Repairs may not be completed successfully
	Repairs may not restore good-as-new condition
Interface and system level issues	
	Failures often occur at interfaces rather than within components
	Failures are often of system level processes
	Failures are often due to component interactions
Human error	
	Design, manufacturing, and assembly errors
	Operations and maintenance errors
	Incorrect operations, maintenance, and repair procedures
Support system failures	
	Power, cooling, and other supporting systems may fail
	The supply chain may fail
Misunderstood environment	
	Difficult or unknown space and planetary environments (meteoroids, lunar dust, microgravity, radiation, Martian soil)
Unknown unknowns	
	Surprise unanticipated system problems
	Unpredictable adverse external system and environmental changes.

The expectation that failures can be repaired using a predetermined number of spares is based on highly questionable assumptions.

1. The component failure rates are accurately known.
2. Component failures account for most system failures.
3. The repair process is completely effective.
4. Interface, system level, and human problems are not significant.
5. The component failures are statistically independent.
6. There will be no surprise failures.

Failure rates are often only initial estimates loosely based on similar systems, these estimated failure rates are often too low due to favorable subjective assumptions, and later, data-based failure rates may not be accurately known due to insufficient testing. Estimated reliability is usually too high because the model ignores most failure causes. All repairs may not be completed successfully.

Failures often occur at interfaces between nominally operating components. Failures are often are seen in system level processes, not component functions, and can be due to component interactions, such as material flow or cascade effects. Human errors can cause incorrect operations, maintenance, and repair that may damage the system.

Many of the objections in Table 1 are based on ways the component failures may not be statistically independent. That is, the failures may have a cause and effect relation or one shared origin, internal or external. Common Cause Failures (CCFs) are essentially failures that are not statistically independent. With CCFs, when one component fails, its replacement is likely to fail in the same way for the same reason. CCFs prevent the usual approach to achieving high reliability by providing sufficient spares.

Design, manufacturing, and assembly errors are also CCFs that can affect most or all the copies of a component. System problems where one component damages another can damage all the replacements in the same way. Harmful operations and maintenance procedures may cause the same problem every time they are used. External systems failures are classic CCFs, damaging components in a repeatable way.

Unknown unknowns are unpredictable and inevitable, but there is a way to cope with them. Systems can be made more robust against surprise disruptions by implementing different, diverse, separate methods to accomplish the same functions or achieve the same goals. This is exactly opposite to providing identical spares to be used in identical systems.

IV. Life support examples of the problems with the sufficient spares assumption

The problems with the adequacy of sufficient spares approach are frequently mentioned and plausible, but the evidence is qualitative and anecdotal. Some of the failures that have been seen in ISS life support systems give some insight into the occurrence and impact of the different failure types considered above. ISS spares are usually Orbital Replacement Units (ORUs).

A. Design oversight Common Cause Failures (CCFs)

The failure and repair data on the ISS Oxygen Generation Assembly (OGA) has been tabulated and analyzed. (Jones, 2016-113) (Jones, 2016-103) The most significant OGA 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, and later new filters were designed and used. 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. Many 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 ACTEX was frequently replaced in subsequent years to prevent cell stack failure. This is a classic case of a sequence of common cause failures that was due to a design oversight and that could only be cured by a significant design change. The cell stack problem accounted for relative few of the maintenance and repair efforts, 16 of 100, but its impact on crew time and operational capability far outweighed all other failure modes. (Jones, 2016-113) (Do, 2016)

B. Unanticipated effects of microgravity

The ISS Carbon Dioxide Removal Assembly (CDRA) has had a continuing high rate of failures. (Jones, 2014-075) A major problem has been damage to the CDRA check valves caused by dust and debris from sorbent pellets, which travelled further than expected in microgravity. The adsorbent beds were redesigned for better containment. Unfortunately this trapped dust, which blocked air flow and required another redesign. (Do, 2016) The ISS Urine Processor Assembly (UPA) distillation assembly failed due to precipitation of calcium sulfate due to unexpectedly high concentrations of calcium in astronaut urine, due to unanticipated astronaut bone loss in microgravity. (Do, 2016)

C. Manufacturing and assembly errors

The ISS UPA Fluids Control and Pump Assembly (FCPA) had multiple failures of the due to manufacturing and assembly errors, including insufficient lubrication, parts misaligned during assembly, and incorrect machining. (Do, 2016)

V. Computing the number of sufficient spares

This section identifies the sufficient number of spares for ISS ORUs using ISS ECLSS failure rates, assuming that the required system failure rate is < 0.001. The sufficient number of spares is computed using earlier estimated ORU MTBFs and later flight data MTBFs. The MTBF is the Mean Time Before Failure in hours. The failure rate is 1/MTBF per hour. Flight failure experience often indicates that the MTBFs are shorter and the failure rates are higher than originally estimated. The higher actual failure rates mean that the number of spares based on lower estimated failure rates is insufficient. If the number of spares is based on lower estimated failure rates, but the failures occur at the higher flight data failure rates, the probability that the system will fail due to insufficient spares can exceed 99 percent.

A. Assumed mission duration

The assumed mission duration corresponds to a typical Mars round trip transit time of 450 days. Conjunction class Mars missions have outbound and return transit times of 200 to 250 days each and Mars surface stays of 400 to 550 days. (Boden and Hoffman, 2000) The total transit time that recycling life support would operate is 400 to 500 days, interrupted by a quiescent period of 400 to 550 days if all the crew is on the surface. The nominal 450 day transit duration is used in the failure probability and sufficient spares calculations below.

B. Calculating the number of spares needed using the Poisson distribution

Suppose that a particular ORU has an MTBF equal to 10 times the mission length L. Then the single system probability of failure during the mission is F = L/MTBF = L/10 L = 0.10. Further suppose that all failures can be fixed by using spare parts and that the system must have less than a 0.001 probability of not having a needed spare. This can be achieved by having one operating ORU and two spares, since the probability that all three will fail is $F^3 = 0.1^3 = 0.001$. In general, suppose an ORU has failure probability F and N redundant units. All the N units, the original system and the N - 1 spare ORUs, must fail for the system to fail. The overall failure probability for N redundant units is F^N.

This approach is usually adequate, especially for small N, but here it overestimates the failure probability by roughly a factor of N. This formula is correct for hot spares, where all the units are operating all the time and all failing at the same rate. Some systems use on-line spares, but not the ISS ECLSS, where the spare ORUs are simply stored. Only one unit is operating and the N - 1 off-line spares have a negligible failure rate. The formula that failure probability equals F^N is correct for hot spares but not for non-operating spares.

The probability of a given number of failures using cold spares is given by the Poisson distribution. It is used here to compute the number of spares needed to have a less than 0.001 probability of not having a needed spare.

The system is required to operate over the mission length L = 450 days. The system has a failure rate, rate f = 1/MTBF, the number of times it is expected to fail per year. 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(t) or fewer failures is the summation of the Poisson pdf from x equals 0 to n(t).

C. The ISS ECLSS ORU MTBFs

This section calculates the number of spares needed to have a less than 0.001 probability of not having a needed spare. It uses both estimated and flight data MTBFs. The estimated MTBF data is from MADS, the ISS Maintenance and Analysis Data Set. (MADS, 2015) The estimates benefited from ground testing but were made before flight. The flight MTBFs are based on the observed lifetimes of systems that have operated on ISS for about eight years. Some

ORUs have had no failures and some have had several. (Bagdigian et al., 2015) In cases where all the ISS ORUs have reached the end of life, the estimated MTBF is the average observed lifetime. In cases where one unit has operated without failure for eight years, the MTBF is estimated as eight years, 70,000 hours. If equipment has operated for a time T with no failures, this suggests that the probability of failure was less than one-half. So the MTBF can be estimated as roughly T. If the MTBF was much less, there would have been many failures. The MTBF could be much longer, but longer run time is needed to prove that. No MTBF is set longer than the successful test time, even where the previously estimated MTBF makes this seem reasonable. In cases where one unit has failed but its replacement continues to operate, the MTBF is estimated as the average of the two observed lifetimes. This produces a strictly flight data based estimate of the MTBFs.

D. Oxygen Generation System (OGS) redundancy and failure probability

Table 2 gives the OGS ORU MTBFs, both initially estimated and based on their observed flight operating lives. The required redundancy and failure probabilities are computed for both estimated and observed MTBFs.

Table 2. OGS ORU MTBFs, failure probabilities, redundancy, and total system failure probability.

Acronym	Full name	MADS estimate MTBFs				Flight data MTBFs				Added spares	Failure probability - MADS spares with flight data MTBFs
		MTBF, hours	ORU failure probability	Needed redundancy	Failure probability with redundancy	MTBF, hours	ORU failure probability	Needed redundancy	Failure probability with redundancy		
Hydrogen ORU	Hydrogen Pressure dome	29,551	0.365	4	0.00056	35,000	0.309	4	0.00030	0	0.00030
Controller	Process Controller	75,677	0.143	4	0.00002	70,000	0.154	4	0.00002	0	0.00002
O ORU	Oxygen Outlet	99,252	0.109	4	0.00001	70,000	0.154	4	0.00002	0	0.00002
~~H2 Sensor ORU~~	~~Hydrogen Sensor~~	~~57,803~~	0.187			~~2,600~~	4.~~154~~				
Pump ORU	Pump	189,433	0.057	3	0.00003	35,000	0.309	4	0.00030	1	0.00389
Inlet DI Bed ORU	Inlet Deionizing Bed	442,487	0.024	3	0.00000	70,000	0.154	4	0.00002	1	0.00055
Nitrogen Purge ORU	Nitrogen Purge	140,195	0.077	3	0.00007	70,000	0.154	4	0.00002	1	0.00055
PSM	Power Supply Module	49,202	0.220	4	0.00008	49,202		4	0.00008	0	0.00008
~~ACTEX~~	~~ACTEX - Recirculation~~					~~11,000~~	~~0.982~~				
~~ACTEX~~	~~ACTEX - By-Pass~~					~~20,000~~	~~0.540~~				
Water ORU	Water	37,885	0.285	4	0.00022	38,000	0.284	4	0.00022	0	0.00022
	Total failure probability				0.00098				0.00097		0.00562

The MADS estimates are from (MADS, 2015) The flight MTBFs are based on (Bagdigian et al., 2015). The MADS MTBF is used for the PSM since no flight data was found. The H2 sensor and ACTEX are limited life items replaced as needed, and so are eliminated from the failure probability calculations.

Six OGA ORUs have different estimated and flight MTBFs. Three of the OGS ORUs have significantly shorter MTBFs and so have significantly higher failure rates. These three each require one additional spare, going from triple to quadruple redundancy. If the original set of spares, based on the originally estimated MTBFs, was used, but the actual flight MTBFs were experienced, The OGA failure probability over 450 days would be 0.00562, 5.6 times higher than the expected 0.00098, based on providing spares for < 0.001. The failure probability using the originally estimated spares is nearly six times higher than expected.

E. Urine Processor Assembly (UPA) redundancy and failure probability

Table 3 gives the UPA ORU MTBFs, both initially estimated and based on their observed flight operating lives. The required redundancy and failure probabilities are computed for estimated and observed MTBFs.

Table 3. UPA ORU MTBFs, failure probabilities, redundancy, and total system failure probability.

Acronym	Full name	MADS estimates				Flight data				Added spares	Failure probability - MADS spares with flight data MTBFs
		MTBF, hours	ORU failure probability	Needed redundancy	Failure probability with redundancy	MTBF, hours	ORU failure probability	Needed redundancy	Failure probability with redundancy		
WSTA	Wastewater Storage Tank Assembly	82,200	0.131	3	0.00034	70,000	0.154	4	0.00002	1	0.00055
FCA	Firmware Controller Assembly	13,453	0.803	6	0.00019	70,000	0.154	4	0.00002	-2	0.00000
SPA	Separator Plumbing Assembly	88,993	0.121	4	0.00001	4,000	2.700	11	0.00012	7	0.28591
PCPA	Pressure Control and Pump Assembly	59,221	0.182	4	0.00004	3,000	3.600	13	0.00010	9	0.48478
DA	Distillation Assembly	41,376	0.261	4	0.00016	5,500	1.964	9	0.00021	5	0.13638
FCPA	Fluids Control and Pump Assembly	22,759	0.475	5	0.00014	1,000	10.800	24	0.00036	19	0.98272
~~RFTA~~	~~Recycle Filter Tank Assembly~~										
	Total failure probability				0.00087				0.00083		0.99451

As before, the MADS estimates are from (MADS, 2015) and the flight MTBFs are based on (Bagdigian et al., 2015).

Six UPA ORUs have different estimated and flight MTBFs. Four of the UPA ORUs have extremely shorter MTBFs and so have failure rates ranging from 7.5 to 22 times higher. These four require 7, 9, 5, and 19 additional spares to achieve failure probability < 0.001. If the original set of spares, based on the originally estimated MTBFs, was used, but the actual flight MTBFs were experienced, the UPA failure probability over 450 days would be 0.99451, 1,000 times higher than the expected 0.00098 < 0.001. The FCPA alone has a probability of insufficient spares of 0.98272. The three SPA, PCPA, and DA ORUs have a combined failure probability of 0.68226. Failure is a near certainty. 99.5 percent probable, using the originally estimated spares. Failure probabilities that are ten, one hundred, or even one thousand times higher than predicted are frequent in life support equipment. (Likens, 1992)

F. Water Processor Assembly (WPA) redundancy and failure probability

Table 4 gives the WPA ORU MTBFs, both initially estimated and based on their observed flight operating lives. The required redundancy and failure probabilities are computed for estimated and observed MTBFs.

Table 4. WPA ORU MTBFs, failure probabilities, redundancy, and total system failure probability.

Acronym	Full name	MADS estimate MTBFs				Flight data MTBFs				Added spares	Failure probability - MADS spares with flight data MTBFs
		MTBF, hours	ORU failure probability	Needed redundancy	Failure probability with redundancy	MTBF, hours	ORU failure probability	Needed redundancy	Failure probability with redundancy		
CR	Catalytic Reactor	27,077	0.399	5	0.00006	14,000	0.771	6	0.00015	1	0.00120
GS	Gas Separator	61,182	0.177	4	0.00004	70,000	0.154	4	0.00002	0	0.00002
IX	Ion Exchange Bed	442,478	0.024	3	0.00000	442,478	0.024	3	0.00000	0	0.00000
MCV	Microbial Check Valve	178,447	0.061	3	0.00004	70,000	0.154	4	0.00002	1	0.00055
MFB	Multifiltration Bed #1	349,650	0.031	3	0.00000	349,650	0.031	3	0.00000	0	0.00000
MFB	Multifiltration Bed #2	349,650	0.031	3	0.00000	349,650	0.031	3	0.00000	0	0.00000
PF	Particulate Filter	560,695	0.019	3	0.00000	560,695	0.019	3	0.00000	0	0.00000
	pH Adjuster										
PC	Process Controller	70,745	0.153	4	0.00002	70,745	0.153	4	0.00002	0	0.00002
PS	Pump Separator	39,429	0.274	4	0.00019	70,000	0.154	6	0.00000	2	0.00002
RHS	Reactor Health Sensor	134,077	0.081	3	0.00008	70,000	0.154	3	0.00055	0	0.00055
S	Sensor	184,618	0.058	3	0.00003	70,000	0.154	4	0.00002	1	0.00055
SF	Separator Filter	642,342	0.017	2	0.00014	642,342	0.017	3	0.00000	1	0.00014
	Start-up Filter	226,850	0.048	3	0.00002	226,850	0.048	3	0.00002	0	0.00002
WD	Water Delivery	81,797	0.132	4	0.00001	70,000	0.154	4	0.00002	0	0.00002
WW	Waste Water	43,669	0.247	4	0.00013	70,000	0.154	4	0.00002	0	0.00002
WS	Water Storage	40,463	0.267	4	0.00017	70,000	0.154	4	0.00002	0	0.00002
	Oxygen Filter	342,548	0.032	3	0.00001	342,548	0.032	3	0.00001	0	0.00001
	External Filter										
	Total failure probability				0.00094				0.00088		0.00314

As before, the MADS estimates are from (MADS, 2015) and the flight MTBFs are based on (Bagdigian et al., 2015). The MADS MTBF is used for the IX, MFB, PF, PC, SF, start-up filter, and oxygen filter since no flight data was found.

Eight WPA ORUs have different estimated and flight MTBFs. Five of the WPA ORUs have shorter MTBFs and have failure rates ranging from 1.9 to 2.6 times higher than the originally estimates. These five require 1, 1, 2, 0, and 1 additional spares. If the original set of spares, based on the originally estimated MTBFs, was used, but the actual flight MTBFs were experienced, the UPA failure probability over 450 days would be 0.00314, about 3 times higher than the originally expected $0.00098 < 0.001$.

G. Summary

The ISS flight data on MTBFs and failure rates make it clear that failure rates are often higher than estimated. Of the 20 ORUs with different estimated and flight MTBFs, 14 had lower actual than estimated MTBFs. 7 of the 20 had failure rates between 1.9 and 6.3 times the original failure rates and usually required one additional spare. Without the additional spare, the probability of insufficient spares in flight would be 3 to 5 times higher than expected. 4 of the 20 ORUs had failure rates from 7.5 to 22 times the estimated failure rates. These were the four UPA ORUs that require 7, 9, 5, and 19 additional spares over the number based on the originally estimated MTBFs. Without these additional spares, having insufficient spares in flight would have more than a 99 percent probability of failure.

H. Why does the ISS life support system have lower reliability than expected?

Some components are very unreliable, and the worst in ISS life support is a pump in the urine processor. This single highly unreliable component increases the system failure rate to a disastrous level, 98%. The source of the problem is clear. Reliability was intentionally sacrificed in the design of space station life support. A review of

space station life support noted that, "A current challenge of operating ECLSS and other vital ISS systems is to make them more efficient." The example of gaining higher efficiency is that, "As we learn to work the systems, they get more robust, and it takes less time for the crews to do repairs in flight." (Goldstein, 2011) Here, increasing efficiency is a more acceptable way of saying we need to improve reliability and maintainability. The review continues, "NASA engineers are also seeking to make WRS (Water Recovery System, WRA and UPA above) equipment more efficient. 'You would think we'd know how to make long-life pumps,' says Badigian (Robert Bagdigian, former designer then life support branch chief at MSFC), but in these applications where we are trying to make these pumps as small as possible … you end up having fairly unique pumps." (Goldstein, 2011) The pumps are unique because they are small, don't have long life (1,000 hours), and fall short in reliability and maintainability. Pumps are rotating machinery, subject to wear and failure. They are usually designed for long life and tested for reliability. But rather than use proven reliable pumps, the space station designed new ones for reduced mass and volume. New unique designs tend to have higher failure rates and undiscovered design errors. Cutting mass and volume ends to reduce robustness. In the shuttle era of very high launch costs, sacrificing reliability for mass was thought necessary. Operational experience showed that, even if needed, sacrificing reliability for mass was highly regrettable. Now that the launch cost has been reduced by roughly a factor of 20, more mass can be flown to achieve higher reliability. "You would think we'd know how to make long-life pumps." We do and we should.

VI. Discussion of the spares model for reliability

There are problems with relying on spares and the spares model to keep a system operating but there are reasons why the model is often used and can be useful.

A. The problem with relying on spares and the spares redundancy reliability model

The problem with relying on spares to keep a system operating is that the expected MTBFs are often too long, the failure rates are too low, and the number of spares provided for the mission is too low. When the actual higher failure rates occur in flight, the spares are likely to be all used before the end of the mission, causing the system to fail. It takes only a few components with much worse than expected MTBFs to make failure almost inevitable.

The problem with the spares redundancy reliability model itself is that is a partial bottom up model with limited scope. It accounts for internal component failures, but there are many other sources of failure that can prevent achieving the desired reliability. Even though spares can replace components damaged by externally caused failures, more will be used than those needed to replace the expected internal failures. All the spares may be used and the system fail.

Unfortunately the spares redundancy reliability model may not be applied sincerely or critically reviewed. There may be a tendency to set the expected failure rates too low based on the optimistic engineering judgment that we can do better than historical experience suggests. It is easy to disregard failure modes not addressed in the spares model, especially those caused by unfortunate human errors. The problems of optimism and willful blindness also cause unrealistic budgeting and scheduling. Favorable performance predictions are welcomed, so the obvious risk of unfavorable events can be overlooked.

B. Why is the spares redundancy reliability model often used?

The component based reliability model has been used since it was invented in the German WWII rocket program. The originator, Lusser, developed the cheaper more reliable V1 cruise missile, while von Braun developed the V2 ballistic rocket. Later in the US, Lusser computed that Apollo was likely to fail. After the Apollo success, reliability analysis was dropped because it was too pessimistic. The shuttle designers did not mitigate obvious risks, for instance by providing crew escape. Risk analysis was reestablished after Challenger.

The spares redundancy reliability model can be very useful if its limitations are recognized. A key limitation is that it accounts for only one of many types of failure, internal component failures. The model can be used to compare systems based on the best case, highest possible reliability. It can eliminate candidate designs that cannot meet the required reliability. The model can guide design for reliability, for instance by selecting reliable components, reducing parts count, and identifying key reliability problems. Component based reliability and redundancy analysis is a fundamental tool in any reliability development program.

VII. A better approach to achieving sufficient reliability

A better approach is needed to achieve sufficient reliability. The method should try to identify, quantify, and mitigate all significant risks. The systems engineering approach includes probabilistic risk analysis, systems engineering trade-offs, and design for reliability.

A. The NASA systems engineering approach

Engineering for reliability should begin with a specific requirement for a certain probability of failure over the mission duration, which defines the maximum acceptable system failure rate per unit time. This requirement is determined by top-down systems PRA. The overall mission first must accept some probability of loss of crew, Pr(LOC). This probability is then allocated to the sequential phases of the mission. For a Mars transit, these include ascent, outbound transit, Mars orbit, return transit, and Earth reentry. For transit, the life support system and other habitat systems would be allocated part of the total Pr(LOC). This sets the limit on the probability of the life support system's failure to perform over the mission duration.

Systems engineering develops alternate architectures and candidate technologies and performs trade-offs considering performance, reliability, and cost. If especially high reliability is needed, simple, flexible, tested, and robust systems should be preferred. Unfortunately this discourages the use of the highest performing, most efficient, and newest technology.

After system selection, the development process should design for reliability from the start. The recent NASA technical standard on reliability and maintainability no longer requires specific reliability activities during the sequential project phases. It requires developing detailed reliability requirements and planning how to implement them to meet overall project reliability objectives. The emphasis is on providing the evidence to show that the reliability requirements are met, rather than on conducting specific prescribed formal reliability activities, such as component based reliability analysis, to meet project milestones. The technical standard defines a comprehensive hierarchy of specific reliability objectives and identifies particular strategies to implement them at each level. The top level objective is defined and then one or more design strategies to implement it are developed immediately, before lower objectives are defined and the strategies developed for them. The new reliability process is aligned with the systems design process and helps ensure that the methods to meet the reliability requirements are built into the design. (NASA-STD-8729.1A, 2017)

B. High risk can be recognized and magnificently overcome

NASA's highest risk human space project was also its most spectacularly successful, Apollo. Understanding and reducing risk was crucial to its success.

The risk was obvious and generally appreciated, but its magnitude was not widely understood. Early in the Apollo program, the "calculation was made by its architecting team, assuming all elements from propulsion to rendezvous and life support were done as well or better than ever before, that 30 astronauts would be lost before 3 were returned safely to the Earth. Even to do that well, launch vehicle failure rates would have to be half those ever achieved and with untried propulsion systems." (Rechtin, 2000)

Everyone in the Apollo program knew the risk was high. Everyone expected others than the Apollo 1 astronauts would die attempting to reach the moon. But this awareness ensured that everyone was intensely focused on reducing risk. The moon landings are a technical and scientific achievement that will forever overshadow the peak achievements of earlier civilizations. "The only possible explanation for the astonishing success – no losses in space and on time – was that every participant at every level in every area far exceeded the norm of human capabilities." (Rechtin, 2000)

Individuals, an organization, a culture that is intensely focused on risk can beat the odds. To achieve high reliability, risk must always be a serious concern. Risk must be identified and minimized as far as possible at every step of a program, from mission planning through systems design, testing, and operations. Everyone in a program should be constantly alert for potential risks, even outside their own area of responsibility. In a safety oriented culture, any anomaly, even a temporary sensor glitch, is traced to its cause, understood, and corrected, no matter who, what, where, or when.

C. Recognizing risk is unpleasant and often avoided

The awareness of risk is an emotional fear response, partly instinctive and partly conditioned. Pilots, astronauts, and soldiers can learn to override fear. Most of us live most of the time in stable and secure circumstances, free of fear and unconscious of risk. But because feeling fear and being aware of risk is very unpleasant, most of us are reluctant to recognize risk. We are often willfully blind or deliberately ignorant of risk. The amazing unconsciousness of risk in space shuttle design and operations was the precise opposite of Apollo awareness.

Risk is inevitable. The future is fundamentally unknowable and unpredictable. The unknown unknowns, the things that we don't know we don't know, cause surprising events. A method to anticipate the unknown is to design with simplicity and conservatism and for robustness and adaptability.

VIII. Conclusion

The model of system reliability based on providing spares to repair failures can provide useful insight into intrinsic maximum system reliability, but using it can cause two serious problems. First, if the component failure rates are too optimistic, the computed number of sufficient spares will be insufficient. More failures will occur than expected, more spares will be needed than expected, and the system is likely to fail for lack of spares. A single component with a much higher than expected failure rate can render the system inoperative. Second, intrinsic component failures are only one of many types of failure modes. New and little tested systems may have requirement or design problems that create common cause failures and require redesigns. If they are prevalent, the disregarded failure modes can completely invalidate the spares model reliability predictions.

The ultimate purpose is to meet the system reliability goal. A supporting objective is to understand the failure modes and the expected reliability. A reliability development approach that is wider and more inclusive can help achieve understanding and reliability. A good systems engineering approach would include risk analysis, trade-offs including reliability, performance, and cost, and a development process that builds reliability into the system. Much more than providing spares is needed.

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.

Boden, D. G., and Hoffman, S. J., "Orbit Selection and Astrodynamics," in Larson, W. K., and Pranke, L. K, eds., *Human Spaceflight: Mission Analysis and Design*, McGraw-Hill, New York, 2000.

Do, S., *Towards Earth Independence - Tradespace Exploration of Long-Duration Crewed Space System Architectures*, Doctoral Thesis, Massachusetts Institute of Technology, 2016.

Goldstein, E., "A green space station," *Aerospace America*, May 2011, pp. 22-4.

Jones, H. W., "Reliability Growth in Space Life Support Systems," ICES-2014-075, 44th International Conference on Environmental Systems,13-17 July 2014, Tucson, Arizona.

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.

Jones, H. W., "We Can't Count on Repairing All Failures Going to Mars," ICES 2016-113, 46th International Conference on Environmental Systems, 10-14 July 2016, Vienna, Austria.

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

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

NASA-STD-8729.1A, *NASA Reliability and Maintainability (R&M) Standard for Spaceflight and Support Systems,* 6/13/2017.

Owens, A. C., and de Weck, O. L., "International Space Station Operational Experience and its Impacts on Future Mission Supportability," ICES-2018-198, 48th International Conference on Environmental Systems, 8-12 July 2018, Albuquerque, New Mexico.

Owens, A. C., and De Weck, O. L., "Limitations of Reliability for Long-Endurance Human Spaceflight," 2016-5308, AIAA Space Conference & Exposition, Long Beach, CA, September 2016.

Owens, A. C., and De Weck, O. L., "Systems Analysis of In-Space Manufacturing Applications for the International Space Station and the Evolvable Mars Campaign," 2016-5394, AIAA Space Conference & Exposition, Long Beach, CA, September 2016.

Owens, A. C., De Weck, O. L., Stromgren, C., Cirillo, W., and Goodliff, K. E., "Supportability Challenges, Metrics, and Key Decisions for Future Human Spaceflight," 2017-5124, AIAA SPACE and Astronautics Forum and Exposition, 12 - 14 Sep 2017, Orlando, FL.

Rechtin, E., *Systems Architecting of Organizations*, CRC Press, Boca Raton, 2000.

Russell, J. F., and Klaus, D. M., "Maintenance, reliability and policies for orbital space station life support systems," Reliability Engineering and System Safety 92 (2007) 808–820.

Stromgren, C., Goodliff, K. E., Cirillo, W., and Owens, A., "The Threat of Uncertainty - Why Using Traditional Approaches for Evaluating Spacecraft Reliability Are Insufficient for Future Human Mars Missions," 2016-5307, AIAA SPACE 2016, Long Beach, CA: American Institute of Aeronautics and Astronautics, 2016.