AN OVERVIEW OF RELIABILITY

John Cooper, CRE, PE
Ops A La Carte, LLC
Santa Clara, CA, USA
johnc@opsalacarte.com

ABSTRACT
An overview of reliability engineering is presented here to serve as an introductory piece for some, or perhaps as a refresher. Much of the discussion here is general in coverage, and may suggest topics for further investigation.

The perspective here is at the electronic product or module level, which typically includes printed circuit board assemblies (PCBAs), power supply modules, software / firmware, product case, and user interface. The analytical and test tools discussed here are generally well suited for application to the electrical module or product hardware.

The intended audience here is the reliability engineer, or engineers wishing to go into that field. Another audience target may be the manager of that reliability engineering function, who wishes to learn more about the field of reliability engineering. The role of these persons is to promote the identification and completion of these reliability engineering tasks. They often act as the interface to other departments such as engineering, quality assurance, marketing, and manufacturing, so as to optimize the project activities related to reliability.

Key words: reliability prediction, reliability engineering, MTBF, reliability test

PURPOSE AND GOALS OF RELIABILITY ENGINEERING

The goals of reliability engineering is to add value to the organization, reduce warranty costs, and satisfy customers, through the enhancement of product life. Activities include product analysis, product testing, failure analysis and reporting with recommendations. Additional activities would include explaining the reliability activities so that others understand it better. All activities should add bottom line value to the organization.

By utilizing the tools of reliability engineering, one can support the goals of faster time-to-market, and reducing warranty returns, field support costs and have a more satisfied customer.

Some examples of reliability tools include the FMEA, HALT testing, circuit stress analysis, failure analysis, among others.

These and other reliability tools are used in each phase of the product development cycle, from design concept through engineering prototype, manufacturing prototype and throughout the product life.

The focus of reliability engineering often extends beyond what is typically done in traditional design engineering, manufacturing engineering or quality assurance. Much of the efforts in these groups are aimed at product design verification or environmental qualification. Reliability engineering is aimed at increasing product life, with an eye on warranty returns and customer satisfaction.

Consequently, the reliability activities are often not understood from the traditional perspective. For example, in reliability testing, we often stress the product well above the operational specifications. In fact, in HALT testing, the product is stressed until it fails. Then, the failure is analyzed, and corrective actions taken to improve the product. By exercising the product to the point where some problem occurs, we are learning about failure modes that will likely occur in the field. Some failure modes are “soft” - the produce may recover fully when the stress factor is reduced. Other failure modes are destructive - the product will never fully recover, or it may start to exhibit intermittent failures. So, as long as the stress levels are not so severe as to change the nature of the failure mechanism, we can observe failure modes that are representative of what will likely occur in the field, and we can take steps in the development cycle to design those failure modes out of the product. This is key to the basic concepts in reliability engineering.

One traditional definition of product reliability is given as follows: the probability that the product will survive (meet requirements) over a given period of time (say, 1 to 5 years), under a number of stated environmental factors (often ambient temperature and / or humidity), with a stated confidence. In contrast, a common goal of the Quality department is to insure that the product operates as specified, when is it shipped, and that the outgoing defective rate is within some given level (the “AQL” - Acceptable Quality Level).

With working with others on the project team, there are many cases where people place a lot of focus on conformance to specification. The role of the reliability engineer includes helping them understand the methods and goals of reliability engineering. That in order to make the product more robust, we need to stress the product beyond it’s specifications, and verify performance during this excess stress condition. Many do not get this - they want to verify the product meets specification and then move on the the next step in the product qualification process; there is a lot of pressure from above to "ship the product".

Reliability engineering tools range from analytical tools, to product test, to failure analysis. Let us consider reliability tools in use and the role of the reliability engineer, and the order of typical project steps from the engineering concept phase through to the manufacturing buildup.

DESIGN FOR RELIABILITY TOOLS

Table 1 shows a number of hardware Design for Reliability (DfR) tools by steps in the project phase. In the earlier phases, the tools are more analytical. As we get closer to manufacturing, we do more testing, such as HALT. In the manufacturing process we use HASS, Highly Accelerated Stress Screening, based on HALT, as a screening process to identify assembly and component issues.

Table 1: Hardware DfR Tools by Project Phase

Phase		Activities	Tools
Concept		Define HW reliability requirements	♦Benchmarking ♦Internal Goal Setting ♦Gap Analysis
Design	Architecture & High Level Design	Modeling & Predictions	♦Reliability Modeling ♦HW Failure Predictive Analysis (FMEA & FTA) ♦HW Fault Tolerance ♦Human Factors Analysis
	Low Level Design	Reliability Analysis	♦Human Factors Analysis ♦Derating Analysis ♦Worst Case Analysis
Prototype (first time product is tested)		Detect design defects	♦HALT ♦ALT ♦DOE ♦Multi-variant Testing
Manufacturing		Identify and correct manufacturing process issues	♦RDT ♦HASS ♦HASA
Operations and Maintenance		Continuous assessment of HW reliability	♦ORT

FAILURE MODES EFFECTS ANALYSIS (FMEA)

In the early project phases, the FMEA (Failure Modes Effects Analysis) is used to identify, compare and prioritize the failures that might be seen in the products. Design FMEAs are often conducted at the schematic level, or they can be conducted on the product software/ firmware, or on the manufacturing process. McDermott gives an introduction to FMEAs [4], and a more thorough discussion is given in Carlson [3]. FMEAs are very helpful in the planning for reliability testing; test plans should address failure modes brought out in the FMEA. Note that FMEAs are similar to FMECAs [3].

The FMEA can be conducted during most any aspect of the project. It could be an FMEA on the service strategy, it could be on user errors. Early in the project prototype phase, the System FMEA is conducted to look at system failure modes prior to any schematic design availability of the schematic or other solid design details. After a preliminary schematic has been released, the Design FMEA is conducted at a detailed component level, looking at each component, at the failure mode for each component or signal.

THE RELIABILITY PREDICTION, THE MTBF

Note also that with the preliminary bill of materials and the preliminary schematic you could also create the reliability prediction, the MTBF. The reliability prediction is useful at this point in the project to help identify relative reliability of the various assemblies, and provide some guidance on planning for replacement assembly inventory levels for service.

The reliability prediction is a method of calculating the steady-state failure rate or reliability of a product or piece of the product, from the bottom up by assigning a failure rate to each individual component and then summing all of the failure rates. The prediction serves several purposes:

- help assess the effect of product reliability on the quantity of spare units required, this feeds into the lifecycle cost model.
- Provide necessary inputs to the system level reliability models. Examples include frequency of system outages, expected downtime per year, and system availability.
- Assist in deciding which product to purchase from a list of competing products.
- Needed as input to the analysis of complex systems to know how often different parts of the systems are going to fail, even for redundant components.
- Used to estimate product warranty and lifecycle costs, including field service and preventative maintenance.
- Can drive design trade-off studies: compare a design with many simple devices to a design with fewer devices that are newer but more complex. Note that the unit with fewer devices is usually more reliable.
- Set achievable in-service performance standards against which to judge actual performance and stimulate action.

Fundamentally, the reliability prediction is one of a number of reliability tools to be used in your reliability program. It is a communication tool and is typically developed in the early phase of the project. It can be developed based on design documents, and does not require samples or testing.

One of the arguments against using the reliability prediction is that the database failure rates are obsolete or too pessimistic. It does, however, provide a basic model into which you can introduce factors for temperature, component quality, stress factors, and other factors. You can also utilize manufacturers data for more realistic component failure rates.

Failure rate data used in the prediction comes from three sources: failure rates given in the prediction standards, component failure rates from the manufacturers or testing in the lab, and failure rates based on actual field data statistics. Field data is typically the best - the most accurate source of failure rate information.

Some prediction abbreviations:
MTBF = Mean Time Between Failures for repairable systems,
MTTF = Mean Time To Failure for non-repairable products, such as components or lower cost consumer products.
Note that the Failure Rate is the inverse of the MTBF; Component failure rates can be added and are easier to use. Component manufacturers often supply FIT rates; that is, failures per 10^9 hours; the quoted FIT rate = 10^9 / MTTF for the component.

BATHTUB CURVE OF FAILURE RATES

The bathtub curve is used to discuss failure rates for the product as a function of time since first assembly in manufacturing. This curve, shown in Figure 1, is more of a conceptual discussion tool, rather than an actual, specific model of failure rates. It is comprised of the failure rates due to various failure modes that dominate at different times. In the early phases or times since after first assembly, we have the so-called infant mortality failures. (The curve terminology comes from the life insurance industry). In the middle portion of the curve, is the so-called flat or constant failure rate region. This middle region is what the MTBF is intended to predict. And to the right side of the curve is the are the wearout failures, such as rechargeable batteries, connectors, user interface panels, printed circuit board failures due to mechanical vibration or thermal cycling, etc.

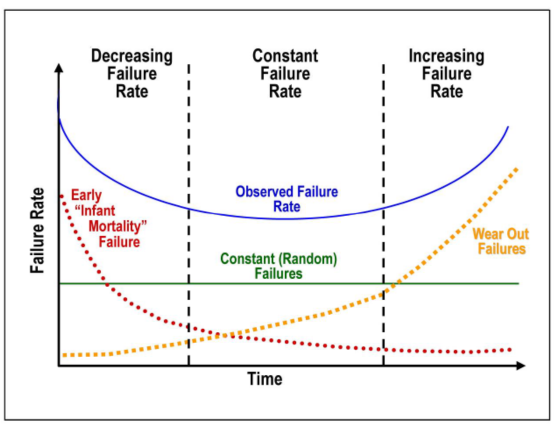

Figure 1: Bathtub Curve

Note then that the MTBF, which refers to the middle portion of this curve, is not really a good tool to predict or measure product life. It does not take into consideration the wearout portion of the curve, nor does it really account for infant mortalities. And yet it is commonly used in the literature perhaps because it seems to offer some relevance when discussing product life. The MTBF is a single number that is mistakenly used to describe what could be completely different product reliability distributions, [5]. More importantly, note that although there is a lot of emphasis put on knowing the MTBF, knowing the MTBF does not, in itself, improve product quality or reliability. It is a major goal of reliability engineering to push for improved product reliability.

An alternate version of the bathtub curve is shown in figure 2, with different regions notated with what reliability tools are used in those areas.

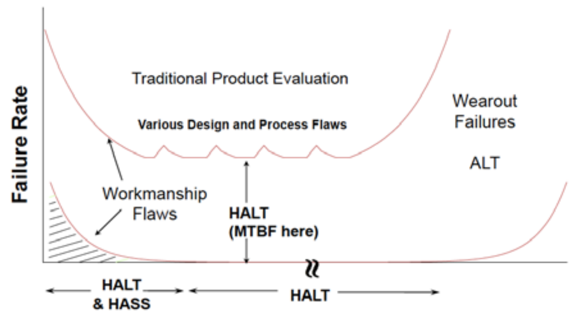

Figure 2: Bathtub Curve with Reliability Tests

In the early product life phase, the HASS process is a manufacturing tool, based on HALT, used to screen out infant mortalities (assembly errors and faulty components). In the flat portion of the curve, HALT test is used to design out failure modes. In the wearout region of the curve, we would use a test like ALT, Accelerated Life Testing, to identify wearout failure modes. ALT could utilize heat chambers or temperature cycling, or it could include button cycling, connector cycling, rechargeable battery cycling, etc. Weibull analysis is used in the wearout region [7,6].

The user interface panel often includes connectors and buttons and display assemblies that are used repeatedly throughout the life of the product. We typically underestimate the usage factors. How many times have we had to call someone on our cell phone multiple times to get through? Yet in the usage models we typically say that a person makes five calls a day or 10 calls a day when in fact they actually use the keyboard to dial out twice or three times the usage model frequency. Usage models should be carefully evaluated; we often underestimate the actual usage, and perhaps should be considering the worst-case or top 10% of worst-case usages. Usage models are often inadequate when considering actual worse-case conditions. There is a certain amount of abuse that should be expected; a moderate amount of abuse is typically considered in safety liability cases. While we may not be able to design the product to withstand all anticipated abuse situations, we should at least consider them, especially in cases of safety or in failures with severe consequences.

Another category of component that often does not get enough attention are the purchased modules, such as a power supply assembly, or a computer board, or complex user interface module. We may not have access to the design details or to the reliability test data, but the manufacturer should, and they should be expected to provide it.

Verification is establishing that the product conforms to requirements or specifications. Validation is establishing that the product meets real customer needs. How many times have we heard that although the product met specifications, it did not satisfy the customer?

ELECTRICAL STRESS ANALYSIS, DERATING ANALYSIS

After the schematic is available, the Electrical Stress Analysis (ESA) and Derating Analysis can be conducted. In this analysis, the circuit schematic is analyzed for electrical stresses that may be present. Examples of stress include voltage stresses on the capacitors, the amount of power dissipated in resistors, and perhaps electrical currents in inductive components and connectors. It may be difficult to obtain the actual current levels or power dissipation but these are key to an accurate stress analysis. These stresses are then compared to the rated component stresses given by the manufacturers, or the recommended stress allowances given in the rating guidelines. For example a typical derating guideline is that the power dissipation in resistors should not exceed 50% of the rated power dissipation. While doing this analysis, components may be identified that are overstressed. That is the electrical stress on the component exceeds the amount recommended by the manufacturer. In these cases it is important that they are carefully analyzed and possibly the the design is corrected. The stress ratings actual stress as a percentage of the rated stress is fed into the reliability prediction to make it more accurate.

Some typical derating guidelines are given in Table 2, from RAC [8].

Table 2: Typical Derating Guidelines

		Environment	
Part Type	Derating Parameter	Severe	Benign
Capacitors	DC Voltage Temp from Max Limit	60% 10° C	90% NA
Diodes	Power Dissipation Max Junction Temp	70% 125° C	90% NA
Lamps	Voltage	94%	94%
Micocircuits	Supply Voltage Fan Out Max Junction Temp	±5% 80% 125° C	±5% 90% NA
Microprocessors	Supply Voltage Fan Out Max Junction Temp	±5% 80% 125° C	±5% 90% NA
Resistors	Power Dissipation Temp from Max Limit	50% 30° C	80% NA

PRODUCT RELIABILITY TESTING

After the construction of electrical prototypes, engineers engineering would typically conduct Design Verification Tests (DVT) to verify compliance to specification over temperature and other parameters, such as power level or voltage. Typically at this point reliability engineering would also get involved and prepare a reliability test plan. That could include testing such as HALT (Highly Accelerated Life Test), RDT (Reliability Demonstration Test) or ALT (Accelerated Life Testing), and other environmental stress tests.

HALT is a stressful quick test of the product subjected to extreme stresses such as temperature step stress, temperature cycling, vibration step stress, and a combination of these stresses. The product is tested to conditions that exceed the specified limits, and is conducted until there are failures reported. Key to the success of HALT is the failure analysis and correction of any failure modes detected. Typically HALT is conducted over a one or two week period and requires several samples. During the testing engineering resources should be available to help analyze the failures and determine corrective steps. The purpose of HALT is to expand the margins, that is the difference between these stresses applied and the specifications how much further beyond the specifications that the product can perform before failure.

HALT traditionally presents stress factors to the product in the form of 5 steps: temperature step stress, first cold then hot, rapid temperature cycling, vibration step stress, and then a combination of all of these. Other forms of stress can and should be used with HALT, such as power supply voltage (both low and high) or frequency, traffic load stress, etc.

Other types of reliability testing include Accelerated Life Testing (ALT), Reliability Demonstration Testing (RDT), Design Verification Testing (DVT), drop or shock testing, humidity, altitude, abrasion, chemical, humidity & salt air, EMC, etc; see Table 3. In most cases, the test is conducted to verify conformance to a specification whether that specification is driven by government, safety, or industry documents, or based on the company's own internal requirements. Additionally that testing may be conducted to test beyond the specifications, to see what margins exist beyond the specified limits, and in some cases this can be used to expand the margins and performance of the product, so as to increase its reliability. In many cases there are no commonly accepted specification requirements and those need to be developed as a result of the testing, and the experience of that company's marketing people and it's customer service people. Tests are listed here to show the variety of stress factors, and to suggest that test planning look beyond the minimal requirements.

Table 3: Variety of Test Stressors used in Reliability or DVT testing

Stressors for Testing	Stressors for Testing
• ESD • Voltage • Current • Power Spikes • Power Cycling • Ripple/Noise • Frequency of circuit crystal oscillators (pulled off nominal) • Drop testing • Bend testing • Water ingress testing • Solar Radiation • Particle contamination • Alpha Particles • Neutron Beam • High Data Rates/ Flooding • Heavy Processor Signal Loading • Pressure • Solar	• Rain • Fungus • Salt/Fog • Sand/Dust • Acceleration • Acoustic Noise • Shock • Pyroshock • Gunfire • Freezing Rain • Ballistic Shock • Data rate – max loading • RF Communications – max data rate transmission • Crush testing (package product) • RF Communications - signal interference • Chemical: skin oil, suntan lotion, etc • Many others

Key to improving the reliability of any product is to introduce test stress factors beyond the minimum requirements, to conduct failure analysis, and consider corrective action for all failures found. In the rush to ship the product, we sometimes fall back on postponing correction, promising to deal with it later. That may be necessary, but the reliability engineer should be investigating and reporting on the consequences of failure modes uncovered, and what potential consequences exist, and follow-up recommendations.

A number of additional reliability product tests are commonly used to verify and validate product performance and reliability. Roughly in order, the following tests may be conducted:

Design Verification Testing (DVT) is conducted during the various product breadboard and prototype phases; it is used to verify product performance against all product specifications with variation of input parameters, such as temperature, power supply voltage, and other usage states.

"The methods used to quantify reliability are the mathematics of probability and statistics" [2]. The topics of probability and statistics used in reliability engineering are well covered in a number of references such as Practical Reliability Engineering by O'Connor and Kleyner [2]. We will make some references to key subjects here but generally the presentation of statistics is beyond the scope of this paper.

HASS, Highly Accelerated Stress Screening, is based on HALT and is used in the manufacturing process to screen out assembly defects and faulty components. Occasionally it will find design defects, so failure analysis is needed to monitor and optimize the process. HASS is nondestructive, although this needs to be verified during the HASS development process. HASS also needs to be confirmed as being effective, i.e., that it can in fact detect failures. If the stresses of HASS are too severe, it can wear out the product. And if the stresses are too mild, it will not detect defects. How many people have experienced a traditional burn in process, where no defects are found? Any test method to be considered should be validated that it is nondestructive (no more than intended), and is effective to detect its intended purpose situations.

Ongoing Reliability Test (ORT), is performed in production on a sampling of outgoing product, while applying a limited amount of stress (typically elevated temperature) that is non-destructive. ORT is useful to demonstrate MTBF, but is not intended to demonstrate failure modes.

FURTHER EXPLORATION

If your goals include getting actively involved in reliability engineering, you should earn the Certified Reliability Engineer (CRE) certification from American Society of Quality (ASQ). The materials to prepare for that class from Quality Council of Indiana are comprehensive, and cover a broad range of topics [9]. While some of the materials covered in the CRE may seem dated, the topics are not, and it is up to the individual practitioner to figure what is relevant and become knowledgeable in the area of interest to them.

There are a number of good books and websites on reliability engineering; one of my favorites is O'Connor and Kleyner, [2]. For a less technical discussion of reliability engineering, see "How Reliable is Your Product?" [1], or for the management of reliability, "Improving Product Reliability", [10].

CONCLUSION

A high-level discussion of reliability engineering is given here, aimed at the engineer or manager seeking an overview of the subject. Discussion is generally focused on hardware reliability of products, systems or printed circuit board assemblies, and discusses many of the concepts and test methods used, and some of the issues and concerns encountered in the field.

Some of the challenges in achieving product liability are discussed, such as the common misunderstanding of stress testing the product beyond its specifications or even beyond its environmental expected use conditions. However these are the steps needed in order to efficiently identify failure modes, and to allow corrections to the design to be made.

The reliability engineering functions when carried out should provide bottom-line financial value to the company, with reduced warranty cost, faster time to market, and improved customer satisfaction.

ACKNOWLEDGEMENTS

The author thanks Jay Muns and staff of Ops a La Carte for their support, and acknowledges the long-term guidance of Mike Silverman.

REFERENCES

[1] M. Silverman, A. Bahret, "How Reliable is Your Product?", SuperStar Press, 2nd Ed., 2016.
[2] P. O'Connor & A. Kleyner, Practical Reliability Engineering, John Wiley & Sons, Inc, 5th Ed., 2012.
[3] C. Carlson, Effective FMEAs - Achieving Safe, Reliable, and Economical Products and Processes using Failure Mode and Effects Analysis, John Wiley & Sons, Inc, 2012
[4] R. McDermott, R. Mikulak & M. Beauregard, The Basics of FMEA, 2nd Ed., CRC Press, 2009
[5] K. Gray, J. Paschkewitz, Next Generation HALT and HASS, Robust Design of Electronics and Systems, John Wiley & Sons, Inc, 2016.
[6] R. Abernethy, The New Weibull handbook, 5th Ed., 2007.
[7] J. McLinn, Practical Weibull Analysis Techniques, Reliability Division of ASQ, 2010
[8] RAC Blueprints for Product Reliability, RBPR-3, 1996
[9] R. Dovich, W. Wortman, CRE Primer, Quality Council of Indiana, 2009.
[10] M Levin, T. Kalal, "Improving Product Reliability" , John Wiley & Sons, Inc, 2003.