Published in: IEEE Design & Test of Computers, 2011

Long-term Thermal Overstressing of Computers

Kirk A. Gray, *Member, IEEE,* and Michael Pecht, *Fellow, IEEE*

***Abstract*— Significant opportunities exist to reduce costs in the design, manufacture and operation of systems by allowing higher use specification temperatures. This paper shares the findings and observations of an experimental study where operating computers were subjected to high steady state temperatures and thermal cycling well beyond their design specifications. The results suggest that significant cost savings can be realized without compromising reliability by abandoning the fundamentally flawed yet long held perception and belief that reliability is mainly temperature driven. This result could enable significant cost savings obtained by extending operating and environmental conditions for both testing and end-use.**

Keywords: **long-term thermal overstressing; computer; reliability; high temperature; cost saving**

I. Introduction

SINCE the early days of solid state electronics, reliability engineers have been taught that the dominant cause of hardware unreliability comes from component failures and that reliability of components can be as much as doubled for each 10ºC reduction in temperature. This belief was a fundamental tenet of the U.S. Military Handbook 217 (MIL-HDBK 217)[1], the first document on reliability prediction of electronic components [1]. While there was no empirical data to support this belief, the concept has been continued and made its way into other reliability prediction handbooks, such as Telcordia SR-332 (formerly Bellcore), PRISM, FIDES, CNET/RDF (European), and the Chinese GJB-299. Furthermore, these prediction methods relied on the analysis of insufficient failure data collected from the field and it was assumed that the components of a system have inherent constant failure rates that are derived from the collected data. It was further assumed that such constant failure rates could be tailored by independent "modifiers" to account for variations in manufacturing quality, operating conditions, and temperature.

In the 1990s, with a host of studies conducted by the National Institute of Standards and Technology (NIST) [2], the U.S. Army [3], AT&T [4], and others [5], it became clear that the approach propagated by these handbooks has been damaging to the industry and a change was needed. The consensus is now that these methods and this type of approach should never be used because they are inaccurate for predicting actual field failures and they provide highly misleading predictions, which can result in poor designs and poor logistics decisions [6]. Although most of these handbooks have been discontinued, and are no longer used by the US military, a few manufacturers of electronic components, printed wiring and circuit boards, and electronic equipment and systems even today still subscribe to the traditional reliability prediction techniques (e.g. MIL HDBK-217 and progeny) in some manner; although sometimes unknowingly.

Because of the use of predictions based on MIL HDBK 217, cooling solutions to reduce the temperature of components were thought to be the most effective way to increase reliability [7][8]. Component manufacturer's data sheets provide absolute maximum ratings (AMR) and recommended operating conditions (ROC). AMRs are provided as the limit for which the part can be operated reliably even though it may not meet electrical function specifications. For example, Motorola states that between recommended and the AMR operation limits, the part may not meet electrical specifications, but that physical failure or adverse effects on reliability are not expected. Yet, system designers are still adding thermal-mechanical complexity, and possibly reducing reliability of a system by attempting to cool below the AMR threshold level of the components in the continuing belief that it will improve reliability. In fact a study by Boeing showed that the use of cooling solutions to drive down system temperatures adds cost and complexity with little hardware reliability benefit [9]. A more recent study by Google [10] on a large population of hard disk drives also showed the failure rate of systems at higher temperatures was acceptable and much lower than expected.

II. The Steady State Elevated Thermal Experiment

Most of the historical basis of traditional reliability theory has been based on the fundamental perception that electronic component failure rates are driven by time at temperature. In this portion of the study, a long-duration experiment was performed with computers to observe the effect of long-term steady state high temperature on hardware reliability. The objective of the experiment was to observe the general effects of running common computer systems at the maximum

Manuscript received October 18, 2010.

K. Gray is with Accelerated Reliability Solutions, L.L.C, 21403 Greenridge Dr, Hutto, TX , US , 78634. Phone: 512-554-3111; Email: kirk@acceleratedreliabilitysolutions.com

M. Pecht is with the Center for Advanced Life Cycle Engineering, University of Maryland, College Park, MD 20742. Phone: (301) 405-5278; e-mail: pecht@calce.umd.edu

*Corresponding author: pecht@calce.umd.edu, phone: 301-405-5323, fax: 301-314-9269.

[1] The last version of Mil-Hdbk 217 was revision "F", in 1995. Since then the document has been cancelled and not updated. Nevertheless, although the predictions are inaccurate and misleading, the document continues to play an influential role in the field of reliability engineering and a few companies still use the results (2).

operational temperature, including the electro-mechanical systems of fans and the hard disk drives (HDD). The goal was to determine a baseline of operational time at elevated temperature of a computer system and observe if there was a common failure mechanism, component or subsystem. Failure was defined as the point at which the system could not power and operate at ambient conditions, what could be called a "hard" failure. Operational failures that were recoverable upon a reboot were documented but not considered failures for this experiment

In May 2008, five state of the art, high end computers with high levels of quality control, were placed into environmental test chambers with an ambient temperature setting of 65-70ºC. An additional computer was also operated in the same conditions from January 2009 until January 2010. The manufacturers end use specifications included: operating temperature of 10C to 35C, and storage temperature of -40C to +65C. Each unit had previously undergone standard reliability development tests (approximately 200 hours total test time) and were considered 'good' samples. Thermocouples were attached to record the actual motherboard and CPU chipset ("Northbridge" and "Southbridge") temperatures and the CPU temperature was the internal temperature diode temperature read via test software (see Table 1).

Two chambers were used for the six systems. Initially, testing was to be conducted at 70ºC but one system was unable to operate without an occasional thermal shutdown of the CPU at 70ºC. That is, the upper maximum operational temperature at the chassis assembly level was limited by the CPU over temperature protection (OTP) circuit set by the CPU manufacturer. As a result, the units were operated at 65ºC environmental conditions. All of the units under test (UUTs) were running applications to functionally stress their hardware using the commercially available stress software (PASSMARK Burn-In) and the graphics intensive gaming demo (Quake Game). The two software applications were run at different times on all the systems. Neither application made an observable difference on the rate of operational reliability.

TABLE 1: MEASURED COMPONENT TEMPERATURES

Component	Temperature (ºC)
Unit #1 Desktop – Northbridge	94
Unit# 1 Desktop – Southbridge	124
Unit# 1 Desktop Motherboard	81
Unit# 3 Desktop Motherboard	76
Unit #2 Desktop – Northbridge	96
Unit #2 Desktop – Southbridge	114
Unit #1 Desktop CPU internal temperature	94
Unit #2 Desktop CPU internal temperature	90
Unit #3 Desktop CPU internal temperature	94
Unit #4 Desktop CPU internal temperature	97

Most of the systems under test did at some time have recoverable intermittent operational failures or "soft failures" occasionally under the stress conditions. During the experiment, these operational failures were documented and the systems rebooted to resume operation. Although this experiment focused on catastrophic hardware failure, thermal stress can "skew" electrical signal propagation and can induce low margin conditions leading to intermittent failures [11].

All systems were tested continuously with different software test configurations changed on a irregular schedule, to observe if any application could rapidly induced a "hard failure. At least one system was allowed to complete over 200,000 reboots during the test period before changing test applications, and most systems ran over 50,000 reboots during the test period. All units were running applications to verify operation during the test period, including an internal diagnostics program, a graphics intensive PC game in a looping demo mode, and a commercially available comprehensive PC stress application.

Non-recoverable failures occurred from 2,300 hours to 8,000 hours. Several of the failing systems were repaired in order to observe secondary failure mechanisms. All the systems had Optical Disk Drive (ODD) failures. Commands to open and close the tray, and the front operation lights were operational but the load trays and mechanical transport of the ODD's were damaged due to plastic deformation from high temperature. Once this occurred, no further analysis on the failures was performed. Soft failure errors rates were negligible except in the one thermal cycling unit having POST failure issues. This exceptional case was not considered further in this experiment. The results of the steady state high temperature test and results of FA (Failure Analysis) are given in Table 2.

TABLE 2: FAILURES OF THE DESKTOP UNITS IN 65ºC AMBIENT CONDITIONS

Unit	1st Hardware Failure	Stress Operational Hours at 65ºC Before Failure	Failure Analysis
UUT#1	HDD	First: 2,600 hours Second: 7,600 hours	First: No FA Second: PSU – no FA
UUT#5	Graphics output	First: 4,400 hours Second: 12,700 hours	First: 4 Aluminum Electrolytic capacitors failure on 1.25 volt circuit Second: PSU Fan
UUT#2	PSU	7,300 hours	PSU – destroyed during FA
UUT#3	PSU	7,300 hours	PSU fan failed
UUT#6	Front fan	8,000 hours	Front fan failure

III. THE THERMAL CYCLING EXPERIMENT

For the thermal cycling experiment four samples of the same model of PCs, also previously used for testing during product development, were placed in one environmental chamber and connected to one KVM (keyboard, video, and mouse) switch to monitor the units with one keyboard, video display and mouse. The chamber was programmed to ramp at a rate of 1ºC per minute to +65ºC, dwell for 30 minutes, then ramp down at a rate of 1ºC per minute and dwell for 30 minutes before beginning the next cycle. The 30 minute dwell times were set based on the time observed for the motherboards to reach thermal equilibrium. The test ran 24 hours a day 7 days a week, with occasional interruptions for chamber maintenance or long holiday periods. The units were checked for operation daily except on weekends and holiday periods. Thermocouples were placed in the center of the motherboards to record the unit's response during cycling. A photo of the thermocouple location is shown in Fig. 1. A total of 2,750 thermal cycles from -10ºC to +65ºC were completed on 3 of the 4 units. All the units also

subsequently passed 168 hours of operation running the commercial test application without failure.

There were no confirmed hardware failures for the computers in the thermal cycling environment although fatigue damage was observed in material cross-sectioning performed post test. One system out of the four was removed after approximately 1,900 cycles due to frequent lock-ups at the initial Power On Self Test (POST) during rebooting cycles. At that time, the unit was removed from the chamber and operated at room conditions for 48 hours and then placed back in the chamber and run under a steady state 65°C for an additional 48 hours rebooting under the same test script without failure. The lock-up condition only occurred during thermal cycling and could not be replicated in steady thermal conditions.

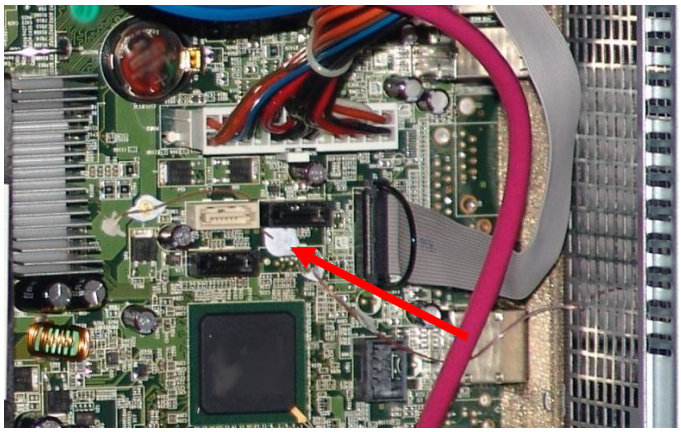

Fig.1: Location of thermal measurement on motherboard.

Operational failures such as lock ups and memory access errors, occurred throughout the experiment, but were recoverable with a reboot. It is likely most end-users would reboot after such operational failures, or that the operating system would do it automatically.

The first suspected hardware failure was a 425W power supply unit (PSU). The unit was unable to power up. A Built-in-Self-Test (BIST) function on the PSU indicated a failure. The PSU was replaced and the unit was then placed back in the thermal cycling chamber. The manufacturer of the PSU took the unit for FA and reported that they were unable to identify a problem because the unit passed all specifications.

After completion of the 2,700 thermal cycles the experiment was ended due to the extensive stress beyond any anticipated end-use stress conditions, and limited resources to support the experiment. All four of the systems, including the one unit that had POST failures in thermal transitional conditions during thermal cycling, were run at ambient conditions for 168 hours with the same general stress application as during the stress without failure. One system was subjected to destructive analysis for circuit card and solder joint degradation. An identical system from the same vintage that had no thermal cycling was used for a comparison reference. The time zero samples and the post 2,700 cycle's samples were both observed to have solder anomalies: lack of solder fill in plated thru-holes (PTH) and thermal header solder attachment voiding in power FET's (see Fig. 2 and 3). The anomalies in these cases were not observed to cause failures.

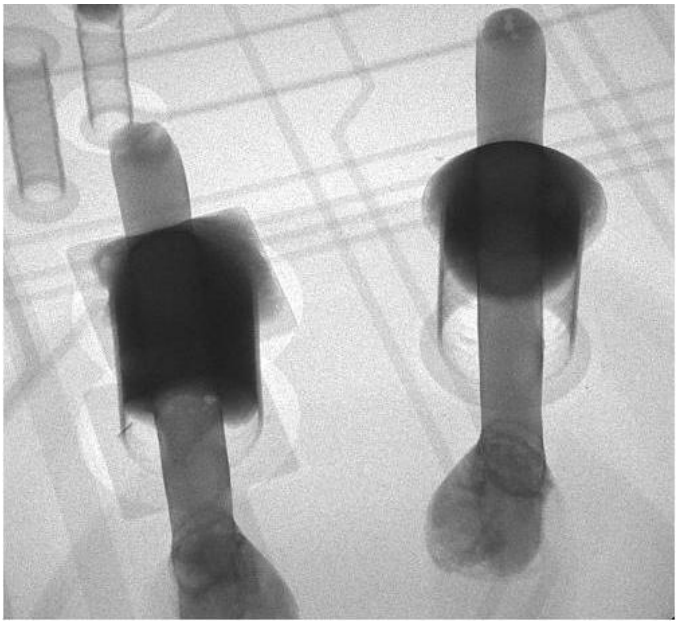

Fig. 2: Anomalies in E Caps PTH for the Thermal Cycle Units.

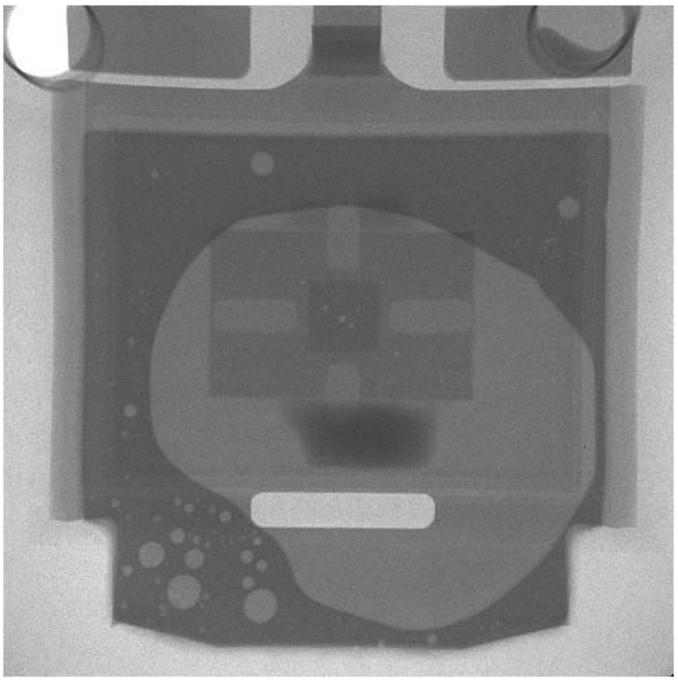

Fig. 3: Power FET Thermal Header Solder Void on Thermally Cycled Units.

Fig. 4: 100% Fracture of BGA Solder Bump on Outer Row.

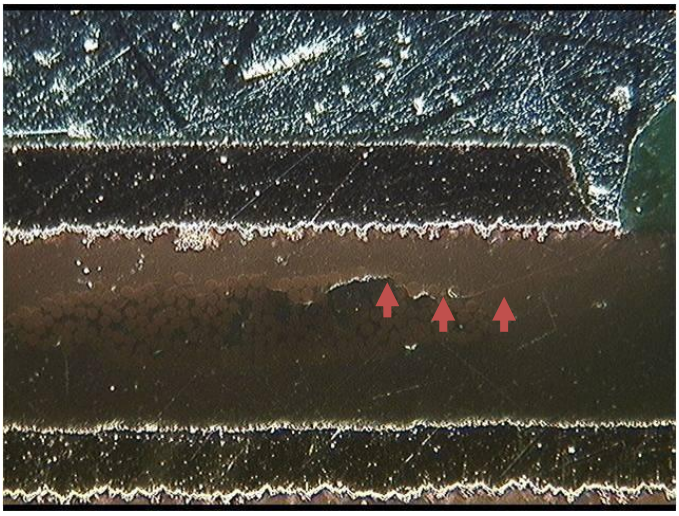

Fig. 5: BGA Pad Cratering and Laminate Separation.

IV. Summary and Conclusions

This study investigated computer hardware reliability under two thermal stress conditions, steady state and transitional thermal environments. Although the sample size and software applications used in this experiment were limited, the results strongly suggest that operating in a higher temperature condition will not shorten the useful life of these systems. While thermal management is necessary for today's computer systems, the set points must not be based on the assumption of the Arrhenius relation. In fact, when this assumption is relaxed over a 30C range from the highest specification value, there was no decrease in reliability. By reducing the stringent component temperature requirements for reliability, designers and manufactures have several potential cost savings opportunities and it may be possible to reduce power consumption from cooling, as well the noise from mechanical cooling systems such as fans. For example, if we assume that the 2,000 hours of operation is equivalent to a standard 40 hour work-week for 50 weeks (one year operation), then the computers tested in the steady state condition operated up to 4 years at 30°C above the manufacturers end-use specifications (excluding the ODD) before catastrophic hardware failure occurred. The systems in thermal cycling survived thermo mechanical stress conditions and fatigue damage well beyond any desktop PC could experience outside of experimentation. For certain fatigue damage did occur in the systems as seen in the detailed post-test material analysis shows, but did not result in significant operational failures. In all of the thermal test conditions, no microelectronic components failed, and thermal transition and steady state high temperature experiments, well outside the system manufacturer's temperature specifications, were needed to observe the first confirmed hardware failure. The results show that our conclusions are valid for high performance systems in the less than 10 gigahertz range.

The electronics industry is still paying the costs of misdirection from MIL-HDBK 217 in thermal solution requirements and in the efficiency and effectiveness reliability development testing. By abandoning the fundamentally flawed yet long held perception and belief that reliability is mainly driven by temperature below the absolute maximum rating, designers and manufactures have several potential cost savings opportunities and it may be possible to reduce power consumption from cooling, as well the noise from mechanical cooling systems such as fans.

The fast pace of today's product development schedules demand shorter and more effective tests for reliability. Most systems when tested to operational limits will have wider inherent empirical thermal capability than the designers expect. In thermal testing, the higher the stress during test, the faster a latent defect or marginal operation can be discovered. Thermal transitions increase the probability of rapidly precipitating many manufacturing latent defects in materials. The greater the range of operational temperatures the greater the thermomechanical stresses on a system. Using the systems full inherent operational range provides a benefit in that the fewest thermal cycles are needed to find defects. The thousands of thermal cycles performed in this experiment demonstrates that using ten to a hundred thermal cycles to detect latent defects during the design or manufacturing phase carries little risk of decreasing the inherent life entitlement of a defect-free system, so large numbers can be safely and rapidly thermally stress tested and shipped to a customer as new. The most effective reliability development processes will be based on capitalizing on the full empirical thermal operational thermal range and long stress-life entitlement inherent in today's electronics systems and materials.

References

[1] U.S. Department of Defense. (1965) MIL-HDBK 217: *Military Handbook for Reliability Prediction of Electronic Equipment, Version A.*

[2] J. Kopanski, D.L. Blackburn, G.G. Harman, D.W. Berning, "Assessment of reliability concerns for wide-temperature operation of semiconductor device circuits," Transactions of the First International High Temperature Electronics Conference, Albuquerque, NM, pp. 137-142,1991.

[3] M.J. Cushing. "U.S. army reliability standardization improvement policy and its impact," IEEE Transactions on Components, Packaging, and Manufacturing Technology, Vol 19, pp.277-278, 1996.

[4] M. G. Pecht and F. R. Nash, "Predicting the reliability of electronic equipment," Proceedings of the IEEE, Vol. 82, pp. 992–1004, July 1994.
[5] M. G. Pecht, "Issues affecting early affordable access to leading electronics technologies by the us military and government," Circuit World, Vol. 22, No. 2, pp. 7-15, 1996.
[6] MJ. Cushing, DE. Mortin, TJ. Stadterman, and A. Malhotra, "Comparison of electronics-reliability assessment approaches," IEEE Transactions on Reliability, Vol. 42, No. 4, pp. 542-546, 1993.
[7] S. P. Park, K. Kang, and K. Roy, "Reliability implications of bias-temperature instability in digital ICs," IEEE Design and Test of Computers, Vol. 26, No. 6, pp. 8-17, 2009.
[8] J. M. Carulli Jr., and T.J. Anderson, "The impact of multiple failure modes on estimating product field reliability," IEEE Design and Test of Computers, Vol. 23, No. 2, pp. 118-126, 2006.
[9] C.T. Leonard, "Mechanical engineering issues and electronic equipment reliability: incurred costs without compensating beneifts," IEEE Transactions on Components, Hybrids, and Manufacuturing Technology,Vol 13. No. 4, pp. 895-902,1990.
[10] E. Pinheiro, W.Weber and L.A. Barroso, "Failure trends in a large disk drive population," Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST'07), February 2007.
[11] H. Qi, S. Ganesan, and M. Pecht, "No-fault-found and intermittent failures in electronic products," Microelectronics Reliability, Vol. 48, Issue 5, pp. 663-674, May 2008.

Kirk A. Gray is currently Principal Consultant at Accelerated Reliability Solutions, L.L.C. He began his career in semiconductor manufacturing. 18 years ago he began consulting on HALT and HASS. He previously worked as a Sr. Reliability Engineer for Dell Inc. His interests are in physics of failure and accelerated stress testing for reliability development based on empirical operational environmental stress limits. He earned a BSEE from the University of Texas at Austin, 1982. He is an IEEE member and member of the IEEE/CPMT Technical Committee on ASTR.

Michael Pecht (M'83–SM'90–F'92) received the M.S. degree in electrical engineering and the M.S. and Ph.D. degrees in engineering mechanics from the University of Wisconsin–Madison. He is the founder of the Center for Advanced Life Cycle Engineering (CALCE), University of Maryland, College Park, which is funded by over 150 of the world's leading electronics companies. He is also a Chair Professor of Mechanical Engineering and a Professor of Applied Mathematics at the University of Maryland. He has written more than 20 books on electronic products development, use and supply chain management and over 400 technical articles. He has consulted for over 100 major international electronics companies, providing expertise in strategic planning, design, test, prognostics, IP and risk assessment of electronic products and systems. Prof. Pecht is a Professional Engineer, an ASME Fellow, an IMAPS Fellow, and an SAE Fellow. He was awarded the highest reliability honor, the IEEE Reliability Society's Lifetime Achievement Award in 2008. He has previously received the European Micro and Nano-Reliability Award for outstanding contributions to reliability research, the 3M Research Award for electronics packaging, and the IMAPS William D. Ashman Memorial Achievement Award for his contributions in electronics reliability analysis. He served as Chief Editor of the IEEE TRANSACTIONS ON RELIABILITY for eight years and on the Advisory Board of IEEE Spectrum. He is Chief Editor for Microelectronics Reliability and an Associate Editor for the IEEE TRANSACTIONS ON COMPONENTS AND PACKAGING TECHNOLOGY.