Engineering Reliability for the NEPTUNE Observatory

Paul N. BOWERMAN [1], Harold KIRKHAM [1], George FOX [1], Philip LANCASTER [2], Stephen LENTZ [3], Richard P. KEMSKI [1], Bruce M. HOWE [5]

1) Jet Propulsion Laboratory, California Institute of Technology
2) Alcatel Submarine Networks
3) 360networks
4) Applied Physics Laboratory, University of Washington

ABSTRACT

This paper reviews aspects of Engineering for Reliability of the proposed NEPTUNE observatory. NEPTUNE has an equipment life requirement of around 30 years, and (because of the local weather) presents limited opportunity for repair. It is necessary to use methods and parts that assure the system is sufficiently reliable and maintainable. Among the factors that influence construction are appropriate parts selection criteria, quality assurance inspections, thorough testing of hardware, understanding and accommodating the shock environment of deployment and maintenance cycles, assuring that there will not be material compatibility issues, minimizing the risk that software will result in system failure, analyzing the hardware designs to assure that they should still work correctly after 30 years and that failures do not propagate, and tracking all anomalies to make sure they are appropriately resolved.

1. INTRODUCTION

On July 3, 2002, the CONTOUR spacecraft blasted into space atop a Delta II rocket. A few weeks later, the comet-chaser broke apart when mission controllers commanded the rocket motor to propel it out of Earth orbit. Space travel is difficult. Such short-lived missions are not unknown in ocean engineering: in September 1998 the H2O observatory encountered difficulties less than a day after deployment when problems were experienced with the power subsystem, and the deployment vessel had to return to recover the instrument.

These two examples highlight *the* essential difference between the space mission and the ocean mission: in space, there is almost no possibility of intervention for repair; in the world of ocean science, it has always been assumed that repairs can be made. It is argued here that this cultural difference must erode if large and permanent installations are to be made on the sea-bed. Considerations that may be new to ocean science must assume prominence. If the NEPTUNE observatory, for example, had a node mean time between failures (MTBF) of 10 years (fairly typical for commercial systems of some complexity), around four repair ship-visits per year may be needed every year for the life of the observatory. This level of repair would tax the available resources.

The potential cost of maintaining a subsea observatory is enough to warrant a significant effort at "designing in" a long life for many of the elements under the water. In particular, the large capital investment required for building (or replacing) the *infrastructure* provides justification. There are issues, collectively known as *mission assurance*, of architecture viability, design approach, fabrication techniques, functional and environmental testing, handling and deployment. In short, just how does one go about making a long-life system?

2. BACKGROUND

Some of what is presented here will be from a culture unfamiliar to the oceanographer accustomed to expedition-based science. In expedition science, equipment lifetimes may be short, and since repair is always possible, failure is tolerated as the occasional cost of doing something difficult. The components used for a short-term measurement as part of a week-long expedition may be of unknown heritage, acquired via a catalog, put together by graduate students, and burn-in may consist only of operation during the deployment cruise. In contrast, the components used in a subsea repeater (or in space) have been *qualified* for the application by a process that begins at the factory where they (the components, that is) are made. The purpose of mission assurance is to understand enough about what can go wrong with a system so that failures can be caught at the design stage, during manufacture, or at least before deployment. Therefore, rigorous quality controls throughout manufacture are expected. Assembly of components into systems is done by qualified personnel in a clean-room environment. Burn-in after system assembly is essential. Extensive analysis, testing and documentation take place throughout the work.

This paper will present the methods being proposed in reliability, quality assurance, electronic part characterization, testing, and reporting for the proposed ocean observatory called NEPTUNE. The material should be capable of being generalized: our intent is to describe how the tasks of the scientist and the engineer are changed by considerations of reliability engineering as they take a concept for an Ocean Observatory, or an instrument to be connected to one, through the stages of proposing, designing, building, testing, and deploying.

A word about semantics: in spite of its title, this paper is not just about *reliability* in the sense that mission assurance people use it. Reliability arithmetic is sometimes made to appear as if it constitutes the entire solution to the problem of making something that can be counted on to work. This is far from true: it is akin to asserting that all that is needed for effective policing is good statistics.

The necessary processes in fact begin with management. Management at all levels from the top down must be involved, setting up appropriate levels of reporting, review and configuration control. Engineering aspects involve manufacture, quality control, maintenance and something called "best practices." Analysis techniques include fault-trees, test designs, and cost trade-offs. These topics, and their importance in the various phases of design and development, are described.

Not all ocean observatories are created equal, nor are the instruments attached to them. The infrastructure of a large-scale deep-ocean observatory requires a much higher standard of reliability than a shallow-water observatory, or else the need for maintenance could consume more budget than is available. One instrument in a suite of measurements that are part of an array may sometimes be allowed to fail without serious impact. A unique instrument in a solo situation must not.

Recognizing this, the reader may decide that some of the labor of reliability engineering described here is inappropriate to his or her application. That is understandable. Our intent is not to dictate what must *always* be done, rather it is to indicate what *could* be done in the striving for ocean systems that can be counted on to work.

3. MANAGEMENT

To make a system that will do its job without maintenance for a long time, everything has to work right, from the beginning. The hardware has to be first-rate, the software has to be debugged to an unusual degree, and the processes by which these things come together has to function effectively. Integration of this kind is a management responsibility: without effective management, mission success is far from assured. The many aspects of Mission Assurance are factors that risk management, a top level responsibility, brings together.

What does this mean in practice? It means that the people working on the project understand the issues, and that management has in place mechanisms to deal with them. According to Dev Raheja 1), 85% of quality problems can be attributed to management. Without management involvement, the developers of subsystems do not discuss the trade-offs between (for example) performance, tolerances, operability, hardware, software, and reliability. Without adequate management, changes can be made in one subsystem that might adversely affect others.

Therefore, management must provide an environment in which reporting is routine, reviews are expected, and configuration management is strict. But system developers cannot work in a vacuum: unless the system requirements are carefully spelled out (a process that must involve the scientists), subsystem performance runs the risk of far exceeding needs, or (worse) of not meeting them. It is a management responsibility to ensure this does not happen.

The design and construction of an ocean observatory is engineering, not science, though scientists must be major participants in the work. Engineering management is needed. One way to inculcate an appropriate management approach into ocean science would be to introduce a Responsibility Matrix (a formal way of saying who does what, who reports to whom, who must approve what and so on), or a Reporting Plan. A Sign Off Matrix for trade-offs is a tool that helps keep subsystem interactions under control.

4. QUALITY ASSURANCE

Included in the overall discipline of QA are factors such as metrics, handling and transport, in-process manufacturing inspections, materiel evaluations and more. And of course, many design issues. It is axiomatic that you cannot *test* reliability into a piece of hardware or software. It has to be designed in by appropriate choice of parts, derating, redundancy and so on.

4.1. Parts Selection

The process begins early, with procurement. In-process inspections may be performed at the supplier's production facility, and then throughout the work, to check that the final product will be satisfactory.

Part selection itself is a major engineering task, involving many steps. For any given part, the QA people will want to:

- Evaluate the design
- Study the reliability history
- Analyze its construction
- Examine vendor's production and QA capability

For NEPTUNE, the budget does not allow all these steps to be carried out for all aspects of the design. The only solution was to design with parts that are already on a list of parts qualified for use in submarine cable systems or in space (both fields require high reliability and long life components). Vendors offering such parts make lists available, and indicate the kind of qualification (military, space, etc).

For the engineer or scientist whose organization is not strong in the area, a search on the Web for "qualified parts list" returns many entries for QPL suppliers. Not only can highly reliable parts be found, their failure rate estimates can be found and used to assess an overall design.

In the case of NEPTUNE, the power delivery system was partly re-designed in recognition of the need for the ultimate in lifetime. The backbone cable, segmented to allow cable faults to be isolated, will be controlled by autonomous hardware designed using known high-reliability parts from the submarine repeater as far as possible. (See a companion paper 2).)

By segregating the backbone cable in this way, the estimate for the number of repairs to the backbone is reduced considerably: it is now estimated that there is a 90% probability that the entire backbone will require less than 2 repairs in the 30 year life, not taking external aggression into account. The more complex science node, on the other hand, cannot reasonably be made to the same level of reliability. It is estimated that if as many as 40 science nodes are deployed on NEPTUNE's backbone, a few (perhaps 3 or 4) will require service by a UNOLS vessel every year.

4.2. Derating

If you look at the data sheet for an electronic part, you will see the *absolute maximum ratings* for the part. These are the values of voltage or current or power or temperature that must not be exceeded. Operation of the part at lower values is called *derating*. Derating is done to extend the lifetime of the component.

Of course, there is no simple relation between the amount by which a part is derated and the effect on its lifetime. Further, the amount of derating that can be used depends on the kind of part. You can derate a light bulb by 50% and make it last a very long time: but it will not be very useful! It is clear then that it is beyond the scope of this paper to explore derating in any depth. The reader is referred to the texts (for example, the book by Fuqua 3)), and the various specifications (7 – 18).

An example from NEPTUNE can be given. In the case of the dc/dc power converter, the design constraints are unusual. The converter has to take an input voltage of 10 kV and convert it down to 400 V. The power rating must be 10 kW.

Now in the case of the typical shore equipment for submarine telecommunications, the situation is similar, except that the conversion is up not down, from 48 V to 10 kV. Such a converter

typically occupies a rack of equipment about 2 m high. NEPTUNE's converter had to be made much smaller if it was to be practical. The designer chose to use a large number of small converters (so the frequency could be high and hence parts size could be small) with their inputs in series. As we will see below, putting things in series like this has the potential to reduce the reliability, for if one part fails, the others may have to be taken out of service. With such a configuration forced on us by the size constraint, the only option was to design for life by derating.

Every part in the dc/dc converter is operated at only 50% of its maximum voltage. Semiconductors used in switching are operated with their junction temperatures below 110 C, whereas the allowable temperature is typically 125 C. By these means, the estimated lifetime of a complete stack of 48 converters is about 11 years, depending somewhat on whose method of estimating is used. If the design can be shown to be such that not all failures necessitate a repair, the estimate becomes much larger.

4.3. Redundancy

Redundancy is an obvious way to improve the availability of some function. Some ways of providing redundancy are simple: a single diode can be replaced by four, in series-parallel. In such an arrangement, any diode can fail open or short without affecting the circuit. Other ways are more complex, and require additional hardware to accomplish the changeover from a failed unit to a good one. In NEPTUNE, for example, there will be a redundant dc/dc converter in the science node, so that operation can continue following a converter fault.

There are in general, two major categories of redundancy: active and standby. At the subsystem level, it may make good sense to have the "spare" in a de-energized mode until the main subsystem fails. Most hi-rel systems have lower failure rates when they are turned off than when they are on (though many owners of antique radios will attest that the failure rate is not zero in the "off" state.)

If NEPTUNE uses three converters at every node (the decision is not yet made), the failure rate is calculated to be so low that replacement may never be needed at a given node in the 30-year planned lifetime of the observatory.

A decision such as this is ideally made by considering the capital cost of the subsystems and the maintenance cost. Maintenance cost will itself be determine by the replacement philosophy: if there are three converters in a node, is it reasonable to repair or replace every failed converter, or should the repair wait until the node is down to its last functioning one, or should repairs be delayed until the node is "off the air"? This sort of decision (addressed again below) can really only be made in collaboration with the funding agency, but it will affect the subsystem design.

4.4. Best Practices

"Best practices" is a term used to cover methods that have been found in practice to result in high quality. For spacecraft, the list includes selection of parts from a pre-approved list, testing and certification of all personnel handling hardware, and assembly in a clean-room. Whether or not these practices are needed for ocean observatories has yet to be proven: but it should be borne in mind that very similar practices are used in the subsea telecom business, where the requirement for long life with no maintenance is also paramount.

4.5. Reviews

Design reviews are an integral part of the Quality Assurance effort. The results of these reviews becomes useful information for the designers and operators of the system. At a minimum, the following reviews should be planned:

- Preliminary Design Review (PDR)
- Critical Design Review (CDR)
- Manufacturing Readiness Review (MRR)
- Deployment Readiness Review
- Safety review

NEPTUNE's first review was before any of these – it was the Concept Design Review for the power subsystem. Subsystem reviews such as this may be done *in addition* to the system reviews. For example, a number of detail design reviews may take place between the PDR and the CDR.

A few words about the PDR and CDR are in order.

4.5.1. Preliminary Design Review

When the PDR is scheduled, the overall system architecture has been established, system and subsystem specification will be known, and some of the subsystems may be well into the design process, preliminary reliability estimates will have been made, and the operations concept will be known.

From the Mission Assurance viewpoint, the PDR can concern itself with issues of fault tolerance, modularity, derating, testing, and so on. Enough about the design is fixed that it is usually not possible to insert new technology into the process after the PDR, ie, this step represents a "technology freeze."

A positive outcome from the PDR means approval to complete the design process.

4.5.2. Critical Design Review

By the time the CDR is held, fabrication drawings will be completed, simulations and prototypes will have been produced, and the verification approach will have been established. From the Mission Assurance viewpoint, the CDR will follow up on the issues from the PDR: fault tolerance, modularity, derating, testing, etc.

A positive outcome from the CDR means approval to build the system.

5. RELIABILITY

5.1. Predictions

Reliability predictions are needed at the start of the project, and they begin in concert with the high level design concepts. Reliability Block Diagrams (RBDs) are developed to illustrate the impact of failures of assemblies on the operation of the overall system. In an RBD, assemblies that are all dependent on each other to operate to provide some higher level functionality are shown as a series of blocks (horizontally) – see Figure 1a; assemblies that are redundant of each other, or independently provide a functionality are shown as stacked (vertically) – see Figure 1b. Sometimes a minimum number of independent assemblies must operate; in this case a k-out-of-n designation is used to indicate the minimum number k that must operate out of the total number n – see Figure 1c.

a) Series assemblies. A and B must function for the system to function

b) Parallel assemblies. Either one (or both) of A and B must function in order for the system to function

c) k-of-n Assemblies - Any group of k (or more) of A_1 through A_n must function in order for the system to function
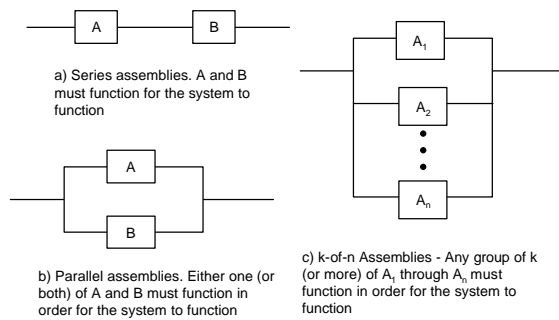
Figure 1. Examples of basic RBDs

Generally, RBDs are much more complicated than these examples, incorporating series, parallel, and *k*-out-of-*n* redundant assemblies in various combinations to model the system. Estimates are used for the reliability (usually expressed as Mean Time Between Failure – MTBF, or failure rate – lambda) of each assembly.

Sometimes the design is influenced by reliability estimates down to the part level, and sometimes at a functional level. For example, as we said above, a single diode can be replaced by four. At the functional level, NEPTUNE's dc/dc converter is the *sine qua non* of the science node. At least one redundant converter will be designed into the node.

When the detailed design has not yet been performed on the assemblies, the MTBFs are generally estimated from similar equipment or from the anticipated complexity of the hardware. At the early stage of the design, there may be a significant error in the estimated reliability of the system, but nevertheless early predictions often point out where designs need significant improvement or redundancy to meet programmatic goals, where designs are marginal and require further investigation as the design matures, and where the design approach is robust to the reliability challenges of the system and further attention can be minimized as the project progresses. For example, in the design of the NEPTUNE node power system, it was clear from the beginning that if the dc/dc converter failed, no science could be done. The addition of one or more redundant converters was therefore planned from the outset. As the design matured, and the estimated MTBF figure for the converter was firmed up, it became evident that a single extra converter in cold standby would enable the power system to meet requirements.

Starting early in the design phase and continuing throughout the design, reliability predictions are an excellent tool to support design trade-off studies. When comparing alternative design concepts, one or several may not be acceptable because of their reliability impact. In this effort, the overall accuracy of the predictions may not be as important as the consistency of the values used. Values propounded by one vendor may be based on different assumptions from those used by others, and reported values may be either optimistic or pessimistic. For example, in the design of the NEPTUNE node, one version of the communication system has reliability data based on fielded hardware, and an alternative version has been modeled using MIL-HDBK-217F, since the design is new. The difference caused by the sources of the reliability data will have to be accounted for in the trade-off, and both designs put on an equal basis. In the case just mentioned, one would have to perform a MIL-HDBK-217F estimate for the first design, since field data would not be available for the competing design.

5.2. Redundancy

For the NEPTUNE system, the inclusion of redundancy can be a two-edged sword. Including redundant units improves the reliability of a single function – going from none redundant to one redundant, may allow for a 50% to 100% improvement in Mean Time Between Critical Failures (MTBCF) for a particular function. However, the number of assemblies failing increases by a factor of 1.5 to 2. If the repair philosophy is to fix the unit on the first failure and reduce the number of critical failures further, then the overall number of ship visits to repair these failures would *increase* by the factor of 1.5.to 2. On the other hand, if repairs are scheduled only after the second failure, then failures requiring repair missions will tend to be clustered later in the system life. Of course, if a node repair was being scheduled for some other reason, the failure causing loss of redundancy should be fixed as well (this may offset the concern about clustered failures). Both the level of redundancy and the repair philosophy need to be well thought out and modeled for NEPTUNE.

5.3. Analyses

But reliability is only one of a suite of Mission Assurance activities performed to assure low failure rates. For instance, if a specific part is applied in a high stress condition, an unusually large number of failures of that part will likely be experienced. In the Reliability Assurance discipline, some activities that help eliminate unanticipated sources of failures include Fault Tree Analysis (FTA); Failure Modes, Effects, and Criticality Analysis (FMECA); Parts Stress Analysis (PSA); and Worst Case Analysis (WCA).

5.3.1. Fault Tree

Fault Tree Analysis (FTA) is a top down analysis that hypothesizes faults at the system level, and using expert knowledge of the system, delves into possible causes that could have resulted in each fault. These possible causes are iteratively checked to determine what could have caused them until root level causes are reached. This information is usually collected in a diagram with AND and OR gates combining the root causes to reach the system level faults. A very useful tool that can be used in concert with the FTA is a Fault Tree Matrix; the Fault Tree Matrix shows each root level cause and the mitigating activities that could be performed to reduce the risk of that cause. Performing the FTA helps ferret out previously unrecognized failure sources and their potential consequences. The Fault Tree Matrix helps prioritize resource allocation to assure that each credible root cause has some mitigating activity taken, or at least considered to be taken, against it.

5.3.2. FMECA

Failure Modes, Effects, and Criticality Analyses (FMECAs) complement the FTAs. A FMECA is a bottoms-up analysis and is typically performed at the electronic part level where there are interfaces, particularly where redundancy is implemented. Parts whose failures could impact the interface are evaluated. Each failure mode of each part is assessed to determine what, and how critical, the effect is on the assembly, subsystem, and system. Where failure modes are found that propagate and result in secondary hardware failures or result in loss of both sides of a redundant design, design modifications are sought to eliminate that possibility, or at least reduce its likelihood. Results of the FMECA can also be incorporated into the Reliability Prediction to model the failure probability better. NEPTUNE needed redundancy to meet reliability goals: this in turn dictated the need to perform FMECAs to verify that the designs will work as desired.

5.3.3. Parts Stress

Parts Stress Analysis (PSA) systematically reviews each electronic part in a design to verify that sources of stress are well below manufacturer's allowable ratings. Parts last longer if they are derated (used at a reduced stress level) relative to manufacturer's requirements. Part operating conditions that can influence lifetime include power dissipation, voltage, current, and internal temperature. The conditions that are checked depend on the type of part being assessed. Broadly, a goal of operating all parts at less than 50% of their rating is conservative. Detailed guidelines for derating levels are to be found in MIL-STD-454, 975 and MIL-PRF-38534. Since NEPTUNE is designed to have a 30-year life, PSA reviews will be needed to verify that parts, and therefore assemblies, do not fail prematurely due to overstress conditions.

5.3.4. Worst Case

As the name suggests, Worst Case Analysis (WCA) assures that circuit designs will continue to operate under the worst case conditions. Part parameters vary due to differences in initial tolerances and drift with time and temperature. In WCA models are generated for the necessary functions that each assembly provides; these models are developed to the constituent part parameters. Part parameter variations due to initial tolerance, time, and temperature are combined. A sensitivity analysis is performed on the circuit model to determine the effect of individual part parameter changes on the modeled circuit performance. Each parameter is set to its worst case value relative to the model, and ability of the circuit to meet its requirement is verified. With NEPTUNE's required life of 30 years, it will be necessary to assure that designs will not be drifting out of specification before the end of its required life.

5.4. Failure Reporting

Monitoring and assuring that anomalies are fixed is another activity that has a great impact on the reliability of the fielded system. Expensive satellite systems have been lost as a result of failures that occurred in flight even though anomalous symptoms occurred in ground test. In those cases repair was not possible (e.g., Mars Polar Lander), or was at least extremely expensive (e.g., Hubble Space Telescope). For NEPTUNE, failures are likely to be very expensive to repair and could potentially take several years (if the problem was built into each node). Therefore, it is important to catalog all anomalies when they occur (so pertinent information about the anomaly is not lost). Anomalies need to be investigated expeditiously, root causes verified, secondary impacts understood, and solutions confirmed by passing the test or other event that originally resulted in the anomaly.

6. TESTING

Testing serves two purposes, in system engineering terminology Validation and Verification. Validation is the process of ensuring that the requirements have been met; the right system has been designed. Verification is the process of ensuring that the system as designed has been built right. Long-life systems require both sorts of test, and the tests must be planned carefully. For the NEPTUNE power system, the Test Plan can be viewed as just one of a series of documents, as shown in Figure 2.
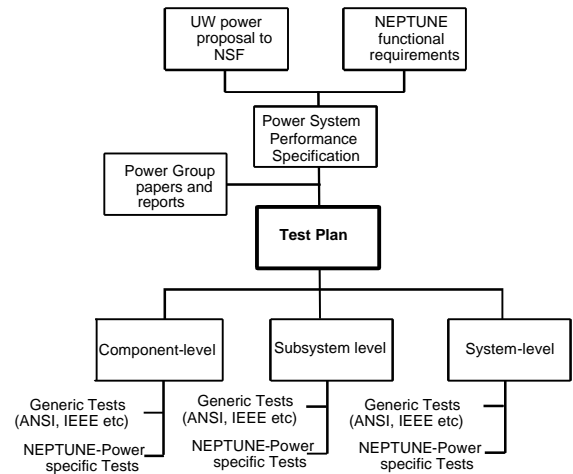

Figure 2 Test Specification Tree

The details of what was promised to the sponsor are contained in the proposal. These, combined with the functional requirements (derived from the science requirements) allow the performance specification to be written. It describes what the power subsystem will do. Voltages, currents (in and out) are specified, as well as conducted and radiated emission limits, controllability, and environmental limits. The Test Plan is then written to ensure that it performs as required.

The plan describes tests to be done at various levels, as shown. In addition, it spells out the resources required, the roles and responsibilities of the participants, and the methods to be used for documentation.

The Test Plan (at present in draft form) and other documents produced by the NEPTUNE Power Group can be seen on the Web site at http://neptunepower.apl.washington.edu/

7. CONCLUSIONS

With the appropriate level of effort in engineering for reliability, a subsea observatory can be designed with acceptable initial cost and acceptable requirements for maintenance. By examining and trading off architectures and designs in the light of reliability engineering, an orderly progress to a successful observatory is made more likely. The design approach, fabrication techniques, functional and environmental testing, handling and deployment are all affected. At times the process may seem overwhelming, but experience in the deepest water and the furthest reaches of space underscore the benefits.

Collectively, these are issues of reliability engineering, or mission assurance, aspects of engineering for a long life that are familiar to the space exploration community and the submarine cable community. They must become as familiar to the world of ocean science, the ultimate users of large observatories, if such observatories are to succeed.

REFERENCES

1) Raheja, D.G.: Product Assurance Technologies, Design For Competitivesness, Inc, Laurel MD, 2002

2) Kirkham, H., Lancaster, P., Liu C-C., El-Sharkawi, M., Howe, B.M.: The NEPTUNE power system: design from fundamentals, a companion paper presented at SSC '03, Tokyo, Japan, June 2003

3) Several reliability prediction models are compared at http://www.relexsoftware.com/reliability/predictionmodels.asp

4) A list of websites related to mission assurance topics is given at http://rac.alionscience.com/rac/jsp/websites/relatedwebsite .jsp

5) Telcordia information can be found at http://www.telcordia.com/products_services/spotlight/systemreliability.html

6) Fuqua, N.B., Reliability Engineering for Electronic Design, Marcel Dekker, New York, NY, 1986

 Some relevant MIL Standards are listed below:

7) MIL-HDBK-217, Reliability Prediction of Electronic Equipment

8) PRISM Reliability Assessment Software superceding HDBK-217, see http://rac.alionscience.com/prism/

9) See also the RELEX implementation of PRISM at http://www.relexsoftware.com/products/prism.asp

10) MIL-HDBK-338, Electronic Reliability Design Hand-book

11) MIL-STD-454, Standard General Requirements for Electronic Equipments

12) MIL-STD-781, Reliability Testing for Engineering Development, Qualification and Production

13) MIL-STD-785, Reliability Program for Systems and Equipment, Development and Production

14) MIL-STD-883, Test Methods and Procedures for Microelectronics

15) MIL-STD975(NASA), NASA Standard (EEE) Part List, Appendix A, Standard Parts Derating Guidelines

16) MIL-STD-1629, Procedures for Performing a Failure Mode, Effects and Criticality Analysis

17) MIL-S-19500, General Specifications for Semiconductor Devices, (deals with high-grade discretes)

18) MIL-M-38510, General Specifications for Microcircuits, (deals with high-grade parts)