Proposal for Applying Thermal HALT as a Tool to Detect Marginal Signal Integrity and Software Faults Arising in Digital Systems from Hardware Parametric Variability

Kirk Gray, *IEEE Senior Life Member*[*]

Abstract

While software failures are often attributed to coding errors or system-level design flaws, an underexamined but critical source lies in physical-layer variability introduced during component manufacturing, board assembly, and system integration. These variations can lead to marginal parametric performance in electronic circuits, which, although within specification, can cause subtle signal-integrity degradation that interacts unpredictably with software execution. This paper explores the hypothesis that such hardware-induced marginality, amplified by complex hardware-software interactions, can manifest as intermittent, non-deterministic software faults. We propose using thermal Highly Accelerated Life Testing (HALT) as a proactive method to skew circuit-level parametrics and expose these hidden faults with limited samples during product development. By subjecting electronic assemblies to empirical thermal operational limits and rapid temperature cycling, HALT rapidly induces timing shifts, signal threshold violations, and noise margin reductions that help uncover weak signal paths and interface instabilities. Case history results demonstrate that thermal HALT increases the probability of triggering failure modes associated with marginal hardware behavior, including those that lead directly to software malfunctions such as execution errors, data corruption, and race conditions. This work highlights the need for cross-domain validation strategies and provides a practical, low-infrastructure method to detect latent hardware-software failure modes that are often invisible to standard test and verification procedures.

Keywords

HALT, Reliability development, Signal Integrity, software reliability

1. Introduction

HALT (Highly Accelerated Life Test) is a reliability development tool that discovers the strength of operating electronics by applying measured and increasing stress to an empirical (observable, not theoretical) stress operational and destructive limits to find and improve the stability and robustness of the product during the design phase. It has been well established that HALT can rapidly precipitate and detect many weaknesses and latent defects in electronic and electromechanical hardware using thermal stepped stress, vibration, voltage margining, and other stresses to operational and destructive limits. A more useful working definition of HALT would be a highly accelerated “limit” test (as that is what it fundamentally discovers) of the empirical stress operational and destruction limits.

Historically, HALT has been used to identify mechanical weaknesses in soldiers, interconnects, and component packaging, thereby increasing product strength and robustness by applying thermal, vibration, voltage, and clock-frequency step-stress to the level of empirical operational limits. Thermal cycling of circuit boards and assemblies causes materials to expand and contract, and the different thermal coefficients of expansion (TCEs) in the material bonds and interfaces result in fatigue damage with each cycle. HASS (Highly Accelerated Stress Screens) are applied

before shipment to precipitate and detect latent defects (hidden defects that will cause failure over time) that may result from manufacturing errors.

When Greg Hobbs, PhD., P.E. created the HALT and HASS methods in the 1980s, digital systems were not as prevalent, and bus speeds were much slower than in today's electronics. By comparison, today's electronics operate at clock and bus speeds up to 1000 times faster than 20 years ago; however, despite these advancements, the materials used in circuit board design have not changed significantly over the same timeframe. FR4 is still being used for PWBA (Printed Wiring Board Assembly) at much higher speeds than were ever expected when it was initially developed. As data bus speeds increase, error-causing effects that were not significant at the time (when bus frequencies were in the megahertz range) have become dominant at the gigahertz frequencies of today's digital systems. Interconnect resistance, capacitance, and inductance are frequency-dependent, and as bus speeds increase and geometries continue to shrink, these variables may prove difficult, if not impossible, to model accurately [1,2].

The materials and methods used to fabricate components, circuit boards, and system assemblies affect the quality and propagation speed of digital signals in a system [3]. Sometimes, this may lead to race conditions in signal transmissions. A race condition (or race hazard) is software behavior in a digital system in which the output depends on the sequence or timing of instructions to logic devices. The term originates from two signals racing each other to influence production first. A race condition becomes a software bug when events do not happen in the order the programmer intended. If the signal quality or propagation is sufficiently skewed by hardware variations, solder defects, or fluctuating temperatures, it may result in timing errors or false binary logic states.

The continual decrease in metallization dimensions and increase in bus frequencies will increase sensitivity to fabrication variations [1]. Studies have demonstrated that crosstalk in a 0.8 µm CMOS device can increase the circuit delay by 100% from the mean due to process variations [2]. "Modern bus designs have become so fast that the designer must calculate the voltage and timing numbers to a resolution as small as a few millivolts and picoseconds. This degree of resolution was unheard of in computer designs just a few years ago" [1]. These new high-frequency bus designs will lead to more SI errors in manufacturing processes when process variations occur across all tiers of assembly. Marginal signal integrity (SI) in a digital system can increase bit error rates (BER), which in turn can cause intermittent software failures or degraded performance due to error-correction processes. The returned hardware may be frequently considered good when tested on the bench due to the sporadic nature of the failure. This increases the cost of sending out new circuit boards and subsystem parts to replace returned parts with no detectable failures.

Higher processor and bus clock frequencies, the higher density of ICs and systems, and their impact on operational reliability pose challenges in quickly finding and correcting marginal SI issues during the development phase before market release. As digital clocks and bus speeds increase and circuit features get smaller, thermal HALT offers potentially significant benefits for discovering marginal SI issues during new product development. Thermal stress skews the electronic parametric operational conditions of components. It enables the components and PWBA to detect and observe a low operational margin arising from future intrinsic parametric variation that may occur during mass manufacturing or over time from aging mechanisms. There is currently a shortage of published data or studies on how the variations in manufacturing materials and methods at all tiers of the circuit board assembly affect SI and software failures. Therefore, thermal HALT is an excellent tool and opportunity for the electronics industry to discover issues with hardware and SI interactions that lead to marginal software reliability earlier during product development by using thermal stepped stress to skew the parametrics of components and PWBA's to empirical operation limits. This paper illustrates how HALT can be used to find empirical high and low-temperature operational limits and potentially discover marginal SI reliability issues. By stressing systems to thermal operational limits, the signal timing and propagation weaknesses are more likely to be found before market release. This suggests thermal HALT is highly effective in discovering operational reliability as digital data transmission speeds increase.

2. Basics of Digital Signal Propagation

It is widely known that the primary function of digital systems is to transmit binary information using electrical signals that represent a 1 or a 0. Ideally, this involves sending and receiving a square electrical wave, with the higher voltage

representing a “1” and the lower voltage representing a “0.” As the speed of signal transmission increases, the ideal square wave becomes trapezoidal with broad boundaries due to signal timing variations, which cause skew and jitter.

Every conductor has a capacitance, inductance, and frequency-dependent resistance. When the frequency is high enough, none of these things is negligible. Thus, a wire is no longer a distributed parasitic element with a delay and a transient impedance profile that can cause distortions and glitches to manifest on the waveform propagating from the driving chip to the receiving chip. The wire is now coupled to everything around it, including power, ground structures, and other traces. The signal is not confined entirely to the conductor but encompasses the local electric and magnetic fields around it. The signals on one interconnect will be affected by the signals on another. Furthermore, complex interactions occur at high frequencies among components within the same interconnect, such as packages, connectors, vias, and bends. All these high-speed effects tend to produce strange, distorted waveforms that give the designer a different view of high-speed logic signals [1].

A digital signal can be observed on a logic analyzer. The signal shown in Figure 1 is referred to as an eye diagram, so named because the spacing between high and low signal levels resembles the shape of a human eye [5].

Eye diagrams are the oldest and most widely used compliance methodology for high-speed serial links. The appearance of the eye diagram allows us to visually assess the performance of a channel and quickly evaluate key parameters of the signal, such as the amount of ISI, noise, and jitter. The more open an eye is, the lower the probability that any bit in the data stream was received in error. An eye diagram plot represents the signal amplitude on the vertical axis and the time on the horizontal axis.

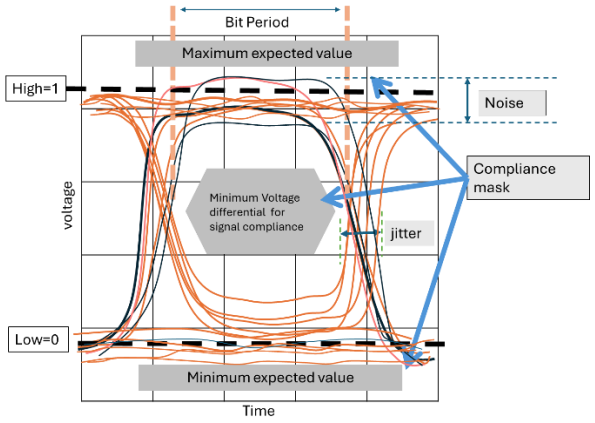

Figure 1. The eye diagram with Compliance Mask showing the minimum voltage differential to transmit a 1 or 0 bit, noise, and jitter. The larger the center eye of the diagram, the less BER [adapted, [5]].

A compliance mask overlay on the signal, also shown in Figure 1, indicates the area that the signal must not cross, or the receiver may misread the one or zero being transmitted. The compliance mask is shown in Figure 1. As the gray areas within this diagram. In the center is the typically diamond- or hexagon-shaped “eye,” hence the name. The eye in the center and the maximum and minimum amplitude values specify the acceptable boundaries for signal transitions. If any part of the signal waveform enters this mask, it indicates a violation of the standard's requirements, suggesting potential issues with timing, voltage levels, or noise margins [2]. Each high-speed digital signaling standard, such as PCIe, USB, and SATA, defines its own compliance mask, tailored to its bit rate, encoding scheme, and channel characteristics. The more open and clean the eye is on the outside of the mask, the better the signal quality and the reduction in Bit Error Rates (BER) . Non-compliant signals result in an increase in BER. BER is the number of bit errors per unit time. High BER can lead to faulty data, processor hangups, reduction in throughput or system crashes [5].

3. Thermal Effects on PWBA's

Thermal stress affects electrical characteristics and the speed of signal propagation in materials used for circuit boards and systems. As the temperature of electrical conductors increases, the thermal vibration of atoms in the material increases, thereby reducing charge-carrier mobility and increasing electrical impedance. Doped semiconductors have a more complex temperature response but generally increase resistance as temperature rises under typical operating conditions [4]. Temperature can be a significant factor in the energy band gap, current density, leakage current, interconnect resistance, and mobility of charge carriers, and is further complicated by the complex interplay of scattering parameters arising from surface roughness, phonons, bulk charge, and Coulombic scattering [4,5,6,16]. In conductors such as aluminum and copper, the wire resistance can change as much as 72% and 77% (respectively) over a military-specified temperature range of -55°C to 125°C [4].

Variations in material quality, semiconductor, and other component fabrication processes can shift the effect of temperature on signal propagation throughout mass-producing circuits and systems during the production or manufacturing life cycle, causing robust SI margins to decline to a level that causes errors in data transmission.

4. Thermal stress and Signal Integrity

4.1 Insertion loss: Insertion loss in digital signaling refers to the attenuation of signal power as it travels through a transmission medium, such as a trace on a printed circuit board (PCB), cable, or connector. A study on the thermal impact of microstrip and stripline traces along with PCB vias found that ~18% of the loss variation can be due to temperature variation
It's a critical parameter in high-speed digital design and signal integrity analysis. It has been shown that there can be an increase from ~ 22 % loss at 20° C to ~ 27 % loss at 100 ° C in stripline traces, depending on the trace and dielectric thickness, along with the spacing between traces [6]

4.2 Conductor Resistance: Stripline and Microstrip conductor resistance increases as the temperature of the components and PWB increases, reducing the signal amplitude and decreasing the eye diagram opening. [6,7].

Parameter	High Temperature Effect	Low Temperature Effect
Vertical Eye Opening	Reduced due to noise and power drop	increases if the noise drops
Horizontal Eye Opening	Narrowed (timing jitter)	Skew-induced asymmetry
Rise/Fall Times	Slower	Faster, but with a potential mismatch leading to a skewed or closed eye opening
Jitter	Increased (thermal noise)	Increased skew and metastability (unstable states of 1 or 0)

Table 1 lists how temperature impacts the eye diagram and the resulting BER

4.2 Degraded Timing Margins: Higher temperatures result in transistors' switch speed slowing down due to the decrease in carrier mobility, causing rise and fall times to slow down. The result is greater timing jitter and narrower eye openings [6]

4.3 Temperature effects on SI and BER Temperatures throughout the digital systems can be significant, and the effects on SI in components, PWBA, transmission lines, vias, and connectors are very complex and can be difficult to model with the potential manufacturing and end-use applications and environmental conditions. A summary of the

impact of thermal stress, both high and low, on the eye diagram that results in increased BER is shown in Table 1 [7,8,9].

5. Soft Failures and the NFF (No Fault Found)

PROBLEM

There are many causes of electronic systems that, when returned to the manufacturer, seem to have no defects or are determined to have No Fault Found (NFF). There are many underlying causes of NFF conditions, including errors in test equipment, misapplication of the product, and a lack of retesting under field-stress conditions. Another contributing factor is the simultaneous replacement of many subsystems, in which only one has failed, with the goal being to return equipment to operation as quickly as possible and not spend time isolating the specific device failure [10]. Another cause of intermittent or marginal operational reliability is poor SI and the errors it creates in binary data. Poor quality of digital signal transmissions or SI can lead to bit errors. Variations in materials, manufacturing processes, and environmental conditions such as temperature, voltage, and humidity affect signal quality, timing, skew and jitter, and electronic noise levels [8,9,10]. Software failures due to these variations are usually "soft failures," where a system can be reset and typically operated. Depending on the frequency of these operational failure events, the user may or may not tolerate their occurrence. When a customer determines that the resettable fault occurs too frequently, it may result in the return of the device or product to the manufacturer.

When a system is returned to the manufacturer due to field failure, it may be disassembled, and each subsystem will be sent for failure analysis. If the system is returned due to reported soft failures due to SI marginality, each subsystem may be tested as functional. The marginal SI may be due to tolerance stack-up of cabling, connectors, or thermo-mechanical stresses on the circuit board that may be included in the failure analysis testing. The returned parts that test suitable may be used for warranty repair and may be less marginal in another system. This can result in the manufacturer never discovering the cause of a high NFF rate in warranty returns for a product, resulting in an expensive cycle of functionally good parts being replaced with other good parts.

Most new electronic systems produced today are digital. In the manufacturing phase of digital electronic systems, many electrical parameters are affected by the accumulation of manufacturing process variations. There are process variations across the many steps of silicon die fabrication that will affect parametric performance from die to die on the same wafer. In deep-submicron processes, variation in transistor threshold voltage can produce over 30% sheet resistance, and variation in polysilicon resistors can reach 40% in some technologies [8,14].

In manufacturing CMOS semiconductors, the primary sources of variation include gate oxide thickness (which may be single-digit atoms thick), random doping fluctuations, device geometry from lithography in the nanometer regime, and transistor threshold voltage. Variations can range from 100% threshold-voltage across a chip, 30% speed variation across a wafer, and 100% leakage-current variation across a wafer manufactured with 130 nm transmission line widths [14]. There is a distribution of values within the minimum and maximum required electrical performance parameters for semiconductor devices because IC fabrication is an imprecise process. For some semiconductor devices, the measured parametric deviations from process variations can be used for product 'binning.' Binning is the selection process for IC die in which IC manufacturers characterize finished products for different markets by measuring thermal and frequency differences and using specific algorithms for an IC's performance. The IC manufacturer does not generally disclose the parametric deviation or margin between the units in each bin category, so that the system manufacturer will not know where the IC is in the distribution of the binning algorithm of selected parametrics.

During PWBA (printed wiring board assembly) manufacturing, variations in the metallization thickness between the lament layer in the substrate and surface roughness of the interconnect, among other variables, will exist. Dimensional variations in PWB can affect impedance crosstalk, noise, and EMI issues in the system. Expansion and contraction of the structures and material interfaces in a PWBA induce thermomechanical fatigue damage during thermal cycling, which has been a primary focus of HALT and HASS methodology. However, the dimensional variations impact SI quality as well.

We know from the SPC (statistical process control) that reducing manufacturing variation is the path to making a defect-free product. Computer simulation models such as SPICE are used to analyze the electrical performance when a complex high-speed digital electronics system is designed. SPICE and other simulation software are helpful tools to verify essential functionality; however, the models are limited in that they cannot include the effect of impact on the operation of potential variations in component manufacturing, circuit board fabrication, solder quality, and second sources of components over the production period, field use, and fatigue damage over time.

Finding marginal SI in a circuit that leads to operational failures during early product development is challenging. Early samples of a new electronics product are typically expensive and are shared by all development teams. Since early prototype samples are built with components from duplicate production lots and use the same production lines, the variation between samples will be slight compared with mass production. Using only a few samples during development makes it difficult to determine whether the parametric variation and SI in the actual manufactured system are marginal. Most companies that have performed thermal HALT to operational limits on digital electronics almost always find/discover an operational limit only, as finding a destruct limit beyond that would come from an irrelevant field failure mode, such as solder being reflowed. Beyond operational levels, it is rare to see a thermal destruct level in digital systems, such as IT (information technology) hardware. The change in material states, such as reflow of solder or melting of wire insulation, beyond operational limits, is usually not a relevant failure in operational field reliability (although it might be in severe non-operational storage or shipping conditions). Hot and cold thermal stress causes signal propagation shifts in conductors and semiconductors, resulting in the "skewing" of signals throughout the system. This is likely why thermal HALT on most digital systems finds an operational limit, not a destructive one. Thermal HALT operation limits in digital systems often come from a failure in SI, and a lock-up or shutdown occurs, but it can easily be reset when the stress is removed.

Process variations exist in all components and PWBs during high-volume manufacture, which producers try to control and reduce with SPC (Statistical Process Control). Marginal operational failures may be observed later in the field from the worst-case stack-up of parametric variations in a smaller percentage of products when thousands or millions are produced. As manufacturing volumes increase from mass production, a wider distribution of parametric variations and end-use stress conditions may cause a percentage of units produced to be near or over the required parametric operational requirements for the SI to be compliant during normal operation at ambient conditions.

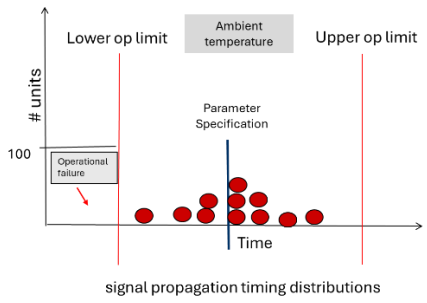

Figure. 2a

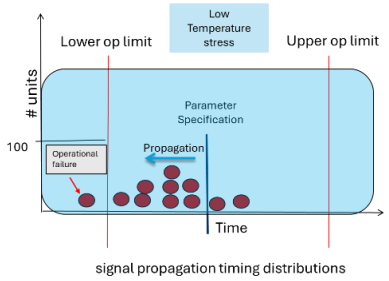

Figure. 2b

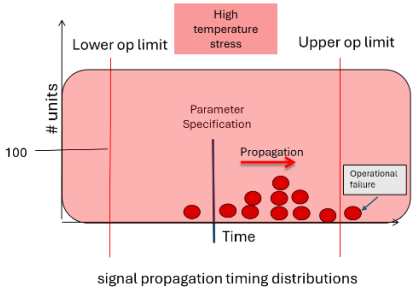

Figure 2c

Fig. 2. Graphical illustration of propagation timing relative to thermal stress. (a) shows the timing deviations from specifications at ambient temperature, (b) Low-temperature HALT increases conductivity and signal propagation, (c) illustrates the decrease in conductivity and signal propagation. (adapted, [10])

Figure 2. Illustrates parametric timing differences between samples resulting from the intrinsic variations in fabrication. Figure 2a shows distributions between samples at ambient operational conditions. Figure 2b. Shows how applying cold thermal stress to empirical operational limits stimulates most marginal samples to slow and fail timing requirements to function. Figure 2c. shows how high temperature stress to empirical operational limits increases timing, resulting in operational failure of the highest speed samples.

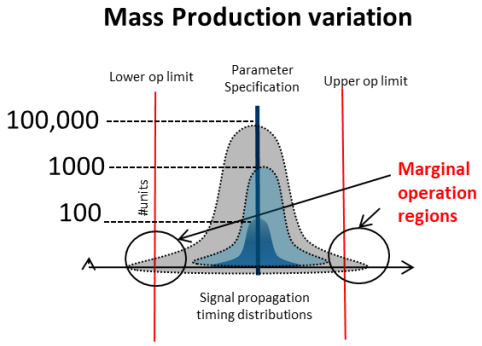

Figure 3. Illustration of the increase in distribution of critical parametric deviations as the production volume increases (adapted [10]).

By applying cold and hot thermal stress to small sample of units in an environmental chamber to empirical operational limits, thermal HALT forces an increase or decrease in noise level, dielectric leakage, and a shift in signal propagation speeds, which results in the samples with the worst case stack up of manufacturing deviations can force operational failures in a short time, and allow determination of low SI margins in PWBA's and systems that result in SI failures.

Figure 3. Is a graphical representation of the potential stack up of parametric deviations that can result in some of the units having marginal operation in the tails of the distributions in larger production volumes

The lot-to-lot and second source of component parametric variation in the larger mass manufacturing population is mixed with high and low-speed distributions. Temperature affects the impedance of electronic conductors and semiconductors, which, in turn, affects the speed of signal propagation. Figure 4 shows an example of the relationship between temperature and signal propagation in a semiconductor. The graph shows the low-to-high propagation delay versus case temperature in Fairchild Octal buffer MM74HC244N (rated for -40 to 85 °C) [12].

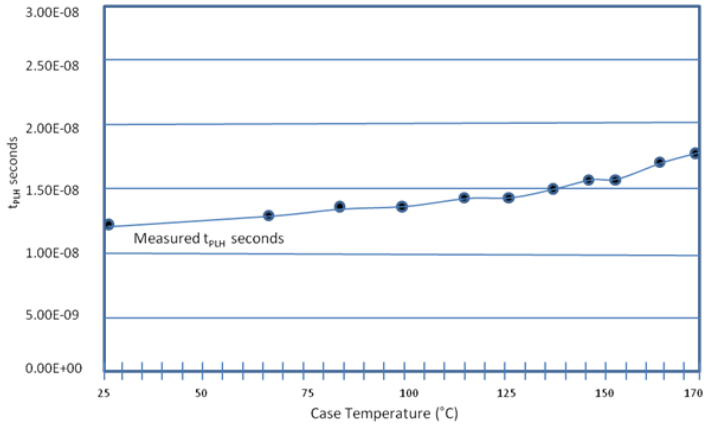

Figure 4. Low-to-high propagation delay versus case temperature in Fairchild Octal buffer (adapted [12])

Another benefit of rapid thermal cycling stress from 40-60 °C/minute using liquid nitrogen cooled HALT chambers with high velocity air circulation is that it helps discover more potential timing variations by creating large temperature gradients across a PWBA and the system. Low-mass components have higher thermal transition rates than larger-mass or high-wattage components, resulting in varying temperatures and thermal gradients across a PWBA. The temperature gradient during thermal transitions skews the impedance and mechanical stress of devices and interconnects at different rates across an operating circuit assembly. In thermal HALT, heating and cooling a single active component effectively isolates the element with a thermal-induced parametric shift that causes an operational limit. A graphical illustration of the synergistic effects in rapid thermal transitions and the creation of thermal gradients, mechanical stresses, and skewing of signal propagation throughout a PWB assembly is shown in Figure 5. The double arrows in the figure represent the thermo-mechanical stress vectors induced by thermal gradients throughout the circuit board assembly.

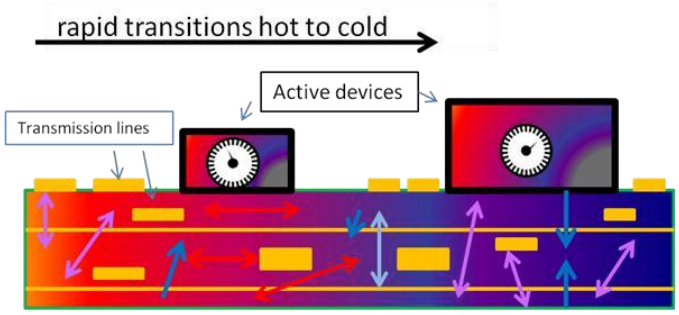

Figure 5. The creation of thermal gradients on a PWBA in rapid HALT thermal cycling that simultaneously skews the components' parametrics differentially during the forced rapid air cooling or heating[10]

6. Case Histories

Manufacturers rarely publish their most effective reliability testing methods to avoid giving competitors new, effective test methods for their products. Examples of the benefits of HALT techniques for finding software issues are rare, but one example case study has been documented by Allied Telesis (formerly Allied Telesyn). Donovan Johnson and Ken Franks of Allied Telesis wrote and published a white paper several years ago on how using HALT has helped them discover software-related reliability issues [15]. Some of the reliability issues Allied Telesis found were:

6.1 Abnormal LED activity

One fault found during cold step stress HALT at minus ten °C was a failure that caused random LED (Light Emitting Diode) activity when the system powered up. This fault occurred only after hard power cycling and not when a soft system reset was done. It was attributed to the reset pulse timing inside the PLD (Programmable Logic Device). They isolated the fault by applying heat to the suspected component. Under a hard power cycle, the PLD reset pulse duration was insufficient. After the code update, the system had no errors at temperatures as low as minus 50°C. [15]

6.2. System Crash

A product that had six months of field use was used in the HALT lab to find its thermal operating margins. The first iteration of HALT found the Upper Operating Limit (UOL) to be 70°C, resulting in a system crash. By changing a register setting for a memory interface inside the boot code, the UOL was extended to over 100°C. [15]. Power Up Sequencing Electronics systems often use onboard logic to control the power sequence of the voltage rails. One product had various components that failed after a power cycle at minus 20°C. By applying heat to the suspected, they could isolate an unreliable PLD component at a lower temperature. Again, the code was modified in the PLD to allow the product to reach temperatures lower than minus 50°C without failures. [15]

In all these cases, the solution was due to a change in program coding, not a hardware component, that significantly extended the thermal operational margins. Figure 6 shows the breakdown of the percentage of software and hardware issues found with HALT at Allied Telesis.

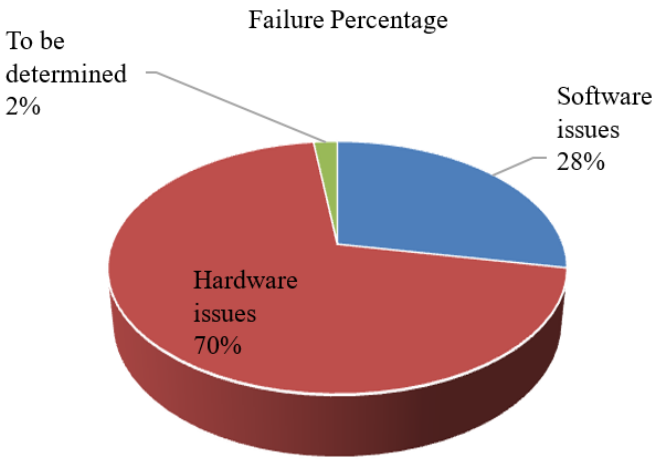

Figure 6. The percentage of hardware and software failures found in HALT at Allied Telesis [adapted [15]].

7. Conclusion

The benefits of HALT in finding latent defects rapidly using thermal, vibration, voltage, and frequency stress to precipitate and detect weaknesses that will result in catastrophic hardware failures in electronics assemblies have been well established over the last several decades. As the speed and density of electronics continue to increase, soft failures due to hardware SI will become more sensitive to parametric variations in the many tiers of the manufacturing processes. The stack-up of parametric deviations during the manufacturing process of components, PWBA, and interconnections can degrade SI margins. Low SI compliance can result in degraded operation or intermittent software failures that are difficult to detect in a limited number of samples during product development. The critical deviations and stack-ups in these systems that lead to poor operational performance or failures may not be detected in initial prototypes during product development. When the deviations become large enough in the mass production and distribution of systems in a variety of use environments, it can impact the reliable operation of some percentage of units. Thermal HALT, along with voltage and frequency margining to empirical operational limits, can be used to detect potential low margins in SI early in product development. Applying thermal stress skews the timing parameters so that the worst-case SI deviations are detected, and low margins can be improved before market release, and helps ensure the reliability of the systems during mass production and throughout their use.

Abbreviations

HALT: Highly Accelerated Life or Limit Test
PWBA: Printed Wiring Board Assembly
TCE: thermal coefficients of expansion
UOL: Upper Operation (thermal) Limit
LOL: Lower Operational (thermal) Limit
LED (Light Emitting Diode)
PLD: Programmable Logic Device
BER: Bit Error Rate
NFF: No Fault Found

Author Contributions

Kirk Gray is the sole author. The author read and approved the final manuscript.

Conflicts of Interest

The author declares no conflicts of interest

References

[1] Hall, Stephen H., and Howard L. Heck. *Advanced signal integrity for high-speed digital designs*. John Wiley & Sons, 2011.

[2] Natarajan, Suriyaprakash, Melvin A. Breuer, and Sandeep K. Gupta. "Process variations and their impact on circuit operation." In *Defect and Fault Tolerance in VLSI Systems, 1998. Proceedings., 1998 IEEE International Symposium on*, pp. 73-81. IEEE, 1998.

[3] Jim Lai and Tristan Lin, "System Operating Environment Effect on PCB Material Electrical Property", 2016 Asia-Pacific International Symposium on Electromagnetic Compatibility (APEMC), 2016, pp. 314- 316: DOI: 10.1109/EPEPS51341.2021.9609215

[4] Wolpert, David, and Paul Ampadu. "Temperature effects in semiconductors." In Managing Temperature Effects in Nanoscale Adaptive Systems, pp. 15-33. Springer New York, 2012.

[5] ONSEMI, "ON Semiconductor," 2016. [Online]. Available: http://www.onsemi.com/pub_link/Collateral/AND9075-D.PDF. [Accessed 8 7 2016].

[6] S. Pathania, M. Vasa, B. Mutnury and R. Sharma, "Thermal Impact on High Speed PCB Interconnects," 2019 IEEE 28th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS), Montreal, QC, Canada, 2019, pp. 1-3, DOI: 10.1109/EPEPS47316.2019.193226.

[7] 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," in *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 12, no. 9, pp. 1542-1556, Sept. 2022, doi: 10.1109/TCPMT.2022.3201673.

[8] A. J. S. Ricketts, Towards minimizing the adverse effects of temperature on high performance digital systems, Ph.D. dissertation, Dept. of Computer Science and Engineering, The Pennsylvania State University, University Park, PA, 2010.

[9] K. A. Gray and J. J. Paschkewitz, *Next Generation HALT and HASS: Robust Design of Electronics and Systems*, John Wiley and Sons, 2016, pp. 190-205.

[10] Qi, Haiyu, Sanka Ganesan, and Michael Pecht. "No-fault-found and intermittent failures in electronic products." *Microelectronics Reliability* 48, no. 5 (2008): 663-674. https://doi.org/10.1016/j.microrel.2008.02.003

[11] Kim, MS., Yoo, SK., Kang, J. et al. Impact of temperature fluctuations on SRAM- based physically unclonable functions. Int. J. Inf. Secur. 24, 113 (2025). https://doi.org/10.1007/s10207-025-01031-1

[12] Condra, Lloyd, Diganta Das, Neeraj Pendse, and Michael G. Pecht. "Junction temperature considerations in evaluating electronic parts for use outside manufacturers-specified temperature ranges." IEEE Transactions on Components and Packaging Technologies 24, no. 4 (2001): 721-728.

[13] Melikyan, V., A. Durgaryan, Abraham H. Balabanyan, Eduard H. Babayan, Milena Stanojlović, and Ashot G. Harutyunyan. "Process-voltage-temperature variation detection and cancelation using on-chip phase-locked loop." In *56th Conference for electronics, telecommunications, computers, automation, and nuclear engineering–ETRAN, Zlatibor*. 2012."

[14] Patel, Janak. "CMOS process variations: A critical operation point hypothesis." In an *Online Presentation*. 2008.

[15] Johnson, Donovan and Ken Franks. “Software fault isolation using HALT and HASS.” Allied Telesyn. 2005.

[16] 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, article no. 5440927, Institute of Electrical and Electronics Engineers, May 2010.

Biography

Kirk Gray has over 45 years of experience in the electronics manufacturing industry. He has been consulting, teaching, and applying HALT and HASS methodology to various design and manufacturing companies since 1992. He co-authored the book *Next Generation HALT and HASS: Robust Design of Electronics and Systems*, published by Wiley & Sons in 2016. He holds a BSEE from the University of Texas at Austin, is a Senior Life Member of the IEEE, and a Senior Collaborator with the CALCE Consortium at the University of Maryland. He is the Principal Consultant at Accelerated Reliability Solutions, L.L.C., www.acceleratedreliabilitysolutions.com.