

Ultra Reliable Space Life Support Systems

Harry Jones

NASA Ames Research Center

ABSTRACT

Ultra reliable space life support systems can be built with small additional mass for direct material supply or about twice the minimum mass for recycling equipment. The required direct supply of a material such as oxygen, water, or food for space life support can be provided in some number “ r ” of identical packages. If only one of the r packages fails, the life support system fails. But by providing $n > r$ packages, so that there are $n - r$ spare packages to make up for failures, the reliability of direct material supply can be greatly increased. Ultra reliability can be achieved if the required direct supply is provided in 10 to 100 or more packages with 1 or 2 spare packages, so the additional mass required for ultra reliable direct life support is only a few percent. Recycling equipment for oxygen or water recovery can be designed to have a large number of subsystems and components so that the failure modes are localized, separated, and contained, and so that the failure rate of each component is small because the failures are distributed over a large number of components. Ultra reliability can be achieved if a sufficient number of spare components is provided. A simple model of this process, a worst-case because of a uniform distribution of fault sources, shows that ultra reliability can be achieved by providing a single spare for each component if the number of components is large enough. Ultra reliability for recycling equipment can be achieved by only doubling the equipment mass. Mathematical analysis and example data expand and confirm this surprising result.

INTRODUCTION

Space life support design can be optimized by trading off performance, cost, and risk. Performance requirements are usually known but can change, cost is usually measured by the Equivalent System Mass (ESM), and the risk of failure is usually ignored in preliminary systems design.

The life support system requirements are given in terms of the mass of oxygen, water, and food per crewmember per day. ESM includes the mass of the material or equipment supplied, plus the mass of all packaging,

volume, power, and cooling required to provide life support. ESM has long been the main focus of life support analysis, diverting attention from other equally important aspects. Analysis of system reliability and failure rate is a key element of systems engineering that has been neglected in advanced life support systems analysis.

The two contrasting methods of providing space life support are direct material supply of oxygen, water, and food, or instead recycling these materials. Direct supply has a fixed mass cost per day and is suitable for short missions, but requires prohibitive mass for long missions. Recycling requires a large mass of equipment but only a relatively small daily supply of materials and spares, so it has less mass cost than direct supply for long missions. The breakeven date between direct supply and recycling occurs when their mass costs are equal. Direct supply is less costly for short missions, and recycling is less costly for longer missions.

Direct supply can meet exactly the life support mass requirements over the mission duration. Recycling can be more flexible, providing excess or emergency capacity, operating in degraded modes, and allowing extended mission duration.

Direct supply, providing oxygen, water, and food stored in tanks and packages, has intrinsically high reliability and low risk of failure. Recycling equipment has a much higher basic failure rate and will need maintenance and repair using spare parts. Either approach must be designed to achieve the required mission reliability. This will increase the total mass of recycling proportionately more than for direct supply. For some designated life support reliability, the mass cost of recycling will increase more than that of direct supply, and the breakeven date will be pushed out to longer mission durations.

The question addressed here is, “How does the difference in intrinsic failure rates between direct supply and recycling change their relative mass costs?” The answer is that achieving ultra reliable space life support requires very small additional mass if using direct material supplies and roughly twice the minimum equipment mass when recycling.

Reliability has not been a major focus of life support analysis. Reliability was an important qualitative factor in the selection of space station life support technology, but reliability was not estimated quantitatively. (Carrasquillo et al.) Recycling life support equipment is expected to be less reliable than direct supply. Active components such as pumps and valves tend to wear out. (Wieland) The result that ultra reliability can be achieved at reasonable mass cost is very encouraging.

MASS COST, FAILURE RISK, AND THE MINIMUM COST AND RISK FRONTIER

The life support system provides oxygen, water, and food to spacecraft crewmembers. This report considers how the expected cost and the allowable risk together determine the best life support system design. The system cost is measured by the required mission mass. The allowed risk is measured by expected failure rate.

Each life support requirement must be met by a constant mass flow input, m , measured in kilograms per crewmember per day (kg/CM-d). The total mass to provide each requirement increases directly with the number, N , of crewmembers and the mission duration, D , in days. The total mass to provide a particular requirement over the mission is $m N D$ kg. The failure risk is the probability that the required the mass flow will not be provided throughout the entire mission duration. The mass and the failure risk both increase with longer mission duration.

DIRECT SUPPLY AND STORAGE - For crewed space missions lasting from a few days to a few weeks, the required life support mass is most economically provided by directly supplying and storing the needed material. Oxygen and water are launched in tanks and food in meal-sized packages. The mass cost for direct supply is equal to the total life support mass multiplied by some tankage factor, t , which could be roughly 1.5. The total mass for direct supply of a life support requirement is then equal to $t m N D$ kg. The failure risk is measured by the probability that the needed life support material flow is not provided throughout the mission duration. This may occur if one or more packages fails to provide the expected material, perhaps due to leakage, spoilage, or poor preparation. The total direct supply failure probability, $pr(ds \text{ fails}, D)$, is typically low, and is assumed to be 1 in 10,000 or a mission probability of 10^{-4} .

RECYCLING - For missions lasting months or years, oxygen can be recovered from crew-produced carbon dioxide. The water in respiration, perspiration, and urine can be purified by recycling. Food might be grown on very long missions, and the crop growth chambers would also convert carbon dioxide to water and purify water by transpiration. Suppose that the mission duration D is sufficiently long that use of some recycling equipment could reduce the total delivered mission mass. A particular life support mass flow requirement can be met by using recycling equipment with a flow rate capacity of $m N$ kg/d. The recycling equipment has a mass proportional to the flow rate, $p m N$ kg, where p is a proportionality factor, assumed to be of the order of 100. The equipment proportionality factor p should include mass penalties for the volume, power, and cooling that the recycling equipment uses. The mass cost is then the total mass delivered to provide the life support material mass flow requirement. If the proportionality factor, p , is 100, the equipment mass is equal to the mass of 100 days output. The recycling equipment also requires a logistics mass flow of make-up life support material and recycling supplies and equipment spares. The logistics mass flow is assumed to be a logistics factor, f , times the life support requirement mass flow rate. This is $f m N$ kg/d, where f is assumed of the order of 0.1.

BREAKEVEN MISSION DURATION - Recycling saves mass and is cost effective only if D is sufficiently long that the required total mass for direct supply, $t m N D$ kg, is greater than the total mass of the recycling equipment and its continuing logistics flow, $p m N + f m N D$ kg. This defines the breakeven mission duration, D_{be} .

$$t m N D_{be} > p m N + f m N D_{be}$$

$$D_{be} > p/(t-f)$$

Recycling breaks even if the mission duration, D , is greater than the equipment proportionality factor, p , divided by the tankage factor, t , minus the logistics factor, f . For the rough estimates that $p = 100$, $t = 1.5$, and $f = 0.1$, break even requires $D > 100/1.4 \sim 70$ days. Since t is usually slightly greater than one and f is much less than one, the breakeven mission duration, D_{be} , is approximately equal to the equipment proportionality factor, p .

$$D_{be} \sim p$$

The recycling equipment has some failure risk, the probability of not providing the required life support material flow over the mission duration. The recycling failure probability, $\text{pr}(\text{rc fails}, D)$, is expected to be higher than the direct supply failure probability, $\text{pr}(\text{ds fails}, D)$, and is assumed to be 1 in 100 or a probability of 10^{-2} . Recycling saves mass and cost for a sufficiently long mission, but a straightforward implementation leads to a higher risk of failure.

RECYCLING EQUIPMENT REDUNDANCY TO IMPROVE RELIABILITY - In general, we can improve system reliability if we are willing to add system mass and increase cost. A direct brute force approach is to provide a second complete recycling system. If we do, the recycling approach does not fail unless the two redundant recycling systems both fail. The probability that both recycling systems fail is $\text{pr}(\text{rc fails}, D) * \text{pr}(\text{rc fails}, D) = 10^{-2} * 10^{-2} = 10^{-4}$. This is equal to the value for $\text{pr}(\text{ds fails}, D)$, the direct supply failure probability. We have regained the low failure rate of direct supply, but we have doubled the recycling equipment mass and therefore doubled the breakeven mission duration to 140 days. We have also, with 99 percent probability, lavishly provided exactly twice the required life support material flow. Providing spare parts and backup life support material would seem more cost effective than implementing full redundancy.

MINIMUM COST AND RISK FRONTIER - The optimum life support design lies on a minimum life support cost and risk frontier. This boundary is defined by the lowest possible cost and mass for a variable allowed risk of failure. Equivalently, it is determined by the lowest achievable failure rate for a given mass cost. This paper begins to identify the minimum cost and risk frontier for space life support. It will also consider ways to achieve extremely high reliability, the ultra quality necessary when failure may cause loss of life.

Identifying the optimum design for some planned mission goes beyond the scope of life support system design. The optimum life support design is located on the life support minimum cost-risk frontier, but the appropriate mass and failure rate point is determined by the overall mission design. The mission managers allocate total mission mass to life support and other elements to maximize the probability of mission success. For a given mission performance and total mass, the optimum mass allocation is such that the last kilogram added anywhere buys the same improvement in probability of mission success. Extremely high life support reliability is cost effective if it is the least expensive way to reduce overall mission risk.

Usually the life support design process proceeds in the opposite direction. The performance, mission duration, and reliability are defined as requirements, and the minimum cost design is developed in response. An unexpectedly high cost can result in significant requirements cuts, but the design approach is rarely changed. The minimum cost-risk chart can help identify when requirements changes should result in design changes.

Figure 1 shows the minimum cost-risk frontier corresponding to the suggested parameter values in the examples above, $t = 1.5$, $p = 100$, and $f = 0.1$. The other example parameters are $m = 1 \text{ kg/CM-day}$, $N = 4 \text{ CM}$, and $D = 200 \text{ days}$. The exact value of m is unimportant, since it is a scaling parameter that appears in all the mass formulas.




Figure 1. Minimum mass and failure risk frontier.

The direct supply mass of $t m N D$ is 1,200 kg, and the recycling failure probability $pr(ds \text{ fails}, D)$ is 10^{-4} . Full dual redundant direct supply would have mass of 2,400 kg and a failure probability of 10^{-8} . This is extremely high reliability, at a high mass cost, but a dual redundant supply is not unreasonable for a long term base (unless recycling is more cost-effective, as in this example). One spare supply tank would be available on the first expedition, in case the primary tank failed, and would be used on the second expedition and replaced by a new spare. If there is no spare tank, the probabilities of failure add for the successive expeditions, so that a $pr(ds \text{ fails}, D)$ of 10^{-4} becomes 10^{-3} , 10 times larger after 10 expeditions. However, in this example, recycling has lower mass for the same failure probability, so that recycling defines the minimum mass risk frontier.

The recycling mass of $p m N + f m N D$ is 480 kg, and the recycling failure probability, $pr(rc \text{ fails}, D)$, is 10^{-2} . Full dual redundant recycling has twice the mass at 960 kg and a failure probability of 10^{-4} . The figure also shows triple redundant recycling with mass of 1,440 kg and failure probability of 10^{-6} , and quadruple redundant recycling with mass of 1,920 kg and failure probability of 10^{-8} .

The nonredundant recycling mass of 480 kg is much less than the minimum direct supply mass of 1,200 kg, but nonredundant recycling has much higher failure probability. Equalizing the failure probabilities to 10^{-4} by using dual redundant recycling systems increases the recycling mass to 960 kg, using most of the mass savings. Similarly, the break even mission duration is much longer for equal failure probabilities, but it is still less than the assumed mission duration. Comparing direct supply to recycling should consider failure probabilities as well as mass costs.

Figure 1 shows that, when we use redundancy, the system mass for both direct supply and recycling increases linearly as the logarithm of the probability of failure becomes smaller. The exact slope depends on the single system failure probability. The ratio of the mass increase, or the redundancy factor, is $\log(\text{required probability of failure}) / \log(\text{single system probability of failure})$. To go from 10^{-2} for a single system to 10^{-4} for a redundant system, the mass increase is $-4/-2 = 2$, or dual redundancy. To go from 10^{-4} to 10^{-6} , the mass increase is $-6/-4 = 1.5$, which is not achievable with full unit redundancy. These mass ratios can be used to estimate the mass needed to achieve some required failure probability, or to equalize the failure probabilities of different design approaches.

The analysis of redundancy helps to define the minimum mass and risk frontier. We can cut mass in half if we increase the probability of failure to its square root, say from 10^{-2} to 10^{-1} . To square the probability of failure, say from 10^{-2} to 10^{-4} , we must double the mass by using dual redundancy.

We will consider other approaches to increasing reliability besides full redundancy. To do this, we need some reliability analysis background. This will also clarify the above results.

STATISTICAL ANALYSIS OF RELIABILITY

Standard reliability techniques can help investigate life support system failure probabilities. We will consider failure rates per unit time, the accumulation of failure probability over time, the number of failures expected in an interval, and how connecting components in series or parallel determines the overall system failure rate.

FAILURE RATE ASSUMED CONSTANT - The failure rate, λ , is the number of times a component or system is expected to fail per unit time, given that it is currently still operating. We will assume that life support systems have a constant failure rate with time. A constant failure rate is sometimes given in failures per thousand or per million hours.

Assuming a constant failure rate is common, simplifies analysis, and seems justified for life support equipment. Some components and systems have a failure rate that decreases with time at first, remains constant during a period of useful life, and finally increases due to wear out. This is called "the bathtub curve." The initially decreasing failure rate is due to burn-in, to failure of defective components, and possibly to detection and repair of minor faults. (Some components and systems are repairable and others are not.) 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, p. 143) Shuttle maintenance data show a rapid decrease in the problem rate to a roughly constant baseline, but no later increase due to wear-out effects. (Shishko, p. 93) If space life support systems are well tested, burned in, and designed to last beyond the mission duration, their failure rates can be expected to be constant.

During a mission, a life support component or system is either operating as required or it has failed. The component may be practically unrepairable, such as a punctured tank that has lost its contents, or it may be at least potentially repairable, such as a clogged recycling system. A repairable component or system will be down

for some time while being repaired and then return to operation. We assume that the repair brings the component or system back to its original condition and failure rate.

RELIABILITY AND CUMULATIVE FAILURE PROBABILITY - The reliability function, $R(t)$, is the probability a component or system survives and operates to time t . The reliability is a probability between 0 and 1. The time is reset to 0 if the system is repaired and returned to service.

The cumulative failure probability, $F(t)$, of a component or system is the probability that it fails one or more times by time t . Since a unit either survives or fails, the failure rate, $F(t)$, is one minus the reliability function.

$$F(t) = 1 - R(t)$$

EXPONENTIAL FAILURE PROBABILITY DENSITY FUNCTION (PDF) - The failure rate, λ , is the number of failures per unit time period, given that the system is operating at the beginning of the period. The reliability, $R(t)$, is the probability that a 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.

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

This is the exponential probability distribution, which is tabulated and available in Excel. 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 failure probability density, $f(t)$, is the failure rate, λ , times the reliability, $R(t)$. The failure rate, λ , assumes that the system is still operating, while the reliability, $R(t)$, is the probability that the system is still operating. Their product is the probability of failure at any given time.

The cumulative failure probability, $F(t)$, is the integral of the exponential failure pdf, $f(t)$, from time 0 to t .

$$F(t) = \int_0^t \lambda e^{-\lambda \tau} d\tau = 1 - e^{-\lambda t}$$

The reliability, R(t), is the integral of the failure pdf from time t to infinity.

$$R(t) = \int_t^\infty \lambda e^{-\lambda \tau} d\tau = e^{-\lambda t} = 1 - F(t)$$

The Mean Time Before Failure (MTBF) is 1/λ, the mean value. The probability that the equipment will not fail before the MTBF is R(1/λ) = e⁻¹ = 0.63.

The exponential model is the simplest of all life distribution models. The constant failure rate, λ, is the only unknown. 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. This justifies the common use of the exponential distribution.

The above reliability relationships are shown in Table 1.

Table 1. Reliability relationships and formulas for constant failure rate.

Name	Symbol	Derivations	Relationships	Formulas
failure rate	λ	constant	f(t)/R(t)	λ
reliability	R(t)	$\int_t^\infty \lambda e^{-\lambda \tau} d\tau$	1 - F(t)	$e^{-\lambda t}$
cumulative failure probability	F(t)	$\int_0^t \lambda e^{-\lambda \tau} d\tau$	1 - F(t)	$1 - e^{-\lambda t}$
failure probability density	f(t)	- dR(t)/dt	λ R(t)	$\lambda e^{-\lambda t}$
Mean Time Before Failure	MTBF	Mean of f(t)	1/λ	1/λ

THE POISSON PROBABILITY DENSITY FUNCTION (PDF) FOR NUMBER OF FAILURES - Suppose the constant rate, λ, of failures is relatively high, such that the MTBF = 1/λ < D, the mission duration. Assuming that the equipment is repairable, we would expect several failures and repairs during the mission.

If the failures (and subsequent repairs) are independent and occur at the constant rate, λ, their times are distributed according to the exponential distribution given above.

The cumulative distribution function of the waiting time between failures is the failure rate, F(t), corresponding to the exponential pdf.

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

n(t) is the cumulative number of failures from time 0 to time t. n(t) has a Poisson distribution. The Poisson distribution gives the probability, for failure rate λ, that there will be exactly n(t) = x failures in time t.

$$\text{Poisson pdf } (n(t) = x) = (\lambda t)^x e^{-\lambda t} / x!$$

The Poisson distribution mean and variance are both λ t. The expected number of failures during the mission duration is λ D. The expected time between failures or MTBF is 1/λ, as before.

The probability of n(t) or fewer failures is the summation of the Poisson pdf from x equals 0 to n(t). The probability of n(t) or more failures is the summation of the pdf from n(t) to infinity. The Poisson model, like the exponential, assumes a constant failure rate and fits the flat portion of the bathtub curve.

MTBF, MTTR, AND AVAILABILITY - The mean time before failure (MTBF) is the mean time expected until the first failure. For repairable systems, MTBF can be read as mean time between failures, the expected time span from a failure to the next failure. If the failure rate, λ, is constant over time, MTBF is the inverse of the failure rate.

For repairable systems, the mean time to repair (MTTR) is the total amount of time spent on repairs divided by the total number of repairs. It is the expected time span from a failure until the return to service.

The availability, A, is the fraction of time a repairable component or system is expected to be in operation. The availability is constant for constant failure rate, λ.

$$A = \text{MTBF} / (\text{MTBF} + \text{MTTR})$$

RELIABILITY DECREASE FOR LONGER DURATION - Increasing mission duration, D, increases the expected number of failures. If the probability of failure is small, the number of failures increases directly with the duration.

To show this, assume a constant failure rate, λ. The expected number of failures in a mission of duration D is λ D. As a crude approximation for very small λ and small λ D << 1, the probability of a failure during the mission is λ D. In this case, doubling the duration, D, would double the expected number of failures and the cumulative failure probability, F(t). The reliability R(t) = 1 – F(t). For this approximation, doubling the duration would double the difference between R(D) and 1. Going from D = 100 to D = 200, with λ = 0.0001, F(D) goes from F(D) = λ D = 0.01 up to 0.02, and R(D) goes from 0.99 down to 0.98.