Reliability Growth in Space Life Support Systems

Harry W. Jones[1]

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

A hardware system's failure rate often increases over time due to wear and aging, but not always. Some systems instead show reliability growth, a decreasing failure rate with time, due to effective failure analysis and remedial hardware upgrades. Reliability grows when failure causes are removed by improved design. A mathematical reliability growth model allows the reliability growth rate to be computed from the failure data. The space shuttle was extensively maintained, refurbished, and upgraded after each flight and it experienced significant reliability growth during its operational life. In contrast, the International Space Station (ISS) is much more difficult to maintain and upgrade and its failure rate has been constant over time. The ISS Carbon Dioxide Removal Assembly (CDRA) reliability has slightly decreased. Failures on ISS and with the ISS CDRA continue to be a challenge.

Nomenclature

CDRA	=	Carbon Dioxide Removal Assembly
ECLS	=	Environmental Control and Life Support
ISS	=	International Space Station
MTBF	=	Mean Time Before Failure
NHPP	=	Non-Homogeneous Poisson Process
WHC	=	Waste and Hygiene Compartment

I. Introduction

THIS paper discusses the process and mathematical modeling of reliability growth and considers its applicability to space life support systems. In failure analysis, it is often assumed a system's failure rate first decreases with time, then remains constant during its normal operating life, and then finally increases due to wear and aging. This is represented by the well known "bathtub curve." Actually most systems do not have an increasing failure rate due to aging, and some have a decreasing failure rate, called reliability growth, due to failure-preventing maintenance and failure-cause-removing upgrades. Reliability will grow if maintenance and redesign eliminate failure causes. Failure sources that are later removed by an improved design can be considered oversights in the original design. Since a design error affects all the units built to that design, it is usually categorized as a common cause failure.

The Duane reliability growth model plots the cumulative failure rate, the total number of failures divided by total test time, versus the test time. The number of failures per unit time often declines with time, and can be approximately plotted by a straight line sloping down on a log-log graph. A formula for the failure rate can be derived by assuming the failures are produced by a time varying random Poisson process. Then the reliability growth slope can be computed from the failure data.

Failure rate data and the possibility of reliability growth are considered for the space shuttle, the International Space Station (ISS), and the ISS Carbon Dioxide Removal Assembly (CDRA). The space shuttle demonstrated significant reliability growth. The shuttle was extensively maintained, refurbished, and upgraded after each flight. The ISS failure rate has been approximately constant over time. The ISS CDRA reliability has been slightly decreasing. Reliability growth requires an active program to discover and analyze failures and make improvements to fix the discovered design deficiencies. The space shuttle program achieved reliability growth. Failures on ISS and with the CDRA have not decreased.

---

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

II. Reliability growth concepts

This section describes how failure rates may vary over time and how a decreasing failure rate, reliability growth, may be achieved. Reliability growth models and their use are explained.

A. System failure rate changes over time

The failure rate is the number of failures per unit time. The usual depiction of a system's changing failure rate over time is the "bathtub curve." This is seen when a system has a failure rate that first decreases with time, then remains constant during the system's useful life, and finally increases due to component wear out. The initial high "infant mortality" failure rate is due to burn-in, to failure of defective components, and to detection and correction of design faults. The failures during useful life are usually assumed to be random events caused internal degradation. The failure rate increase at end-of-life can be caused by mechanical wear or aging related to chemical or thermal activity. The bathtub curve is shown in Figure 1.
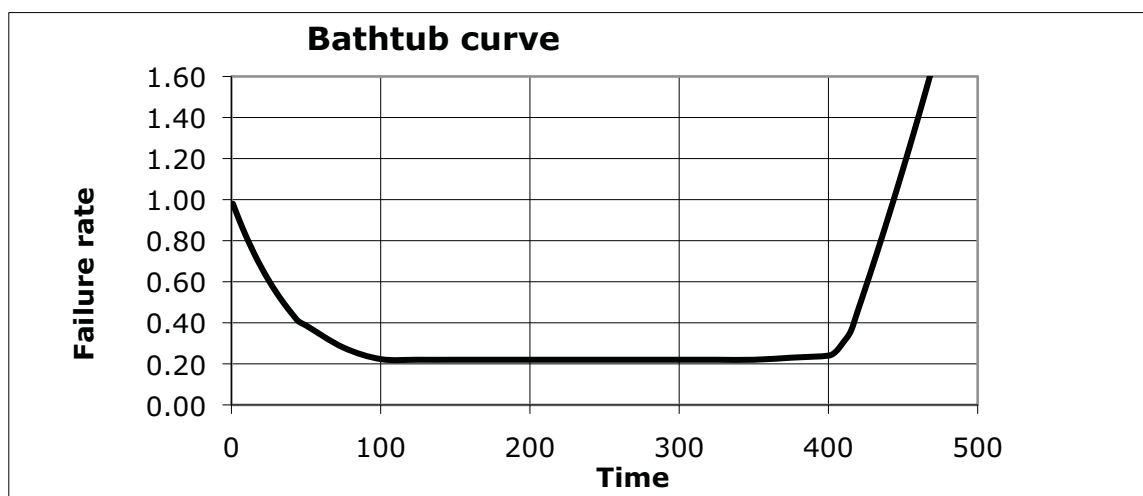

Figure 1. The bathtub failure rate curve, arbitrary time units.

More than two-thirds of all systems show infant mortality and then a constant failure rate, but no final aging period. (Hansen, p. 143) Individual shuttle maintenance data show a rapid decrease in the problem rate to a baseline that then decreases slightly, indicating reliability growth. The shuttle had no failure rate increase due to end-of-life effects, probably because of extensive maintenance and refurbishment. (Shishko, p. 93)

Standard reliability analysis attempts to estimate the failure rate of a system during its useful life, between burn-in and wear-out. A constant failure rate during useful life is usually assumed. However, reliability growth is often observed.

B. Reliability growth concept

Reliability can grow when effective efforts are made to reduce failures. Different kinds of systems have different failure modes that require different fault testing and remediation.

1. High volume production hardware

Commercial mass-produced hardware of a given design and model year, such as an automobile, appliance, or machine, does not show reliability growth during operational life. Failure rates follow the standard bathtub curve, although maintenance can postpone wear-out. Except for the occasional recall, the failures that occur are accepted as normal and repaired without a redesign. The production system's reliability is the result of design and development trade-offs with cost, schedule, and performance.

The reliability of commercial hardware usually increases from one year to the next, due to advances in technology and design techniques. Longer operating time does not itself increase reliability. More testing does not itself increase reliability. Testing and operating a system produce failures. Redesigning systems to reduce or eliminate failure modes increases reliability.

2. Software

Software is different from hardware. Hardware has inherently limited reliability since it always fails sooner or later due to some component or material failure. Software implements a logical structure that always does exactly

International Conference on Environmental Systems

what it is programmed to do. Simple, well-designed, thoroughly tested software can achieve perfection and work flawlessly forever.

Software often fails by producing unintended results or crashing. Rebooting and rerunning the same software always produces the same results, as long as the hardware is correct. Software failure modes are built into the code. Software failures are due to design errors. Software is developed, tested, integrated, and tested again, debugged, to detect and repair design errors, bugs. It is often not possible to test moderately complicated software enough to discover all flaws because of its very large number of different states and paths. An error in software may first appear after years of operation.

Debugging during development and operations usually produces reliability growth. It may happen over time that the original software designers leave and less knowledgeable programmers make bug fixes that introduce new flaws, "software rot."

3. Development hardware

Development hardware is produced in only a few units, unlike production hardware. Development hardware may be preproduction or perhaps only one or a few units may be needed, as in space applications. The typical development process is to design, test, investigate failures, and then redesign to eliminate the failure causes. This process is intended to reduce the failure rate and cause reliability growth. More failures will occur in high stress accelerated life tests.

If the development process could remove all design flaws, a hardware system would have only unavoidable random failures, usually modeled by a constant failure rate, and no further reliability growth would be possible. In well-designed hardware, the probability of failure due to design flaws is usually much lower than the probability of random failure due to unavoidable physical or chemical deterioration.

4. Design flaws lead to common cause failures

System faults are either removable or non-removable. Failure sources that can be removed by redesign are design errors or oversights. All software faults are design errors. Since all copies of a particular software program will produce the same failures, these failures could be called common cause failures, although the term is usually used only for hardware.

A useful definition of a common cause failure is one that cannot be cured by redundancy. Suppose a hardware system has two redundant components in a parallel path. If one fails the other can be used. But if both components are identical and have the same design flaw, both will fail, defeating redundancy. A design error can produce a common cause failure, perhaps better termed a "potential design improvement" or "reliability growth opportunity."

Common cause failures are caused by design errors, failures in human understanding and implementation. Common cause failures are usually identified by a failure analysis that determines that a design change is needed.

In poorly designed systems, common cause failures may account for half the failures. The typical rate of common cause failures is ten percent but some systems have much lower rate. (Jones, 2012-3602)

5. Different types of systems

Failure testing varies for different types of systems. Some hardware operates continuously so failures are distributed over clock time. Many units and long test times are needed to detect infrequent failures and verify high reliability. Machines and automobiles are operated occasionally and failures occur with operating time or miles driven. One-shot systems such as rockets and explosive bolts work or fail on each test or flight. On-demand systems, such as alarms and back-up generators are usually not in operation but can be tested repeatedly. Software is timeless and always gives same result if run the same way, but many permutations cannot be tested when test time is limited.

III. Reliability growth models

The common Duane model of reliability growth is described first. The Duane model is then given a statistical basis and reformulated by assuming the failure data is produced by a time varying random Poisson process.

A. The Duane model

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. It is assumed in basic reliability analysis that the failure rate, λ, is constant during the useful life of a system. The Mean Time Before Failure (MTBF) is the inverse of the failure rate, λ.

$$MTBF = 1/\lambda$$

Reliability growth, a decreasing failure rate over time, can occur when design improvements are made or failure modes removed. If the failure rate varies with time, it is $\lambda(t)$, the instantaneous failure rate.

The Duane model is the earliest reliability growth model and is still commonly used. Duane observed in 1964 that a plot of the measured cumulative failure rate, [$N(t)/t$], versus the cumulative test time, t, often closely follows a

straight line when plotted on log-log graph paper. This occurred when failures were fixed by redesigns that improved reliability.

Duane's log-log plot is described by the equation

$$\log [N(t)/t] = \log k - \alpha \log t$$

Here t is the total test time, N(t) is the cumulative number of failures until time t, k is a constant scale factor greater than zero, and α is the failure growth rate parameter that varies between zero and one.

Taking antilogarithms

$$N(t)/t = k\, t^{-\alpha}$$

Given the two parameters of Duane's model, k and α, the cumulative number of failures, N(t), and the cumulative failure rate, [N(t)/t], can be calculated for any time t. It is possible to solve this equation to find the test time it will take to achieve a specific failure rate, assuming that the reliability growth rate remains constant. Measured values of α, the growth rate parameter, average about 0.4 and usually vary from 0.2 to 0.6. (Yamada and Osaki) (MIL-HDBK-189C)

B. A Duane model reliability growth analysis

The Duane model is a graph-based analysis of reliability growth data. A data set of 56 failures occurring over 400 hours is used to plot the data points of Figure 2. The data set is the list of failure times. These are used to compute the cumulative number of failures, N(t) and the cumulative failure rate, N(t)/t. (MIL-HDBK-189C, Table XV, p. 114)
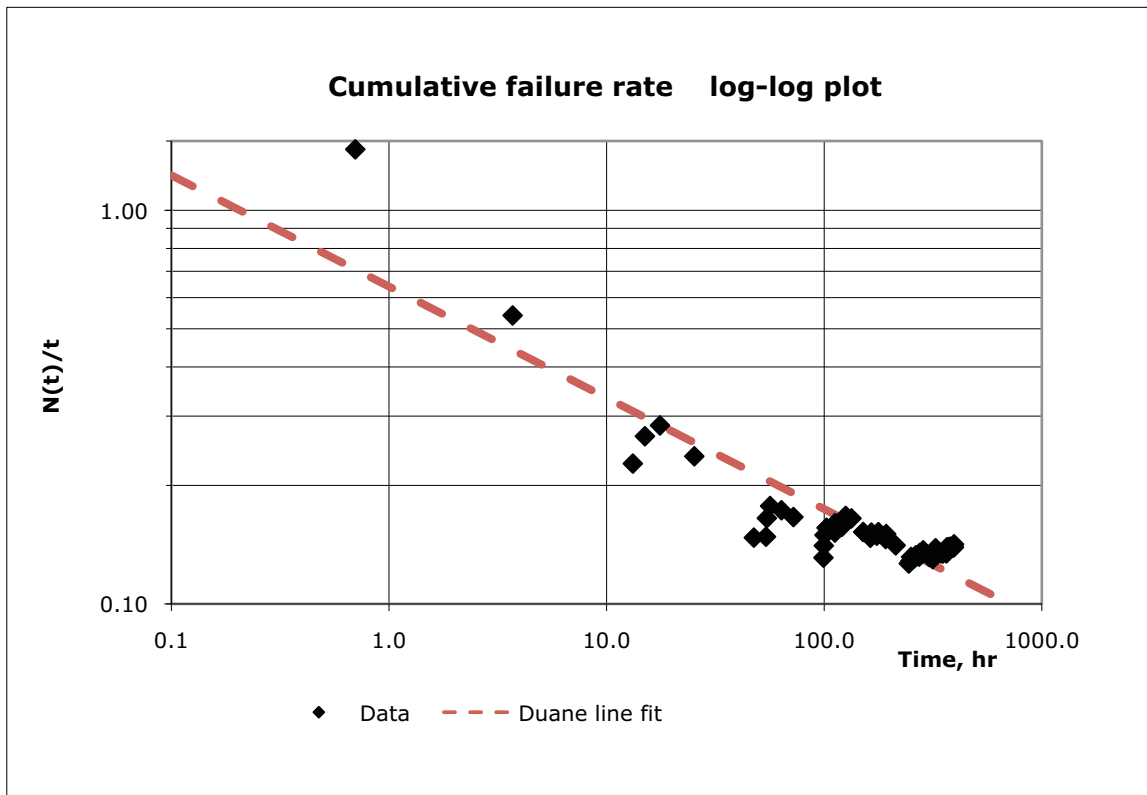

Figure 2. Cumulative failure rate N(t)/t plotted versus time, t, in a Duane log-log graph.

As expected, the data fall near a straight line on the log-log graph. The best Duane straight line fit has the equation,

$$\log [N(t)/t] = 0.640 - 0.283 \log t$$

So $\alpha = 0.283$ and $k = 0.640$, and the Duane line fit is also plotted in Figure 2.

$$N(t)/t = 0.640 \, t^{-0.283}$$

The Duane line fit shows significant reliability growth. If the process of fixing failures continued as during the first 400 hours of test, the cumulative failure rate, $N(t)/t$, would drop below 0.10 at about 700 hours, according to the Duane line fit.

However, it is easy to see that the sharp downward slope, and the high reliability growth parameter, $\alpha = 0.283$, are due to the high early failure rate, infant mortality. The final data points for $N(t)/t$ are above the fitted line, suggesting that reliability may no longer be growing! If the first two failures are ignored, the Duane line downward slope is much shallower, corresponding to $\alpha = 0.171$.

We will consider reliability growth for this data set again, but first the Duane model will be reformulated by assuming the failure data is produced by a time varying random Poisson process. This statistical model can provide a better estimate of the reliability growth.

C. Crow's formulation of Duane's model using the non-homogeneous Poisson process

Statistical reliability growth models usually assume that system failures are caused by a non-homogeneous Poisson process (NHPP). The NHPP is non-homogeneous in time, meaning the statistics and resulting failure rate change over time.

A Poisson process counts the number of events, $N(t)$, that occur during the time interval $(0, t)$. The time varying failure rate is $\lambda(t)$. A Poisson process occurs if:

(1) the numbers of events (failures) in non-overlapping intervals are stochastically independent;

(2) the probability that exactly one event occurs in the brief interval $(t, t+\Delta t)$ is approximately $\lambda(t) \, \Delta t$; and,

(3) the probability that more than one event occurs in an interval of length Δt approaches zero as Δt approaches zero.

If these requirements are met, then the number of events occurring by time t, $N(t)$, has a Poisson distribution with a mean value, $m(t)$, which is equal to the integral from 0 to t of $\lambda(t)$, $m(t) = \int \lambda(s) \, ds$. $N(t)$ is the actual number of failures and $m(t)$ is the expected number of failures occurring in the interval $(0, t)$. (Yamada and Osaki) (MIL-HDBK-189C)

1. Crow's formulation of Duane's model

Crow in the 1970's reinterpreted the Duane model. He assumed that the failures of a system during development testing occur according to an NHPP with a power law mean value function, $m(t)$, and the Weibull distribution failure rate. The mean number of failures is assumed to be

$$m(t) = k \, t^{\beta}$$

where k is positive and β is between zero and one.

The instantaneous failure rate is the time derivative of the number of failures.

$$\lambda(t) = d[m(t)]/dt = k \, \beta \, t^{\beta-1}$$

This is known as the Weibull distribution failure rate. The resulting reliability growth model is equivalent to the Duane model. Its mathematically expected cumulative failure rate is given by

$$\text{Expected } [N(t)/t\} = m(t)/t = k \, t^{\beta-1}$$

The Duane model α is equal to $1 - \beta$ in Crow's reformulated NHPP based model. The parameter k is the same as in the Duane model. The parameter β is the ratio of the current instantaneous failure rate, $\lambda(t)$, to the average cumulative failure rate, $m(t)/t$.

$$\beta = \lambda(t)/[m(t)/t] = k \, \beta \, t^{\beta-1}/k \, t^{\beta-1}$$

A β less than one corresponds to a decreasing failure rate and positive reliability growth. (Yamada and Osaki) (MIL-HDBK-189C)

2. Reliability growth parameter estimation

The assumption that failures occur according to an NHPP with a Weibull distribution failure rate allows the reliability growth model parameters k and β to be calculated from the failure data.

Suppose that N failures are observed during the test time (0, T), and that they occur sequentially at times s1, s2, … , sN. The maximum likelihood estimate of β is

$$\beta^* = N \; / \; \sum \ln (T/si)$$

where ln is the natural logarithm and the summation $\sum$ is over i = 1 to N. The maximum likelihood estimate of k is

$$k^* = N \; / \; T^{\beta^*}$$

(Yamada and Osaki) (MIL-HDBK-189C)

3. An NHPP model reliability growth analysis

The same data set of 56 failures occurring over 400 hours used in Figure 2 is used again in Figure 3.

Figure 3. Cumulative failure rate N(t)/t plotted versus time, t, in an NHPP log-log graph.

The computed value of β = 0.927 and the computed k = 0.217. This β corresponds to an α = 1 - β = 0.073, which is much less than the α = 0.283 found using the Duane method. The fitted line is

$$\text{Expected } [N(t)/t] = m(t)/t = k \, t^{\beta-1} = 0.217 \, t^{-0.073} = 0.217 \, t^{-\alpha}$$

As can be seen in Figures 2 and 3, the α which corresponds to the downward slope of the fitted line and measures the estimated reliability growth is much less for the NHPP model than for the Duane model. The NHPP model is much less influenced by the early infant mortality data and gives a more realistic projection of reliability growth. If the process of fixing failures continued as during the first 400 hours of test, the cumulative failure rate, N(t)/t, would be 0.131 at 1,000 hours in the NHPP model. This is a significantly higher failure rate than in the Duane model line fit.

4. Does this data show continuing reliability growth?

Crow used this data set to illustrate the NHPP model method and it was analyzed in (MIL-HDBK-189C, pp. 113-6). The same values of k and β were obtained there. It was noted that, "While growth is small, hypothesis testing indicates it is significantly different from 0. Thus growth is occurring and the failure intensity (failure rate) is decreasing." MIL-HDBK-189C, p. 115)

Figure 4 shows the cumulative failure rate, N(t)/t plotted versus time, t, but in actual values, not in the usual Duane log-log graph.

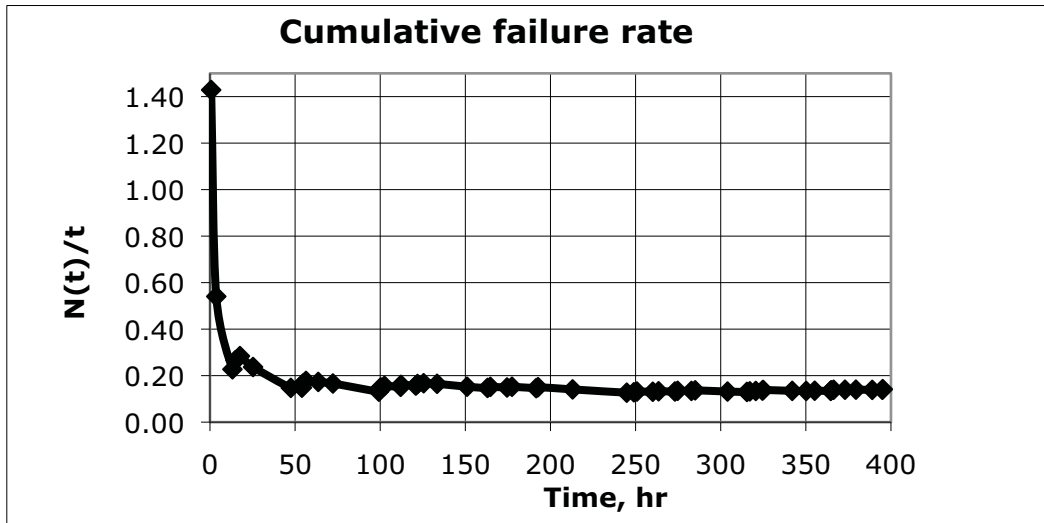Figure 4. Cumulative actual N(t)/t plotted versus time, t, not log-log.

Each data point is a failure occurrence, so when the next failure is long delayed, the cumulative failure rate falls. After the high initial failure rate is smoothed out over time, the cumulative failure rate in Figure 4 looks nearly constant. It actually increases slightly after 350 hours. Eliminating only the first failure data point in the computation of β increases β about ten percent, so that then the value of β is greater than one and α is negative, indicating that there is no reliability growth. The entire data set does show reliability growth, but it is apparent only in the early part of the test.

The cumulative average failure rate from time zero is shown in Figure 4. As the number of failures increases to 56, the cumulative average converges to a steady state constant number and additional failures change it only slightly. Figure 5 shows the average recent failure rate based on the last 20 failures.
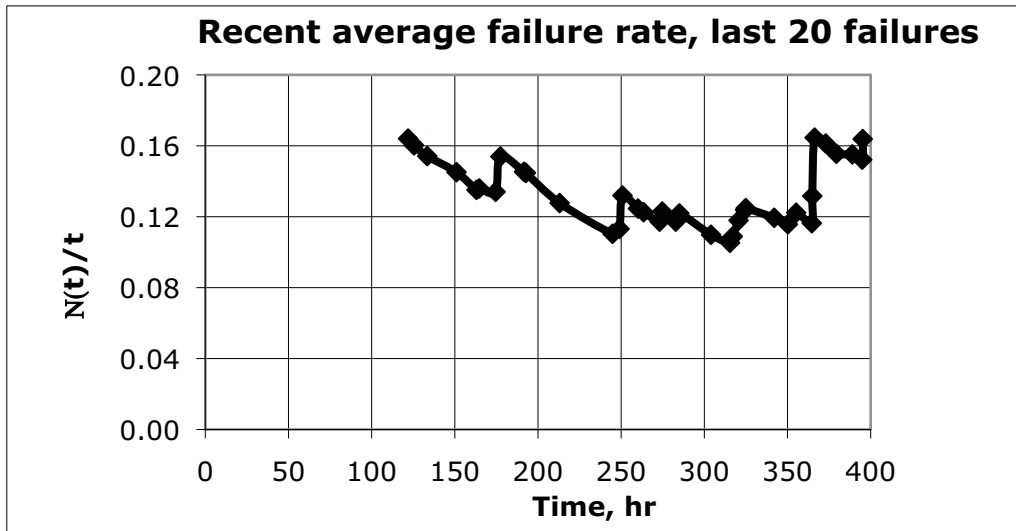

Figure 5. The average recent failure rate based on the last 20 failures.

The first point at 122 hours is the average failure rate based on the first 20 failures. Each new failure is added and the earliest one is dropped, to compute a new moving time window average after each subsequent failure. The recent average failure rate decreases until 250 or 320 hours, but then it increases, first slightly and then with a jump. It could be judged that the average recent failure rate is roughly constant from 220 to 365 hours. Three failures close together at about 365 hours cause a sharp increase in the recent average failure rate.

The statistical analysis based on the NHPP model shows a small positive reliability growth. The final increase in the recent failure rate could then be attributed to random variation. However, the final increase is large enough to

cancel all the previous reliability growth after the infant mortality phase. Clearly, the reliability growth process is not producing significant failure rate decrease in this textbook example.

It appears that the Duane log-log plot, and even the NHPP mathematical line fit can exaggerate the long term reliability growth because they are strongly influenced by the very common early "infant mortality' failures. The recent failure rate and trend are the best predictors of the near future failure rate. And many systems do exhibit an "end-of-life" increase in failure rate.

5. Projecting future reliability.

Reliability growth is driven by an active program to test a system, discover its failures, and then make improvements and fix design deficiencies. The Duane and NHPP reliability growth models attempt to characterize this process. The NHPP model does a better job of fitting the later stages of reliability growth.

The reliability growth models can be used to project the future reliability growth if the past test and upgrade process is continued into the future. The NHPP model indicates that N(t)/t could reach 0.131 at 1,000 hours.

If the test and fix process is terminated and the system put into operation, the best estimate of the future failure rate is the final failure rate at the end of the reliability growth effort. For the NHPP model, the final cumulative failure rate, N(t)/t, is 0.140. The final recent average failure rate, based on the last 20 failures is 0.164, higher tha expected.

IV. Space systems reliability growth

Failure rate data and reliability growth are considered for the International Space Station (ISS), the space shuttle, and the ISS Carbon Dioxide Removal Assembly (CDRA).

A. International Space Station (ISS)

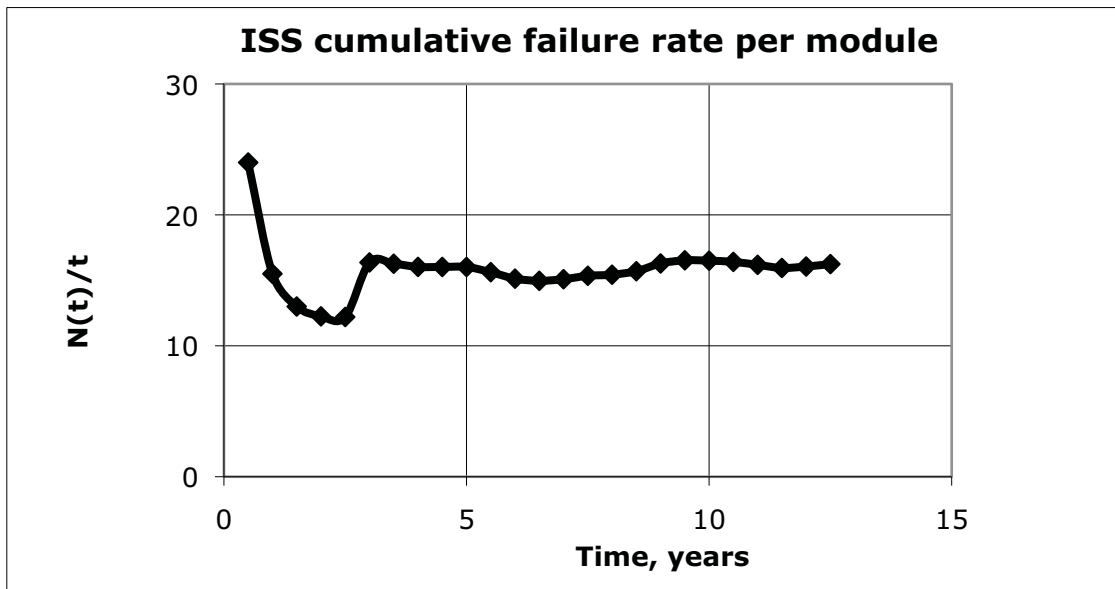The cumulative failure rate versus time for the ISS is shown in Figure 6.

Figure 6. The cumulative failure rate for the ISS.

The failure rate has been essentially constant over time. The ISS failure rate data is based on the number of unscheduled maintenance actions from February 1999 to February 2011. (Cirillo et al.) The number of pressurized modules grew from two to fourteen over this period, and the data is normalized as failure rate per pressurized module. Since the ISS was not permanently inhabited for the first two years, the number of unscheduled maintenance actions was much lower. The same ISS data is plotted in Figure 7 in a log-log plot with HNPP best fit line.
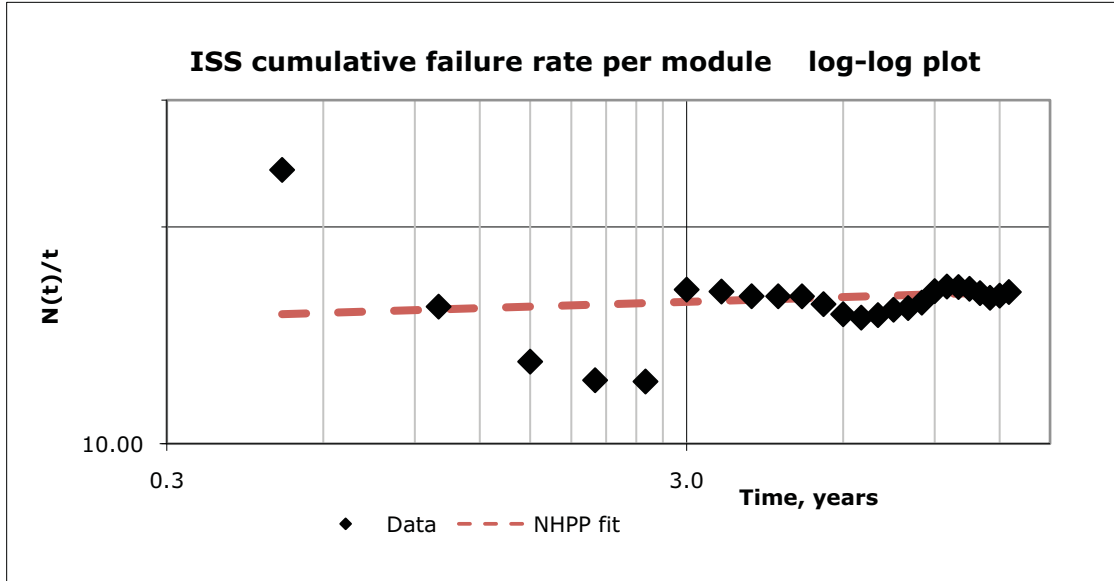
Figure 7. Cumulative failure rate for ISS in a log-log graph with an NHPP line fit.

The NHPP fit line has a β = 1.022, α = - 0.022, roughly corresponding to constant or slightly decreasing reliability. To achieve reliability growth, failures must be analyzed and the design changed to prevent the problems recurring. The ISS is difficult operational environment. Developing and implementing improvements requires a long time cycle, even compared to its long operational life.

Russell and Klaus found that Environmental Control and Life Support (ECLS) repair was a major component of ISS maintenance. They also found that the ISS ECLS maintenance load was constant over an 865 day period. (Russell and Klaus) (Russell et al.) The ISS and ISS ECLS failure data do not show reliability growth.

The ISS failure data was provided by (Cirillo et al.) as the number of maintenance actions related to failures per month. The formula given above for computing the NHPP fit line uses the individual failure time data. Estimating the NHPP line fit for grouped data, such as failures per month, requires a trial and error procedure using a nonlinear equation that is described in (MIL-HDBK-189C, p. 85). The space shuttle and CDRA failure data below are also grouped.

B. Space shuttle

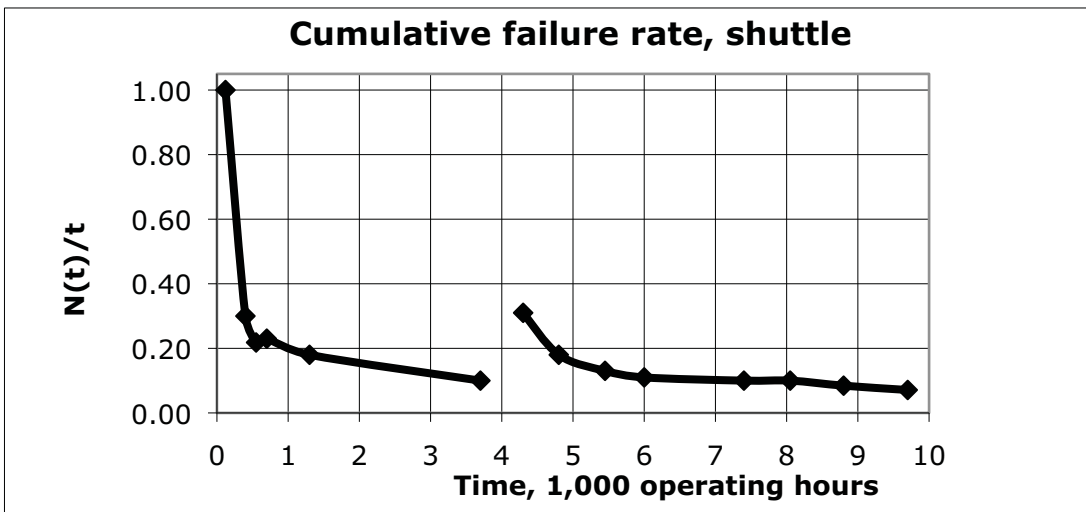The cumulative failure rate data versus time for the space shuttle is shown in Figure 8.

Figure 8. The cumulative failure rate for the space shuttle.

International Conference on Environmental Systems

The data is for the space shuttle Columbia for its first 10,000 operating hours and was reported in this form. The discontinuity at about 4,000 operating hours is due to the interruption of shuttle flights after the 1986 Challenger accident. The other shuttles show similar failure rate declines. (Shishko, SP-6105) The same Columbia space shuttle data is plotted in Figure 9 in a log-log plot with the NHPP best fit lines.
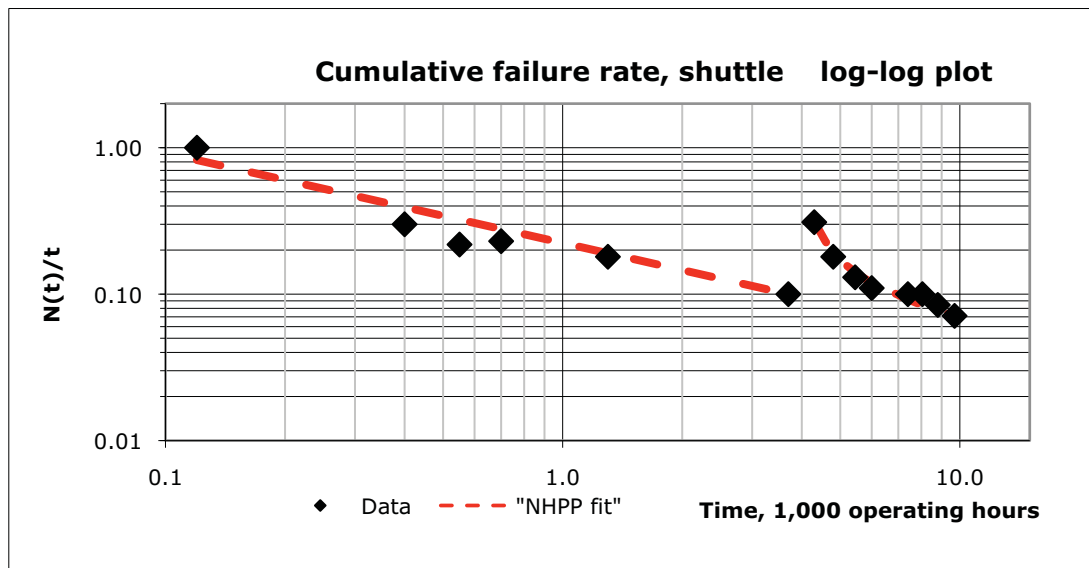

Figure 9. Cumulative failure rate for the space shuttle in a log-log graph with NHPP lines fitted.

The first NHPP line has $\beta = 0.385$, $\alpha = 0.615$, and the second line has $\beta = 0.496$, $\alpha = 0.504$, both corresponding to significant reliability growth. The shuttle was extensively refurbished, maintained, and upgraded after each flight. A large staff had the time it needed in ideal hangar conditions, which is a significant contrast to the ISS in continual orbital flight. The space shuttle failure history is a convincing demonstration of reliability growth in a space system.

C. ISS Carbon Dioxide Removal Assembly (CDRA)

The ISS Carbon Dioxide Removal Assembly (CDRA) has been operating on board ISS since 2001. The number of failures per year was determined from yearly ISS life support system status papers given at ICES. (Reuter, 2000-01-2248, Reuter and Reysa, 2001-01-2386, Gentry et al., 2002-01-2495, Williams et al., 2003-01-2589, Williams and Gentry, 2004-01-2382, Williams and Gentry, 2005-01-2777, Williams and Gentry, 2006-01-2055, Williams and Gentry, 2007-01-3098, Williams and Gentry, 2008-01-2131, Williams and Gentry, 2009-01-2415, Williams et al., 2010-6180, Williams et al., 2012-3612) Two units were on board during the last year reported and the cumulative failure rate is given per unit. The cumulative failure rate for the ISS CDRA is shown in Figure 10.
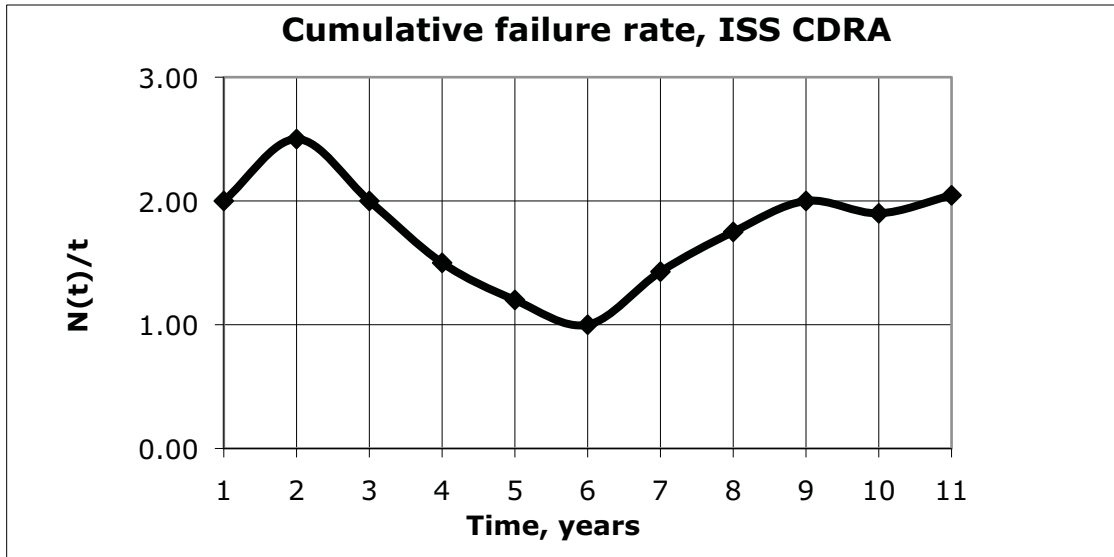
Figure 10. The cumulative failure rate for the ISS CDRA.

The data was gathered in the form of the number of failures per year. No failures were reported in years 4, 5, and 6, which were 2004, 2005, and 2006. This gives an impression of initial reliability growth, but many failures occurred in the following years. The same the ISS CDRA data is plotted in Figure 11 in a log-log plot with HNPP best fit line.
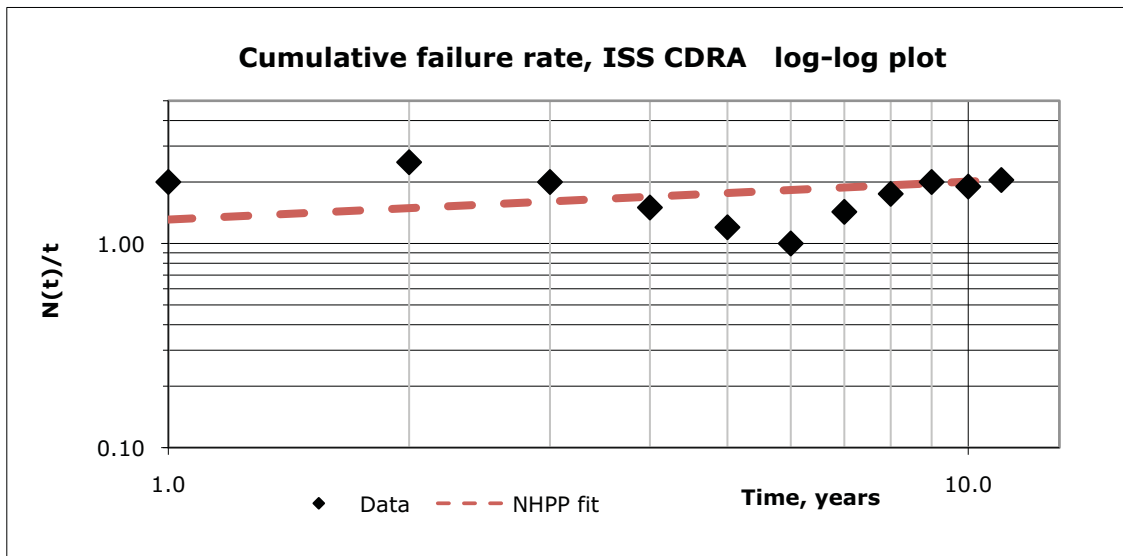

Figure 11. Cumulative failure rate for the ISS CDRA in a log-log graph with an NHPP line fit.

The NHPP line has β = 1. 186, α = -0.186, corresponding to decreasing reliability. The NHPP line fit procedure for grouped data was used. The ISS CDRA, like the ISS ECLSS and overall ISS, does not show reliability growth. The continuing poor reliability of ISS life support is acknowledged and well known. "It continues to be a challenge to provide the functionality necessary to support the six crew members on ISS with all of the problems that have occurred with the new regenerative ECLS racks, WHC, and CDRA." (Williams et al., 2012-3612)

V. Conclusion

The motivation for this work was a hope that the expected reliability growth in space life support would lead to the high reliability systems needed for future missions to deep space, an asteroid or Mars. Although the ISS ECLS is

effectively supporting the crew and providing valuable knowledge and experience, it is clearly not evolving into the ultra reliable system needed for Mars.

As a practical matter, the great complexity of analyzing failures and implementing upgrades on the orbiting ISS makes it very difficult to improve ISS reliability. The ISS maintenance approach accepts the possibility of failures and relies on the resupply and onboard storage of replacement units, which is effective in Earth orbit but not in deep space. And, ISS life support reliability does not seem to be significantly improving.

Logically, systems engineering and design is about making the best practical trade-offs. Reliability is very important, but it is only one objective that must be balanced against many others, including launch mass, crew time for operations, maintenance, and repair, crew time and equipment for science, and the hardware development and total mission costs. The best overall life support system design for a space station, if it was possible to invest development funds to cut total mission cost, would be more reliable and initially expensive than the current system, but would still be much less reliable and expensive than the life support system needed for Mars. If life support fails on the space station, parts and materials for repairs can be provided or the crew can return to Earth. If life support fails on the way to Mars, the crew may be lost. The expense of the ultra reliability needed for Mars is not justified for an orbital space station. The reliability of life support for a space station, even for the Moon or anywhere within the Earth-Moon system, will never grow through the expected process of failure analysis and upgrade, to reach the ultra reliability needed for deep space and Mars.

References

Cirillo, W., Goodliff, K., Aaseng, G., Stromgren, C., and Andrew Maxwell, A., "Supportability for Beyond Low Earth Orbit Missions," AIAA 2011-7231, AIAA SPACE 2011 Conference & Exposition, 27 - 29 September 2011, Long Beach, California.

Duane, J. T., "Learning curve approach to reliability monitoring," IEEE Trans. Aerospace, pp. 563-566, 1964.

Gentry, G. J., Reysa, R. P., and Lewis, J. F., "International Space Station (ISS) Environmental Control and Life Support (ECLS) System Equipment Failures, Causes, and Solutions February 2001 - February 2002." SAE 2002 - 01 - 2495; 32nd International Conference on Environmental Systems; San Antonio, Texas, July 2002.

Hansen, R. C., Overall Equipment Effectiveness, Industrial Press Inc, New York, New York, 2001.

Jones, H. W., "Common Cause Failures and Ultra Reliability," AIAA 2012-3602, 42nd International Conference on Environmental Systems, 15 - 19 July 2012, San Diego, California.

MIL-HDBK-189C, Department of Defense Handbook Reliability Growth Management, 14 June 2011.

Reuter, J. L., "International Space Station Environmental Control and Life Support Status: 1999 – 2000," SAE 2000-01-2248, 30th International Conference on Environmental Systems; Toulouse, France, July 2000.

Reuter, J. L., and Reysa, R. P., "International Space Station Environmental Control and Life Support Status: 2000 – 2001," SAE 2001-01-2386; 31st International Conference on Environmental Systems; Orlando, Florida, July 2001.

Russell, J. F. and Klaus, D. M., "Maintenance, Reliability and Policies for Orbital Space Station Life Support Systems," Reliability Engineering and System Safety, Vol. 92, No. 6, 2006, pp. 808-820.

Russell, J. F., Klaus, D. M., Mosher, T., "Applying analysis of international space station crew-time utilization to mission design," J Spacecraft Rockets, 2006;43(1):130–6.

Shishko, R., NASA Systems Engineering Handbook, NASA-SP-6105, June 1995.

Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2003 – 2004," SAE 2004-01-2382; 34th International Conference on Environmental Systems; Colorado Springs, Colorado, July 2004.

Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2004 – 2005," SAE 2005-01-2777; 35th International Conference on Environmental Systems; Rome, Italy, July 2005.

Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2005 – 2006," SAE 2006-01-2055; 36th International Conference on Environmental Systems; Norfolk, Virginia, July 2006.

Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2006 – 2007," SAE 2007-01-3098; 37th International Conference on Environmental Systems; Chicago, Illinois, July 2007.

Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2007 – 2008." SAE 2008-01-2131; 38th International Conference on Environmental Systems; San Francisco, California, July 2008.

Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2008 – 2009," SAE 2009-01-2415, 39th International Conference on Environmental Systems; Savannah, Georgia, July 2009.

Williams, D. E., Dake, J., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2009 – 2010," AIAA 2010-6180, 40th International Conference on Environmental Systems; Barcelona, Spain, July 2010.

Williams, D. E., Dake, J., and Gentry, G., "International Space Station Environmental Control and Life Support System Status for the Prior Year: 2010-2011, AIAA 2012-3612, 42nd International Conference on Environmental Systems, 15 - 19 July 2012, San Diego, California.

Williams, D. E., Lewis, J. F., and Gentry, G., "International Space Station Environmental Control and Life Support Status: 2002 – 2003," SAE 2003-01-2589; 33rd International Conference on Environmental Systems; Vancouver, British Columbia, Canada, July 2003.

Yamada, S., and Osaki, S, "Reliability growth models for hardware and software systems based on nonhomogeneous Poisson processes: a survey," Microelectronics Reliability, Vol. 23, No. I, pp. 91-112, 1983.