From Prediction to Discovery: The Changing Paradigm in Electronic Reliability

Kirk A. Gray
Accelerated Reliability Solutions, L.L.C.
Revised August 2026

Abstract

Electronic reliability engineering developed during an era when vacuum tubes, discrete components, and early solid-state devices exhibited prominent intrinsic failure and wear-out mechanisms. Reliability programs therefore emphasized component failure-rate prediction, burn-in, and qualification to specified environmental limits. Modern electronic systems are different. Highly integrated devices, complex assemblies, global supply chains, rapid component turnover, firmware-software interactions, and manufacturing variation have shifted much of the practical reliability challenge from estimating a single failure rate toward discovering design weaknesses, process-induced defects, and insufficient operating margins. This article argues for a discovery-based reliability strategy centered on stress-to-limit testing, root-cause analysis, physics of failure, Highly Accelerated Life Testing (HALT), and Highly Accelerated Stress Screening (HASS). The objective is not to discard quantitative reliability analysis, but to use prediction only where its assumptions and evidence are appropriate, while placing greater engineering emphasis on discovering how a product actually fails and increasing the distance between normal use stresses and the product's true limits. This distinction is important: empirically validated life models can be useful when a specific physical failure mechanism is known, the relevant stresses can be measured, and the model is supported by representative empirical data. That is fundamentally different from assigning statistically derived component failure rates to a complex solid-state electronic system and treating the resulting number as a prediction of future field reliability.

Keywords: electronic reliability, HALT, HASS, reliability prediction, design margin, physics of failure, manufacturing variation, accelerated testing

1. Why the Reliability Paradigm Is Changing

For many decades, electronic hardware was comparatively fragile. Vacuum tubes contained glass structures and filaments with recognizable wear-out mechanisms. Early solid-state devices were vulnerable to contamination, material instability, packaging limitations, and other mechanisms that could produce time-dependent failures. Because many early defects were accelerated by elevated temperature, burn-in became an effective method for removing some early-life failures. Large discrete components also imposed mechanical limitations that made vibration and shock survival a major design concern.

In that environment, reliability testing was often considered successful when equipment survived its specified shipping, storage, and operating environments. Reliability engineering also developed a strong quantitative tradition: component failure rates were estimated, summed, and converted into system-level metrics such as failure rate and MTBF.

Those methods provided a common language for comparison and acquisition decisions, but the electronics industry has changed dramatically. Today's systems contain highly integrated devices, dense interconnects, high-speed interfaces, complex power-management networks, firmware, software, and components produced through rapidly changing global manufacturing processes. The resulting field reliability is influenced not only by intrinsic component mechanisms, but also by materials, geometry, assembly, process control, mission profile, environmental loading, design margin, and interactions among subsystems.

The important question for the practicing reliability engineer is therefore no longer simply, "What failure rate does a handbook predict?" It is also, and often more importantly, "Where are the weak links, what stresses expose them, and how much margin does the design actually have?"

2. Prediction Has a Role - but It Is Not the Same as Discovery

Traditional reliability prediction methods remain useful for some purposes, particularly when they are used as consistent comparative tools and when their assumptions, data sources, uncertainty, and intended decision are made explicit. IEEE 1413, for example, emphasizes that a credible hardware reliability prediction should identify its inputs, assumptions, data sources, methodologies, and uncertainties so that the risk of using the result can be evaluated [1]. IEC 61709 similarly provides stress-conversion models for component failure-rate data under defined reference conditions, while explicitly limiting such conversion to specified component operating limits [2].

The difficulty arises when a calculated failure rate is treated as though it were a direct measurement of the future field reliability of a complex electronic system. The National Academies has documented important limitations of MIL-HDBK-217-type approaches, including their lack of explicit treatment of materials, geometry, failure mechanisms, and realistic loading histories, and has contrasted those limitations with physics-of-failure approaches [3]. More recent peer-reviewed work has raised similar concerns about constant-failure-rate assumptions and the use of handbook-style predictions as a basis for reliability improvement decisions [4].

MIL-HDBK-217F was validated by the U.S. Department of Defense for continued use in acquisition in June 2022 [5]. However, the 2022 action was a validation notice, not a substantive technical revision of the handbook. The current technical methodology remains MIL-HDBK-217F, Notice 2, dated February 28, 1995. The distinction matters: administrative validation does not demonstrate that models and component data developed for an earlier generation of electronics accurately represent the dominant failure mechanisms or field reliability of modern electronic systems.

For many modern electronic products, intrinsic component wear-out is not the dominant cause of failures that occur within the product's technologically useful service life. In practice, failures are often associated with design margin deficiencies, manufacturing variation, assembly defects, supplier changes, intermittent interconnects, thermal or mechanical interactions, contamination, and unforeseen combinations of stresses. These mechanisms are difficult to discover by summing generic component failure rates.

3. A Discovery-Based Reliability Strategy

A discovery-based reliability program starts with a different objective. Instead of asking only whether a product passes a specified environmental test, the engineer deliberately explores the product's actual operating and destruct limits. Stress is increased in controlled steps, failures are investigated, and the physics behind each weak link is understood. Corrective action is then directed at increasing margin rather than merely demonstrating compliance.

This philosophy is consistent with the qualitative purpose described in IEST-RP-PR003, which characterizes HALT and HASS as methods for improving product reliability and quality rather than as quantitative methods for calculating field life [6]. HALT can reveal whether a design requires corrective action to become more robust; HASS can then use knowledge gained during HALT to detect adverse process or component changes in production.

The distinction is fundamental: qualification asks whether a product meets a requirement; HALT asks how far the product can be pushed, what fails first, why it fails, and whether that limit should be improved.

4. Margin, Variation, and the Weakest Units

Every volume-manufactured electronic product has a distribution of inherent strength. Component tolerances, solder joints, connector engagement, material properties, assembly processes, cooling effectiveness, firmware settings, and supplier variation all contribute to unit-to-unit differences. The stresses experienced in the field also form a distribution. Reliability problems occur when the low-strength tail of the product population overlaps the high-stress tail of the use environment.

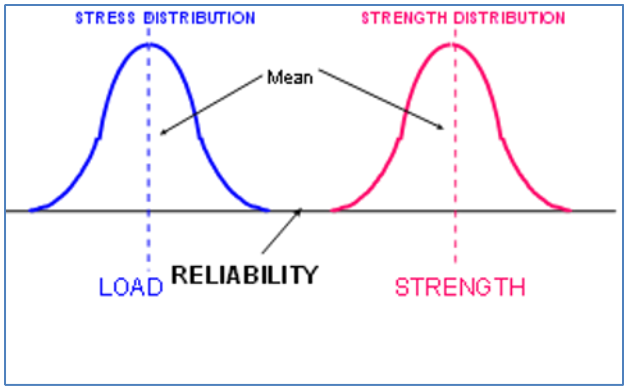

Figure 1. Conceptual separation between fatigue-life consumption and the inherent strength distribution of electronic assemblies. Reproduced from the author's original article.

A robust design maintains substantial separation between these distributions. Manufacturing changes can reduce that separation without producing an obvious immediate failure. An engineering change, operator change, new supplier lot, process drift, altered material, or several small changes acting together may shift the strength distribution downward. The result can be a population of latent defects that passes ordinary production tests but fails under real-world stress.

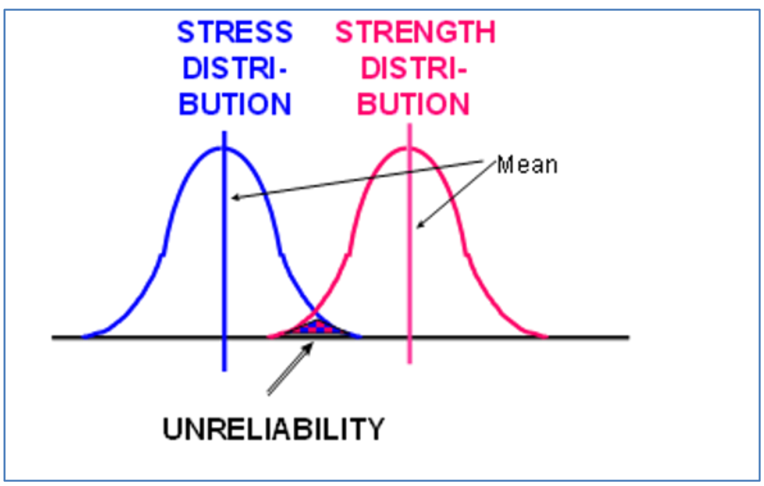

Figure 2. Conceptual field-failure region created when the low-strength tail of the product population overlaps the applied-stress distribution. Reproduced from the author's original article.

The important engineering response is not to attempt to predict every future combination of events. It is to discover the weak links, understand what controls their strength, increase the available margin, and establish production controls that can detect significant shifts before affected units reach customers.

5. Case History 1 - From Burn-In to a One-Hour HASS

One client organization, referred to here as Company A, had a product with an annualized field return rate of approximately 5%. Its existing reliability screen relied on an elevated-temperature burn-in room operating at approximately 90 °F. A HASS process was developed using knowledge of the product's margins and failure behavior.

Six months after implementation of a one-hour HASS process, the reported annualized field return rate had fallen to approximately 0.5%. The important lesson is not that every HASS program will produce the same numerical improvement. Rather, the case illustrates what can happen when a production screen is based on experimentally discovered product capability and is designed to expose latent manufacturing defects rapidly instead of simply accumulating time at a moderately elevated temperature.

This type of result also explains why detailed industrial HALT/HASS results are not commonly published. The combination of lower warranty cost, faster screening, reduced rework, and improved customer experience can create a meaningful competitive advantage.

6. HALT: Finding the Real Limits

HALT is conceptually straightforward, although effective execution requires disciplined instrumentation, functional monitoring, failure analysis, and engineering judgment. Stress is increased beyond normal specification limits in steps while the unit under test remains functionally monitored. The objective is to identify operational limits, determine whether recovery occurs when stress is reduced, continue toward destruct limits when appropriate, and investigate the root cause of each failure.

A practical sequence often begins with low temperature, followed by high temperature, random multi-axis vibration, and then combinations of stresses. The exact sequence should be adapted to the product and the information sought. Individual stresses can simplify root-cause isolation, while combined stresses can expose interactions and are especially valuable when developing production screens.

Three to five samples can provide substantial engineering information during an exploratory HALT, although more units may be required when the objective is to characterize the distribution of a particular failure mechanism. Statistical analysis remains valuable, but it should not displace the central HALT objective: identify the mechanism, understand the physics, and increase the margin.

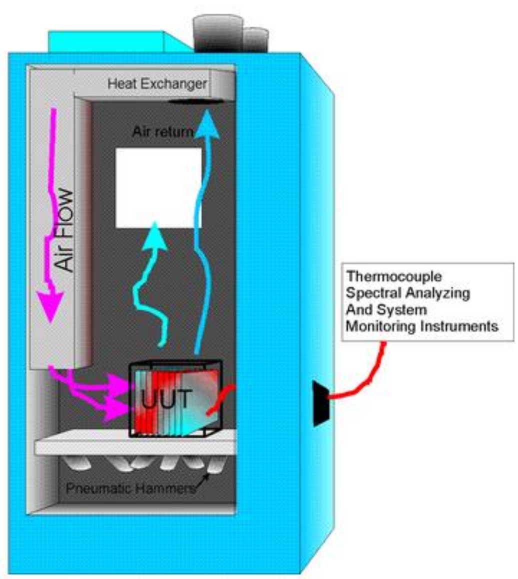

Figure 3. Typical HALT/HASS chamber configuration showing rapid thermal airflow, pneumatic vibration excitation, and functional monitoring. Reproduced from the author's original article.

7. Case History 2 - A 30 °C Margin Improvement from a Small Design Change

A second case demonstrates why failures beyond the product specification can be valuable engineering information. During step-stress testing of an RF power supply, the unit under test reached an upper operating limit of approximately 60 °C. Investigation traced the limitation to a small regulating diode in the +15 Vdc auxiliary power supply. The diode was positioned next to a heat sink but did not make thermal contact with it.

The component was repositioned so that it contacted the heat sink. That small, essentially no-cost change increased the observed upper operating limit by approximately 30 °C, from 60 °C to 90 °C.

The original product already met its specified operating requirement. A conventional pass/fail qualification test therefore would have provided little reason to investigate this component. Step-stress testing exposed the true limit, root-cause analysis identified the mechanism, and a minor design change produced a major increase in thermal margin.

8. Small Physical Changes Can Produce Large Reliability Gains

The same principle applies to mechanical design. Component mass, lead length, support, orientation, mounting method, and local board stiffness can have a large effect on vibration response. Relatively inexpensive layout or attachment changes made early in development can therefore provide substantial improvement in mechanical stress margin.

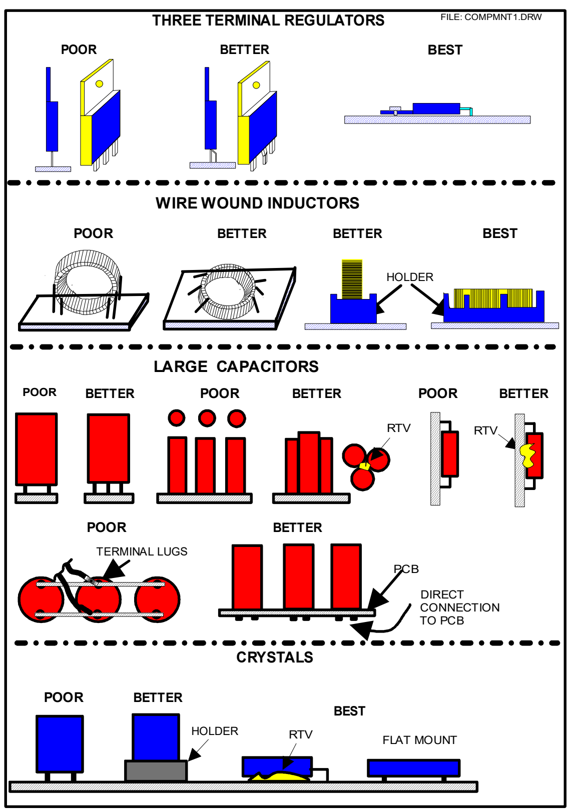

Figure 4. Examples of component-placement and support practices that can improve vibration margin. Original figure courtesy of C. Felkins, Storage Technology; reproduced from the author's original article.

The value of stress-to-limit testing is that these weaknesses become visible while changes are still inexpensive. Reliability improvement is usually far more economical during development than after field returns, redesign, requalification, production disruption, and customer dissatisfaction have begun.

9. HASS: Using Demonstrated Margin to Control Manufacturing

Once HALT has established meaningful operating and destruct margins, that knowledge can be used to develop a HASS process. HASS is not simply 'HALT on production units.' Its purpose is to expose manufacturing defects and process shifts quickly without consuming an unacceptable fraction of the useful life of conforming hardware.

The screen must therefore be both effective and safe. Gregg Hobbs emphasized the importance of Safety of HASS because an overly severe screen can itself reduce usable field life [7]. Appropriate HASS development requires knowledge of the product margins, monitoring of relevant functions, verification that the screen precipitates defects, and evidence that good units are not being damaged.

When properly developed, a higher-stress combined-environment screen can often reveal latent defects much faster than a long-duration burn-in. The engineering advantage is not simply shorter test time; it is that the screen is tied to the actual capability of the product rather than to an arbitrary time-temperature recipe.

10. Physics of Failure Complements Stress-to-Limit Testing

The argument for discovery-based reliability should not be interpreted as an argument against modeling. Physics-of-failure (PoF) methods use knowledge of materials, geometry, life-cycle loading, potential failure sites, and failure mechanisms to assess reliability [8]. When the relevant mechanism and loading are known, such models can provide valuable quantitative insight.

The most productive relationship between testing and modeling is iterative. HALT and failure analysis can reveal the weak mechanism; PoF can help explain its sensitivity to stress and design variables; design changes can increase margin; and subsequent testing can verify the improvement. This is fundamentally different from treating a generic component failure rate as a complete representation of system reliability.

10.1 The Mission-Profile Problem: You Cannot Model Stresses You Do Not Know

More sophisticated reliability-prediction methodologies attempt to address some limitations of traditional component-count approaches by incorporating mission profiles, environmental stresses, supplier quality, and manufacturing-process factors. FIDES, for example, explicitly requires characterization of a product's mission profile and recognizes that local equipment conditions - such as internal temperature rise and the damping or amplification of vibration - must be derived from the system-level environment [9]. This is an important improvement in framing, but it exposes a more fundamental problem: the accuracy of the prediction depends on the accuracy and completeness of the mission profile supplied to the model.

The age of the underlying methodology compounds the mission-profile problem. The current technical content of MIL-HDBK-217F has not received a substantive update since Notice 2 in February 1995 [5]. In the more than 30 years since then, electronic materials, semiconductor technologies, packaging, interconnect density, printed-circuit-board construction, power density, operating speeds, manufacturing processes, software-hardware interactions, and global supply chains have changed dramatically. A methodology rooted in component models and data from that earlier technological era therefore faces two separate questions: whether its component models remain representative of modern hardware, and whether the mission-profile assumptions supplied to those models adequately represent the stresses that today's products actually experience.

For the practicing engineer, relying on such an old methodology without strong validation against relevant modern field data can be analogous to using a medical textbook from a distant era to diagnose a modern patient. Some underlying principles may remain valid, but the materials, technologies, diagnostic knowledge, and available evidence have changed. The concern is not merely that the prediction may be imprecise. A prediction that appears quantitative and authoritative can create false confidence, direct engineering resources toward satisfying a calculated number, and divert attention from discovering the actual weak links, manufacturing variations, and interacting stresses that cause field failures. In that sense, an unvalidated prediction can do more harm than good.

For many complex electronic products, manufacturers do not possess solid, population-wide distributions of the stresses that customers will actually impose over the complete life cycle. A portable consumer product such as a smartphone may experience widely different combinations of ambient temperature, solar heating in a parked vehicle, internal self-heating, charging temperature and frequency, mechanical shock, bending, vibration, humidity, contamination, software workload, and duty cycle. An automobile presents an even broader problem: electronics behind the windshield, inside the passenger compartment, in a door, near a wheel, or under the hood may experience very different local thermal, mechanical, electrical, and humidity histories even though they belong to the same vehicle.

The external environment is only the first level of uncertainty. Reliability-relevant stress must ultimately be known at the location of the potential failure mechanism. The field environment must be translated into the local stress experienced by a particular circuit board, solder joint, connector, semiconductor package, capacitor, battery, or interconnect. Thermal paths, structural resonances, shielding, airflow, mounting, electrical loading, and neighboring components can attenuate or amplify the external stress. Multiple stresses can also interact. Consequently, a nominal system-level mission profile is not necessarily the mission profile of the failure site that determines system reliability.

FIDES itself acknowledges this input problem. Its 2022 Guide states that mission-profile definition is crucial because it determines prediction accuracy, requires assumptions to be documented and validated, and notes that the available input data are generally insufficient to create a mission profile without supplementing them with assumptions [9]. Thus, a sophisticated model does not remove uncertainty when the field-stress distributions, local transfer relationships, and interactions among stresses are not known with sufficient fidelity.

This limitation is especially important when a calculated failure rate is presented as though it were a measurement of future field performance. An independent 2023 assessment of FIDES 2022 concluded that its constant-failure-rate foundation, omission of critical physical attributes such as materials and geometry, and subjective model inputs make it inadequate for producing an accurate reliability prediction. The authors further cautioned against using FIDES predictions as inputs for product reliability improvement, reliability-test design, or support and sustainment decisions [4].

Stress-to-limit testing approaches the problem from the opposite direction. Instead of assuming that every important future stress distribution and interaction can be characterized accurately, HALT experimentally explores the system until the relevant weak links and operating limits become observable. Failure analysis then identifies the mechanism and the local conditions that produced it. The resulting knowledge can be used to increase design margin and to develop production screens. The objective is not to reproduce every possible customer environment; it is to discover which stresses and combinations expose the actual weaknesses of the design before those weaknesses are encountered in the field.

11. What Reliability Engineers Should Measure

For a discovery-based program, the most actionable measures are often not predicted MTBF values but engineering margins and demonstrated process capability. Useful questions include:

- What are the lower and upper operating limits of the complete system?
- What are the destruct limits, and what physical mechanisms define them?

- How much margin exists between normal use and those limits?
- Which design parameters move the limits most strongly?
- What manufacturing variations reduce the margin?
- What is the deviation distribution of margins among multiple samples and mixed builds?
- How do the new products' margins compare with predecessor products and their field failure histories?
- Can the relevant functions be monitored continuously while stress is applied?
- What production screen will reveal significant defects quickly without damaging good hardware?
- Do field-return root causes match the weaknesses discovered during development and production testing?

12. Conclusion - Reliability as Discovery and Margin Improvement

The central challenge of modern electronic reliability is not simply to calculate a number that represents future failure. It is to discover the mechanisms that can cause failure, determine how close the design is to those mechanisms, and increase the margin before customers encounter them.

Traditional reliability prediction methods such as MIL-HDBK-217 continue to be used in acquisition programs and as standardized methods for comparing calculated reliability values between designs. However, their continued use should not be interpreted as empirical evidence that the resulting predictions correlate with the actual field reliability of modern solid-state electronic systems. The author has found little publicly available field data demonstrating such a correlation. A standardized calculation may provide a consistent number for comparison, but consistency of calculation is not evidence of predictive accuracy.

This criticism should not be extended indiscriminately to all reliability and life modeling. Mechanical systems and structures can experience identifiable fatigue, wear, corrosion, creep, fracture, and other cumulative damage mechanisms. Electronic assemblies can also contain materials, interconnects, connectors, solder joints, batteries, capacitors, and other elements with identifiable degradation or wear-out mechanisms. When a physical failure mechanism is known, the relevant stresses can be measured, and the model is supported and validated by empirical data representative of the actual application, life modeling can provide valuable engineering information. The objection is to assigning a statistically derived failure rate to complex solid-state electronic systems and treating that number as evidence of their future field reliability without demonstrated physical and empirical correlation.

HALT changes the frame of reference from testing only to specification toward discovering actual operating and destruct limits. Root-cause analysis and physics of failure turn those limits into engineering understanding. HASS then uses the design's demonstrated strength to detect manufacturing defects and process shifts rapidly.

The most useful reliability question is therefore not merely, "What failure rate can we predict?" It is, "What is the weakest link, why is it weak, and how much margin can we create before the product reaches the customer?"

AI Assistance Disclosure: The author used OpenAI's ChatGPT to assist with editorial revision, organization, language refinement, and identification of potential reference sources. The technical arguments, engineering interpretations, figures, case histories, and conclusions are those of the author. The author reviewed and approved the final content and is responsible for its accuracy.

References

[1] IEEE Std 1413-2010, IEEE Standard Framework for Reliability Prediction of Hardware. Source

[2] IEC 61709:2017, Electric components - Reliability - Reference conditions for failure rates and stress models for conversion. Source

[3] National Research Council, Reliability Growth: Enhancing Defense System Reliability, National Academies Press, 2015. See discussion comparing MIL-HDBK-217 with physics-of-failure approaches. Source

[4] A. Gaonkar, R. B. Patil, D. Das, M. H. Azarian, B. Sood, and M. Pecht, "Assessment of the FIDES Guide 2022 Electrical, Electronic, and Electromechanical Reliability Prediction Methodology," e-Prime - Advances in Electrical Engineering, Electronics and Energy, 2023, 100353. Source

[5] U.S. Department of Defense ASSIST, MIL-HDBK-217, Reliability Prediction of Electronic Equipment, current status listing. Source

[6] Institute of Environmental Sciences and Technology (IEST), IEST-RP-PR003, HALT and HASS. Source

[7] G. K. Hobbs, "Pitfalls to Avoid in HALT and HASS," Hobbs Engineering. Source

[8] M. Pecht et al., Physics-of-Failure approaches to electronics reliability and prognostics; overview in Prognostics and Health Management of Electronics, Wiley. Source

[9] FIDES, FIDES Guide 2022 Edition A (July 2023), sections on mission profiles and reliability prediction methodology. Source

[10] K. A. Gray, Next Generation HALT and HASS for Electronics, Wiley, 2016. Includes HALT/HASS case histories and applications. Source

Author's Note

The two case histories and the original figures in this article are based on the author's prior engineering and consulting experience and were preserved from the May 7, 2026 version of this article. Published references are provided to distinguish the broader industry and academic literature from the author's experience-based observations.