Redundancy: How Many Unreliable Spares are Needed for High Reliability and Confidence on a Time Limited Mission?

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

This paper investigates the number of redundant units needed to achieve high reliability with high confidence. The approach applies to the case where the unit failure rate is too high for a single unit to provide the required reliability over the mission duration. To achieve high reliability, the design then uses N redundant units, one operating unit and N – 1 spares. If the unit failure rate is f, the mission length is L, and f * L is small (not the case assumed here), the unit failure probability over the mission duration is F1 = f * L << 1. In this case, the probability that all N units will fail is FN = $F1^N$, and the needed N = LN(FN)/LN(F1). For the case of large f * L assumed here, F1 = f * L > 1, and F1 is the expected number of failures during the mission. The needed redundancy, N, to achieve the specified N unit reliability, FN, can be computed using the cumulative Poisson distribution with mean equal to F1. The number of spares, N - 1, is increased until the probability - that the total number of failures will be less than N -1 - achieves the required reliability. The confidence that this reliability can be achieved can be computed using the cumulative Poisson distribution or the chi-square distribution. Since the measured unit failure rate, f, has some uncertainty, the confidence that the rate is not lower than the actual failure rate and the required reliability is not overestimated is about 50%. Adding more redundant units increases the confidence that the required reliability, FN, will be achieved. For a fixed number of redundant units, the expected reliability and confidence can be traded off, since lower reliability goals have higher confidence in being achieved. Both the required reliability and confidence can be specified initially and the needed number of redundant units computed using the measured failure rate.

The unit failure rate is determined by initial reliability growth testing to remove design errors and to better estimate the final constant failure rate. Reducing the failure rate and reducing its variance both reduce the number of redundant units needed for the required reliability and confidence. Since the total cost is the sum of the costs of the units and of the testing, there is an optimum test time that produces minimum cost.

Nomenclature

a	=	a constant (abcd model)
b	=	reliability growth rate (abcd model)
c	=	the constant uncorrected failure rate (abcd model)
d	=	an additional constant failure rate due to correctable but uncorrected failure modes (abcd model)
f	=	single unit failure rate
F1	=	single unit failure probability over the mission duration, F1 = f * L
FN	=	the probability that all N redundant units will fail over the mission duration
Fx(t)	=	the increased number of failures for a single unit over the mission length L, Fx(t) = λx(t) L
k	=	a constant (Duane model)
L	=	mission length
m	=	Poisson mean, equal to variance
M	=	number of test units
MTBF	=	Mean Time Before Failure, MTBF = 1/f
mu	=	upper bound on the mean number of failures
N	=	Number of redundant units

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

$n(t)$ = number of failures up to time t
P = probability of $n(t)$ or fewer failures
t = time
t_d = time of end of reliability growth
x = 1 – confidence
y = 1 - reliability
α = reliability growth rate (Duane model)
$\lambda(t)$ = measured time varying failure rate, $\lambda(t) = n(t)/t$
$\lambda x(t)$ = upper bound on failure rate for confidence 1 - x

I. Introduction

THIS paper determines how much testing time is cost effective in computing the number of redundant units needed to achieve high reliability with high confidence. This requires a two-phase test program, an initial period of testing to provide reliability growth followed by life testing to more accurately determine the final failure rate achieved by reliability growth. Newly designed systems usually have high initial failure rates which are often reduced by finding the high rate failure modes and removing them by redesign until the system achieves an acceptable failure rate. This failure rate will be constant if there are no further redesigns or wear out. To plan for high reliability with high confidence using redundancy, longer test time more accurately determines the failure rate and reduces the number of redundant units needed for high reliability with high confidence.

The measured failure rate decreases throughout the reliability growth period as failure modes are removed. After the reliability growth effort is terminated, the first few failures during life testing provide an estimate of the final failure rate. Given the failure rate and the desired reliability using redundancy, the number of spares can be determined and the confidence in the reliability computed. As long as there are only a few failures, the failure rate estimate has a wide variance. There is a 50% chance that the actual final failure rate is higher, and it could be much higher. If a too low estimated failure rate is used to compute the redundancy needed to achieve the required reliability, the number of spare units provided will be too low. Using the estimated failure rate gives only a 50% confidence that the failure rate and number of spares are not too low.

In the current approach, both the redundant reliability and the confidence level are initial requirements. The confidence that the actual redundant reliability is not too low can be increased by increasing the number of spares. When there are only a few failures, the variance of the failure rate is high and the failure rate and number of spares must be increased greatly to reach high expected reliability with high confidence. The higher number of spares increases the cost of redundancy. Longer test time reduces the variance in the failure rate and reduces the number of spares needed to increase confidence.

As the test time is increased, the test cost increases linearly but the number of needed spares drops, at first exponentially. The total cost is the sum of the cost of the redundant units and of the test time. There is an optimum test time that produces the minimum total cost for the system failure rate, mission length, reliability, and confidence level. If the total cost must be reduced, the reliability, confidence level, or both must be reduced. The required reliability and confidence identify and justify a minimum total cost for redundant units and testing. Determining the optimum test time that minimizes the cost to meet requirements can help ensure sufficient testing.

II. Background

Many systems must operate continuously with high reliability. If the required reliability cannot be provided by one single operating unit, several redundant systems are often used. Spare systems can be stored and replaced as they fail. If spares cannot be readily replaced because the operating location is remote or the spares require a long lead time, the on-hand stock of spares must be sufficient for the duration of the operation.

NASA has recycling life support systems on the International Space Station (ISS) and plans to use similar systems for the moon and Mars. The ISS life support was tested only briefly before flight and some systems have had unexpectedly high failure rates. The crew on ISS and even the moon can return or be resupplied with spares in a few days, but a Mars transit requires more than a year without any opportunity for the crew to return or obtain spares. Planning a deep space mission requires knowing how many spares are needed to provide high reliability with high confidence.

MIL-HDBK-781A discusses reliability growth and failure rate confidence bounds. Its major purpose is planning failure testing to minimize mistaken acceptence or rejection of hardware due to an inaccurate failure rate.[1] Owens and de Weck previously investigated how much testing is needed to determine the number of spares needed for deep space missions.[2] Their approach is similar to that used in this paper, but both they and the military have a significantly different goal. Owens and de Weck, like MIL-HDBK-781A, consider the problems due to overestimating the failure rate as well as underestimating it. In the military case, overestimating the failure rate leads to improper rejection of hardware that meets requirements. Beyond LEO, overestimating the failure rate leads to taking too many spares and increasing the mass and launch cost. Balancing the effects of the two opposing errors of too high and too low failure rates, Owens and de Weck conclude, "In the end, there is no simple answer to the question of how much testing is required to manage supportability risks for beyond-LEO missions." Instead, it seemed better to search for a quantitative answer defining how much testing should be done. The major current problem in testing is that not enough testing is done. A good justification would lead to the correct level of testing. Given the great impact of the loss of a crew due to insufficient spares, the major risk is that a too low failure rate will lead to insufficient spares. Too many spares would actually increase reliability and confidence, while their additional cost seems relatively trivial.

The approach of the current paper is to set the required reliability and confidence in consideration of both the mission cost and the safety of the crew. Two previous papers by this author describe the method to identify how much testing time is cost effective to determine the number of redundant units needed to achieve high reliability with high confidence.[3] [4] These papers use a somewhat similar method, but the first[3] one uses a nonstandard approach and develops an unfamiliar terminology. The second paper[4] uses more standard analysis and notation, as does the current paper.

III. Approach

The next section IV on reliability growth describes the abcd model of reliability growth and shows a data set with the model and the upper 90% confidence bound on the failure rate. The upper 90% confidence bound is used later to compute the number of redundant units needed for 90% reliability with 90% confidence. Results on the number of redundant units, the test cost, and the total cost are given in the next section but derived later.

The following section V describes how to determine the upper confidence bound on the measured failure rate. The two methods are generating and scanning a table of the cumulative Poisson distribution or using the inverse chi-square distribution. The third section VI shows the computation of the required number of redundant units based on the upper confidence bound on the measured failure rate and the required reliability. The Poisson Distribution with its mean equal to the single unit failure probability is used to determine the minimum number of units for the required reliability. The fourth following section VII estimates the cost of test time in units of system cost per time unit and determines the optimum test time for minimum cost. For a significant test cost, the optimum test time is not much longer than the reliability growth period.

The next section VIII presents equations for estimating the number of units required for a particular failure rate, confidence, and reliability. The approximate equations fit the data. Section IX shows the trade-off between confidence and reliability for a particular failure rate, mission length, and the number of units needed for reliability = confidence = 0.90. Section X shows the increase in the number of units needed as the required reliability = confidence increase. Section XI discusses practical issues and XII is the conclusion.

IV. Reliability Growth

Newly designed systems often have high initial failure rates (called "infant mortality") which can be reduced by successive redesign until the system achieves an acceptable constant failure rate. Reliability grows rapidly if the causes of all the failures that occur are identified and removed without introducing new failure modes. The time varying failure rate is $\lambda(t) = n(t)/t$, where n(t) is the total number of failures up to time t. Suppose only one failure occurs, its cause is corrected, and testing continues without further failures. Then $\lambda(t) = 1/t$ and declines as t^{-1}. This is the most rapid possible failure rate decline. The commonly used Duane-Crow reliability growth model represents such exponential decline, with the main issue being the time exponent of the failure rate decline.[1 5 6]

The Duane reliability growth model is

$$n(t)/t = k\, t^{-\alpha} \tag{1}$$

The reliability growth rate is α, the downward slope of n(t)/t versus t. It usually varies from 0.2 to 0.6. k is a constant.[1 4] Crow used a 56-failure data set to illustrate reliability growth.[1]

A graphical Duane model fit to this data gives

$$n(t)/t = 0.640\ t^{-0.283} \qquad (2)$$

A problem in using the Duane-Crow reliability growth model is that it assumes that reliability growth continues and the failure rate decreases throughout the test and operations period. It is more usual that reliability growth stops when the failure rated is low enough. Growth testing is often followed by testing with a low constant failure rate due to rare or uncorrectable failure modes. This is the flat portion of the "bathtub curve." As more and more low constant rate acceptable failures accumulate after the period of reliability growth, the reliability growth time exponent α decreases toward zero. It is more accurate to model a period of initial reliability growth followed by testing to more accurately determine the final constant failure rate. This is done in the abcd model.

$$n(t)/t = a\ t^{-b} + c \text{ from } t = 0 \text{ to td} \qquad (3a)$$

$$= c + d \text{ after td, where } d = a\ td^{-b} \qquad (3b)$$

The term a t^{-b} describes the continuous reliability growth that continues out to time t_d and c is the constant uncorrected failure rate. The parameter d represents an additional constant failure rate due to correctable but uncorrected failure modes. After the reliability growth process is terminated, the failure rate n(t)/t = c + d.[1 5]

The abcd model was computed for the Crow reliability growth data set, with reliability growth out to t_d =100. The abcd model gives a better fit than the Duane-Crow model.

$$\text{Failure rate} = 1.37\ t^{-0.99} + 0.14 \text{ from } t = 0 \text{ to td} = 100 \qquad (4a)$$

$$= 0.14 \text{ after td} = 100 \qquad (4b)$$

The constant uncorrected failure rate, c, is assumed to be 0. Figure 1 shows the Crow failure rate data, n(t)/t, the abcd model for that data, and the 90% confidence bound on the measured n(t)/t.

Figure 1. Crow failure rate data n(t)/t, the abcd model, and the 90% confidence bound.

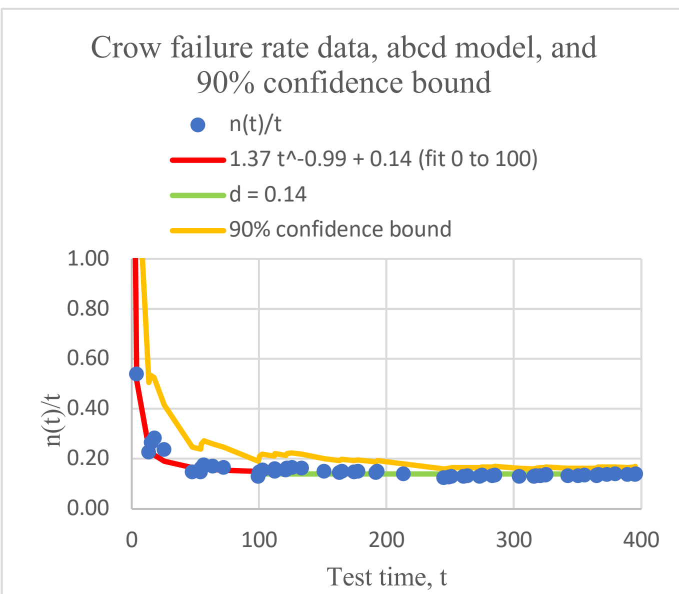

Redundant units, spares, are needed when the final failure rate λ(t max) = n(tmax)/tmax, equal to c + d in the abcd model, is relatively high or the mission length, L, is relatively long. For the example of Table 1, L = 10. Since λ(t = 395.2) = 0.14, The expected number of failures during the mission is 0.14 * 10 = 1.4 and spares are needed to prevent mission failure.

The failure rate n(t)/t and the 90 90% upper confidence bound drop rapidly during reliability growth. After reliability growth ends at td = 100, even after t = 50, the measured failure rate is constant while the 90% confidence bound continues a slow decrease due to the increased number of failures reducing variance of the upper bound.

Table 1 shows the Crow failure rate data in full up to t = 100 and with most failure counts eliminated for t = 100 to 400.

Table 1. Data, redundancy, and cost results for the Crow data set.

n	Time t	λ(t)	λ0.9(t)	Units N	Test cost	Total cost
1	0.7	1.43	5.56			
2	3.7	0.54	1.44	20	1.1	21.1
3	13.2	0.23	0.51	9	1.2	10.2
4	15.0	0.27	0.53	9	1.2	10.2
5	17.6	0.28	0.53	9	1.3	10.3
6	25.3	0.24	0.42	8	1.4	9.4
7	47.5	0.15	0.25	6	1.7	7.7
8	54.0	0.15	0.24	6	1.8	7.8
9	54.5	0.17	0.26	6	1.8	7.8
10	56.4	0.18	0.27	6	1.8	7.8
11	63.6	0.17	0.26	5	2.0	8.0
12	72.2	0.17	0.25	5	2.1	8.1
13	99.2	0.13	0.19	4	2.5	7.5
14	99.6	0.14	0.20	4	2.5	7.5
15	100.3	0.15	0.21	4	2.5	7.5
23	151.0	0.15	0.20	5	3.3	8.3
30	213.0	0.14	0.18	5	4.2	9.2
40	304.0	0.13	0.16	4	5.6	9.6
56	395.2	0.14	0.17	4	6.9	10.9

The data in Table 1 includes the failure count, the failure times, the decreasing failure rate λ(t) = n(t)/t, the increased failure rate for confidence = 0.9, λ0.9(t), the required number of units N for reliability = 0.9 and confidence = 0.9, the increasing test cost for testing one unit, and the total cost. The cost unit is the cost of developing one system. The test cost is assumed to be 0.015 per unit time. The data in the last three columns is shown in Figure 1. The computations are described below.

V. Upper Confidence Bounds

The testing that occurs after the reliability growth period is not accompanied by redesign to remove the discovered failure causes, so the failure rate does not continue to decrease. The transition from reliability growth testing occurs at about t = 50 to 100 for the Crow data in Figure 1 and Table 1.

The reason to continue testing is to more accurately determine the final failure rate achieved by reliability growth. The data-based failure rate λ(t) = n(t)/t is an average. The fewer the failures, the wider the measured λ(t) may vary. If testing stops after only a few failures, it is possible that the true λ(t) would have produced many more failures than have actually occurred. If the measured λ(t) is much lower than the true λ(t), the calculated number of redundant units N will be too few.

Using the measured λ(t) gives only about a 50% confidence that the λ (t) is not too low, which would cause the number of spares to be too low. If higher confidence is needed that there are sufficient spares, say 90%, this requires using an increased failure rate λ(t) that would produce fewer than the measured number of failures only 10% of the time. The test failure rate should be increased, for example, to the 90% confidence failure rate. Using this increased failure rate, λ0.9(t), there is a 90% confidence that λ0.9(t) is not lower than the actual failure rate and the number of spares is not too low for 90% confidence in the predicted reliability. λ0.9(t) is referred to as the 90% upper confidence bound on λ(t) and is shown in Figure 1 and Table 1.

The upper confidence bound on λ(t) can be determined using the cumulative Poisson distribution.

$$\text{Cumulative Poisson} = \sum m^x e^{-m}/x!,\ x = 0, 1, 2, \ldots M \quad (5)$$

The mean number of failures m = λ(t) L, where L is mission length. The summation is over the number x = n of failures included, 0 to M.

The upper bound of λx(t) for a confidence of x corresponds to an upper bound on the mean number of failures in the mission, mu = λx(t) L. The correct value of mu satisfies the equation

$$\sum mu^x e^{-mu}/x! > P,\ x = 0, 1, 2, \ldots n(t) \quad (6)$$

The upper bound on the mean number of failures, mu, must be high enough that the corresponding probability of n(t) or fewer failures is P. Stated conversely, given n(t) failures, the probability that the true mean value m = λ(t) L is less than mu is P. The value mu determines the upper bound on failure rate for confidence x, λx(t) = mu/L. Using this approach requires generating and examining tables of the cumulative Poisson distribution, which is cumbersome.[7] [3]

The upper confidence bound on λ(t) can also be determined using the chi-square distribution, since the cumulative Poisson and chi-square distributions are related.

$$\sum m^x e^{-m}/x!,\ x = 0, 1, 2, \ldots M = \text{Probability } \chi 2\ (x, f) > 2m \quad (7)$$

The x in $\chi 2$ (x, f) is not the counting index x = n(t). The x in $\chi 2$ (x, f) is the fraction of the distribution summed, equal to the cumulative probability of being below the upper bound when the mean of the Poisson distribution is mu. And f is the number of degrees of freedom. Here f = 2 (n(t)+1) where n(t) is the number of failures.[7]

Like the cumulative Poisson distribution, the chi square distribution is included in the standard spreadsheet, and tables can be generated and examined as when using the Poisson distribution. The standard spreadsheet also has a function that iteratively computes the inverse of the chi-square distribution.

The inverse chi-square distribution computes the inverse of the left-tailed probability of the chi-square distribution for probability P and f degrees of freedom. We use

$$\lambda x(t) = \text{Inverse chi-square } (x, 2*(n(t)+1)/(2\ t) \quad (8)$$

The same results are produced by examining tables of the cumulative Poisson distribution and using the inverse chi-square distribution.[3] The 90% confidence bounds equal to $\lambda 0.9(t)$ are shown in Figure 1 and Table 1. As expected, the distance between the upper 0.9 probability confidence bound and its measured value decrease as test time t increases.

VI. The Increased Number of Redundant Units N Based on Confidence Bounds

The number of spares required for any particular mission reliability is determined by the Poisson distribution for a constant failure rate. In reliability analysis, it is usually assumed that the expected number of failures in a given time interval is constant and that the numbers of failures in different time intervals are independent. The Poisson distribution is used to determine the expected number of failures in a given time interval. The number of spares needed is equal to the expected number of failures and the total number of redundant units, N, is equal the number of spares plus one, the initial operating unit.

The Poisson distribution gives the probability (Pr) that the number n events will occur in an interval, given that the expected or mean number of events is m.

$$\text{Poisson } (x, m) = \Pr (n = x) = m^x e^{-m}/x! \quad (9)$$

The number of redundant units, N, must be sufficient that the probability of N-1 failures is less than the required reliability y. N-1 is determined by the cumulative Poisson distribution, which is the sum of the probabilities of n =0, 1, 2, …N-1 failures occurring.

$$\text{Cumulative Poisson} = \sum m^x e^{-m}/x!,\ x = 0, 1, 2, \ldots M \quad \text{(5 repeated)}$$

The number of redundant units, N, can be determined by inspection of tables of the cumulative Poisson distribution, which can be computed in the standard spreadsheet.

The upper 0.9 probability confidence bound, $\lambda 0.9(t)$, is used for the failure rate. The increased number of failures for a single unit over the mission length L is $Fx(t) = \lambda x(t)\ L$. The values for $\lambda 0.9(t)$are listed in Table 1. The required number of redundant units can be computed for confidence = 0.9 and any reliability. Here a reliability of 0.9 is chosen. As shown in Table 2, the cumulative Poisson distribution is scanned down for an increasing number of failures and an increasing probability that the # of failures will be less than indicated.

Table 2. Using the cumulative Poisson table to determine N for reliability y = 0 9 and confidence x = 0.9.

	Time, t	15.0	25.3	56.4	120.9	260.1
	F0.9(t)	5.33	4.16	2.73	2.14	1.64
N = 1 + # failures	# failures	Cumulative Poisson				
1	0	0.00	0.02	0.07	0.12	0.19
2	1	0.03	0.08	0.24	0.37	0.51
3	2	0.10	0.22	0.49	0.64	0.77
4	3	0.22	0.40	0.71	0.83	**0.92**
5	4	0.38	0.60	0.86	**0.93**	0.97
6	5	0.56	0.76	**0.94**	0.98	0.99
7	6	0.71	0.87	0.98	0.99	1.00
8	7	0.83	**0.94**	0.99	1.00	1.00
9	8	**0.91**	0.97	1.00	1.00	1.00
10	9	0.95	0.99	1.00	1.00	1.00

Table 1 shows that as the test time t increases, the increased failure rate for confidence = 0.9, $\lambda 0.9(t)$, decreases from 5.56 to 0.17. F0.9(t) = $\lambda 0.9(t)$ L is the increased single unit failure probability over the mission length L = 10 that is required to achieve confidence x. Table 2 plots a segment of time horizontally, from 15.0 to 260.1, and F0.9(t) decreases from 5.33 to 1.64. The single unit failure probability for confidence 0.9, F0.9(t), is used as the mean of the Poisson distribution. The cumulative Poisson is tabulated for 0 to 9 failures. Scanning down the table shows the # failures that must be replaced using spares for the redundancy reliability of 0.9 or more to be achieved. To always have an operating system, the number of redundant units $N \geq 1 + \#$ failures. The numbers in bold are the smallest reliability that exceeds the required reliability of 0.9. The number of required units drops from 20 (Table 1) to 4 as test time increases. The redesign that occurs during reliability growth reduces the number of redundant units needed to achieve a particular mission reliability with a particular confidence. Longer testing produces slower reduction in N.

VII. Minimizing the Total System Development Cost

The overall mission cost is equal to the cost of developing N redundant units plus M test units and the estimated cost of testing the M test units. Suppose the cost of testing is the fraction 0.015 of the unit development cost per hour. The test cost is

$$\text{Test cost} = M + 0.015\, M\, t \quad (10)$$

The test cost starts from a fixed value and increases linearly with time. The total cost is the cost of developing the N redundant units plus the test cost.

$$\text{Total cost} = N + M + 0.015\, M\, t \quad (11)$$

The increasing test time reduces the upper confidence bound on the system failure probability so that the number of redundant units N decreases with time. The initial decrease during the reliability growth period is exponential and the later decrease during long term testing is incremental. Table 1 and Figure 2 show the number of redundant units N, the test cost for M = 1 test unit, and the total cost. The selected cost metric is the cost of producing a single unit, so the cost of N units is N.

Figure 2. N, test cost, and total cost for increasing test time.

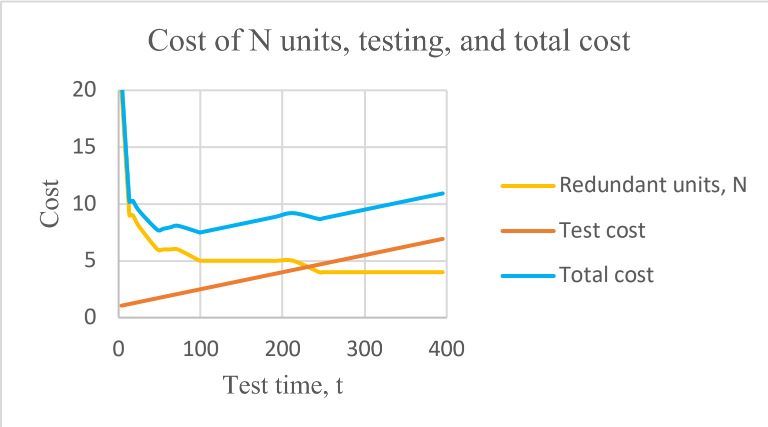

The minimum total cost is 7.5 units at t = 100 hours. However, a cost of 7.7 is reached at 50 hours, so the test time for minimum total cost is 50 to 100 hours. Table 1 shows that in this time range, N decreases from 6 to 4, the minimum, while test cost increases from 1.7 to 2.5, The decline in the number of redundant units N is not smooth because N must be an integer and often provides more than the minimum required redundant reliability, as seen in Table 2.

The long-term failure rate is λ = 0.14 from Table 1, so the final Mean Time Before Failure (MTBF) = 1/0. 14 = 7.14 hours. The test time for minimum total cost is equal to 7 to 14 times the MTBF. Testing for only twice the MTBF, t = 15, would give N =9 and total cost =10.2.

The computation of N uses the 0.9 upper confidence bound on the failure rate, λ0.9(t). The number of units, N, for reliability 0.9 and confidence 0.9 is 20 at 3.7 hours, 9 at 15 hours, 6 at about 50 hours, and 4 or 5 beyond 60 hours.

Extended testing reduces the number of units needed for high confidence in high reliability. The long test time shown here is not justified by the reduction in cost or the number of redundant units. For this example, the optimum test time is roughly equal to the extent of the reliability growth period.

VIII. Estimating Versus Computing the Number of Redundant Units, N

Equations to estimate λx(t), Fx(t), and N were developed by fitting mathematical formulas to the data.[3] The redundancy reliability is 1 -y. The failure rate estimation confidence is 1 – x. λx(t) is the upper bound failure rate for confidence 1 - x.

$$\lambda x(t) = 1.285\ \lambda(t) + (-\ 1.24\ LN(x) + 0.19)/t \qquad (12)$$

LN is the natural logarithm. Fx(t) is the corresponding single unit failure probability over the mission length L that is required to achieve an upper confidence probability of 1 - x, the probability that the failure rate is not underestimated.

$$Fx(t) = \lambda x(t)\ L \qquad (13)$$

N is the increased number of redundant units required to increase the probability of having sufficient spares to 1 - y.

$$N = -\ (0.305\ LN(Fx(t)) + 0.86)\ LN(y) + Fx(t) \qquad (14)$$

N depends on the measured system failure rate, λ(t), the mission length, L, the required redundancy reliability 1 -y, and the required confidence in the failure rate estimation upper bound, 1 – x. The value y is the probability that all redundant units fail. The value x is the probability that the upper bound failure rate is too low. The confidence is (1 – x) 100 percent that the probability that the upper bound is not too low. The approximation for is close for λ(t) L > 0.3, where λ(t) L is the failure probability of a single system over the mission length. The exact N for any case can be calculated as shown above and previously.[3] Substituting for the full equation for N,

$$N = -\ (0.305\ LN((1.285\ \lambda(t) +(-\ 1.24\ LN(x) + 0.19)\ t)\ L) + 0.86)\ LN(y)$$

$$+\ (1.285\ \lambda(t) +(-\ 1.24\ LN(x) + 0.19)/t)\ L \qquad (15)$$

Figure 3 compares the calculated and estimated numbers of needed redundant units for the Crow data set considered here.

Figure 3. Calculated and estimated number of units, N.

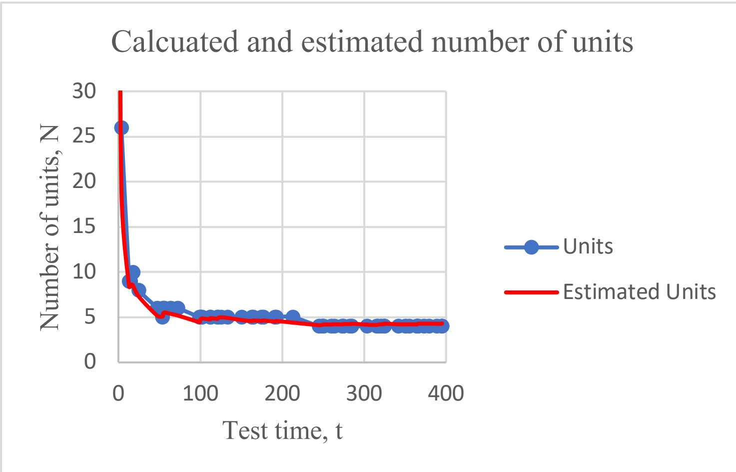

In Figure 3, the computed units are in integers while the estimated units are decimal and would be rounded up in practice. The estimation equations are easier to use than the statistical calculations and can be used to check the complicated direct statistical process.[3]

IX. The Trade-Off Between Confidence, 1 -x, and Reliability, 1 - y

For any given failure time data set, we can compute the minimum cost and optimum test time for any x and y. There is a defined trade-off between confidence and reliability. For the finally chosen redundancy, N, a higher reliability will be met with lower confidence and a higher confidence can be achieved with a lower reliability estimate.

For the Duane data set of Table 1, we consider the result for N at time t = 50 with failure rate $\lambda(t) = 0.15$ and mission length L =10. For reliability and confidence both equal to 0.90, the computed integer N is 6, and the N estimated from equation 15 is 5.17. The trade-off between reliability and confidence is investigated by picking either one and then finding the other so that the estimated N is again 5.17. The trade-off is shown in Table 3 and Figure 4.

Table 3. Reliability and confidence.

Reliability, 1 - y	Confidence, 1 – x
0.95	0.0400
0.94	0.4400
0.93	0.6600
0.92	0.7800
0.91	0.8550
0.90	**0.9000**
0.85	0.9785
0.80	0.9935
0.70	0.9989
0.60	0.9997
0.50	0.9999

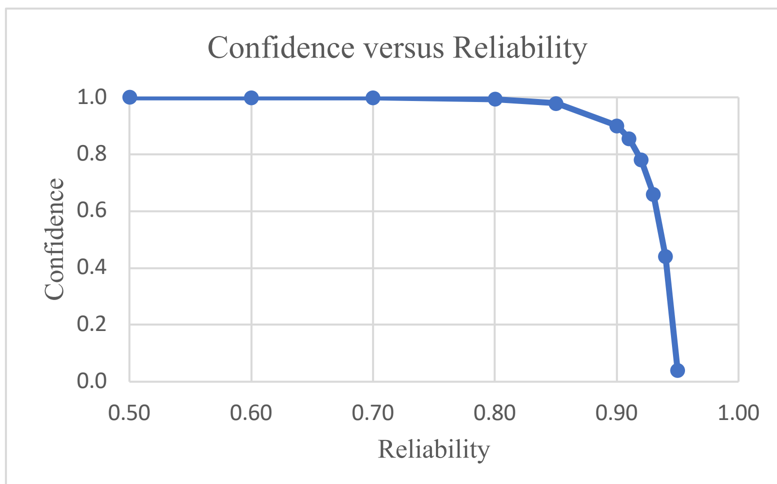

Figure 4. Confidence versus reliability.

Confidence and reliability are varied around the design point where both reliability and confidence are equal to 0.90. Increasing the reliability requirement from 0.90 to 0.91, 0.92, … 0.95 causes the confidence that the requirement will be met to drop rapidly toward zero. Reducing the reliability requirement in steps from 0.90 to 0.50 increases confidence form one 9 to two, three, and even four 9's.

X. The Increase in N as Reliability and Confidence Increase.

We again consider the result for N at time t = 50 with failure rate λ(t) = 0.15 and mission length L =10, in the case where reliability and confidence both equal but requirements other than 0.90. The increase of the estimated N for increasing reliability and confidence is shown in Table 4 and Figure 5.

Table 4. Increasing N for higher reliability and confidence

Reliability = confidence	N
0.995	9.76
0.99	8.66
0.98	7.58
0.95	6.20
0.93	5.69
0.91	5.32
0.90	**5.17**
0.85	4.58
0.80	4.17
0.70	3.60
0.60	3.20
0.50	2.89

Figure 5. N versus increasing reliability and confidence.

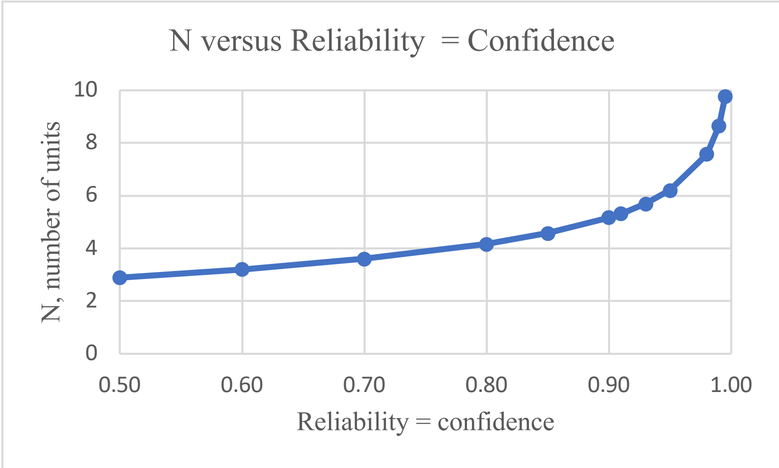

Although the estimated N is decimal, the actual number of units should be the next higher integer. As the requirements for the reliability equal to confidence increase from 0.50 to 0.995, the required number of redundant integer units N increases from 3 to 10.

XI. Practical Issues

Often failures occur at a decreasing rate during reliability growth and later occur at a constant rate during long term resting. This is usually modeled during reliability growth using the nonhomogeneous or variable rate Poisson process and later according to the homogeneous or constant rate Poisson process.

This assumes that the failures are internally generated and independent of each other, but this may not be true. Many failures are caused by the environment or operator error.[6] These are usually discounted when considering the intrinsic system reliability, but systems should be designed to reduce them.

Common Cause Failures (CCFs) occur when several failures result from a single shared cause, internal or external. If common cause failures are present, the failures are not statistically independent as assumed in reliability analysis. The essential effect of CCFs is that they defeat redundancy and can prevent achieving high reliability using redundant units. The cure for common cause failures is to use diverse redundancy, to have the different redundant units use different designs developed by different organizations. At least one of each design type would have to be tested for reliability growth and to improve the estimate of failure probability, so test cost would increase by the diversity factor.

Even if we attempt to design for reliability, eliminate CCFs, and test as we fly, unanticipated failure modes can appear. However, the designed for failure probability of N redundant units is a useful lower bound on the true failure rate. The required reliability and confidence cannot be achieved unless the number of redundant units is sufficient.

Since large systems usually have many subsystems that must all operate correctly for the overall system to work, the subsystem with the highest failure rate determines the overall system failure rate. It is usually cost effective to improve the reliability of subsystems with the highest failure rates.

The approach described here shows how much reliability and confidence can be achieved for a minimum cost by optimizing test time. Having an exact engineering method to estimate and minimize total cost should help overcome the short term problems that lead to overly brief test times and produce poor reliability. The cost of too short testing is the larger number of redundant units that must be provided to achieve the required reliability and confidence.

XII. Conclusion

A statistical process was developed to compute N, the number of redundant units needed to achieve any required redundant reliability and confidence level. N depends on the measured system failure rate, $\lambda(t)$, the mission length, L, the required redundancy reliability, and the required confidence in that reliability. The first step is to determine the upper confidence bound on the measured failure rate needed to achieve the required confidence using the Poisson or inverse chi-square distribution. This produces $\lambda x(t)$, the increased single unit failure rate over the mission that is used to achieve the required confidence. The second step is to determine N, the increased number of redundant units N required to achieve the required reliability and confidence. This is determined by scanning tables of the cumulative Poisson distribution. Fortunately, the results of this statistical process can be closely approximated by an equation.

It is possible to set the requirements for both the redundant reliability and the confidence level and then test until the time needed to minimize the total cost required to achieve these requirements. The total mission cost is the sum of the redundant units cost and the test time cost. The optimum test time produces the minimum total cost given the unit failure rate, the mission length, and the required reliability and confidence level. Initial testing produces reliability growth, which often has a major impact in reducing the unit failure rate. Testing to better determine the long term failure rate reduces the confidence interval of the failure rate, which seems to have a smaller effect. Extending testing is justified as long as it reduces cost.

References

[1] MIL-HDBK-189C, Department of Defense Handbook Reliability Growth Management, 14 June 2011, p. 114.

[2] Owens, A. C., and de Weck, O. L., "How Much Testing is Needed to Manage Supportability Risks for Beyond-LEO Missions?" ICES-2019-66, 49th International Conference on Environmental Systems, 7-11 July 2019, Boston, Massachusetts.

[3] Jones, H., "Verified Cost-Effective High Reliability for New Deep Space Systems," ICES-2020-221, International Conference on Environmental Systems, ICES 2020.

[4] Jones, H., "Extended Testing Can Provide Cost-Effective Redundancy With High Reliability and High Confidence," submitted to 67th Reliability& Maintainability Symposium (RAMS), Orlando FL, Jan. 2021.

[5] Jones, H., "Achieving Maximum Reliability Growth in Newly Designed Systems," 64th Reliability& Maintainability Symposium (RAMS), Reno, NV, Jan. 2018.

[6] Jones, H., "A Method and Model to Predict Initial Failure Rates," 66th Reliability& Maintainability Symposium (RAMS), Palm Springs, CA, Jan. 2020.

[7] Sahai, H., and Khurshid, A., "Confidence Intervals for the Mean of a Poisson Distribution: A Review," Biometrical Journal, 1993.