TEMPERATURE IS A STRESS, NOT A CLOCK

Replacing Arrhenius-Based Life Assumptions with Thermal Margin, Signal-Integrity, and Hardware-Marginality Discovery

Kirk A. Gray

Accelerated Reliability Solutions, L.L.C.

Central thesis

Temperature should not be treated as a universal clock that consumes electronic life at a calculable rate. It should be used as an empirical engineering stress that shifts electrical, timing, material, and mechanical parameters so that marginalities can be observed, understood, and corrected.

Informative technical article with numbered references
Prepared July 2026

Abstract

Electronic reliability practice has often assigned two incompatible meanings to temperature. In traditional prediction and accelerated-life calculations, temperature is treated as a clock: a 10°C increase is assumed to double failure rate or halve life. In reliability development, temperature is used as a stress: it changes semiconductor switching behavior, conductor resistance, dielectric loss, timing, mechanical strain, and signal-integrity margin so that hidden weaknesses become observable. This article separates these uses. It explains why the Arrhenius equation is valid only for a defined thermally activated mechanism and why the universal 10°C rule is not a physical law. It then integrates evidence from long-term computer overstressing, large field populations, semiconductor research, signal-integrity studies, and HALT case histories. The evidence supports a more useful engineering model: temperature does not have one universal meaning, and it should not be converted automatically into field-life equivalence. Instead, hot, cold, and rapid thermal transitions should be used to discover empirical operating limits, reveal marginal signal paths and hardware-software interactions, identify physical failure mechanisms, and guide design improvement. Arrhenius remains valuable when the mechanism and activation energy are validated. Thermal HALT remains valuable precisely because it need not assume Arrhenius life consumption.

Keywords—Arrhenius equation, 10°C rule, accelerated testing, HALT, signal integrity, timing margin, semiconductor reliability, thermal cycling, intermittent failure, no fault found.

Executive Summary

The familiar statement that electronic failure rates double for every 10°C increase is not a universal consequence of the Arrhenius equation. The calculated multiplier depends on activation energy and absolute temperature, and a complete electronic product contains many competing mechanisms that do not share one activation energy. Treating a system as though it were one thermally activated chemical reaction creates an appearance of precision without demonstrating the physics of failure. [1]-[7]

Yet rejecting the 10°C rule does not reduce the importance of thermal testing. Temperature is exceptionally effective at perturbing a digital system. It changes propagation delay, threshold voltage, leakage, conductor resistance, dielectric behavior, power-distribution performance, clock relationships, mechanical dimensions, and contact forces. These changes can narrow an eye diagram, increase bit-error rate, violate setup-and-hold margins, trigger race conditions, and produce software-visible lockups or data corruption. [17]-[24]

The distinction is simple but consequential: a recoverable system crash at a thermal limit demonstrates inadequate operational margin; it does not by itself demonstrate permanent aging or an Arrhenius-equivalent number of field hours. Conversely, a progressively degrading mechanism may be accelerated by temperature, but only after the same mechanism is shown to occur at stress and use conditions.

The argument in one sentence

Temperature can accelerate a specific physical reaction, but it is often more valuable as a controlled stimulus that exposes marginal operation than as a universal measure of consumed life.

1. Two Meanings of Temperature in Reliability Engineering

Temperature has been given two roles in electronics reliability. The first is predictive: temperature is inserted into a handbook, model, or acceleration equation to estimate failure rate or lifetime. The second is empirical: temperature is deliberately varied while a product operates so that weak margins, latent defects, intermittent interconnections, and material limits can be discovered.

These roles are often conflated. When a unit fails at high temperature, the failure may be described as "accelerated aging" even when it resulted immediately from slowed logic, a shifted threshold, a regulator limit, a thermal-protection circuit, or a material state change. When a unit survives thousands of hours at high temperature, the result may be discounted because a prediction model claims that an enormous amount of life must already have been consumed. Both interpretations can be wrong if the actual mechanism is not identified.

A more useful approach is to ask two separate questions: (1) Did temperature accelerate a defined physical degradation mechanism? and (2) Did temperature perturb the operating system sufficiently to reveal inadequate margin? The first question concerns life extrapolation. The second concerns design discovery. They require different evidence and different models.

2. Why the 10°C Rule Is Not a Physical Law

A common Arrhenius acceleration-factor expression is:

$$AF = \exp[(E_a / k) \times (1/Tuse - 1/Tstress)]$$

Here E_a is the activation energy of the particular process, k is Boltzmann's constant, and temperature is absolute. The equation cannot produce a unique multiplier until a mechanism-specific activation energy and two temperatures are supplied. NIST explicitly treats Arrhenius as a model for particular chemical, diffusion, or migration mechanisms—not for the total failure rate of an arbitrary electronic product. [4]-[6]

Table 1. Calculated acceleration for a 10°C increase

Temperature change	0.3 eV	0.7 eV	1.0 eV
25°C to 35°C	1.46×	2.42×	3.54×
55°C to 65°C	1.36×	2.08×	3.00×
85°C to 95°C	1.27×	1.85×	2.52×
125°C to 135°C	1.20×	1.65×	2.16×

Calculated from the Arrhenius equation. A 10°C increment does not produce a constant multiplier; the result changes with activation energy and starting temperature.

The phrase "Arrhenius is erroneous" is therefore best understood as a criticism of application, not of physical chemistry. The equation can be appropriate for a known mechanism. It becomes erroneous when a customary activation energy is assigned without evidence, when multiple mechanisms are collapsed into one multiplier, or when the result is treated as a prediction of complete-system field reliability. Edward B. Hakim raised this concern directly in his 1990 editorial, "Reliability Prediction: Is Arrhenius Erroneous?" and later publications. [2], [3], [15]

3. A Component or System Does Not Have One Activation Energy

A semiconductor package may contain silicon, gate dielectrics, copper or aluminum conductors, barrier metals, passivation, mold compound, die attach, bond wires, solder, and numerous interfaces. A printed wiring board assembly adds laminates, vias, connectors, cables, regulators, capacitors, oscillators, fans, mechanical supports, and software-controlled operating states. The first failure in use is governed by whichever weakness reaches a critical condition first.

Table 2. Mechanisms that should not be collapsed into one temperature multiplier

Mechanism or response	Dominant stresses or variables	Appropriate interpretation
Diffusion, corrosion, chemical reaction	Temperature, humidity, material chemistry	Arrhenius may apply after mechanism validation
Electromigration	Temperature, current density, geometry	Multi-stress model; not temperature alone
Solder or package fatigue	Temperature range, strain, dwell, cycle count	Creep-fatigue or strain-based model
Fan, bearing, and connector wear	Speed, load, lubricant, vibration, contamination	Mechanical/wear model
Timing or signal-integrity failure	Temperature, voltage, frequency, process variation	Operational-margin discovery; often recoverable
Polymer deformation or solder reflow	Threshold temperature and material state	Limit violation, not normal-life equivalence
Software-visible intermittent failure	Hardware timing, BER, sequencing, environment	Cross-domain failure analysis

The engineering error is not merely numerical. A single multiplier suppresses the question that matters most: What actually failed, and why?

4. Evidence from Long-Term Thermal Overstressing of Computers

Gray and Pecht reported an exploratory study in which operating business computers were subjected to steady-state high temperature and thermal cycling far beyond the manufacturers' specified ranges. The sample size was limited, and the study was not a statistically powered life demonstration. Its value lies in the observed failure pattern and in the contrast between permanent hardware failures and recoverable operational failures. [1]

4.1 Steady-state high-temperature operation

Five computers were placed in environmental chambers in 2008, with a sixth added later. Their specified operating ambient range was approximately 10°C to 35°C. The original target was 70°C, but one CPU reached its protective shutdown limit, so the systems were operated at approximately 65°C ambient while running diagnostics, graphics-intensive software, and comprehensive stress applications. [1]

Measured internal temperatures included motherboard values of 76°C to 81°C, CPU temperatures of 90°C to 97°C, northbridge temperatures of 94°C to 96°C, and southbridge temperatures of 114°C to 124°C. Nonrecoverable failures occurred after approximately 2,300 to 8,000 hours in the first stress operation. The observed failures involved hard-disk drives, power supplies, fans, electrolytic capacitors, and heat-deformed optical-drive mechanisms. No microelectronic component failure was confirmed. [1]

4.2 Thermal cycling and transitional failures

Four computers were cycled from approximately 10°C to 65°C at 1°C per minute with 30-minute dwells. Three completed approximately 2,750 cycles and then passed 168 hours of operation. One unit developed frequent POST lockups during thermal cycling, but the failure could not be reproduced under steady-state ambient or 65°C conditions. Cross-sectioning found fatigue damage and solder anomalies, but these did not produce a confirmed hard failure during the experiment. [1]

What the computer study supports—and what it does not

The study supports the conclusion that extreme temperature did not produce the universal pattern

of rapid semiconductor failure implied by the 10°C rule. It also shows that thermal transitions can expose recoverable operating marginalities and physical damage. It does not prove that temperature never accelerates semiconductor degradation or that all computer designs will respond similarly.

The same experiment therefore connects the two meanings of thermal stress. Constant elevated temperature did not act as a simple semiconductor life clock, yet thermal stress and transitions did reveal system margins, protective limits, intermittent lockups, and weaknesses in electromechanical and power subsystems.

5. Temperature as a Parametric Stress in Digital Systems

A digital system operates through relationships among voltage, time, and logic thresholds. Manufacturing variation, board fabrication, connectors, component tolerances, software sequencing, and environmental conditions all influence these relationships. Thermal stress changes many parameters simultaneously, making it an unusually effective stimulus for discovering low margin. The companion paper on thermal HALT describes this as skewing circuit-level parametrics to empirical operational limits. [17]

Conductors have resistance, inductance, and capacitance. At high edge rates, traces, vias, packages, and connectors form distributed transmission structures. Temperature changes conductor resistance, dielectric properties, semiconductor mobility, leakage, threshold behavior, and propagation delay. The effects accumulate across a path and can alter signal amplitude, rise and fall times, jitter, skew, and noise margin. [18]-[24]

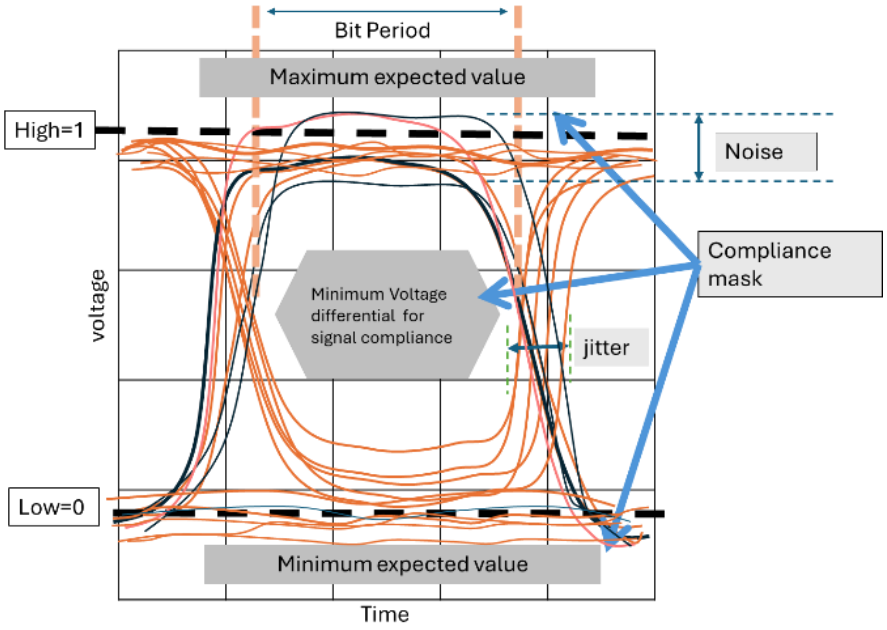

Figure 1. Eye diagram and compliance mask. A smaller vertical or horizontal opening represents reduced voltage or timing margin and a higher probability of bit errors. Adapted from the author's thermal-HALT signal-integrity paper. [17]

An eye diagram overlays many transitions in a serial data stream. The vertical opening represents amplitude and noise margin; the horizontal opening represents timing margin. If a waveform crosses the compliance mask, the receiver may misinterpret a binary state. A degraded eye can therefore appear to software as corrupted data, reduced throughput, retries, lockups, or crashes rather than as a visibly failed component.

6. Thermal Effects on Signal Integrity and Timing

6.1 Conductor resistance and insertion loss

Copper and aluminum resistance increase with temperature. The author's signal-integrity paper cites changes of approximately 77% and 72%, respectively, across -55°C to 125°C, illustrating that wide thermal excursions can significantly alter interconnect behavior. PCB studies have also reported temperature-related insertion-loss changes in microstrip, stripline, and via structures. [17], [20]

6.2 Semiconductor propagation delay

Semiconductor switching delay changes with temperature because carrier mobility, threshold voltage, leakage, and drive current do not remain constant. The direction and magnitude depend on technology and operating point, but the practical consequence is that clock and data paths do not shift identically. A system with adequate room-temperature margin can approach setup, hold, reset, or sequencing limits when hot or cold.

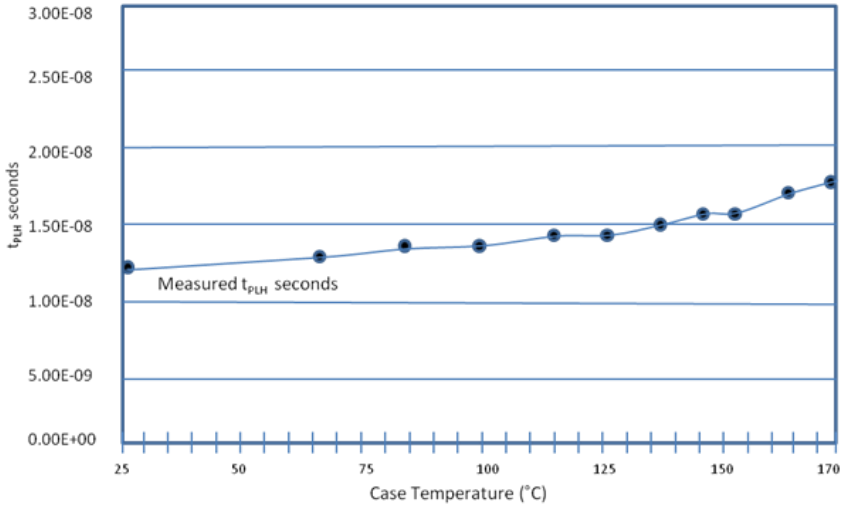

Figure 2. Example of low-to-high propagation delay increasing with case temperature for an octal buffer. Adapted from the signal-integrity paper and manufacturer data. [17], [25]

6.3 Eye opening, jitter, and bit-error rate

Table 3. Typical thermal effects observable in an eye diagram

Signal measure	High-temperature tendency	Low-temperature concern
Vertical eye opening	Can decrease from higher loss, leakage, and power drop	May improve or degrade depending on threshold and noise shifts
Horizontal eye opening	Can narrow as propagation slows and jitter rises	Can become asymmetric when paths speed up differently
Rise and fall times	Often slower	Often faster, but path mismatch can increase skew
Jitter and skew	Thermal noise and unequal path delay may increase	Metastability or sequencing errors may appear
Bit-error rate	Can increase as margin closes	Can increase when thresholds or timing relationships shift

These effects are not equivalent to wear-out. They are immediate operating responses. Their importance is that they expose whether the system has sufficient margin to tolerate production distribution, component substitutions, aging drift, workload variation, and environmental conditions.

7. Manufacturing Variation and the Small-Sample Development Problem

Early prototypes are commonly built from a narrow set of component and fabrication lots. The samples may therefore cluster near the center of the production distribution and fail to represent the tails that will appear when thousands or millions of units are manufactured. Process variation exists in semiconductor geometry, oxide thickness, threshold voltage, resistance, leakage, interconnect dimensions, dielectric properties, surface roughness, solder attachment, connectors, and cables. [17]-[19], [23]

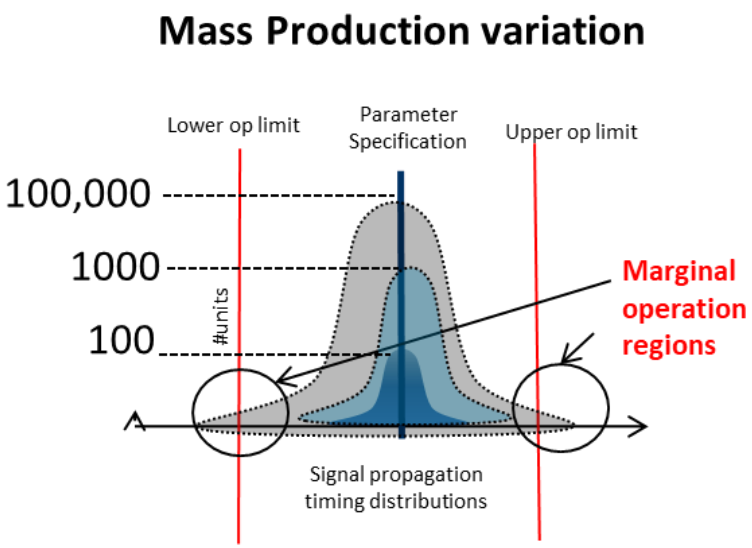

Figure 3. As production volume increases, the tails of the parametric distribution are more likely to approach operational limits. Thermal HALT can help a small sample emulate these worst-case margins. Adapted from the author's signal-integrity paper. [17]

Thermal HALT does not reproduce every manufacturing variation. Instead, it shifts the available samples through a wider parametric range. Hot and cold stress can move propagation delays, thresholds, impedance, noise, and power behavior toward conditions that would otherwise appear only in tail units or adverse field environments. The method therefore increases the probability of revealing an inadequate design margin with limited samples.

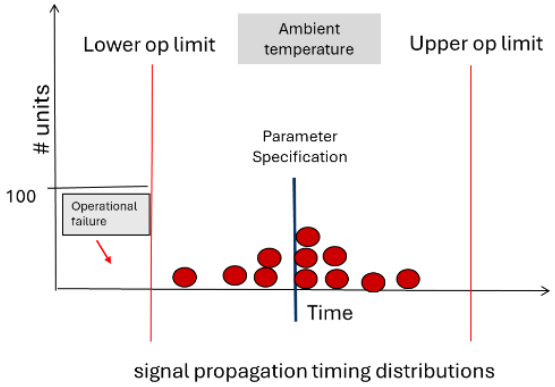

Figure 4a. Parametric timing distribution at ambient temperature. A marginal tail may already approach an operational limit. [17]

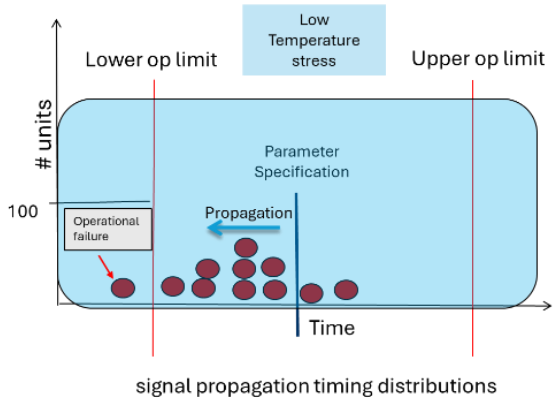

Figure 4b. Low-temperature stress shifts propagation behavior and can expose a different marginal tail. [17]

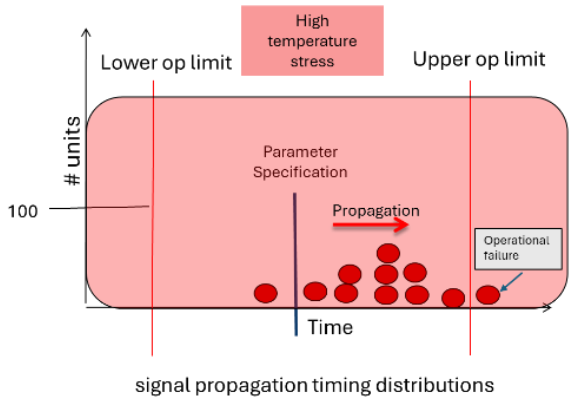

Figure 4c. High-temperature stress shifts propagation behavior toward the opposite operational limit. [17]

8. Rapid Thermal Transitions: More Than Cycle Count

Rapid thermal transitions create temperature gradients because components differ in mass, power dissipation, location, and airflow. A low-mass device can change temperature much faster than a large heat sink, transformer, connector, or board region. During the transition, electrical parameters and mechanical dimensions therefore change at different rates across the system.

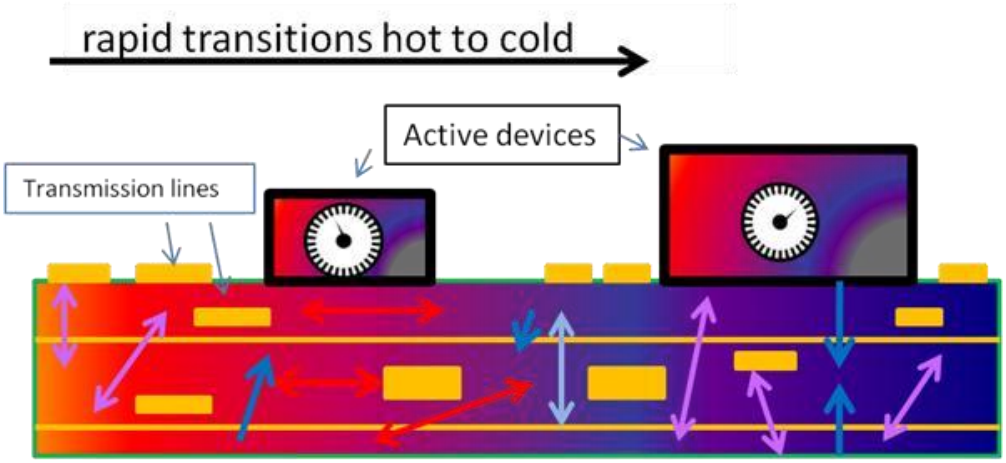

Figure 5. Rapid hot-to-cold transitions create temperature gradients, differential parametric shifts, and thermomechanical stress throughout an operating printed wiring board assembly. [17]

This differential response can expose a failure that is absent at either steady-state extreme. Clock generators, memory, buses, regulators, and reset logic may be temporarily at different temperatures and therefore at different propagation speeds. Simultaneously, board strain can affect solder joints, connectors, and marginal interconnections. Counting cycles alone misses this combined electrical and mechanical discovery mechanism.

Temperature can accelerate a response without accelerating life

A system that fails only during a rapid transition may be revealing a transient timing, gradient, contact, or strain problem. The observation is valuable even when no defensible conversion to equivalent field hours exists.

9. Software-Visible Failures from Hardware Marginality

A marginal digital signal does not announce itself as a hardware defect. It may appear as a corrupted transaction, retry, memory error, boot failure, abnormal LED sequence, processor hang, watchdog reset, or non-deterministic software crash. After rebooting or returning to room temperature, the product may pass every conventional functional test. The result is frequently classified as no fault found. [17], [26]

The no-fault-found cycle is expensive. Complete subsystems may be replaced without isolating the interaction that caused the field symptom. Each removed board can test good by itself because the original failure depended on a particular stack-up of cables, connectors, boards, temperatures, voltages, workloads, and software sequence. The replacement board may simply have more margin.

9.1 Allied Telesis case histories

Johnson and Franks documented HALT examples in which apparent software or system failures were traced to temperature-sensitive timing and sequencing. In one case, random LED activity occurred after hard power cycling near -10°C because the reset pulse generated inside a programmable logic device was too short. Applying local heat helped isolate the device; a code change extended operation to below -50°C. In another case, a system crashed at an upper operating limit near 70°C. A memory-interface register change in boot code extended the upper limit to more than 100°C. [27]

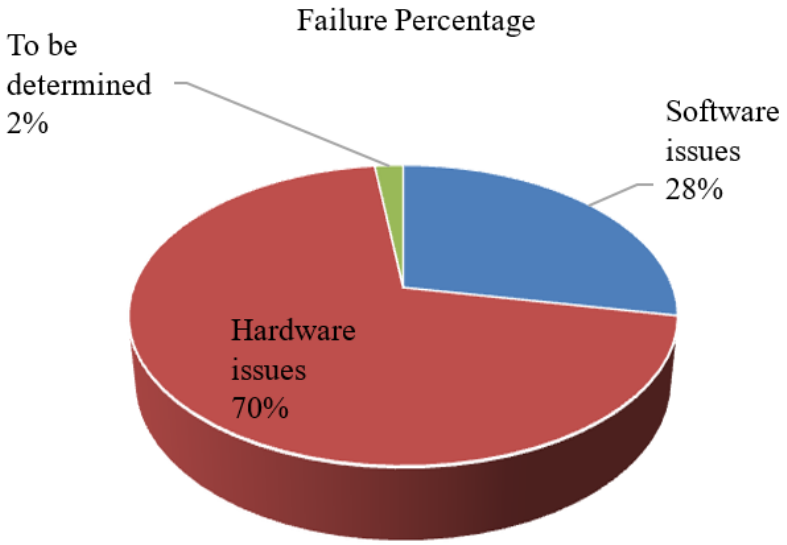

Figure 6. Reported Allied Telesis HALT findings included a significant fraction of software-related issues, illustrating that thermal stress can reveal cross-domain hardware-software marginality. Adapted from [27].

The corrective action in these examples was software or programmable-logic code, but the discovery required a physical stress that shifted hardware timing. This is an important reliability lesson: the category of the corrective action does not necessarily identify the physical origin of the failure symptom.

10. Field Evidence Against a Universal Temperature-Failure Law

Large field studies provide a complementary perspective. Pinheiro, Weber, and Barroso found that disk-drive temperature and activity were much less correlated with failures than commonly believed. Schroeder, Pinheiro, and Weber found a surprisingly small temperature effect on field DRAM error behavior after other variables were considered. A multi-organization data-center study found that many temperature/failure relationships were absent, linear rather than exponential, or more strongly associated with temperature variability than with average temperature. [8]-[10]

These studies do not prove that temperature is unimportant. They show that real products fail through competing mechanisms and that the dominant field population often does not follow the simple exponential relationship assumed by a universal 10°C multiplier. This is consistent with the computer-overstress experiment and with the signal-integrity argument: temperature may have its most immediate effect on operating margin, while long-term field failures are governed by a broader set of causes.

11. Semiconductor Reliability Is Not One Arrhenius Process

Semiconductors are frequently treated as the strongest justification for the 10°C rule, yet modern devices contain many distinct degradation processes. Ultra-thin gate-oxide research has reported non-Arrhenius temperature dependence, voltage-driven defect generation, percolation behavior, and acceleration parameters that vary with field. Bias-temperature instability involves stress, rapid recovery, waveform, measurement delay, charge trapping, and high-k material effects. Electromigration depends strongly on current density and geometry in addition to temperature. [11]-[14]

These examples do not invalidate Arrhenius. They invalidate the assumption that a semiconductor has one intrinsic activation energy governing all failures. A defensible acceleration model begins with the mechanism, not with a preferred multiplier.

Scope limitation

This article does not claim that high junction temperature is harmless or that semiconductor wear-out is temperature independent. It argues that temperature acceleration is mechanism-specific and that universal component or system failure-rate multipliers are unsupported.

12. High Temperature Can Create the Wrong Failure

An accelerated test is useful for life extrapolation only when the same physical failure mechanism occurs at stress and use conditions. Excessive temperature can soften plastics, dry or vent electrolytes, alter lubricants and seals, reflow solder, damage insulation, trigger protective shutdown, or activate reactions that would never be life limiting in normal operation.

The deformed optical-drive mechanisms in the computer study are a direct example. The test found a genuine material limit, but the failure cannot be converted automatically into normal-use life through an Arrhenius factor. Hakim similarly criticized extrapolating plastic-encapsulated microcircuit reliability from high-temperature/high-humidity corrosion conditions when the accelerated mechanism did not represent the intended application. [1], [15]

A high-temperature test may therefore serve three different purposes: discovering a functional limit, precipitating a latent defect, or accelerating a known degradation mechanism. Only the third purpose supports quantitative life extrapolation, and then only after mechanism equivalence is established.

13. A Better Model for Thermal Reliability Development

The traditional model can be summarized as: temperature predicts life. A more productive model is: temperature perturbs the system, observation identifies what the perturbation reveals, and failure analysis determines whether a life model is appropriate.

Table 4. Thermal-test observation and defensible interpretation

Observation	What it may demonstrate	What it does not automatically demonstrate
Immediate lockup at a hot or cold limit	Marginal timing, power, threshold, signal integrity, or software-hardware interaction	Long-term wear-out rate
Failure only during a thermal transition	Gradient, sequencing, interconnect, contact, or transient timing sensitivity	Arrhenius acceleration
Gradual drift during a dwell	Possible thermally activated degradation	Field-life equivalence without mechanism confirmation
Plastic deformation or solder reflow	Material threshold exceeded	Normal-use failure mechanism
Failure after repeated cycles	Possible fatigue accumulation	Constant-temperature acceleration
No semiconductor failure after prolonged overstress	Wide capability under the tested conditions	Proof that semiconductors are never temperature sensitive

14. Recommended Thermal HALT Practice

1. **Define the discovery objective.** Decide whether the test is intended to find operating limits, precipitate latent defects, characterize margin, or accelerate a known degradation mechanism.
2. **Operate the complete system.** Use realistic workloads, interfaces, software states, and power conditions so that cross-domain failures can occur.
3. **Find empirical hot and cold limits.** Increase stress in controlled steps while monitoring functionality, errors, temperatures, voltage rails, clocks, and critical interfaces.
4. **Use rapid transitions deliberately.** Create gradients to expose timing, sequencing, contact, and strain-sensitive failures, but do not treat ramp rate or cycle count as field-life equivalence without a validated model.
5. **Combine relevant margins.** Voltage and clock-frequency margining can distinguish timing and power sensitivity from purely thermal behavior.
6. **Separate recoverable and permanent failures.** Document whether the symptom clears after reboot, local heating or cooling, stress removal, or substitution of a subsystem.
7. **Localize the sensitive element.** Apply controlled local heating or cooling, instrument suspect paths, and reproduce the failure before changing the design.
8. **Perform failure analysis before modeling.** Select Arrhenius, fatigue, electromigration, humidity, wear, or another model only after the mechanism is supported.

9. **Verify corrective action by margin extension.** A repair is convincing when the empirical operating limit moves substantially, and the original failure cannot be reproduced.
10. **Compare with production and field evidence.** Use warranty returns, NFF data, manufacturing variation, and field environments to confirm that the discovered mechanism is relevant.

15. When Arrhenius Is Appropriate

Arrhenius remains appropriate when all of the following are demonstrated: a specific thermally activated failure mechanism has been identified; the same mechanism and failure signature occur at stress and use conditions; data are collected at multiple temperatures; activation energy is measured or justified for the technology and condition; voltage, humidity, current, duty cycle, and other relevant stresses are controlled; no mechanism shift occurs; and extrapolation is limited to a supported range. [4]-[7]

Under those conditions, Arrhenius is not erroneous. It is a mechanism-specific scientific model. The misuse occurs when the model is applied to an unspecified component failure rate, when a generic activation energy is assumed, or when a complex system is assigned an equivalent age without evidence that the life-limiting field mechanism has been reproduced.

16. Limitations and Necessary Cautions

The computer experiment was exploratory and involved a limited number of systems. It cannot provide a statistical estimate of population life. The signal-integrity discussion establishes plausible and measured temperature dependencies, but a thermal HALT failure still requires isolation before it can be attributed to a particular signal path. Large field studies contain confounding variables and do not isolate every intrinsic semiconductor mechanism. The Allied Telesis cases are valuable demonstrations, but manufacturers rarely publish complete HALT data, so the public case-history base remains limited.

Thermal management remains necessary for performance, safety, material limits, leakage control, acoustic and power objectives, and protection of mechanisms known to be temperature sensitive. The argument is not for operating all electronics hotter, although that may be possible without decreasing its usable life before it is technologically obsolete. It is for replacing automatic temperature-failure assumptions with mechanism evidence and empirical margin discovery.

17. Conclusion

Temperature is neither unimportant nor singular in meaning. It can accelerate a defined chemical, diffusion, migration, or degradation process. It can also change the operating state of an electronic system immediately by shifting propagation delay, impedance, leakage, thresholds, power behavior, timing relationships, mechanical strain, and contact conditions.

The universal 10°C rule confuses these effects. It treats a complete product as though it has one activation energy and converts temperature into an assumed rate of consumed life before identifying what fails. Evidence from long-term computer overstressing, large field populations, semiconductor research, signal-integrity studies, and HALT case histories shows that this simplification is not adequate for modern electronic systems.

The better use of temperature is empirical. Hot, cold, and rapid transitions should be used to find operational limits, expose latent defects, reveal marginal signal integrity and hardware-software interactions, and identify the physical mechanism that actually limits reliability. Arrhenius can then be used where the mechanism supports it—not as a substitute for discovery.

Temperature is a stress, not a clock. The engineering question is not how many field hours an oven hour is assumed to represent. The question is what the stressed product revealed, whether the failure mechanism is relevant, and how the design can be made more robust.

Acknowledgment of AI Assistance

This article was developed with assistance from OpenAI's ChatGPT for literature research, organization, drafting, and editorial refinement. The underlying experimental work, technical position, interpretation, and conclusions are those of Kirk A. Gray. All cited sources and technical claims were reviewed by the author, who accepts responsibility for the final article.

References

[1] K. A. Gray and M. Pecht, "Long-Term Thermal Overstressing of Computers," IEEE Design & Test of Computers, vol. 28, no. 6, pp. 58-65, Nov.-Dec. 2011, doi: 10.1109/MDT.2011.127.

[2] E. B. Hakim, "Reliability Prediction: Is Arrhenius Erroneous?" Solid State Technology, vol. 33, no. 8, p. 57, Aug. 1990.

[3] E. B. Hakim, "Microelectronic Reliability/Temperature Independence," Quality and Reliability Engineering International, vol. 7, no. 4, pp. 215-220, 1991, doi: 10.1002/qre.4680070405.

[4] National Institute of Standards and Technology, "Arrhenius," NIST/SEMATECH e-Handbook of Statistical Methods, sec. 8.1.5.1.

[5] National Institute of Standards and Technology, "What Are Some Common Acceleration Models?" NIST/SEMATECH e-Handbook of Statistical Methods, sec. 8.1.5.

[6] National Institute of Standards and Technology, "How Do You Choose an Appropriate Physical Acceleration Model?" NIST/SEMATECH e-Handbook of Statistical Methods, sec. 8.2.4.

[7] P. Lall, "Tutorial: Temperature as an Input to Microelectronics-Reliability Models," IEEE Transactions on Reliability, vol. 45, no. 1, pp. 3-9, Mar. 1996, doi: 10.1109/24.488908.

[8] E. Pinheiro, W.-D. Weber, and L. A. Barroso, "Failure Trends in a Large Disk Drive Population," Proc. 5th USENIX Conference on File and Storage Technologies, pp. 17-29, 2007.

[9] B. Schroeder, E. Pinheiro, and W.-D. Weber, "DRAM Errors in the Wild: A Large-Scale Field Study," ACM SIGMETRICS Performance Evaluation Review, vol. 37, no. 1, pp. 193-204, 2009, doi: 10.1145/2492101.1555372.

[10] N. El-Sayed, I. Stefanovici, G. Amvrosiadis, A. A. Hwang, and B. Schroeder, "Temperature Management in Data Centers: Why Some (Might) Like It Hot," Proc. ACM SIGMETRICS/PERFORMANCE, pp. 163-174, 2012, doi: 10.1145/2254756.2254778.

[11] D. J. DiMaria and J. H. Stathis, "Non-Arrhenius Temperature Dependence of Reliability in Ultrathin Silicon Dioxide Films," Applied Physics Letters, vol. 74, 1999, doi: 10.1063/1.123677.

[12] J. S. Suehle, "Ultrathin Gate Oxide Reliability: Physical Models, Statistics, and Characterization," IEEE Transactions on Electron Devices, vol. 49, no. 6, pp. 958-971, 2002, doi: 10.1109/TED.2002.1003712.

[13] J. P. Campbell, "NBTI: Why Won't This Thing Go Away?" National Institute of Standards and Technology, 2009.

[14] B. Zhu, J. S. Suehle, Y. Chen, and J. B. Bernstein, "Negative Bias Temperature Instability of Deep Sub-Micron p-MOSFETs Under Pulsed Bias Stress," Proc. Integrated Reliability Workshop, 2002.

[15] E. B. Hakim, "Plastic Encapsulated Microcircuit Reliability Prediction: Why?" Microelectronics Reliability, vol. 38, no. 4, pp. 681-683, Apr. 1998, doi: 10.1016/S0026-2714(97)00219-9.

[16] M. J. Cushing, D. E. Mortin, T. J. Stadterman, and A. Malhotra, "Comparison of Electronics-Reliability Assessment Approaches," IEEE Transactions on Reliability, vol. 42, no. 4, pp. 542-545, Dec. 1993, doi: 10.1109/24.273574.

[17] K. A. Gray, "Proposal for Applying Thermal HALT as a Tool to Detect Marginal Signal Integrity and Software Faults Arising in Digital Systems from Hardware Parametric Variability," unpublished manuscript, Accelerated Reliability Solutions, L.L.C., 2026.

[18] S. H. Hall and H. L. Heck, Advanced Signal Integrity for High-Speed Digital Designs. Hoboken, NJ: Wiley, 2011.

[19] S. Natarajan, M. A. Breuer, and S. K. Gupta, "Process Variations and Their Impact on Circuit Operation," Proc. IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems, pp. 73-81, 1998.

[20] S. Pathania, M. Vasa, B. Mutnury, and R. Sharma, "Thermal Impact on High Speed PCB Interconnects," Proc. IEEE Electrical Performance of Electronic Packaging and Systems, pp. 1-3, 2019, doi: 10.1109/EPEPS47316.2019.193226.

[21] K. Son et al., “Thermal and Signal Integrity Co-Design and Verification of Embedded Cooling Structure With Thermal Transmission Line for High Bandwidth Memory Module,” IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 12, no. 9, pp. 1542-1556, Sept. 2022, doi: 10.1109/TCPMT.2022.3201673.

[22] D. Wolpert and P. Ampadu, “Temperature Effects in Semiconductors,” in Managing Temperature Effects in Nanoscale Adaptive Systems. New York: Springer, 2012, pp. 15-33.

[23] A. J. S. Ricketts, “Towards Minimizing the Adverse Effects of Temperature on High Performance Digital Systems,” Ph.D. dissertation, Pennsylvania State University, 2010.

[24] J. Fan et al., “Signal Integrity Design for High-Speed Digital Circuits: Progress and Directions,” IEEE Transactions on Electromagnetic Compatibility, vol. 52, no. 2, pp. 392-400, May 2010.

[25] L. Condra, D. Das, N. Pendse, and M. G. Pecht, “Junction Temperature Considerations in Evaluating Electronic Parts for Use Outside Manufacturers-Specified Temperature Ranges,” IEEE Transactions on Components and Packaging Technologies, vol. 24, no. 4, pp. 721-728, 2001.

[26] H. Qi, S. Ganesan, and M. Pecht, “No-Fault-Found and Intermittent Failures in Electronic Products,” Microelectronics Reliability, vol. 48, no. 5, pp. 663-674, 2008, doi: 10.1016/j.microrel.2008.02.003.

[27] D. Johnson and K. Franks, “Software Fault Isolation Using HALT and HASS,” Allied Telesyn, white paper, 2005.

[28] K. A. Gray and J. J. Paschkewitz, Next Generation HALT and HASS: Robust Design of Electronics and Systems. Hoboken, NJ: Wiley, 2016.

Author Note

Kirk A. Gray is Principal Consultant at Accelerated Reliability Solutions, L.L.C. He has more than 40 years of experience in electronics reliability, environmental stress testing, and reliability development. His work emphasizes empirical discovery of operational limits, HALT and HASS, failure-mechanism identification, and the use of environmental stress to reveal design, manufacturing, signal-integrity, and hardware-software marginalities before field deployment. He is a Senior Life Member of IEEE and coauthor of **Next Generation HALT and HASS: Robust Design of Electronics and Systems.**