Highly Accelerated Life Testing – Testing With a Different Purpose

Neill Doertenbach
QualMark Corporation

Biography

Mr. Doertenbach holds an Electrical Engineering degree from CSU in Fort Collins, Colorado. He has a wide range of industry experience, including digital hardware design, software and Quality Assurance, and is the OVS Technical Sales Manager of QualMark Corporation.

Abstract

This paper describes the technique of HALT – Highly Accelerated Life Testing – and the advantages gained by using the technique. HASS – Highly Accelerated Stress Screening –is also introduced and described. The paper begins with a discussion of the HALT philosophy and how it differs from traditional Design Verification Testing (DVT). The advantages of the technique are highlighted. The process of HALT is described in detail, with emphasis on contrasting HALT with DVT and the logic behind the differences. The discussion of the technique will include preparing for the test, fixturing, the sequence of the applied stresses and the post-test activities. HASS is introduced, including the development of a screen, proof of screen and fixture mapping.

Keywords: HALT, HASS, Accelerated Life Testing, Accelerated Stress Screening, DVT, ESS, Proof of Screen.

Overview and Definitions

In recent years, the test techniques known as HALT (Highly Accelerated Life Testing) and HASS (Highly Accelerated Stress Screening) have been gaining advocates and practitioners. These test methods, quite different from standard life testing, design verification testing and end-of-production testing, are becoming recognized as powerful tools for improving product reliability, reducing warranty costs and increasing customer satisfaction. This paper provides a basic description of these techniques, highlights the differences between these techniques and more conventional testing and provides a guideline for their implementation.

HALT is a test that is performed on a product as part of the design process. Typically it is performed on a product when pre-pilot or pilot run units are available, before the design verification testing begins. During HALT, a product is stressed far beyond its specifications as well as far beyond what the product will encounter in a typical use environment. The actual functional and destruct limits of the product are found and pushed out as far as possible. These limits are used as the basis for the implementation of HASS during the production of the product. HASS is a production screen test, performed on products built as part of the production process. Since HALT is required for the implementation of HASS, HALT will be discussed first.

HALT vs. DVT - the difference is the purpose

When first exposed to the concept of HALT, many design engineers are skeptical of the method. Much of this skepticism stems from the fact that these engineers are used to doing standard life testing and design verification testing, and the HALT methods differ so dramatically from these conventional methods that they seem to be almost at odds with them. The key to understanding the value of HALT lies in understanding the basic difference in the purpose of the testing being done. The basic purpose of Design Verification Testing (DVT) is well understood - it is to demonstrate that the product meets its specifications, and to demonstrate that the product will function in its intended environment. DVT is considered successful when all the tests are passed, with no failures detected.

The purpose of HALT is dramatically different. In HALT, the goal is to over-stress the product and to very quickly induce failures in the product. By applying these stresses in a controlled, stepped fashion, while continuously monitoring the product for failures, the testing results in the exposure of the weakest points in the design. At the completion of HALT, the functional and destruct limits of the product are known, and a “laundry list” of design and process limitations are defined, with corrective actions often defined as well. In short, the goal of HALT is to quickly break the product and learn from the failure modes the product exhibits. The key value of the testing lies in the failure modes that are uncovered and the speed with which they are uncovered. HALT is considered a success when failures are induced, the failure modes are understood, corrective action has been taken, and the limits of the product are clearly defined and pushed out as far as possible. Unlike DVT, HALT is not a pass/fail test. It is a process of discovery and design optimization.

Although these failure modes are induced by stresses in excess of specification, they are typically valid failure modes that would show up in the product in the field. A full failure analysis of all modes found will help confirm this. The important thing to remember is that HALT is finding the weakest parts of the design. These weak links will be the source of warranty problems in the field. The controlled over-stresses applied during the HALT process simply accelerated the precipitation of these failures to allow early detection and correction. The advantage of HALT is that it quickly finds failure modes that would not be brought out in DVT. A typical HALT will take only 3 to 5 days.

Because the purpose of the tests is so clearly different, HALT is not intended to replace DVT. It is true that HALT will find most, if not all, of the failure modes that would show up in DVT (along with many more). However, HALT will not provide you with the documented evidence that you often need to prove that your product meets specification. By doing HALT before DVT is started, you help insure that your DVT will be completed in one pass, with no defects found. This will greatly speed your time to market, avoiding the slow process of repeating DVT until no more failures are precipitated and detected.

Choosing HALT stresses and equipment

The basic concept of HALT can be implemented using many different stresses. However, the stresses most often used are thermal extremes, extreme thermal rates of change, vibration and the combination of thermal and vibration. Other stresses, such as voltage margining, frequency margining, power supply loading and power cycling can also be applied, resulting in additional valid failure modes being exposed.

It is worth remembering that HALT is not intended to demonstrate that a product will function in its intended environment. Consequently, the stresses do not attempt in any way to duplicate those expected in “real life”. Rather, the stresses are specifically designed to quickly bring out failure modes. This logic affects the choice of chamber used to apply the stresses as well as the type of vibration fixturing used and the routing of the air flow through the product. Given that extreme stresses are to be applied, the chamber must be capable of reaching both hot and cold thermal extremes, executing very fast thermal ramps and providing high vibrational energy that will quickly bring out failure modes. This, of course, precludes the use of mechanical refrigeration systems.

The vibration system that has been proven to be the most effective for HALT is a Repetitive Shock (RS) system with a wide frequency and acceleration range and 6 degree-of-freedom vibration. In order to rapidly and effectively bring out failure modes it is important to excite the product at the resonant frequency of all assemblies, sub-assemblies, components and leads and legs of components in the product, regardless of what that resonant frequency, or the orientation of the assembly or component may be. An RS shaker, designed to provide energy from 2 Hz to 10,000 Hz will do this most effectively.

Preparing for HALT and planning the test

Once the purpose of HALT is understood and accepted, the process and stresses used during the testing begin to make more sense. Because the stresses applied are increased until failure occurs, it is not necessary to test a large population of product to insure that a failure mode will be found. A relatively small sample - typically 4 to 6 units - is adequate. This number will allow verification of a failure mode in more than one unit as well as providing for a spare or two in the event of a catastrophic failure of a unit under test. In order to preserve these samples and get as much information as possible from them, the stresses are applied starting with the least destructive and going to the most destructive. For the thermal and vibration stresses, this means starting with cold step stressing, then hot, then rapid thermal extremes, then vibration, followed by a final <u>combined</u> thermal/vibration environment.

If the product being tested is more complex than simply a single board or small system, then one of the first questions to consider is what level of the product to test. In general, the goal of HALT is most effectively met by testing at the lowest possible subassembly. Card cages or other assemblies can dampen vibration and block air flow, reducing the stresses applied to subassemblies inside them. Of course, the trade-offs of functionality and testability must be considered. Also, there will be interconnect circuitry and connections that may not be tested at the subassembly level. An ideal HALT on a complex product would include HALT on all subassemblies, with a final HALT on the upper level assembly as well.

The functional test equipment used during HALT is extremely important. Since the value of HALT is the detection of failure modes induced, it is critical to be able to detect the failures when they happen. This means that the units under test must undergo complete diagnostics while they are being stressed. Much valuable information will be lost if the product is stressed without being monitored, then removed from the stress and tested at ambient. By testing under stress, you will be able to detect "soft" failures that only show up under a particular stress or combination of stresses. These soft failures define the operating limit of your product, and can be the source of troublesome "no defect found" failures when the product reaches the field.

The vibration fixturing used in HALT is very different from that used when testing with typical Electro-Dynamic (ED) or hydraulic shakers. In HALT, the fixture is not designed to mimic the real-life mounting of the product. Instead, it is designed to maximize the transmission of energy into the product to speed the precipitation of failures. This results in simple, inexpensive fixturing with the goal of simply clamping the product to the vibration table as tightly as possible. Figure 1 shows a typical product fixtured in a HALT chamber. To maximize air flow through the product as

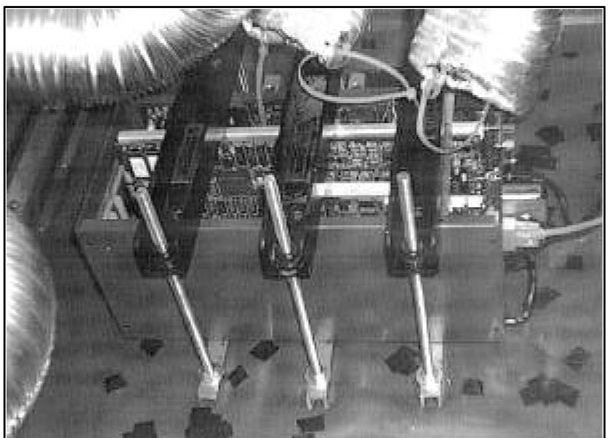

Figure 1: HALT Fixturing Example

well as to improve the transmission of the low frequency energy, the product is set up on aluminum u-channel rather than being placed directly on the table top. The u-channel across the top of the product and the all-thread rod and nuts clamp the product to the table.

Air flow through the product is also planned with the HALT goal in mind. Using flexible air ducts, the air flow is routed to maximize the temperature rate of change on the thermally sensitive parts of the product and to insure that all parts of the product experience maximum temperature extremes. The normal air flow through the product during use is not considered when the ducting is designed. If necessary, holes should be cut in the product's case to allow sufficient air flow across its components.

To aid in failure analysis and to insure that the stresses are being coupled into the product effectively, it is important to instrument the product under test. Thermocouples should be placed at key points on the product, and accelerometers can be placed on boards and subassemblies to evaluate the transmission of energy into the product. However, the actual accelerometer placement should be delayed until after the thermal portion of the stressing is complete, since the accelerometers would be exposed to stress levels that may shorten their life.

A final, important part of the HALT setup is to clearly define what parameters in the product will be monitored, and what constitutes a failure. This fairly obvious step in the test process can be easily missed, making the interpretation of HALT findings more difficult.

Margin Discovery – the core of HALT

With the test set up, the process of Margin Discovery can begin. As mentioned above, HALT will uncover the operational and destruct limits of your product. During testing, the stress is steadily increased in a stepwise fashion, with a complete functional test done at each step. The operational limit is defined as the stress necessary to cause a product to malfunction, but the product returns to normal operation when the stress is removed. Essentially, it is the point of "soft" failure. The destruct limit, as you may guess, is the level of stress necessary to cause a permanent, or "hard" failure to occur. The difference between these limits and your operating specifications is your margin for that particular stress. As the failure modes are found and eliminated that are responsible for these limits, you push the limits further and further out, maximizing your margins and increasing your product's life and reliability.

Figure 2 graphically represents these limits. The stress applied is shown in the X axis, with number of failures shown in the Y axis. The curve drawn around each of the limits represents the distribution of the failure that is responsible for that particular limit. The operating specifications and margins are also shown.

This figure can be helpful in gaining an intuitive understanding of the value of HALT. Consider a failure mode – say, a high ripple on the output of a power supply – that causes a unit under test to fail. If you were able to test hundreds of units, you could see and understand the distribution on that failure mode, as sketched on the graph. However, you do not typically have that luxury. By increasing the stress until the failure is seen, then it doesn't matter where in that distribution the unit under test falls – the failure mode will be detected. If the tail of that distribution happens to fall in the operating specifications, then the failure mode would have been an out-of-box failure mode on some fraction of your products. By doing HALT and stressing to failure, you will find the failure mode without having to hope that your sample size is big enough to exhibit the failure within operating specifications.

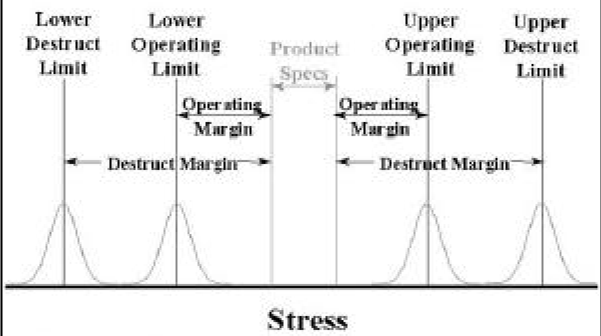

Figure 2: HALT Margin Discovery Curves

But, what if the tails of the distribution are well outside of the product specification, as shown on the graph? Is the high ripple a failure mode that can be ignored? Consider for a moment what happens to this distribution and limits as your product ages in the field. Components fatigue and begin to drift out of specification, power cycles and lightning strikes stress the product, and these limits begin to creep in. If you have chosen to ignore the failure, then you will find that it is one of the first failures to begin showing up in warranty issues. By pushing the stress until the failure occurs, you have effectively accelerated time, precipitating a failure mode in just a few days that could have taken months to come up in the field.

As illustrated in the above example, a failure mode found beyond the operating limits of the product can, indeed, be a "valid" failure mode that could cause warranty problems in the future. However, it is also clear that you may find a failure mode that is completely due to the extreme stress applied, and would never occur in the field. Consider a failure mode precipitated by the softening of a plastic boss at high temperature. A brief failure analysis will reveal that the distribution on this failure mode is clearly understood, will never have a tail that is in the product specification, and will not shift with time and fatigue. Consequently, this failure mode can be safely ignored. Of course, the distribution on most failure modes is not that easily understood. This is one reason why a complete failure analysis is always necessary on HALT failures. In general, it is unusual when a HALT failure can be safely ignored. It is important to resist the urge to ignore a failure mode simply because it happened outside of the specification for the product.

As you test to higher and higher extremes of stress, pushing limits further and further, an obvious question comes up – When do I stop testing? The stopping point will be either the limit of the test equipment, or the fundamental limit of the technology.[1] This fundamental limit is the point where multiple failures begin to occur with small increases in stress. Failure analysis reveals fundamental and catastrophic failures across several devices, with corrective action being prohibitive or impossible. In vibration testing, multiple components are coming off the board.

With this understanding of the margin discovery process, the process of margin discovery can begin. As described earlier, stresses are applied starting with the least destructive and progressing to the most destructive. This helps conserve samples. Cold step is done first.

Cold step testing begins at ambient temperature The temperature is dropped in 5 C^{o} steps. At each step the temperature is allowed to stabilize for 10 minutes. This dwell helps insure that the entire product is stabilized at this temperature, and makes the testing more repeatable. At the end of 10 minutes, a full functional test of the product is done. If the product passes, the temperature is dropped again, and the process repeated. When a failure occurs, the testing is stopped and an investigation into the failure is done. Often, once the failure mode is defined, it is possible to "work around" the failure with a quick patch and continue testing, saving the intensive failure evaluation for later. As described above, this step process is continued until you reach the limits of your test equipment or until you reach the fundamental limit of the technology.

After the cold step is completed, hot step testing is done in a similar manner. Again, testing is started at ambient, then increased in 5 C^{o} steps . The dwell and functioal testing are identical to those done in cold step testing.

The third stress applied in HALT is rapid thermal extremes. Now, the product is functionally tested continuously while the product temperature is changed as rapidly as allowed by the chamber. The upper and lower limits of these ramps are determined by the results of the step stressing, and stay within the operating limits found there (there is no point in repeating failures that were found earlier). If the product cannot tolerate these maximum thermal ramps, then the ramp rate is decreased, and then increased in a stepwise fashion, similar to the thermal step stressing. When failures are encountered, they are addressed in a similar fashion as before.

With the thermal only portion of the testing completed, the product is now exposed to vibration. With accelerometers applied to the product to verify adequate energy transmission to the product, vibration testing is begun at a stress level of 3 to 5 G_{RMS}. Just like in the thermal phase, there is a 10 minute dwell, then a complete functional test of the product is executed. Again, the stress is stepped up, in 3 to 5 G_{RMS} increments, until the chamber limit is reached or you begin to see the catastrophic failures indicative of the fundamental limit of the technology.

The final environment is combined thermal and vibration. Now, the temperature is ramped as it was during the "rapid thermal extremes" portion of the testing, while the vibration is stepped up as it was during the vibration only portion.

It is important to remember that the HALT will be made more effective if additional stresses can be incorporated. By combining more and more stresses, you will bring out failure modes that may occur in the field only under a unique stress situation. This can eliminate a failure mode that could cause a lot of headaches if you were forced to look for it using traditional methods, after the product was released.

At the completion of the step stress testing, you will have found many valuable failure modes for your product. You will have a clear understanding of the margins in your product. You will know not only what your limits are, but WHY they are where they are, giving you a unique understanding of the weaknesses in your product. After doing a root cause failure analysis on all failures found and implementing corrective action, you can do a verification HALT to test your fixes and make sure you have not introduced any new "weak links" in the design with your changes. In the end, you will have optimized the design of your product so that it will last as long as possible in the field.

HASS – maintaining optimization

After your design is ruggedized through HALT and you have completed DVT, you will begin production. As anyone who has seen a product into production knows, the production process can introduce many failure modes that are not related to a faulty design, and the sustaining process can certainly introduce new design problems. HASS is intended to catch these new failure modes more quickly and more effectively than burn-in or other ESS testing done in production.

Once again, an understanding of the purpose of the test is helpful. Burn-in is designed to weed out infant mortality in a product, aging it to induce early life failures before the product ships. HASS has a broader purpose. The goal in HASS is to verify that no new "weak link" has crept into the product since HALT that has shifted either the operational or destruct limits found in HALT.

An important first step to setting up HASS is the completion of HALT on the product. The HASS limits will be set based on the operational and destruct limits found in HALT. Prior to setting up HASS, it is important that corrective action has been implemented on all HALT failures and a verification HALT has been done.

The HASS process and equipment

The equipment used to do HASS is similar to that used in HALT, although often a larger chamber is used to accommodate production quantities. The fixturing can be quite different in HASS, simply to accommodate the production flow. The speed with which product can be fixtured in the chamber becomes important, as well as maximizing the number of products in the chamber. Quick release clamps are often used in lieu of nuts and bolts for securing the product.

An important part of designing a fixture for HASS is the mapping of the fixture. The goal is to insure that the vibration and thermal stresses at each point in the fixture are roughly equal (although precise uniformity is not important). Mapping the fixture involves taking accelerometer and thermocouple readings on a product in each of the fixture locations. It is important the fixture is completely loaded with product for the test, since the load will affect the vibration characteristics. Thermal inconsistencies can be corrected by changing air flow through baffling or other air distribution changes. Vibrational inconsistencies can be corrected through fixturing changes, with the introduction of dampening materials or changes in clamping mechanisms.

During HASS, the stresses are applied simultaneously. Typically, the product is subjected to continuous vibration while the temperature is ramped between its limits, with short dwells at the extremes.

Defining the screen

The levels of the stresses to be applied during the screen are based on the limits found during HALT. There are two parts to the screen.[3] The first part is the Precipitation screen. This screen stresses the product beyond the operational limits and near the destruct limits found in HALT. It is intended to precipitate failures in the product due to latent defects. Because the product is being stressed beyond its operational limit, you do not expect it to function properly, so no testing is done on the product at this point. The product should be powered, however, since applied power can be a significant stress for the product in itself when combined with the other stresses of HASS. The second part of the screen is the Detection screen. During the Detection screen the product is stressed to near the operational limit found in HALT. Now, the product is being functionally tested. Any hard failures induced during the Precipitation screen will be detected, as well as any soft failures that may be induced by the stresses.

Figure 3 can provide an overview of the purpose and

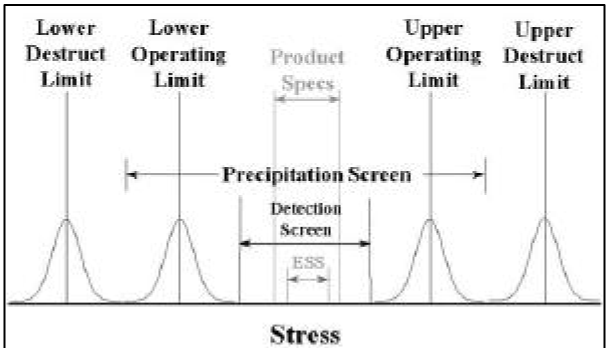

Figure 3: Precipitation and Detection Screen Limits

limits of these screens. It shows the margin discovery curves, overlaid with the Precipitation and Detection screens. The limits on the screens are set so that they are outside of the tails of the distribution of the failure mode(s) that define the operational and destruct limits for the product. Consequently, product which has no new latent failure modes should pass the screen undamaged. Any new failure mode, however, will be exposed. Figure 4 illustrates a typical thermal profile for a HASS screen.

There is one key problem with setting up the limits on the screens from this data – the small sample size used in HALT means that you really have no idea what the distribution looks like on these limits or where the tails may be. Consequently, a more empirical method is

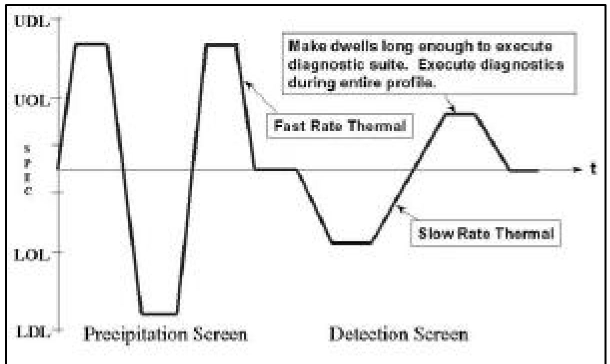

Figure 4: Typical HASS Thermal Profile

used. A baseline for the stresses is derived by guardbanding the limits found in HALT. Typically, vibration is reduced by 50% and thermal excursions are reduced by 20%.[1,2] These limits can be used as a starting point for the Proof of Screen process.

Proof of Screen (PoS) is a critical part of HASS implementation. The goal of PoS is to demonstrate that the screen will reliably find defects without inducing failures or significantly reducing the life of the product. The process of PoS is fairly straightforward. A sample of product – typically a full chamber load – is run through the proposed HASS multiple times. The sample includes some seeded failures – perhaps some "no defect found" failures from field trials. The final configuration of the screen will depend on two factors – the number of cycles through the screen necessary to precipitate the seeded failures, and the number of cycles good product is able to tolerate before exhibiting end-of-life failures.

Figure 5 demonstrates the logic behind PoS. Ideally,

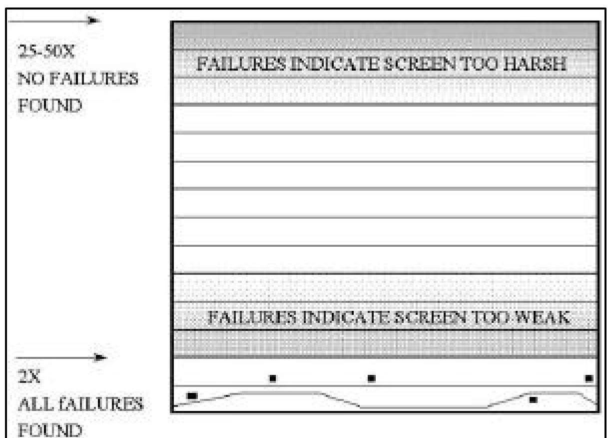

Figure 5: Interpretation of Proof of Screen Results

one or two cycles through the screen will precipitate all the seeded failures. This will yield a short, efficient screen, typically lasting less than 2 hours. As Figure 5 shows, if seeded failures are not precipitated until several passes through the screen, then the severity of the screen should be increased. This part of the PoS verifies that the screen will reliably find defects.

Multiple repetitions of the screen will demonstrate that the screen is not taking an unacceptable amount of life out of the product. Ideally, good product will tolerate 20 to 50 passes through the screen without exhibiting failures. If end-of-life failures are seen before 20 or more cycles are complete, the screen may need to be reduced in severity. A rough estimation can be made of the amount of life being removed from the product by the screen by simply comparing the number of cycles in the proposed production screen to the number of cycles necessary to cause end of life failures to occur. For example, if your production screen consists of 2 passes through the precipitation and detection screens, and your proof of screen showed that 20 cycles through the screen induced no end-of-life failures, then your screen is removing less that 2/20, or 10%, of the useful life of your product.

The stress levels can be adjusted, or the vibration duty cycle can be changed, to achieve the proper balance between the number of cycles necessary to bring out defects versus the amount of life being taken out of the product. If stresses are increased as a result of the PoS, the PoS must be repeated on new, unstressed samples.

In reality, in can often be difficult to seed failures sufficiently to accurately verify that the screen will find defective units. Consequently, it is typically necessary to make a conservative estimate of the number of passes through the screen that are necessary, then tune the screen after a reasonable population of product has been through it. If you find that all of your failures are being precipitated in the first one or two passes through the screen, then no more than two passes should be necessary. Conversely, if you are running 3 passes through the screen and are seeing equal failures in each pass, you should either make the screen more aggressive or increase the number of passes through the screen.

Once your HASS process is defined and proven, it is not necessarily “set in stone”. Product changes can bring acceptable changes in the limits, if they are understood. However, it is always important to base your decisions on a complete failure analysis and a thorough understanding of the impact of the change. Remember that a verification HALT is a useful tool when considering these changes.

Summary

A clear understanding of the unique goals of HALT and HASS provides the basis necessary for introducing the techniques into an R&D and production process. This understanding will also enable you to intelligently make changes in the process. If carefully executed, the end result will be increased product life and reliability, reduced warranty expenses, faster time to market and delighted customers.

---

[1] Hopf, A.M., “Highly Accelerated Life Testing for Design and Process Improvement”, Sound and Vibration, November, 1993, pp. 20-24
[2] McLean, H., “Exceeding the Limits of Traditional Reliability Tests”, Medical Device & Diagnostic Industry, April, 1994
[3] “HALT and HASS, The New Quality and Reliability Paradigm”, G.K.Hobbs, 1996. Published by Hobbs Engineering Corporation, Available Upon Request