AEROSPACE REPORT NO.
TOR-2013(1459)-2

Lifetest of Actel AX2000 Field Programmable Gate Array

April 5, 2013

Susan H. Crain,[1] John R. Scarpulla,[2] and Heidi C. Graziano[3]
[1]Space Science Applications Laboratory
Physical Sciences Laboratories
[2]Reconnaissance Systems
Imagery Programs Division
[3]Systems Engineering
Navigation Division

Prepared for:

Space and Missile Systems Center
Air Force Space Command
483 N. Aviation Blvd.
El Segundo, CA 90245-2808

Contract No. FA8802-09-C-0001

Authorized by: Space Systems Group

AEROSPACE REPORT NO.
TOR-2013(1459)-2

Lifetest of Actel AX2000 Field Programmable Gate Array

April 5, 2013

Susan H. Crain,[1] John R. Scarpulla,[2] and Heidi C. Graziano[3]
[1]Space Science Applications Laboratory
Physical Sciences Laboratories
[2]Reconnaissance Systems
Imagery Programs Division
[3]Systems Engineering
Navigation Division

Prepared for:

Space and Missile Systems Center
Air Force Space Command
483 N. Aviation Blvd.
El Segundo, CA 90245-2808

Contract No. FA8802-09-C-0001

Authorized by: Space Systems Group

Lifetest of Actel AX2000 Field Programmable Gate Array

Approved by:

Dan J. Mabry, Director
Space Instrumentation Department
Space Science Applications Laboratory
Physical Sciences Laboratories

Thomas D. Powell, Systems Director
GNSS Engineering and Technology
Navigation Division
Space Program Operations

Gerald J. Trombley, Systems Director
Reconnaissance Systems
Imagery Programs Division
National Systems Group

Michael J. Baxter, Systems Director
Systems Engineering/Effectiveness
Cross Program Engineering Operatons
MILSATCOM Division
Space Systems Group

Abstract

The radiation-tolerant Axcelerator (RTAX) family of field-programmable gate arrays (FPGAs) manufactured by Microsemi offers high performance in a non-volatile antifuse technology that could be advantageous for space applications. However, the antifuse is not a standard element in normal CMOS integrated circuits, and there is no commonly accepted understanding of their reliability. The Aerospace Corporation designed and executed a life test to determine a failure rate for the antifuse portion of the FPGA by accumulating a large number of device-hours on a large quantity of parts operating in typical spaceflight temperature environments.

Nearly 26.5 million device-hours were collected on 760 AX2000 FPGAs (a commercial-grade Axcelerator part with 2 million gates) divided into three temperature environments: hot (held for the duration of the test at 85°C), cold (held for the duration of the test at −40°C), and cycling (cycled between +80°C and −40°C and back about every 5 hours). No antifuse failures occurred. The data is presented here along with some discussion on applying the data to determine a failures in time (FIT) rate for the RTAX FPGA.

Acknowledgments

The authors want to thank Steve Morgan, Charles Rowe, Brent Morgan, Valerie Lang, Lawrence Greenberg, Joseph Brady, Patrick Riley, Doug Daughaday, and Allyson Yarbrough for their support of this work. The technical oversight and reviews done by Lawrence Harzstark and Steven Moss were very valuable as well. The authors also gratefully acknowledge Dan Mabry for the transfer of knowledge for the 54SXA testing and initial test design, Tom Tsubota and Hung Ngo for their work on the SRAM test, and the help of Dhruv Patel and Tom Gibson in maintaining the environment chambers throughout this test.

Contents

Figures

Tables

1. Introduction

The Actel (now Microsemi) AX2000-FG896I field-programmable gate array (FPGA) is a commercial-grade version of the high-performance device using nonvolatile antifuse technology. This particular FPGA family provides the equivalent of 2 million system gates along with 288 kb RAM. The radiation-tolerant version, the RTAX2000S, employs the same antifuse technology, making it very desirable for space applications. In order to verify the reliability of the antifuses in these devices, a lifetest was undertaken. The lifetest of this FPGA was designed to utilize the same method of operation used by The Aerospace Corporation (hereafter Aerospace) to test Actel's A54SX-A FPGAs. The lifetest consisted of three components: fixed temperature (hot), fixed temperature (cold), and temperature cycling, with a separate set of parts dedicated to each lifetest component.

Because antifuses are not a standard element in normal CMOS integrated circuits, there is no commonly accepted understanding of their reliability. Furthermore, there is a question as to whether low-probability defective antifuses might dominate the reliability, versus whether an antifuse wearout mechanism might exist. Still further, there are no test structures that have been made available to answer these questions more directly with accelerated testing. Even if there were, the acceleration models are subject to question because the physical mechanisms of antifuse failures are not well understood. For these reasons, a "brute force" large-quantity approach to performing reliability tests was warranted: logical test structures were programmed onto FPGAs ,which were then subjected to lifetest. The advantage of this approach is that no specific acceleration model need be applied, and there are no assumptions needed to quote the reliability. A failure rate for the antifuse portion of the FPGA was determined in this test by running a large quantity of parts through typical environments for long durations, thus accumulating a large number of device-hours.

In this technology, a one-time programming is performed by permanently "fusing" or "melting" antifuses that are initially open. (An antifuse is initially an open circuit and becomes a short circuit after programming). The antifuses consist of a metal-amorphous silicon-silicon nitride-metal arrangement where the amorphous silicon fuses through the silicon nitride dielectric to provide a conducting link. There are special high-voltage programming transistors built into the FPGA, normally only active during programming. These form the antifuse links between the top metal layer and the layer beneiath in the technology, defining the desired circuitry. The antifuses are categorized into low-and high-current antifuses, the low-current antifuses only driving high-impedance gate inputs, and the high-current antifuses driving long interconnects and their larger capacitances. A further categorization divides antifuses into dynamic versus static antifuses—dynamic antifuses have a logic signal passing through them, while static antifuses only have a pull-up or pull-down function.

The focus of this effort was twofold. The first was to see whether the antifuse technology in these FPGAs was vulnerable to failure with time and temperature stress. In particular, the low-current antifuses were scrutinized by designing the circuit programmed into the FPGA to maximize their use. Second, the SRAM section of the device was exercised to look for any loss of function in the internal memory.

We performed both SRAM testing and antifuse testing on 760 parts distributed into three environmental control chambers. The first chamber held the devices-under-test (DUTs) at a fixed case temperature of 85°C; the second chamber held additional DUTs at a fixed case temperature of −40°C; and the last chamber cycled the remaining DUTs over the range of −40°C to 85°C. These cycled devices were powered in the same way as the devices in each of the hot and cold chambers, and all devices were continually exercised throughout the lifetest. Testing began in May 2007, and finished in January 2012, resulting in nearly 26.5 million device-hours. This data can be used in determining a failure in time (FIT) rate. The failure rate for any particular device is ultimately dependent on the specific device and operating conditions. It is estimated that for an average part with 185,000 dynamic antifuses and 100% duty cycle that a 60% confidence bound on the failure rate due to the antifuse portion of an RTAX FPGA is 25.88 FITs. To obtain a total RTAX FPGA failure rate, this FIT rate needs to be combined with the failure rate for the CMOS portion of the FPGA and the applicable failure rate for the package. This was not calculated here but can be done by using the MIL-HDBK-217 model or other accepted methodology and adding the result to the FIT rate defined here.

2. Description of the AX2000 Test Article

The AX2000 devices used in this lifetest were all programmed with an identical circuit design. The block diagram for the FPGA DUT (Figure 1) shows the sections dedicated to the testing of the antifuses and the SRAM as well as some of the support functions necessary to ensure that the device was exercised somewhat typically. In this design, more than 97% of the cells used were dedicated to the actual test. The remaining cells were either unused or used for the peripheral support circuits needed to execute the test and ensure that the device operated in a manner typical of an actual FPGA application. For instance, all of the device power supplies needed to operate within the product specifications and be able to handle switching demands on the device. Microsemi recommends in the data sheet that there not be more than eight simultaneously switching outputs (SSOs) in order to minimize switching noise. To this end, I/O cells were configured as 32-bit shift registers that would be loaded and clocked continually during testing to tax the power lines with a heavy switching demand (4 times the recommended limit of the datasheet). The shaded regions in the block diagram mark the circuits that targeted either the SRAM or the antifuse devices of interest.

2.1 Antifuse Test

The FPGA was programmed with several long chains or tracks of specific cell types so that the outputs from each track could be monitored over time to determine performance of the targeted antifuse type. Test cell design incorporated antifuses of specific types as denoted in Table 1. Cells were selected from Microsemi's cell libraries that utilize the K-type antifuse between routed clock track and inputs and the single DB antifuse between local track and inputs. Both of these antifuse types are low-current antifuses. Table 1 shows the totals for the various antifuse types for the final circuit design of the FPGAs under test. The test was devised to measure in-situ routing delay for each of the tracks and used that as a figure of merit to characterize the antifuse degradation. Chains of sequential cells and chains of combinatorial cells were used, and signals moving through these long chains were monitored for changes in propagation times that could be attributed to antifuse failure. All data readout was multiplexed to one I/O pin for monitoring.

The sequential networks shown in Figure 2 were made up of a chain of 280 D-type flip-flops configured as a shift register with the first seven elements (DFE4F cells from the Microsemi cell libraries) preset to a high output ,and the remaining 273 (DFE3C cells) cleared to a low output at the start of the test. The resulting pulse traveling through the shift register had a duty cycle of 2.50%. Behavior of this "7-hot" circulating shift register was monitored with a frequency counter that could measure the period and duty cycle for each track. For the sequential network, the data included a measurement of the clock frequency, the period of the output pulse, and the duty cycle. An antifuse failure would be reflected in an inability of this circuit to circulate the travelling pulse or a change in either the duty cycle or period.

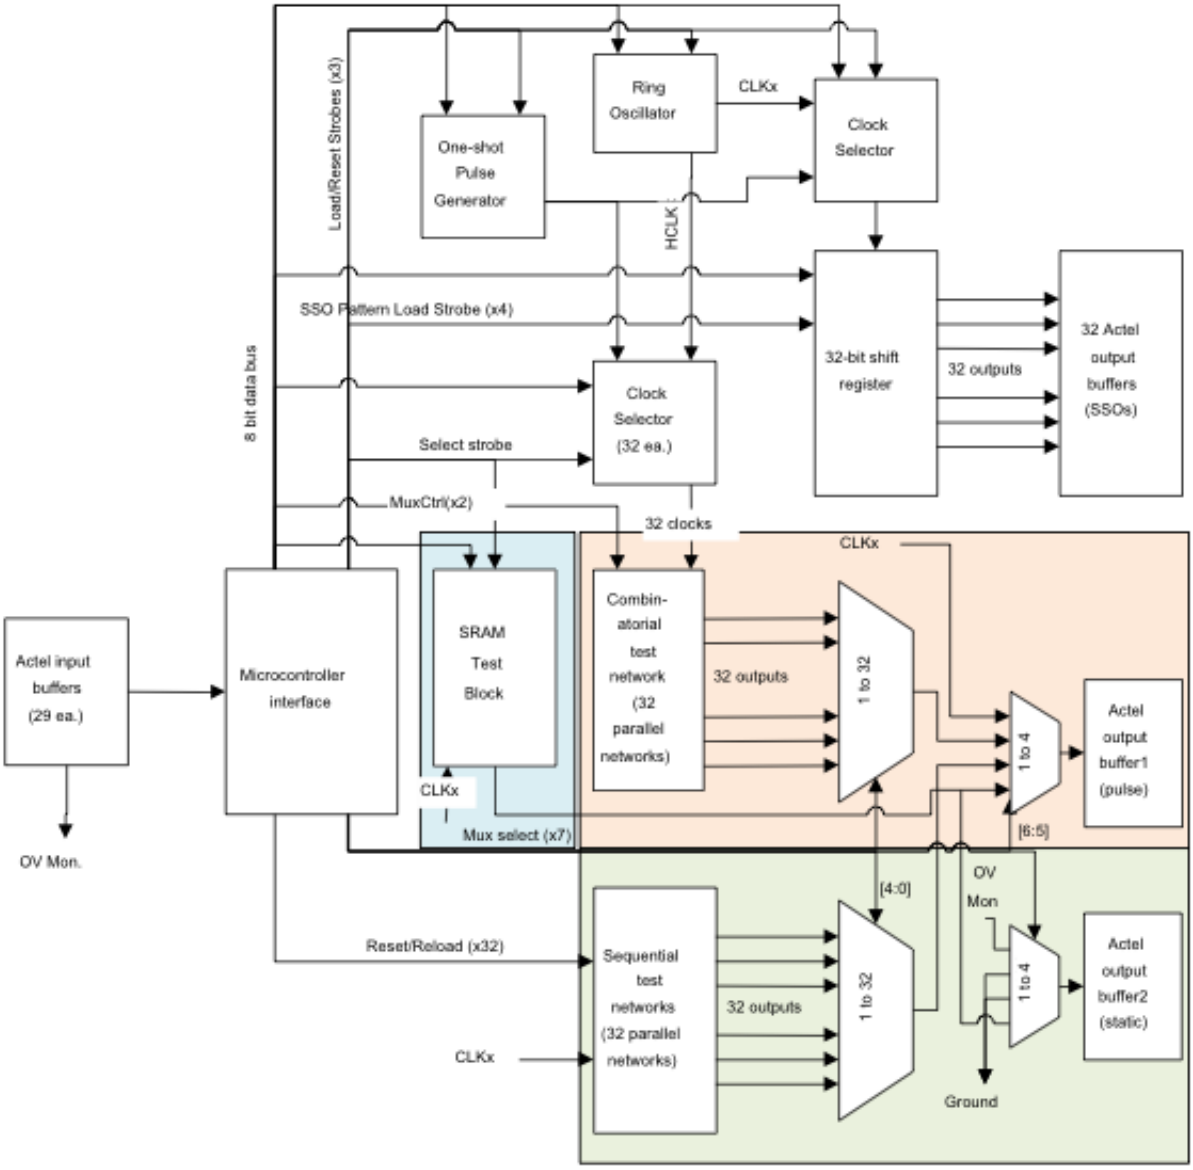

Figure 1. Block Diagram for the DUT showing antifuse test sections in orange (upper-right shaded block) for the Combinational tracks and in green (lower-right shaded block) for the Sequential tracks. The SRAM test block is shown in blue (left shaded block).

Table 1. Antifuse Statistics Report from the Actel Designer Program That Does the Place and Route for the AX2000 FPGA Circuit Designs

Antifuse Statistics Report

Product: Designer
Release: v7.2 SP2
Version: 7.2.3.2
Date: Thu Dec 14 14:23:37 2006
Design Name: ACTELDUT Family: Axcelerator Die: AX2000 Package: 896 FBGA
Design State: Layout

F	=>	29812	Antifuse between output track and vertical segment
V	=>	6625	Antifuse between two vertical segments
X	=>	28899	Antifuse between horizontal segment and vertical segment
H	=>	14831	Antifuse between two horizontal segments
LDH	=>	39	Antifuse between two horizontal long lines
LDV	=>	51	Antifuse between two vertical long lines
LLO	=>	558	Antifuse between long line and local track output
80815	**Total number of High Current Antifuses.**		
CI	=>	0	Antifuse between vertical segment and hclock module input
I	=>	40025	Antifuse between horizontal segment and input
DB	=>	115	Antifuse between local track and input
SD	=>	8899	Antifuse between output track and input (semi-direct)
K	=>	8421	Antifuse between routed clock track and input
LLI	=>	976	Antifuse between long line and local track input
58436	**Total number of Low Current Antifuses.**		
139251	**Total number of Dynamic Antifuses.**		
SIG	=>	0	Silicon signature antifuse
CSR	=>	8272	Antifuse between IO and IO banks
SSR1	=>	10	Silicon signature antifuse
SSR2	=>	0	Silicon signature antifuse
UID	=>	5	User Id antifuse
GBS	=>	0	Global set/reset antifuse
P	=>	1	Programming antifuse
S	=>	0	Security antifuse
J	=>	388948	Antifuse between horizontal NVCC or NGND and input
CJ	=>	3456	Antifuse between vertical NVCC or NGND and hclock module input
Y	=>	0	Antifuse between horizontal segment and vertical NVCC or NGND
Z	=>	0	Antifuse between vertical segment and horizontal NVCC or NGND
LLJ	=>	24605	Antifuse between local track and LDNVCC
425297	**Total number of Static Antifuses**		
564548	**Total number of Antifuses**		
9048	**Total number of Hardwired connections**		
73	**Total number of Single DB Antifuses**		

2.1.1 Sequential Test Network

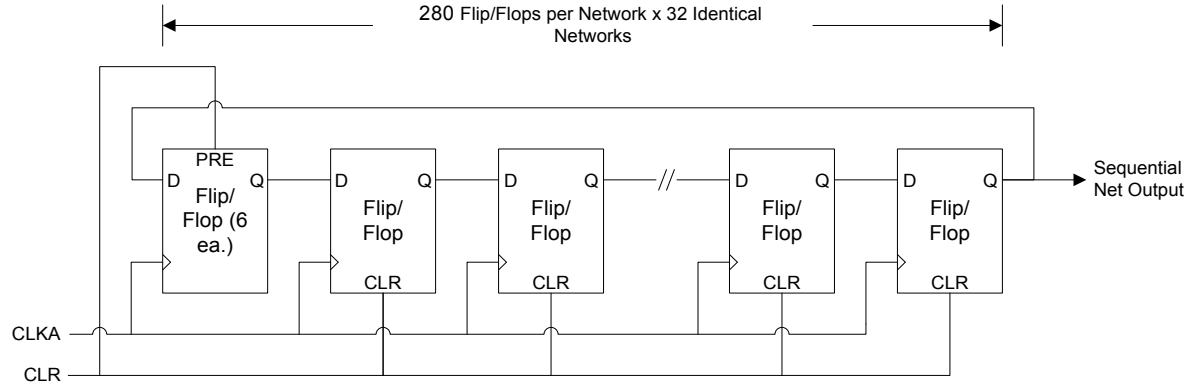

Figure 2. Chain of sequential cells used to exercise some of the low-current antifuses.

2.1.2 Combinatorial Test Network

The combinatorial networks were made up of either the inverter cell (INV cell) or the full combinational module (CMFE cell). The cells were connected end to end in long tracks. A measure of the time it takes a signal to travel from the start of the track to the end was made using a timer. An antifuse failure would be reflected in a change in the time it took to pass the signal through the track. Figure 3 shows the configuration of cells for the INV inverter. A digital pulse signal applied at the first inverter would start the counter, and that signal exiting the last inverter would stop it. In order to maximize the use of K antifuses, half the tracks were implemented with a similar configuration but using a CMFE mux cell in place of the inverter. The CMFE cell is a multiplexor so one of four lines is passed through depending on the state of two select lines. In order to test the antifuses with various loads, four configurations were arranged for the four possible modes of the multiplexor. These tracks were exercised for all four of the possible multiplexed tracks in modes that would route the input to the output in the manner shown in Figure 4. Four tracks would be related to each other so the signal would move through the cells exercising the targeted antifuses. In mode 0, the signal propagated down a single track. In the other three modes, the signal would be directed through sections of the other tracks until ultimately ending up back on the original track for measurement.

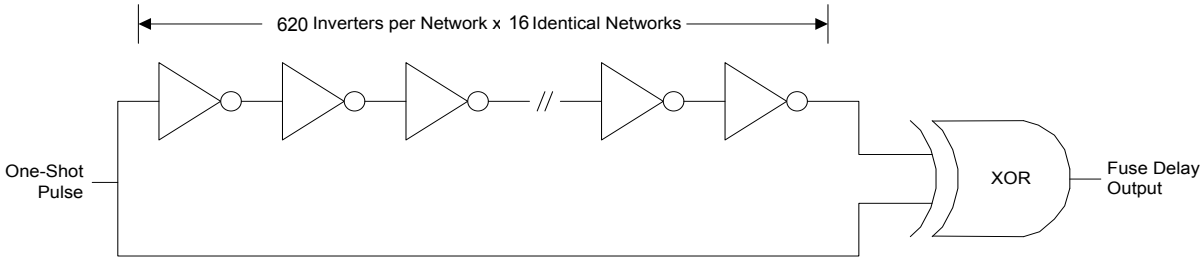

Figure 3. Inverter cells used in a combinational chain for an asynchronous test of the antifuses.

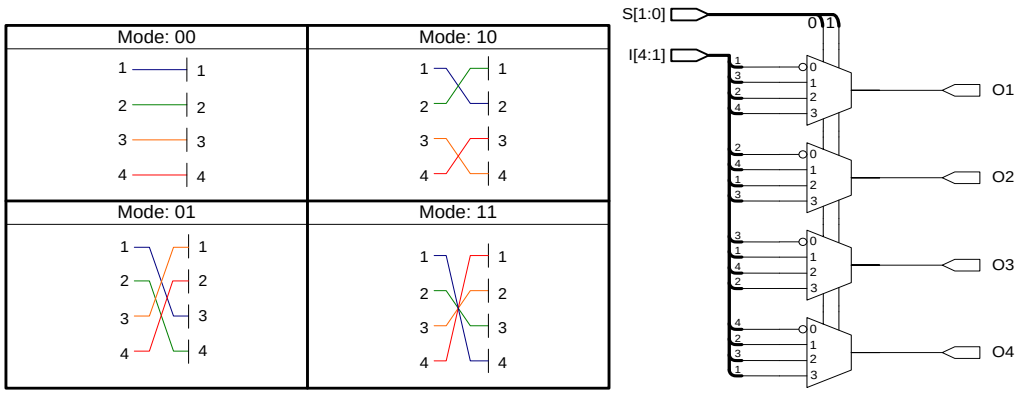

Figure 4. CMFE cell used to maximize the K antifuse count. Since all four modes of the mux needed to be tested, the chains were grouped together.

2.1.3 SRAM Test

The SRAM test was designed to use AX2000 elements for an internal memory test that would generate address and data vectors and compare results to expected values. A flag would be set on a single pin of the FPGA that was monitored as an analog housekeeping point through the DVM. The first and last error and address were stored in registers on the chip to be read at the first opportunity for the operator once the data was reviewed and the flag was seen. These memory test circuits used the internal clock generated for the antifuse test.

The memory test comprised two separate tests: a long-term operational test and a retention and minimum operation voltage test. For the long-term operational test, the presence of an error flag would indicate potential failure that would require investigation. Consistent indication of an error or errors would be a failure ,while non-repeatable errors would be noted for further discussions.

In the retention/minimum operating voltage test, the distribution of minimum operating voltage and minimum retention voltage would be collected and statistically evaluated. Changes in minimum voltage levels over the lifetime would need to be evaluated and discussed while consistent indication of error(s) at the nominal voltage would be a failure. It was planned that this test would be performed if the results of the long-term operational test indicated issues with the SRAM reliability in order to refine the understanding of the performance.

An 8-bit word was used to command the SRAM test. Prior to normal testing, the SRAM in each device was reset. The normal test then followed the Modified Walk 1/0 sequence where the entire memory was written with a single pattern, read out, and verified, and then a second pattern written, read, and verified. This procedure was repeated for 20 patterns that comprised a walking 1 in a field of 10 zeroes followed by a walking 0 in a field of ones. If the reading of a memory location did not match what had been written, the data and address were saved, and an error flag set.

2.1.4 Data Recording

A fully automated test set was used to exercise the devices-under-test (DUTs). The test boards were controlled through a daisy-chained RS232 port from a PC that would select each board, each FPGA, and each track of cells to monitor performance over the lifetest. The circuits designed into the

FPGAs were programmed to produce a known frequency or pulse width depending on the cells used in the subcircuit. These products were monitored with a Universal Counter (from Agilent, 53132A) under GPIB control. Housekeeping for each part was also monitored using a Multimeter (from Agilent, 34401A) to measure the I/O voltage and the core voltage, the current for the I/O and the core, the quiescent current, and a temperature. In addition, this GPIB-controlled meter was used to monitor the input voltages and currents for each of the boards.

Each board was first tested separately from all the others. For the first 75 hours or so data were referred to as High Visibility data since each device could be interrogated about every 45 min. This data was used to calculate the baseline timing for each device track. Once this initial period was over, the board was moved to the long-term Lifetest. With the maximum number of boards (14) in a static oven, it was possible to read and record every track for every part for every board once per day. The cycling chamber had room for only 10 boards, but because of the time taken to transition the temperature, a full data set still took about one full day to read out. Data was recorded for half the boards at the maximum temperature and half the boards at the minimum temperature. Figure 5 shows one full temperature cycle and when the data was read.

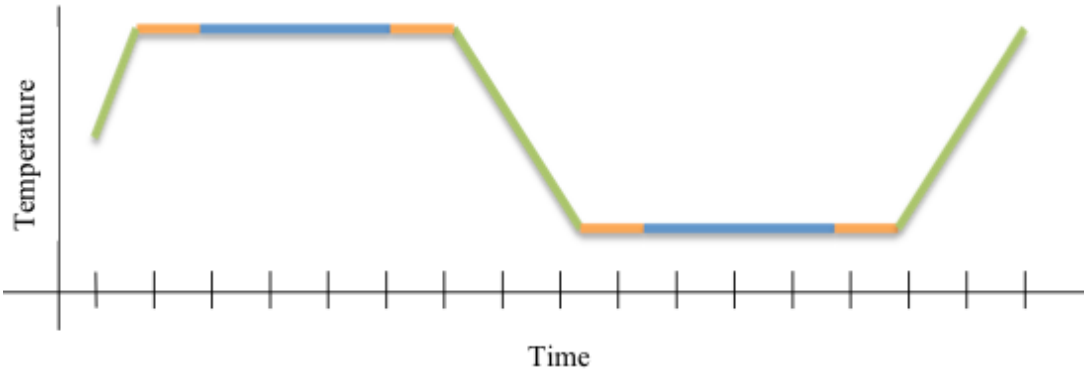

Figure 5. One full temperature cycle. From room temperature, the chamber was taken to +80°C in about 15 min, leaving about 5°C headroom for the devices to warm up but remain within the specifications for the part. The boards were held at this temperature for 15 min (shown in orange). Data for one board was taken in about 45 min (blue) followed by an additional 15-min soak (orange). The chamber then transitioned to −40°C in about 30 min (green). Again, the boards were held in a soak for 15 min (orange). Data was read out in 45 min (blue), soaked for an additional 15 min (orange) then transitioned to +80°C (green) to repeat the cycle.

3. Test Time and Dates

The test started in May 2007. Each board was first run alone in a chamber for three days in order to accumulate concentrated data. Each device would be read out about every 45 min. Once the board had accumulated at least 72 hours on each device, it was transitioned into a second identically configured system that would share power to many boards. In this system, with the accumulation of more boards over time, the devices would be polled less frequently. Ultimately, once each chamber was fully populated, each part was interrogated once a day.

Table 2 shows the total hours of operation for each board in the three chambers. The last board was turned off in January 2012. Over the span of the lifetest, some devices were taken off test, as discussed in Section 7. The total of device hours reflects this.

Table 2. Total Number of Hours Seen by Each Board in Each of the Three Test Chambers Identified by How the Temperature Was Controlled

	Hot Oven			Cold Oven			Temperature Cycling Oven		
PCB No.	**Board Hours**	**DUTs on Test**	**Device Hours**	**Board Hours**	**DUTs on Test**	**Device Hours**	**Board Hours**	**DUTs on Test**	**Device Hours**
1	39,845	20	796,907	37,233	19	709,397	31,001	19	585,157
2	39,840	19	756,955	37,237	20	744,744	31,001	15	523,417
3	39,742	20	794,845	37,165	20	743,297	30,928	16	541,954
4	39,665	20	793,290	37,087	18	667,565	30,500	12	412,732
5	39,596	20	791,923	37,019	20	740,377	30,795	19	585,113
6	39,528	19	751,035	36,951	19	702,072	30,666	19	606,491
7	39,417	20	788,338	36,845	20	736,903	29,745	18	566,252
8	39,356	20	787,127	36,781	20	735,625	30,440	17	592,904
9	39,265	20	785,293	36,631	17	636,205	29,301	18	575,122
10	39,185	20	783,698	36,709	20	734,181	30,018	19	598,876
11	38,066	19	727,338	36,562	20	731,247			
12	38,977	20	779,535	36,497	20	729,937			
13	39,010	20	780,196	36,434	20	728,688			
14	38,897	20	777,940	34,571	19	671,334			
Device Hours			10,894,421			10,011,572			5,588,018
Total Device Hours									**26,494,010**

THIS PAGE INTENTIONALLY LEFT BLANK

4. Description of the Test Facilities

Each of the three setups (hot, cold, and cycled) were made up of nearly identical equipment racks (housing the data collection computer, the housekeeping digital volt meter, the universal counter, the power supplies, an Uninterruptible Power Supply [UPS], and a house power monitoring system) and a thermal chamber. The power monitoring system was designed to recognize a power problem, and after a short time with no correction would initiate a controlled power shutdown of the system in less time than the UPS battery would last. The hot and cold chambers were CSZ ZP-32 temperature chambers large enough to hold the 14-board set and capable of holding a steady temperature at the required high and low temperatures. The thermal cycling chamber was an Espec TSA-71S-(A/W) capable of quickly transitioning from −40°C to +80°C and back. This chamber only had room for 10 boards.

THIS PAGE INTENTIONALLY LEFT BLANK

5. Test Pass/Fail Criterion

The test pass/fail criterion for these devices was that the propagation delay through these long strings of cells rich in the critical antifuse types would not change more than ±2 ns from the baseline measurement. The baseline for each track was an average value of the times measured for the first 100 hours after the device was in the long-term test with all the other boards. The inverter tracks were strings of 620 INV cells resulting in propagation times of approximately 490 ns. The CMFE tracks were made up of 640 elements and saw delays of about 550 ns for mode 0, 440 ns for mode 1, and 460 ns for modes 2 and 3. Each device and each track varied from these numbers some amount based on device variation, local temperature of the device, and the frequency of the device ring oscillator. However, once a baseline for a given track was established, any deviation from the baseline was simple to detect.

THIS PAGE INTENTIONALLY LEFT BLANK

6. Deviations from Procedure

Over time, system problems would arise. The control electronics were designed to recognize problems with the house power and respond with a controlled shutdown of the test after 5 min without resolution. This would guarantee that the devices under test would be turned off safely before the UPS ran out of battery life in the case of a prolonged power outage. Over the course of the five years, each chamber had several controlled shutdowns such as this. In every case, the devices were brought back on line after power was restored, and there was no evidence of antifuse failures.

The cold chamber failed to maintain temperature three times, bringing the devices up to room temperature. The devices continued to be exercised through this temperature change. In these cases, the data was set aside, the parts returned to the proper testing temperature, and evaluated for the pass/fail criteria. No failures occurred; so the hours of operation were included in the totals. On one occasion in the cycling temperature setup, the Universal Counter failed to function. No data was recorded for a short period. The meter was replaced, and the data from the devices evaluated, and no evidence of failure was seen.

The universal counters and the multimeters used in this test to measure the pulse widths of the signals coming from the test articles were calibrated instruments rented from outside companies. The universal counters came from TNS RenTelCo, and the multimeters came from Test Equity, Inc. These instruments fell outside of Aerospace's internal system for calibration assurance. This meant that they remained in the test setup measuring the data output past the two-year period of calibration. However, the system was self calibrating. Since all data from a device was multiplexed to a single I/O, and then all parts were, in turn, multiplexed to a board output that was, in turn, multiplexed with all the other boards on test, any deviation in the measuring instrument would show up in all the data. If there were a false correction of the data, it would have been on every data point. It would be impossible for a change on a single track to be masked out. Still, in order to validate this, data was taken on one board from each chamber at the end of the test with Aerospace equipment that was properly calibrated as stated in Appendix A and compared for differences. None were found. This served the purpose of validating the data taken throughout the test period.

THIS PAGE INTENTIONALLY LEFT BLANK

7. Failure or Problem Reports

Table 3 lists the devices that were taken off test for failure to report data. The DUTs were ball grid array parts, and these large devices resided on large boards that were flow soldered by an outside

Table 3. Devices That Were Removed From the Test Set

Chamber	Board	Device	Hours	Comments
Cold	1	3	1979	Lost power to the core
Cold	4	9	0	No start
Cold	4	10	0	No start
Cold	6	12	0	No start
Cold	9	7	0	No start
Cold	9	10	0	No start
Cold	9	15	13473	Intermittent data out, failing to switch tracks
Cold	14	2	14477	no power to the core
Hot	2	15	0	No start
Hot	6	15	0	No Start
Hot	11	16	4086	SRAM
TC	1	1	24460	No current - lost power to core
TC	2	1	29979	no power to the core
TC	2	9	12902	Intermittent data, half current, solder issues
TC	2	11	0	No start
TC	2	16	8040	No housekeeping data so turned off
TC	2	18	7741	no power to the core
TC	3	8	20884	Increase in current consumption although data remained in spec
TC	3	10	7668	no power to the core
TC	3	16	12504	increase in current consumption, no sequential data - came back on line after reflow with no indication of antifuse failure
TC	3	18	11037	can't set clock, loss of sequential function
TC	4	1	7240	data out is intermittent
TC	4	11	0	No start
TC	4	12	0	No start
TC	4	13	0	No start
TC	4	14	0	No start
TC	4	15	0	No start
TC	4	16	29022	no power to the core
TC	4	17	10470	no power to the core
TC	5	12	0	No start
TC	6	17	23837	Lost all communication to board - solder fatigue. Couldn't recover power to part 17
TC	7	3	4190	SRAM
TC	7	16	26659	no power to the core
TC	8	6	25392	Intermittent output, no sequential data
TC	8	11	26730	no power to the core
TC	8	16	23299	no power to the core
TC	9	5	28207	no power to the core
TC	9	15	19505	no power to the core
TC	10	3	28540	no power to the core

vendor. Out of 760 devices, 14 did not start up due to insufficient solder on the ball grid array (BGA). Over time, there were more problems with solder failure, especially in the taxing temperature cycling environment. It was found that after about a year and a half, the boards needed to have some of the solder on the peripheral support circuitry reflowed. This was repeated at about three years. At the point that the test was ended, the boards were all requiring attention again, and it was determined that it was a good place to stop. Note that the purpose of the test was to evaluate the reliability of the antifuse technology used in these FPGAs, not the packaging method.

The most common problem encountered was that the solder on the BGA failed, and power was lost to the core of the FPGAs. This was seen as a drop in current consumption to zero and a loss of all data from the output port. It was not a simple task to reflow the BGA solder without putting the still-functioning devices at risk, so once this happened, the devices (1C-3, 14C-2, 1TC-1, 2TC-1, 2TC-18, 3TC-10, 4TC-16, 4TC-17, 6TC-17, 7TC-16, 8TC-11, 8TC-16, 9TC-5, 9TC-15, and 10TC-3 where these designators take the form board number followed by chamber letter and then the part number on that board) were removed from the test set. Solder issues also arose for the support chips and caused the removal of one of the devices in the cycling chamber (2TC-16) when the housekeeping could no longer be read out.

Some other devices were taken off line for intermittent performance in the readout of the data. When data was successfully coming out it was correctly within the ±2 ns range. When data was not coming out, it was all the data missing—not just one, or a few of the data. The intermittence was over long periods of time. That is, the device output would fail for a day or two or more and then start to work again with no cycling of the power or other human intervention. This is not the signature of an anti-fuse failure but rather of more solder problems. These devices (9C-15, 2TC-9, 3TC-8, 3TC-16, 3TC-18, 4TC-1, and 8TC-6) were no longer counted in the test.

Two devices were taken off test because of SRAM failure. Careful analysis of the devices was undertaken by both Aerospace and Actel. The first device (7TC-3) was observed to only fail at cold temperature but work at room temperature. The other device (11H-16) seemed to fail at all temperatures but at only two addresses. None of the devices used in this test went through blank burn-in processing nor did they receive any military temperature screening. They only had electrical screening test at room temperature before being shipped to Aerospace. After characterizing the failure modes, Aerospace returned the parts to Actel for more analysis. Actel determined that 7TC-3 failed the manufacturing test (t1_sertest) at cold and passed at hot, as was observed on test at Aerospace. The 11H-16 device passed at hot (this device was being tested at hot) but failed at cold. Eventually Actel performed de-processing on this device to look for defects or damage that correlated with the failure mode. They did not find any and determined the most likely cause of the failure to be CMOS-related issues within the SRAM logic. A possible cause could be micro pin-hole in the gate oxide of the SRAM cell, for instance. The failure would have likely been caught by the dynamic blank burn-in and military temperature testing done as part of the 883B flow for military-grade parts. The failures observed would not be seen in flight-quality devices.

An issue that became apparent a few thousand hours into testing in every chamber was that there was localized pooling of heat in the board stacks. It was found that parts 16, 11, and to a lesser degree parts 17, 18, 6, and 1 on some of the boards in each system were running warmer than they did when the baselines were calculated. Looking at the configuration of the boards in the housing used for the

two larger chambers, the hot and the cold, it can be seen in Figure 6 that these parts were along the top of the boards where the heat of all the other parts would rise and toward the back side of the boards. There was a metal bar across the top at the back that would have impeded the outflow of this heat. Likewise, there was a cross-strap at the back that could change airflow. The effect for this setup was most noticeable on the inner boards, for instance in the cold chamber data, boards 3−10 showed a shift in the times it took a pulse to traverse a delay chain. It was found, too, that shifting the whole board assembly inside the chamber could produce a different temperature profile across the parts.

The frame that held the boards in the cycling chamber was a little different in its construction, but in that chamber the board assembly was only slightly smaller than the cavity itself, which did not give a completely free flow of air. However, in this chamber, the heated or cooled air was forced through by blowers. This caused the highest temperature location to shift to the two boards closest to the door that were surrounded by a mass of cables which impeded flow. The data is shown in Section 8, Test Summary, and shows shifts in the data that reflect the increase in temperature locally for some of the parts.

What is not easily seen in the data as presented is that for each part affected by an increase in temperature, there was a secondary pattern that helped to determine that no antifuse failure was present. In the cases where the local temperature of a part increased, it was observed that the time it took the signal to travel down a track increased for the tracks measured first and decreased or remained close to baseline for tracks measured later. Figure 7 shows the effect reading out errors had on the internal temperature of the parts. The program control of the parts would turn off the internal ring oscillator to make the measurement of the time it took for a signal to traverse the tracks of combinational cells (see Appendix E for software flow). With the oscillator off, the part would cool. If an error was detected, an additional set of steps was taken to record the error and display it on the monitor. This increased the time it took to read out all the tracks. The effect on the internal temperature was that the

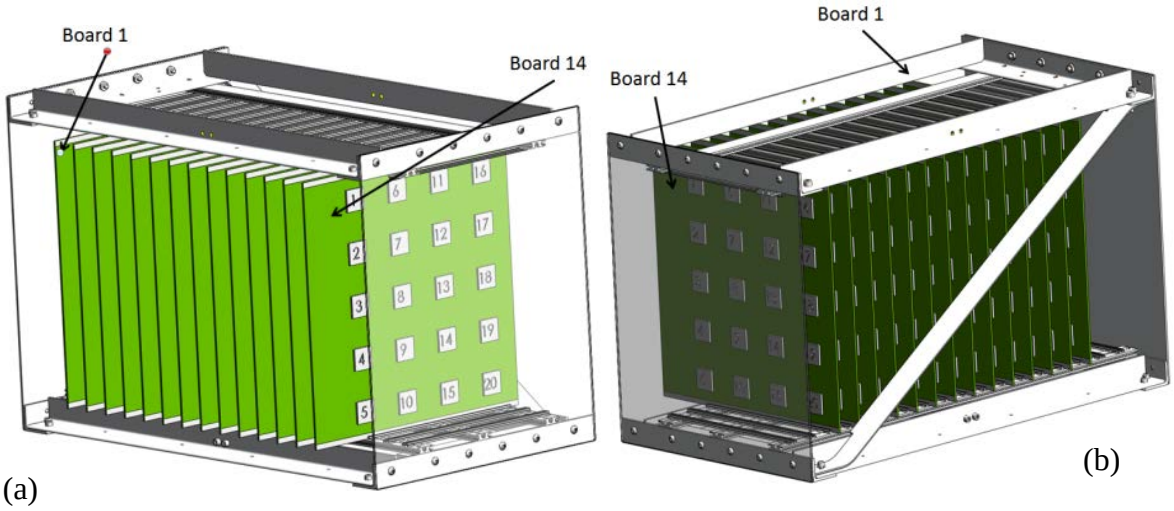

Figure 6. Arrangement of test boards shown from the front (a) and the back (b) and showing the orientation of the parts on the boards. The conditioned air in all chambers originated from the back of the chamber. For the hot and cold chamber, the back of the board set was approximately in the middle of the chamber.

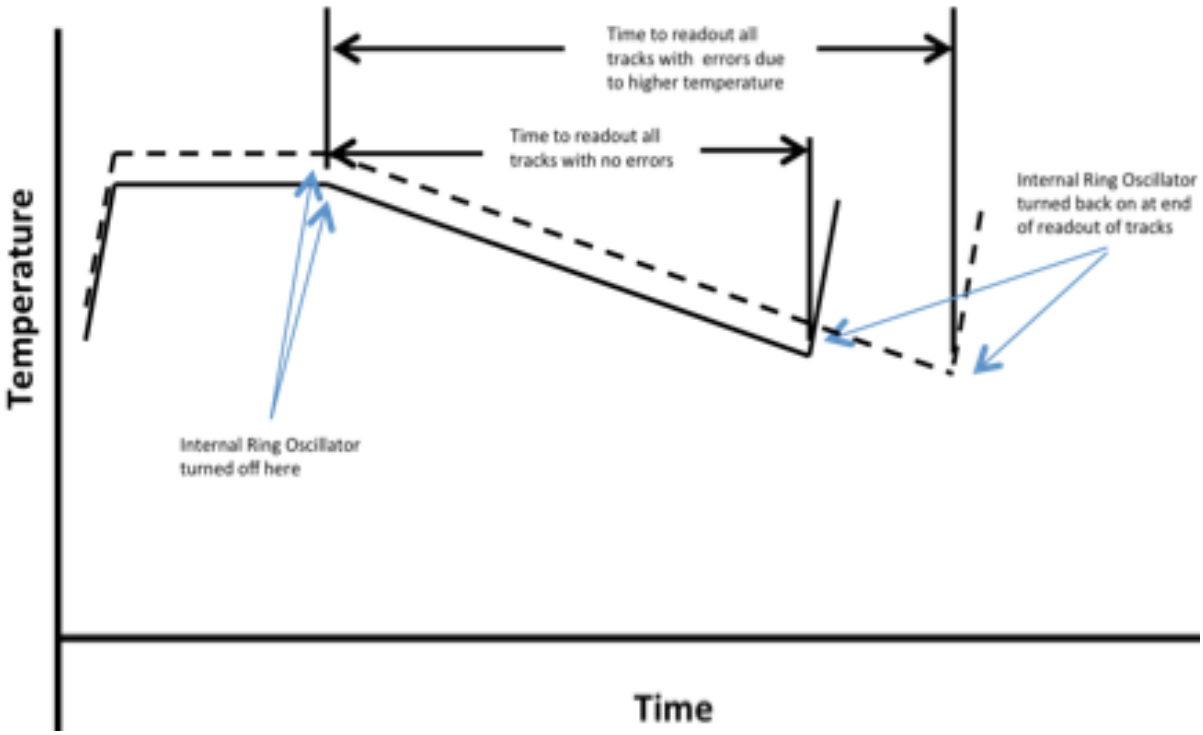

Figure 7. The effect of reading out errors on the internal temperature of the device resulted in deviations from the baseline only in the tracks read out first. With the ring oscillator off and the program taking longer to read each track when recording and displaying errors, the tracks read out later had time to cool to baseline temperatures.

tracks read out earlier were warmer from the ambient temperature being higher, but the longer delay in reading out the errors cooled the part more so that the later tracks are at the baseline temperature. Only the tracks read first have an increase in time taken for the signal to propagate the entire track.

8. Test Summary

The test resulted in nearly 26.5 million device-hours without antifuse failure. This data set became large enough to exceed the capability of the plotting program. It currently cannot plot all of the delta delay distributions from one chamber in a single graph. Therefore the 38 boards are plotted independently in the following figures. The figures are probability plots of the changes in delay from the baseline delays and are plotted at 1000-hour intervals. The probability plots are shown as percentiles of the normal distribution with plotting positions from the median ranks approximated using the formula of Benard and Bosi-Levenbach (1953). The data from board TC1 in the cycling chamber are shown in Figure 8. On average, the parts in the cycling chamber experienced 6350 temperature cycles.

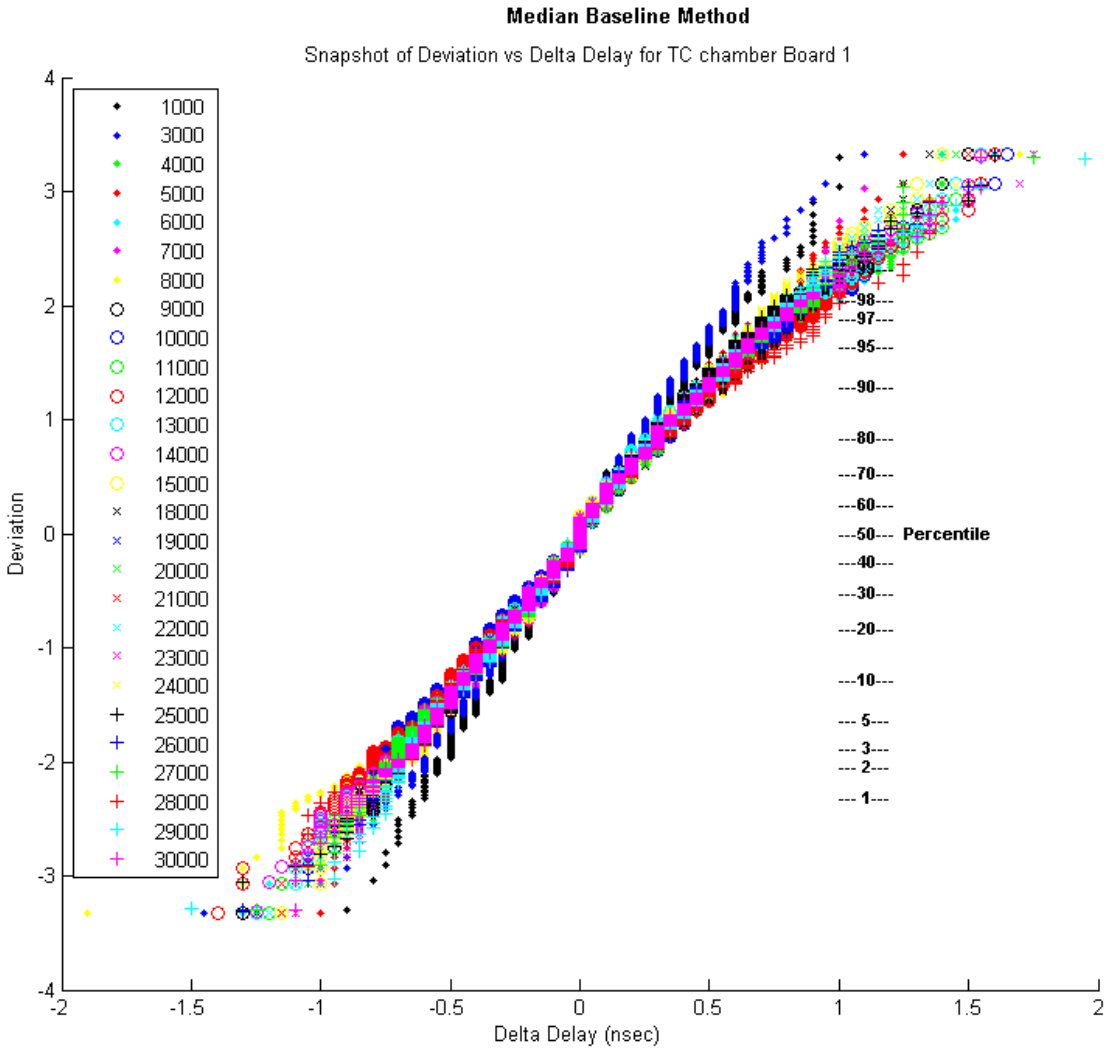

Figure 8. Data from Board 1TC in the temperature cycling chamber.

In Figure 9, the data stays nicely distributed except around 500 hours. As the experiment boards moved from the high-visibility, single-board setup to the combined board lifetest setup, there were some issues with current limiting in the power supply where the limit was set too low. Around 5000 hours into the test for the first boards, the core voltage was pulled lower than the specified voltage. Once returned to the correct voltage, the data came back within the ±2 ns deviations from the baselines.

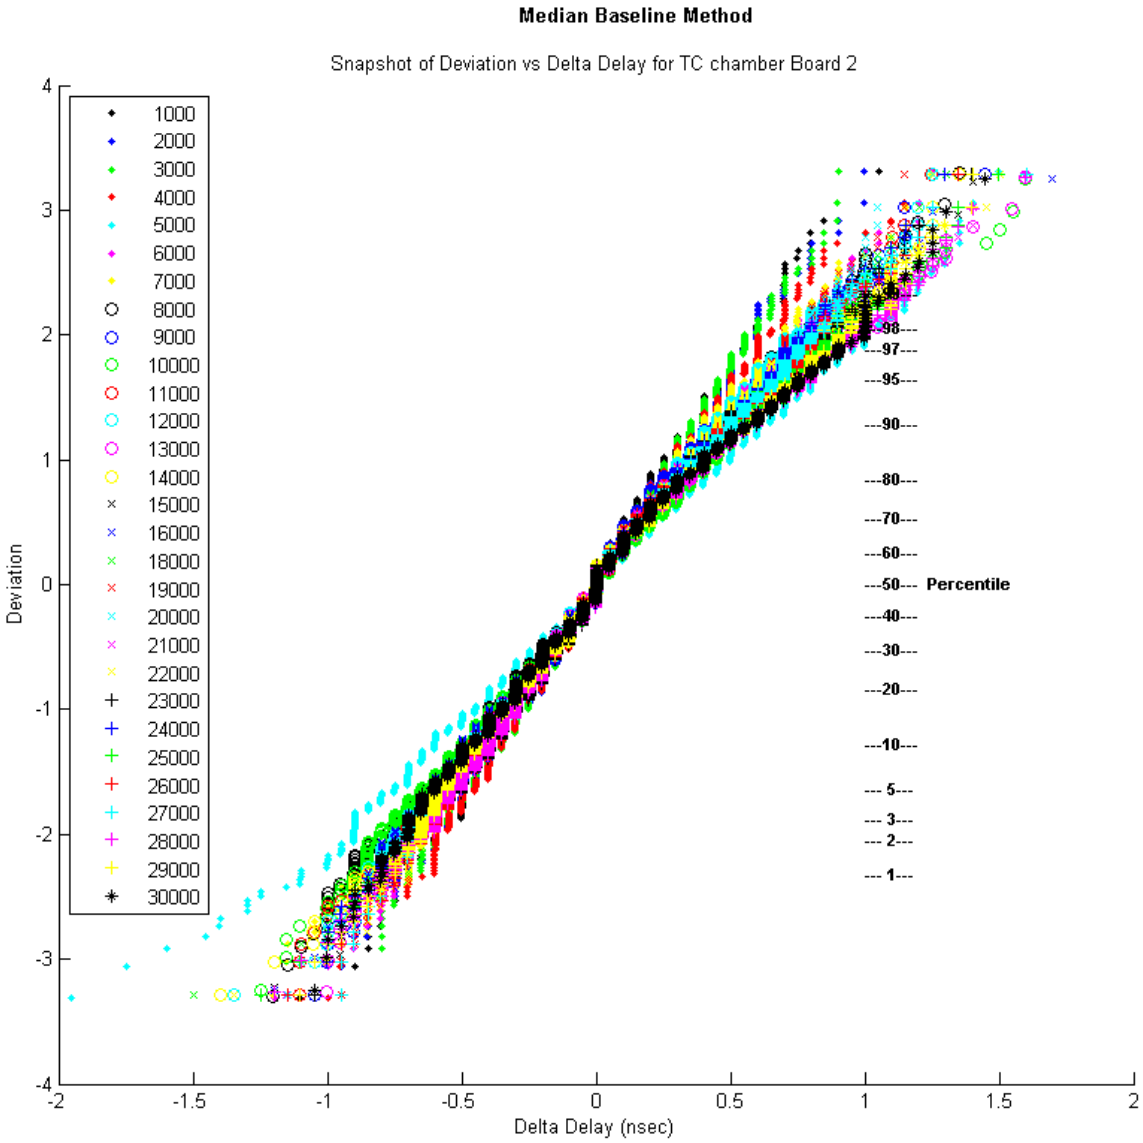

Figure 9. Data from Board 2TC in the temperature cycling chamber. A drop in power to the board around 500 hours shows up in the data as a temporary shift to faster propagation times.

In Figure 10, the data for board 3 of the cycling chamber are plotted every 1000 hours. In the cycling oven, all parts experienced the full temperature cycle. However, for any given part, data was taken at only one temperature in order to be able to make the comparison to a single baseline value. On some occasions, the ambient temperature was not correct. For instance, at about 5000 hours, the chamber experienced a problem with icing and needed to be defrosted. Data was still being taken for some of the time the chamber was malfunctioning. When the chamber was brought back on line, the data fell back into the required range. Therefore, there are some gaps where data were not included in the plotting because they were taken at the wrong temperature.

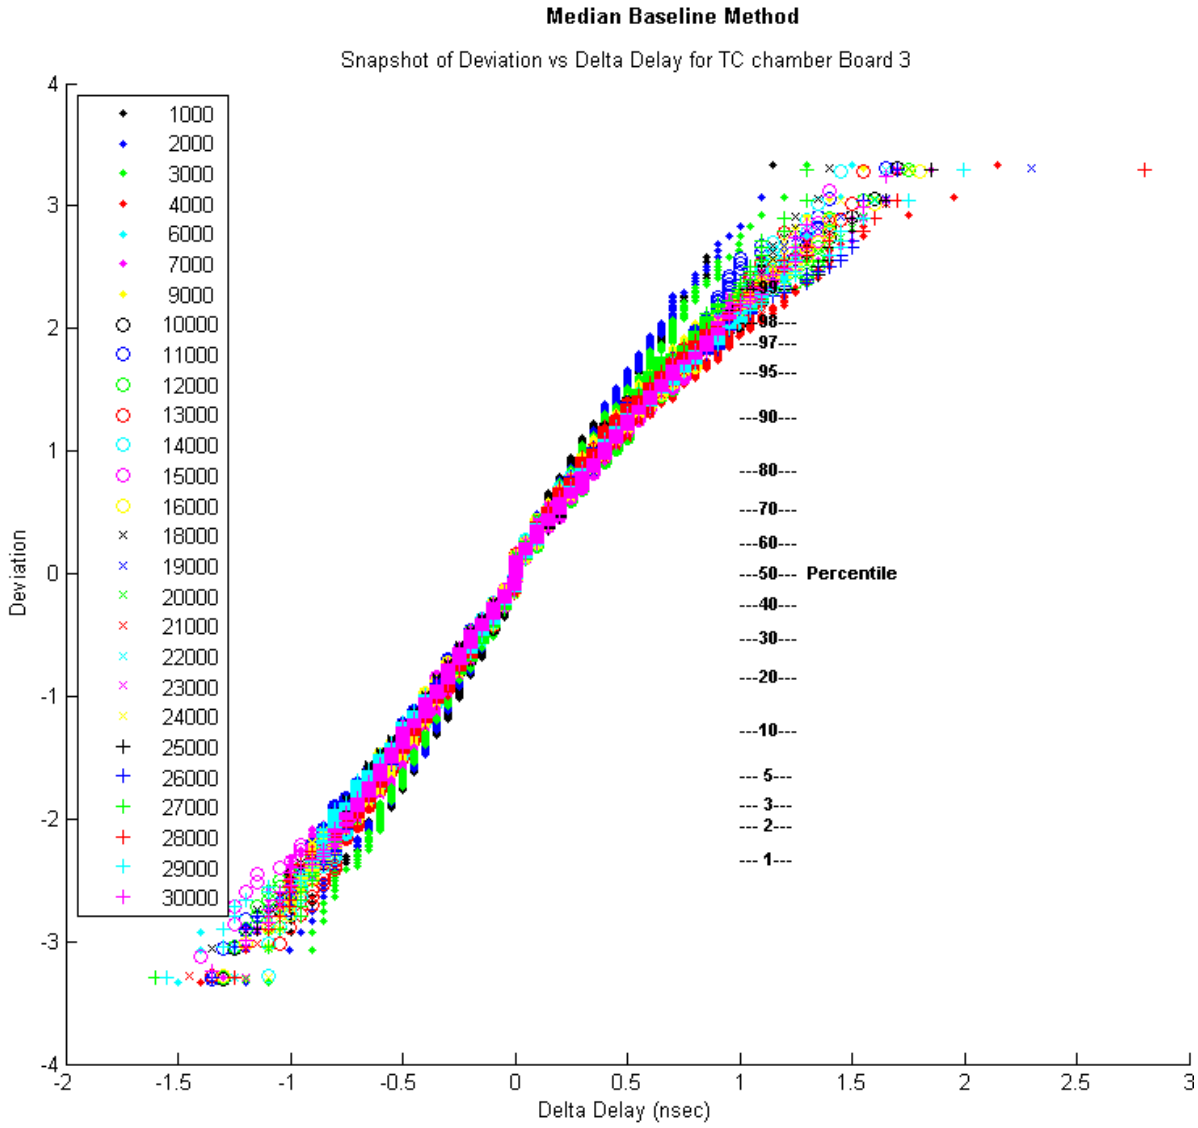

Figure 10. Data from Board 3TC in the temperature cycling chamber. There are 3 data points outside of the +2 ns bound. These outliers were not permanent; that is, data taken at later times did not fall outside the bound. This behavior is not indicative of antifuse failure.

Board 4TC showed a deviation in the data around 6000 hours, as seen in Figure 11. The data that is outside the ±2 ns bounds are all from part 6. The data for this part deviated out of bounds from 5766 hours to 6126 hours. There is no obvious reason for this—the housekeeping showed the temperature was stable, and the power to the part was correct. However, the data came back into range after this time with no indication of antifuse failure. Figure 12 shows the distribution plot with the data from part 6 removed for the interval where the data was outside the bounds. Figure 13 shows all the measured times for all the tracks in part 6 for the entire test. The deviations around 6000 hours never repeated. The data for boards 5TC and 6TC are shown in Figures 14 and 15, respectively.

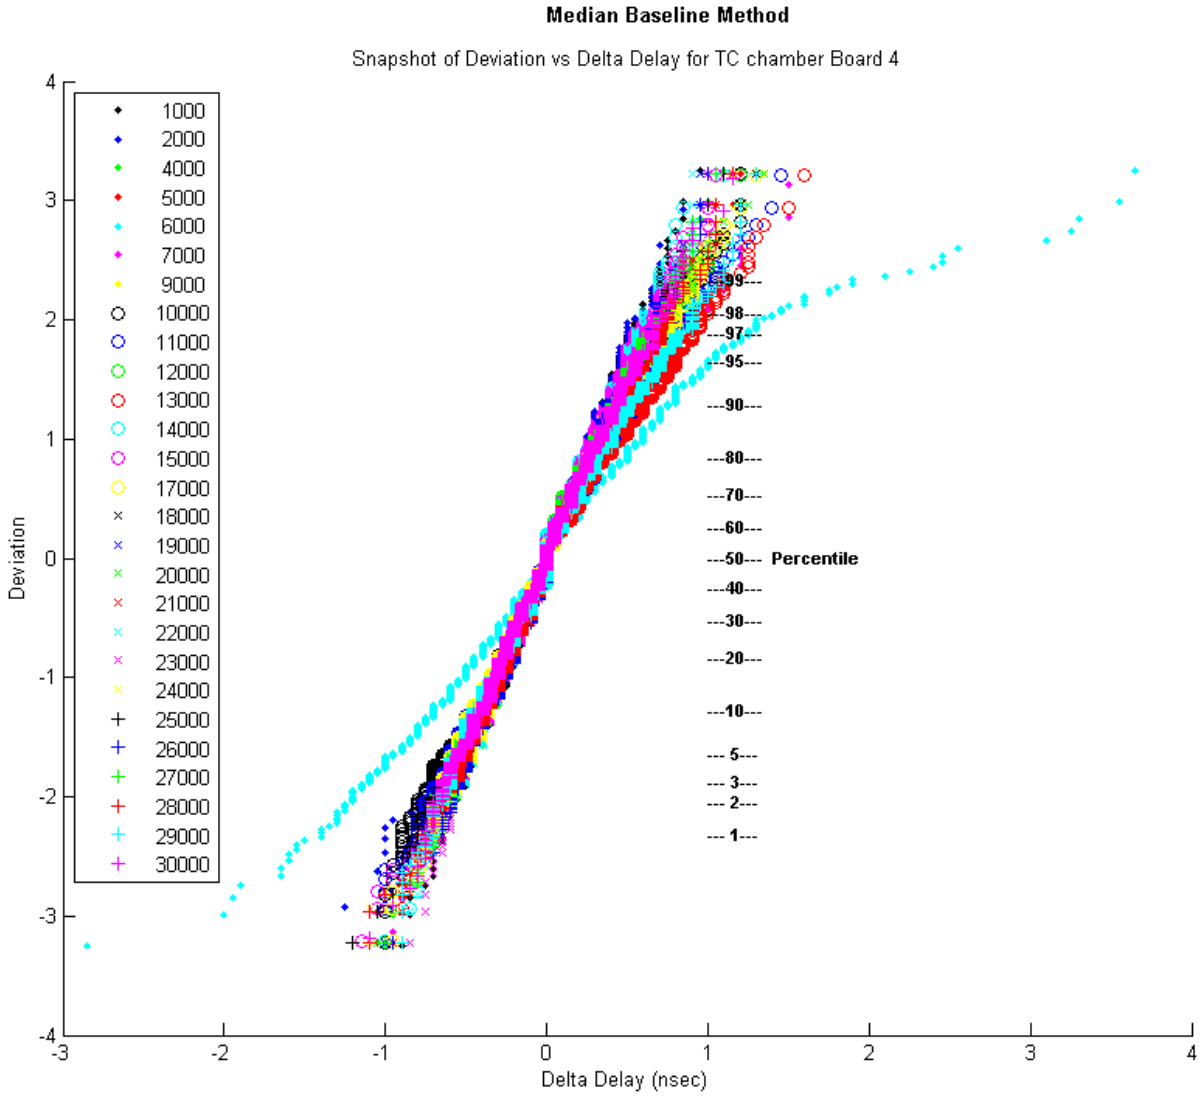

Figure 11. Data from Board 4TC in the temperature cycling chamber. Data outside the ±2 ns bounds occurred around 6000 hours. Because the deviations were not permanent this is not an indication of antifuse failure.

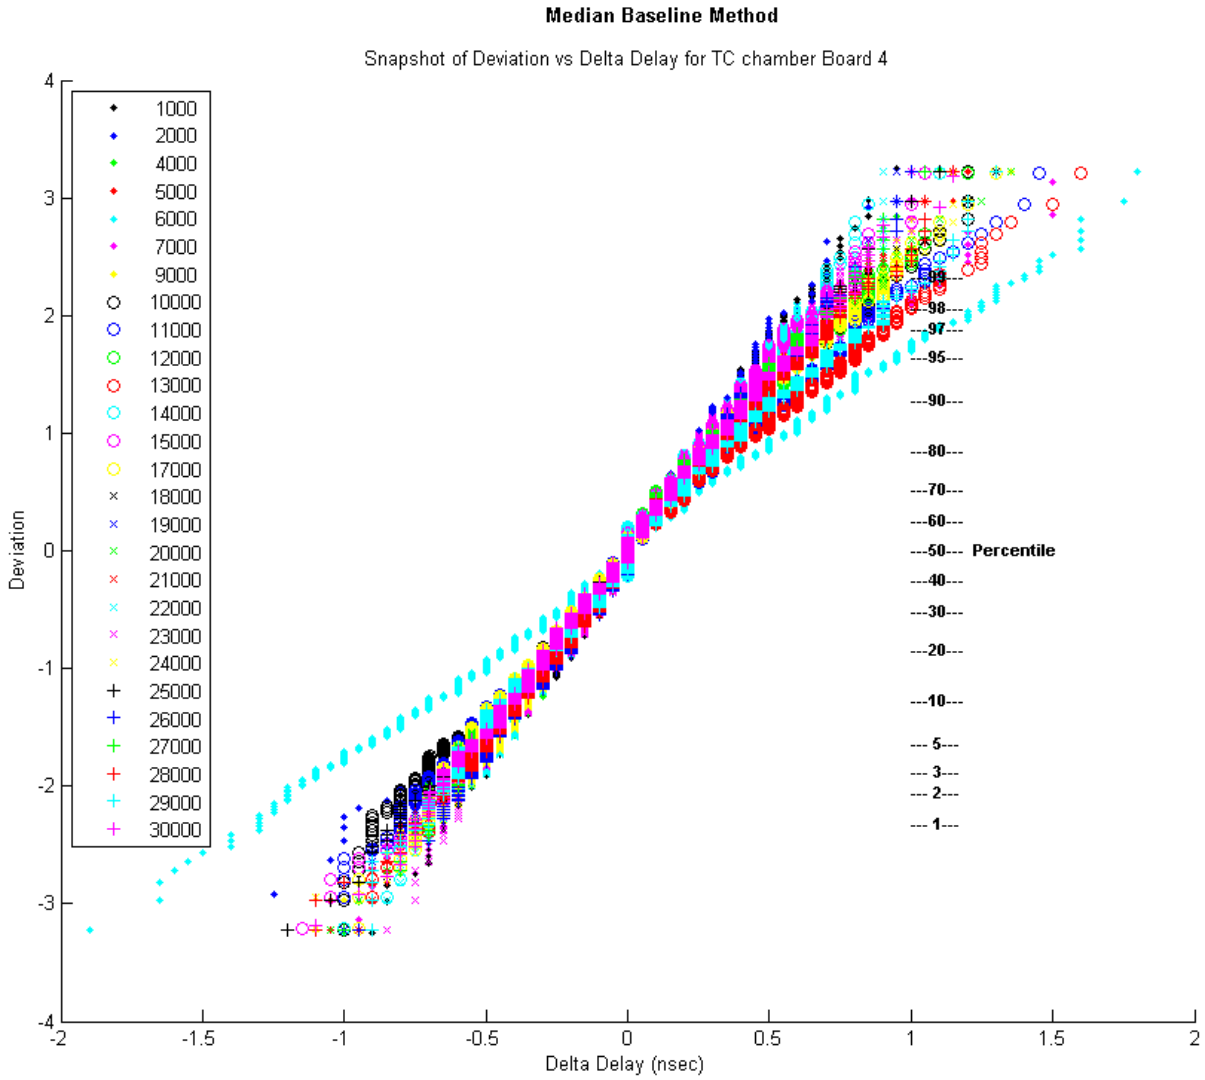

Figure 12. Data from Board 4TC in the temperature cycling chamber with data from part 6 between the hours 5766 and 6126 removed in order to show that the outliers in Figure 10 were entirely from this single part. All the data for this part after this period are still included and show the part functioning with no evidence of anti-fuse failure.

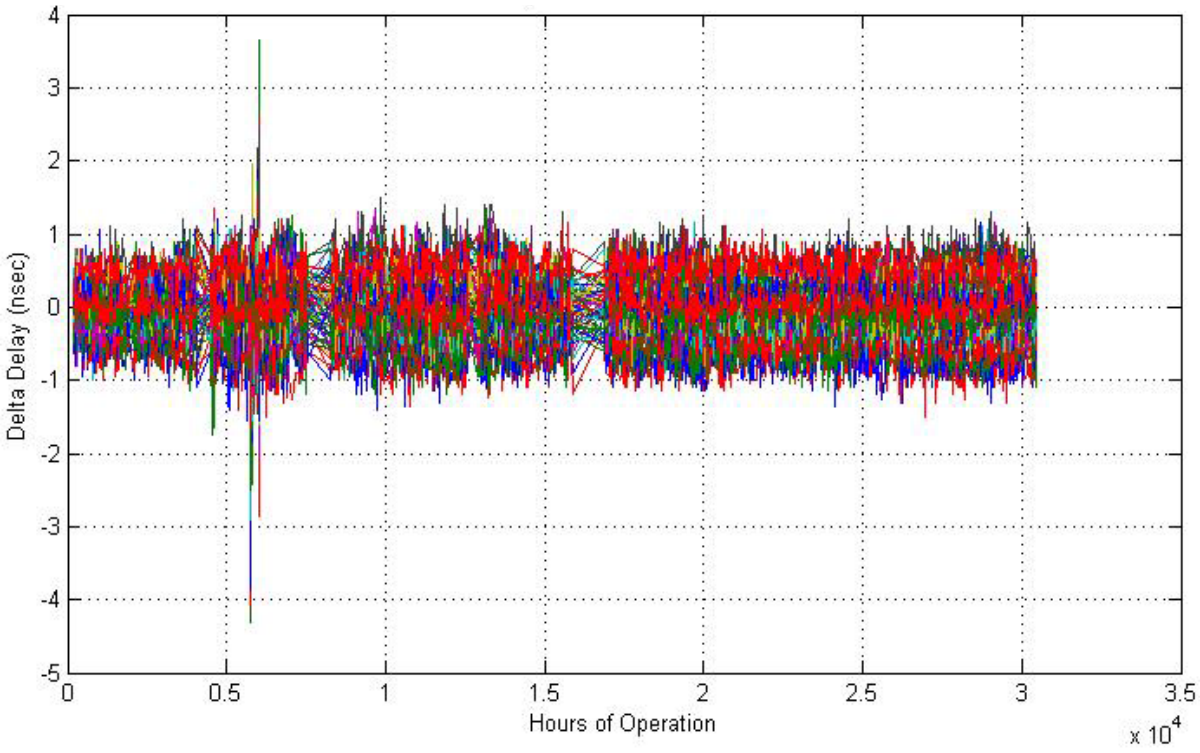

Figure 13. Delta delays for Part6 on board 4TC shown for every track. There are 80 traces overlaid here, one for every track being monitored (16 inverter tracks and 16 CMFE tracks in 4 modes each). The delta is the measured delay minus the base-line delay for that track. Note that the data between the hours 5766 and 6126 shows the deviation from the baseline outside of the ±2 ns limits. However, for the rest of the testing time, the data remains tightly aligned with the baseline values.

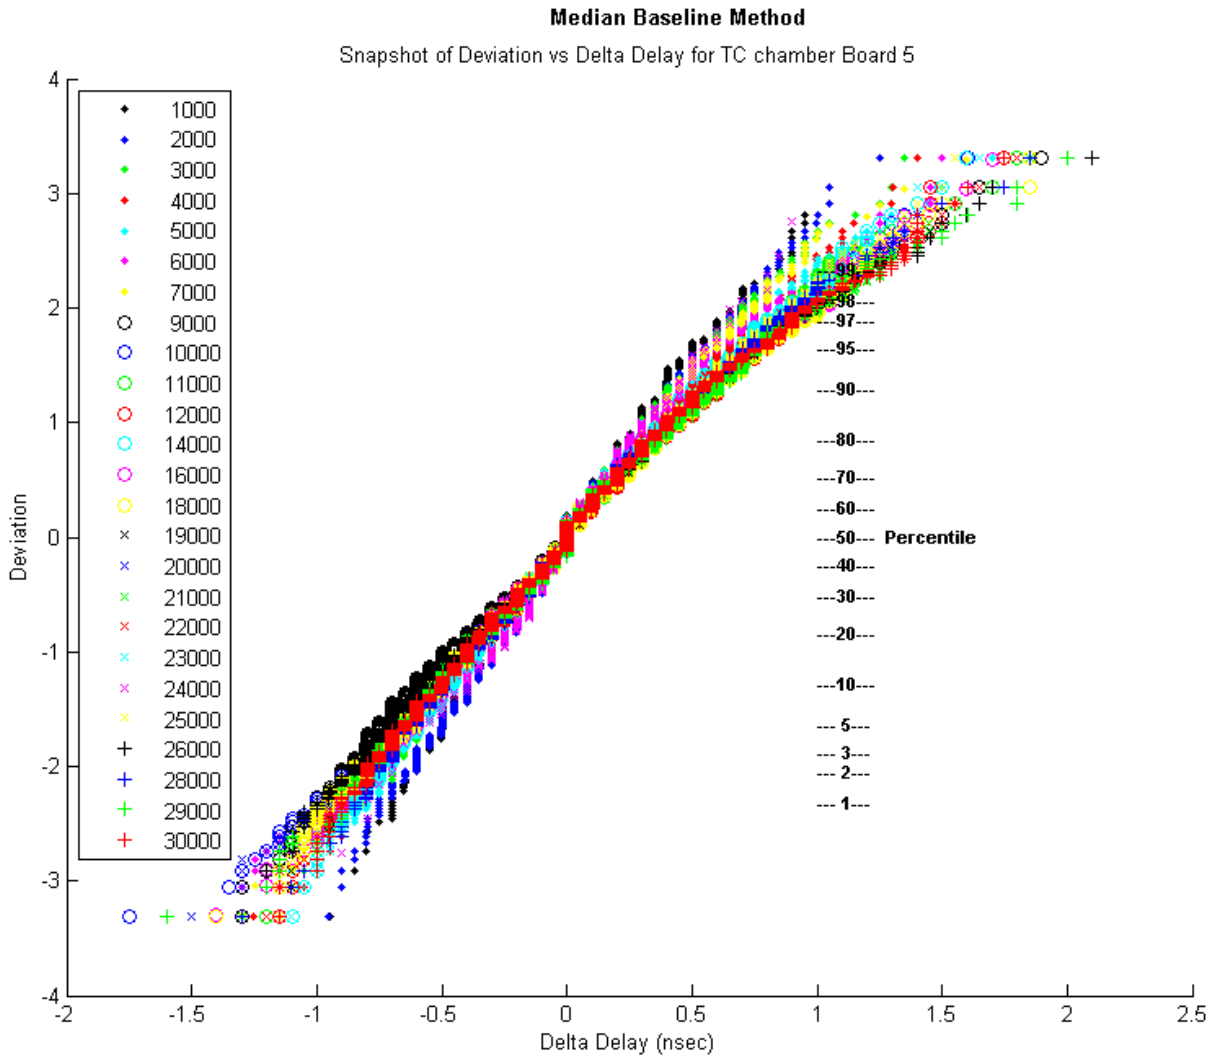

Figure 14. Data from board 5TC in the temperature cycling chamber. The one data point at 26000 hours that exceeds the + 2ns bound was for part 1. Data taken after this showed that the deviation was not permanent and therefore not an antifuse failure.

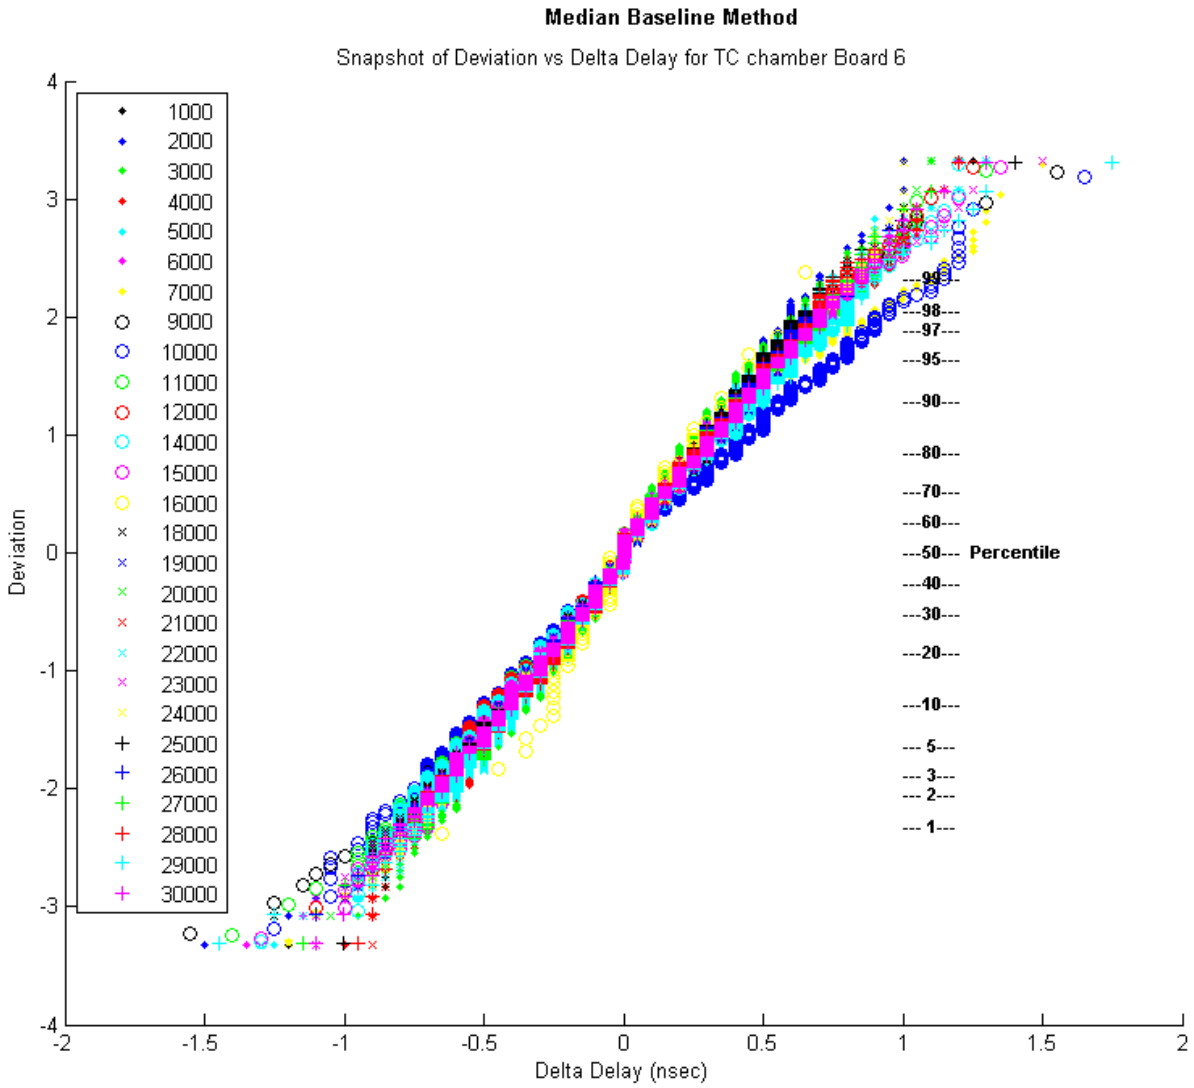

Figure 15. Data from board 6TC in the temperature cycling chamber.

The data shown in Figure 16 for board 7TC shows a shift in the data a few thousand hours into the testing. However, once shifted, the data remained consistently distributed around the new mean and still within the expected limits of ±2 ns of the original baseline. In Figure 17, the data show a similar fanning out of the delay distribution for board 8TC. For these boards, both 7TC and 8TC, these shifts are from changes in the temperature at which the data was taken. This temperature effect is even more pronounced in the data for board 9TC shown in Figure 18.

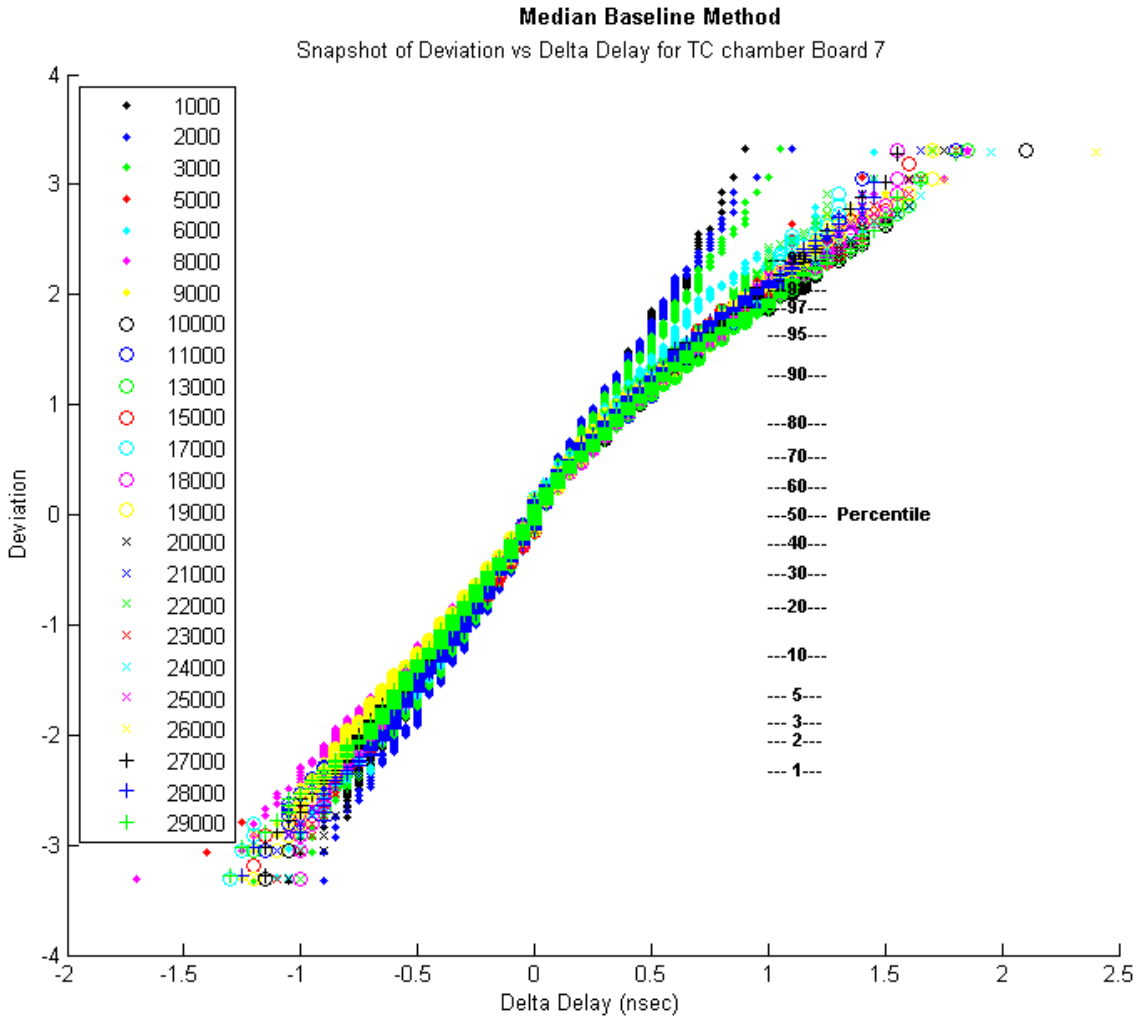

Figure 16. Data from Board 7TC in the temperature cycling chamber. There is a data point at 9000 hours for part 16 and another at 10000 hours for part 11 outside the +2 ns bound. Subsequent data show that these deviations were not lasting and therefore not due to antifuse failure.

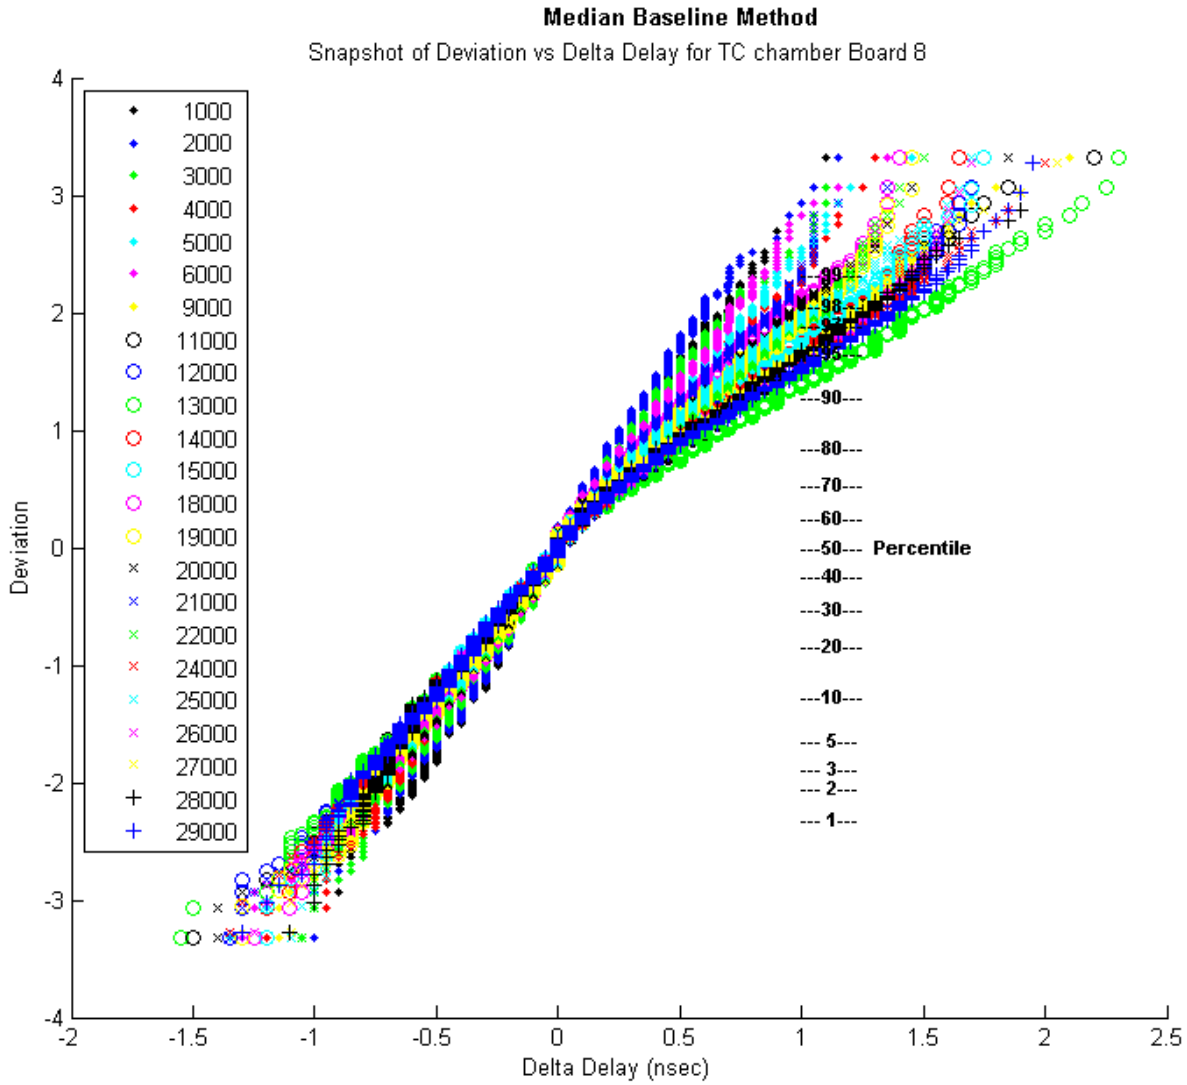

Figure 17. Data from Board 8TC in the temperature cycling chamber. There are 4 data points at 13000 hours and another at 9000 for part 6 that exceed +2 ns bound. Another point at 11000 for part 11 also exceeds this limit. These deviations were not lasting so are not indicative of antifuse failure.

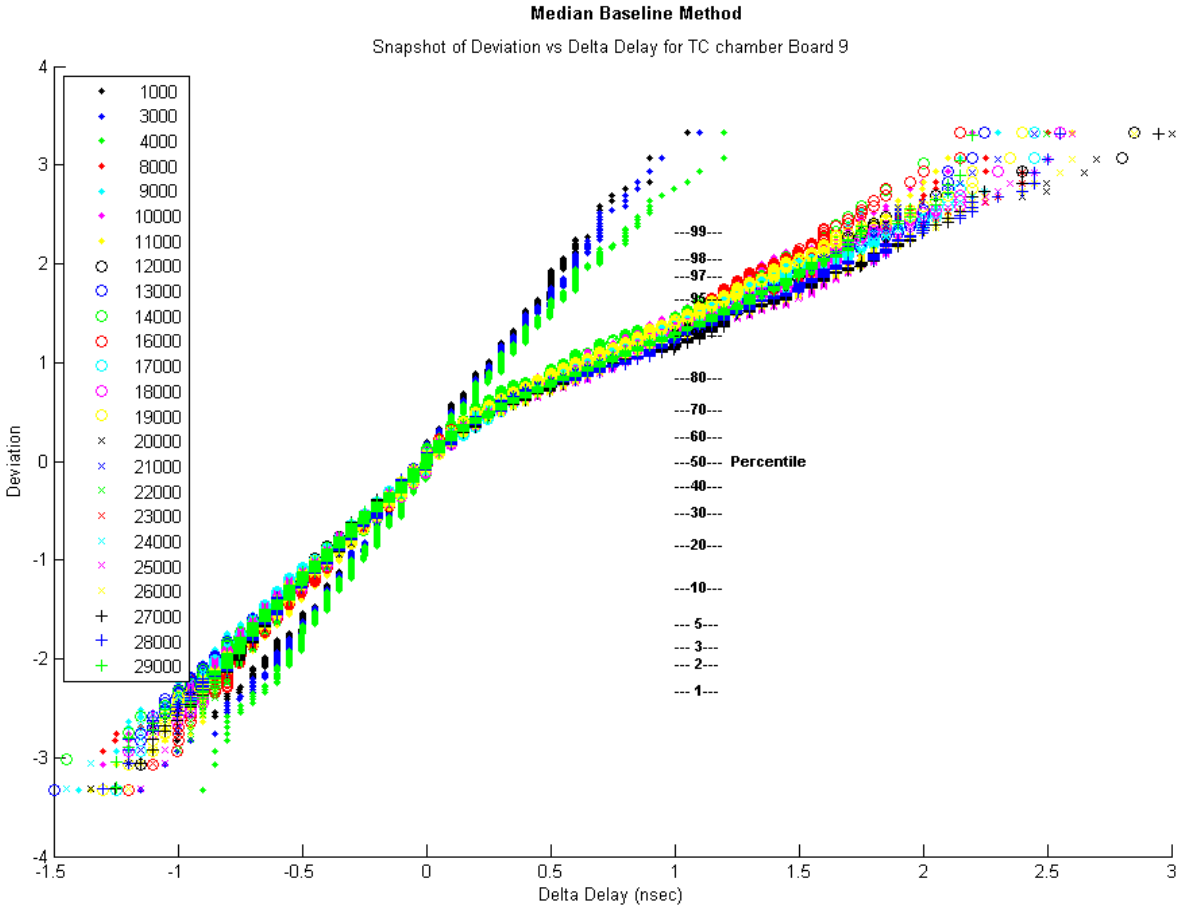

Figure 18. Data from Board 9TC in the temperature cycling chamber. The shift in the data was due to a change in the temperature of all devices at the hot plateau in this chamber. It was determined to not recalculate the baselines a few thousand hours in to the test but rather to monitor the data for any additional shifts from the tight grouping shown here. All of the data outside the +2 ns bound are from parts 11, 16, 17, and 18.

The parts on board 9TC all saw an increase in the temperature at which they were interrogated after about 5000 hours on test. This caused the dramatic shift in the data shown in Figure 18. In early 2008, we were experiencing difficulty with the reliability of the RS232 interface and the currents in the returns of the daisy-chained housekeeping circuit. During this time, we removed and replaced the complete set of test boards in the chamber and rearranged cabling. A large metal bar was added at this time to improve the system ground by shorting out any currents in the return lines to the power supplies. This changed the overall temperature profile of the chamber, leading to higher temperatures on boards 9TC and 10TC. Similarly, the data for board 10TC shown in Figure 19 showed a shift from the baseline. Monitoring over the next few weeks showed that while the data was different from its original baseline, it remained constant with time consistent with the tight grouping of the data in Figure 18. It was determined that the baseline would not be recalculated but that the data would be monitored for additional changes. None occurred.

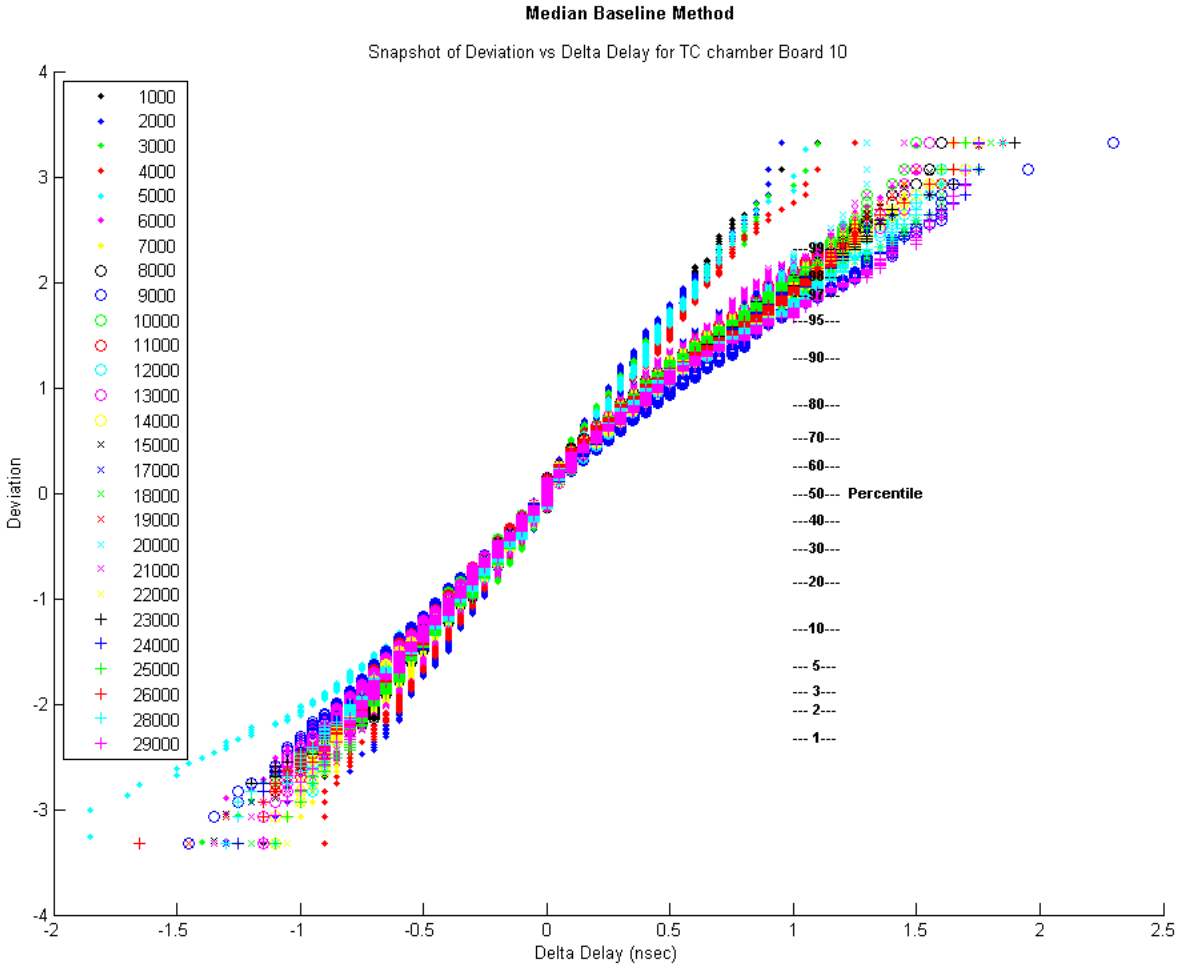

Figure 19. Data from board 10TC in the temperature cycling chamber. Part 16 had a data point outside the +2 ns bound at 9000 hours. This was not a permanent deviation so not indicative of an antifuse failure.

The data shown in Figure 20 is from board 1C in the cold chamber. In this chamber, the effect of local heating for the devices numbered 16, 11, and 6 on the boards begin to show a clear pattern. The boards at the ends of the board stack, boards 1C and 2C, 11C, 12C, 13C, and 14C do not have as dramatic a shift for these devices as the boards in the middle of the stack up, boards 3C through 10C. In Figures 22 through 36, the data is plotted for all the devices. But for these middle boards, the data is then shown again without the data for parts 16, 11, and 6 on the boards. A close look at the data shows the same tight range of delta delays distributed around a different baseline as was seen in the data on board 9TC in the cycling chamber. Figures 21−40 show all the data for boards 2C-14C in the cold chamber.

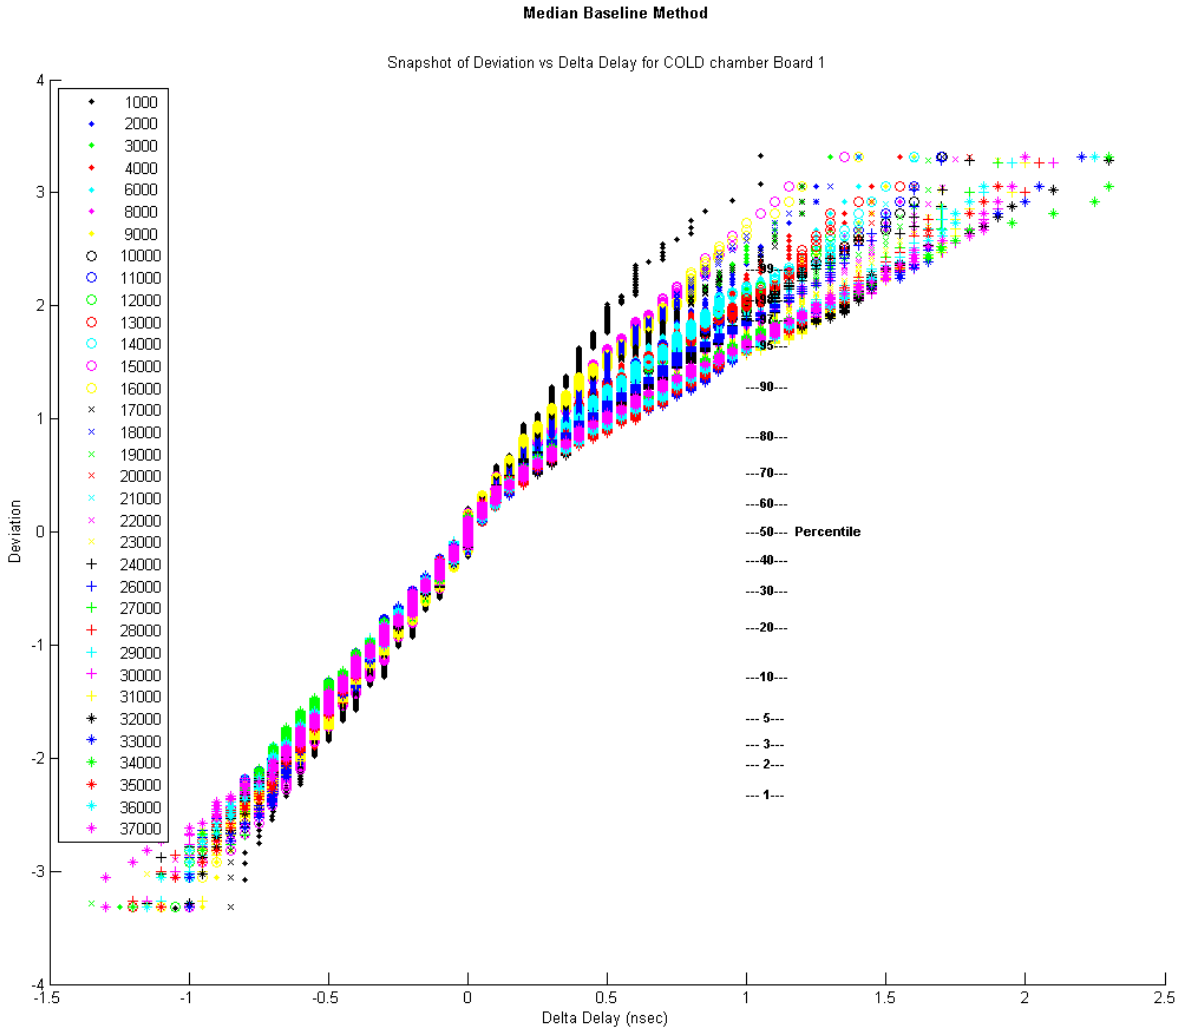

Figure 20. Data from board 1C in the cold chamber. The 11 points that exceed the +2 ns bound are for parts 16, 11, and 6 that experienced in increase in local temperature between 31000 and 35000 hours. Because the deviations were not permanent, there is no evidence of antifuse failure.

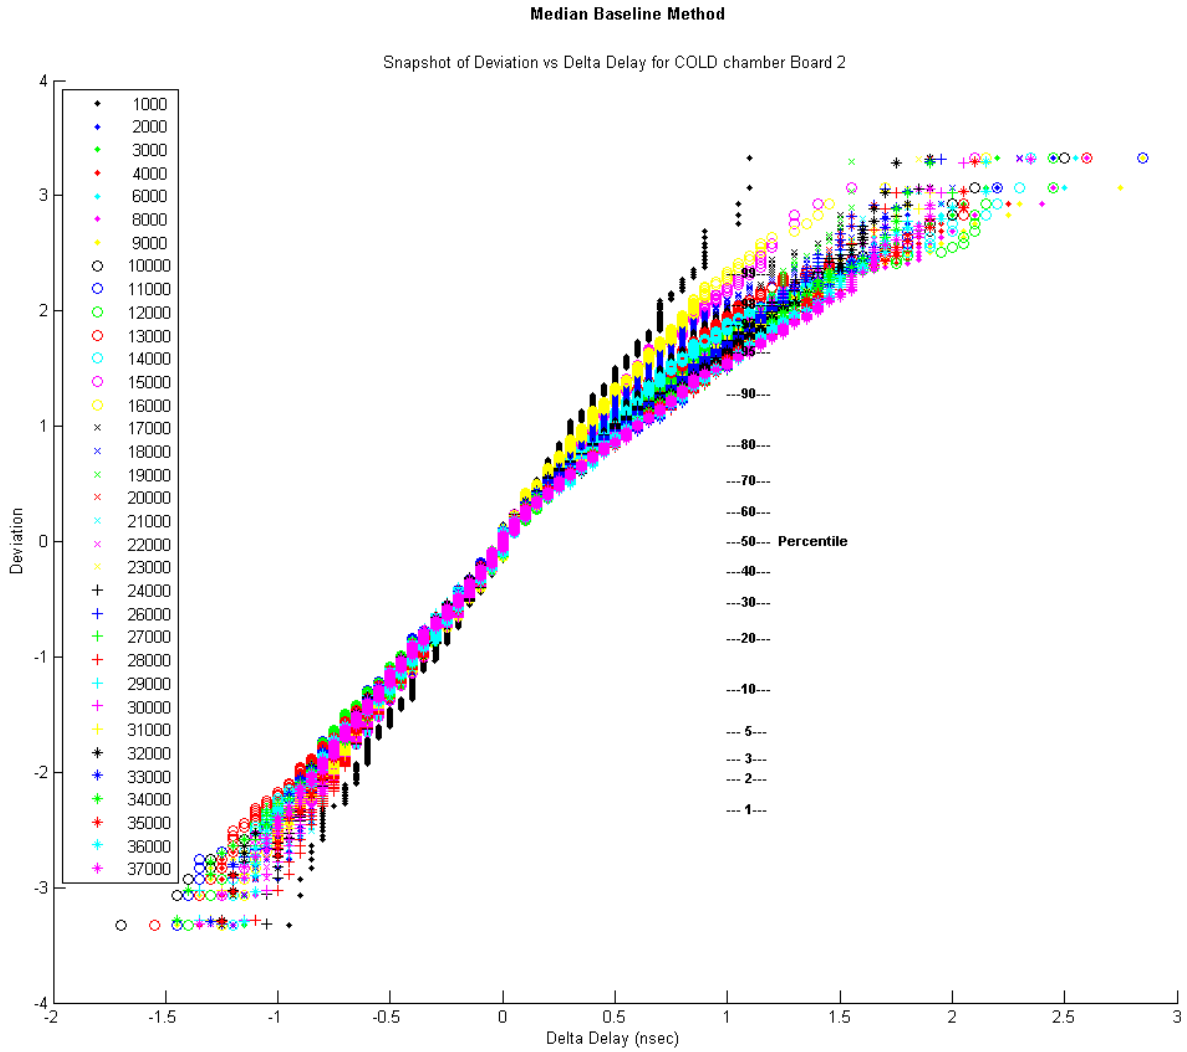

Figure 21. Data from board 2C in cold chamber. Cold chamber held the ambient temperature at −40C. The data is showing the overall temperature shift the parts on the upper, back corner of the board, parts 11 and 16 in particular, experienced in this chamber. The deviations in the data were not permanent and so no antifuse failures were evident.

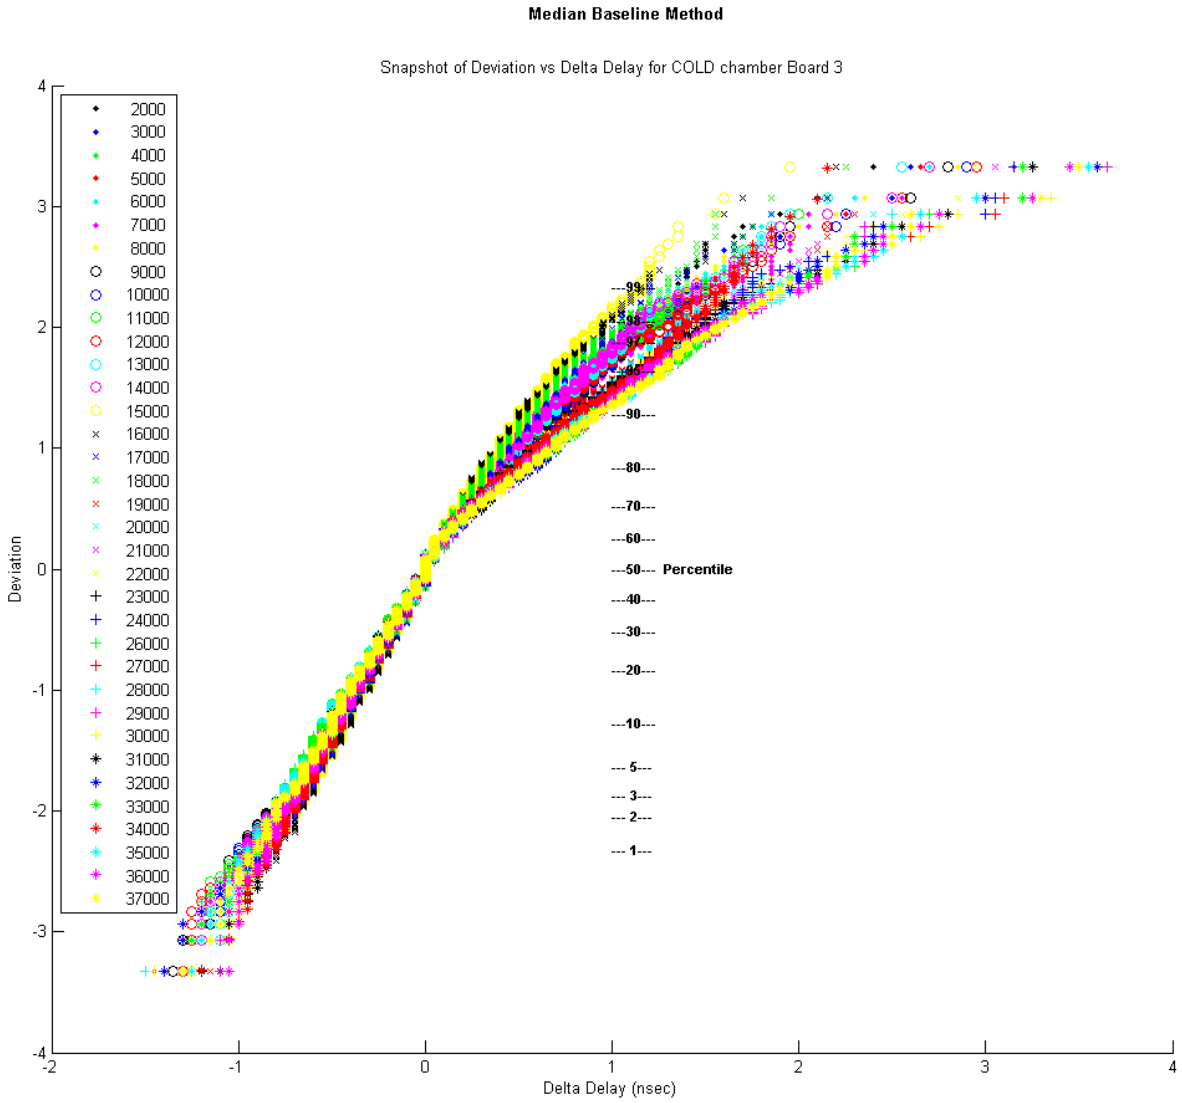

Figure 22. Data from board 3C in the cold chamber. The data shows a greater shift for the local heating of the upper corner devices. This data is replotted in the Figure 23 without devices 16, 11, and 6 to show the effect of the temperature shift for these parts.

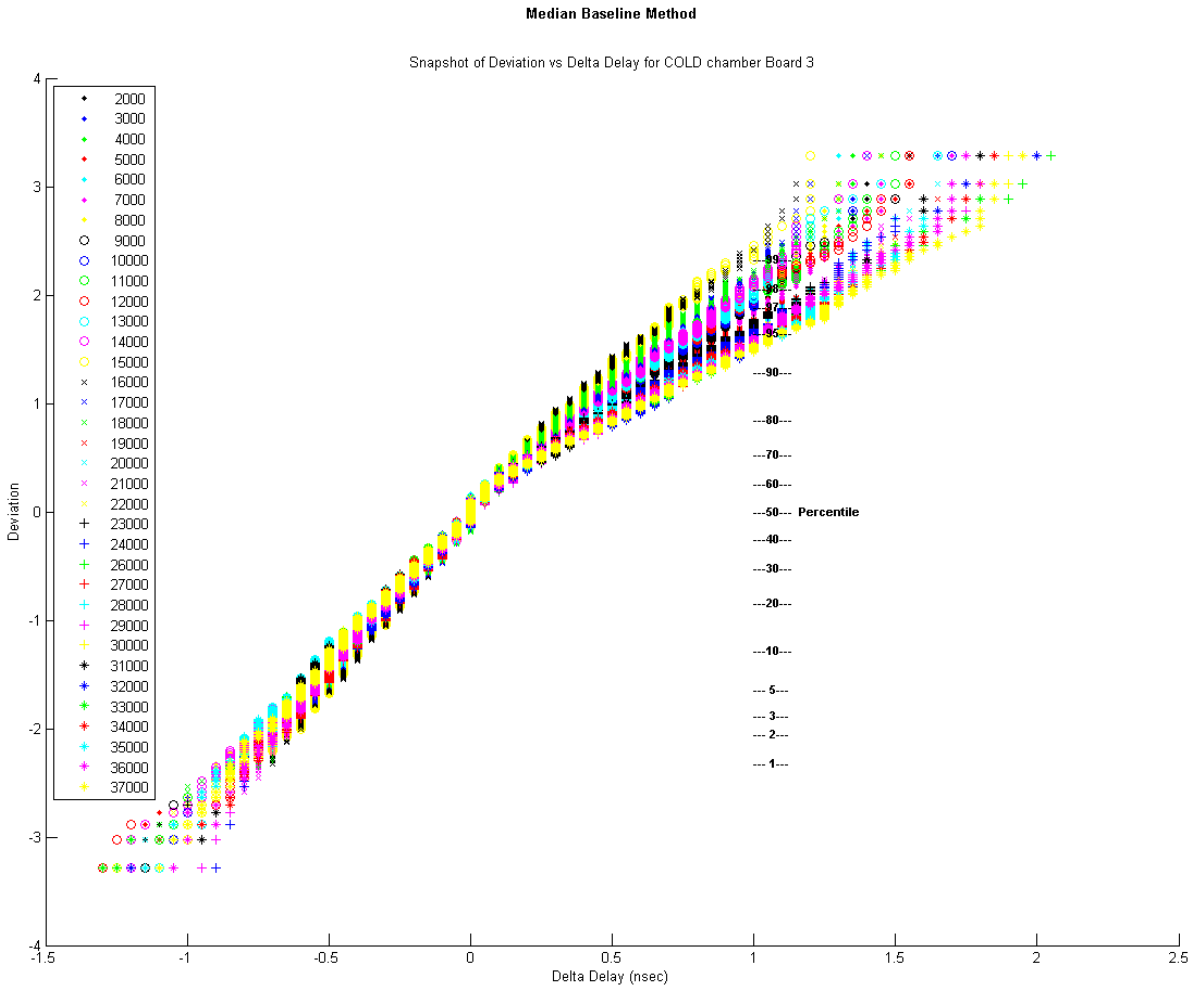

Figure 23. Data from board 3C in the cold chamber with all parts plotted except for 6, 11, and 16. There is one data point outside of the +2 ns bound for part 13 at 26000 hours, but this deviation does not persist so does not indicate an antifuse failure.

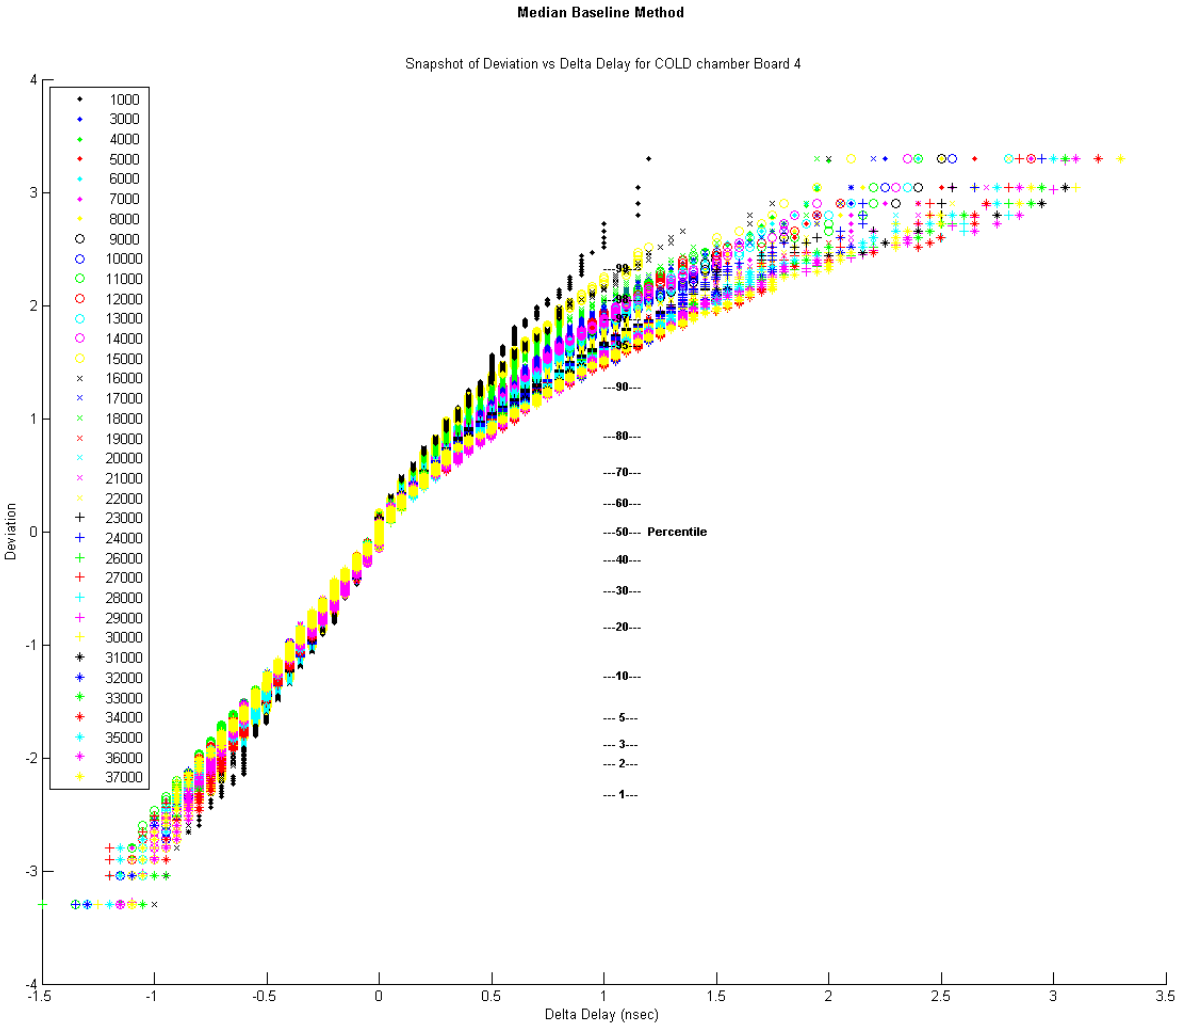

Figure 24. Data from board 4C in the cold chamber. The shift in the data for the parts locally warmed on this board is shown in the spread of the distributions. The data outside the +2 ns bound is essentially all from parts 6, 11, and 16.

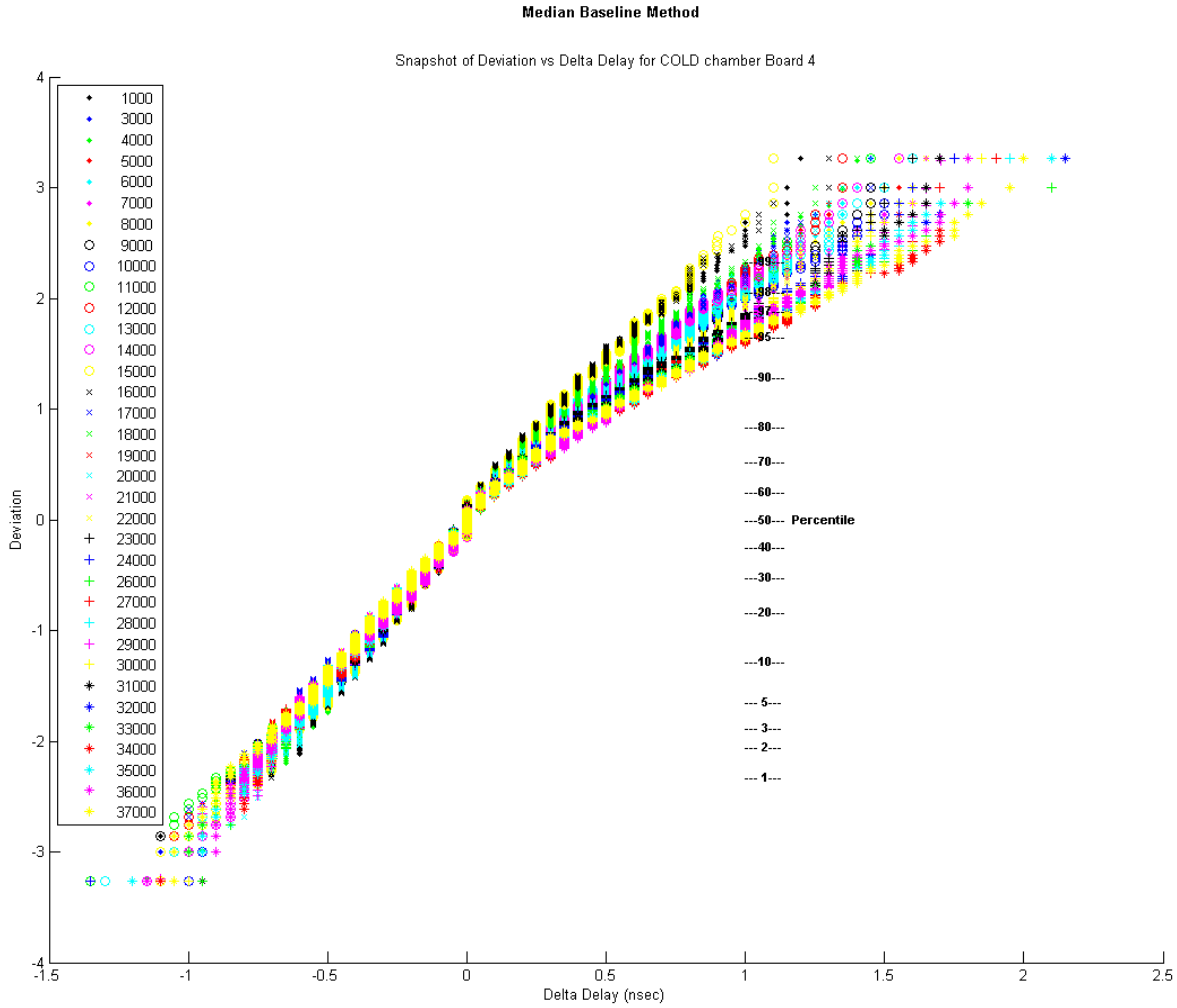

Figure 25. Data from board 4C in the cold chamber with all parts plotted except for 11, and 16. The three points that exceed the +2 ns bound are for part 6 at 26000 and 28000 hours and part 7 at 32000 hours. These deviations do not persist so do not indicate antifuse failure.

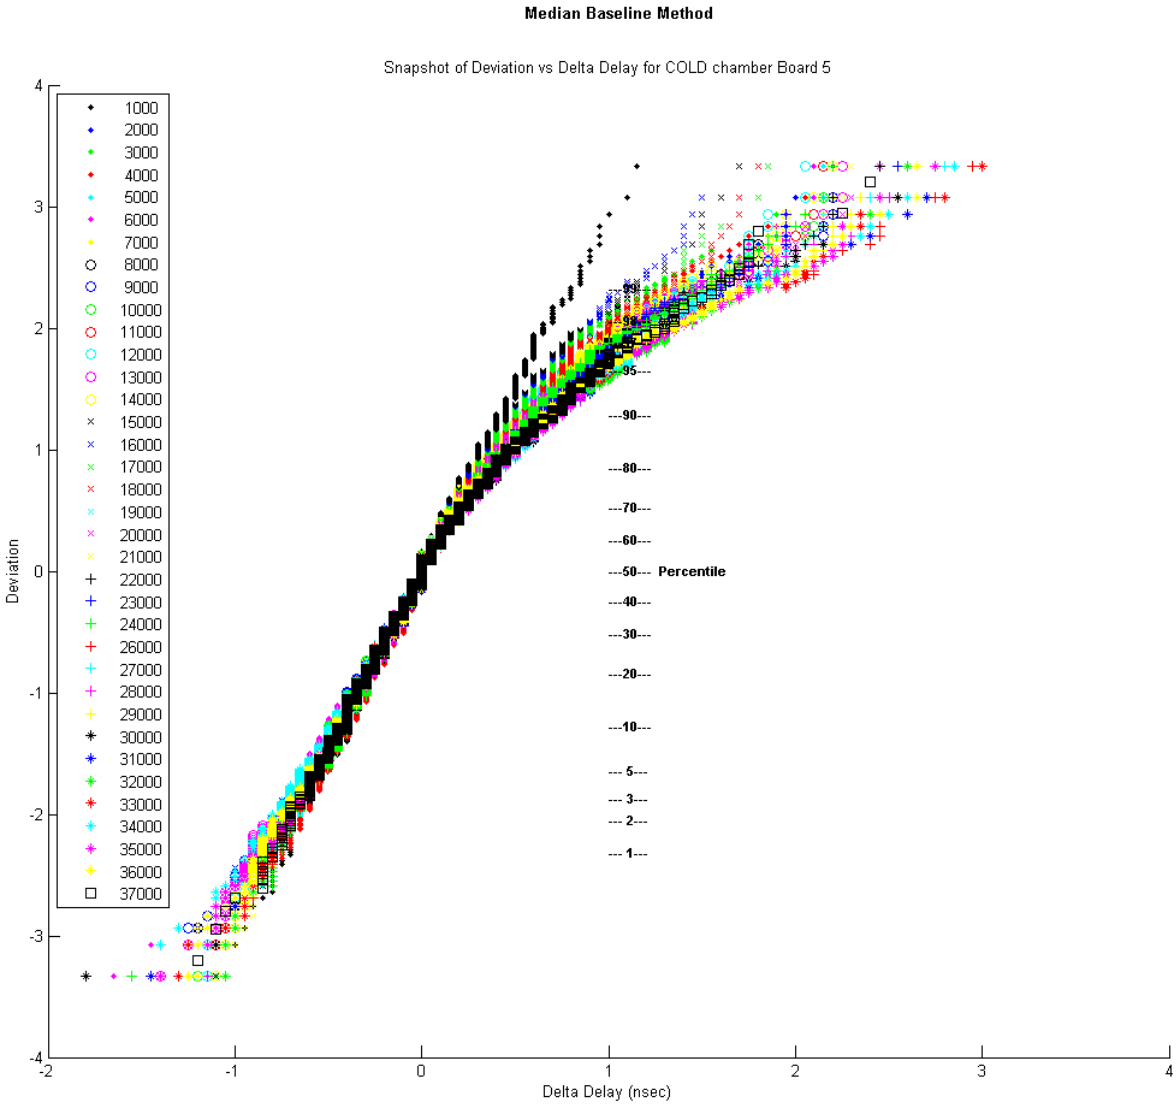

Figure 26. Data from board 5C in the cold chamber. The shift in the data for the parts locally warmed on this board is shown in the spread of the distributions. The data outside the +2 ns bound is essentially all from parts 6, 11, and 16.

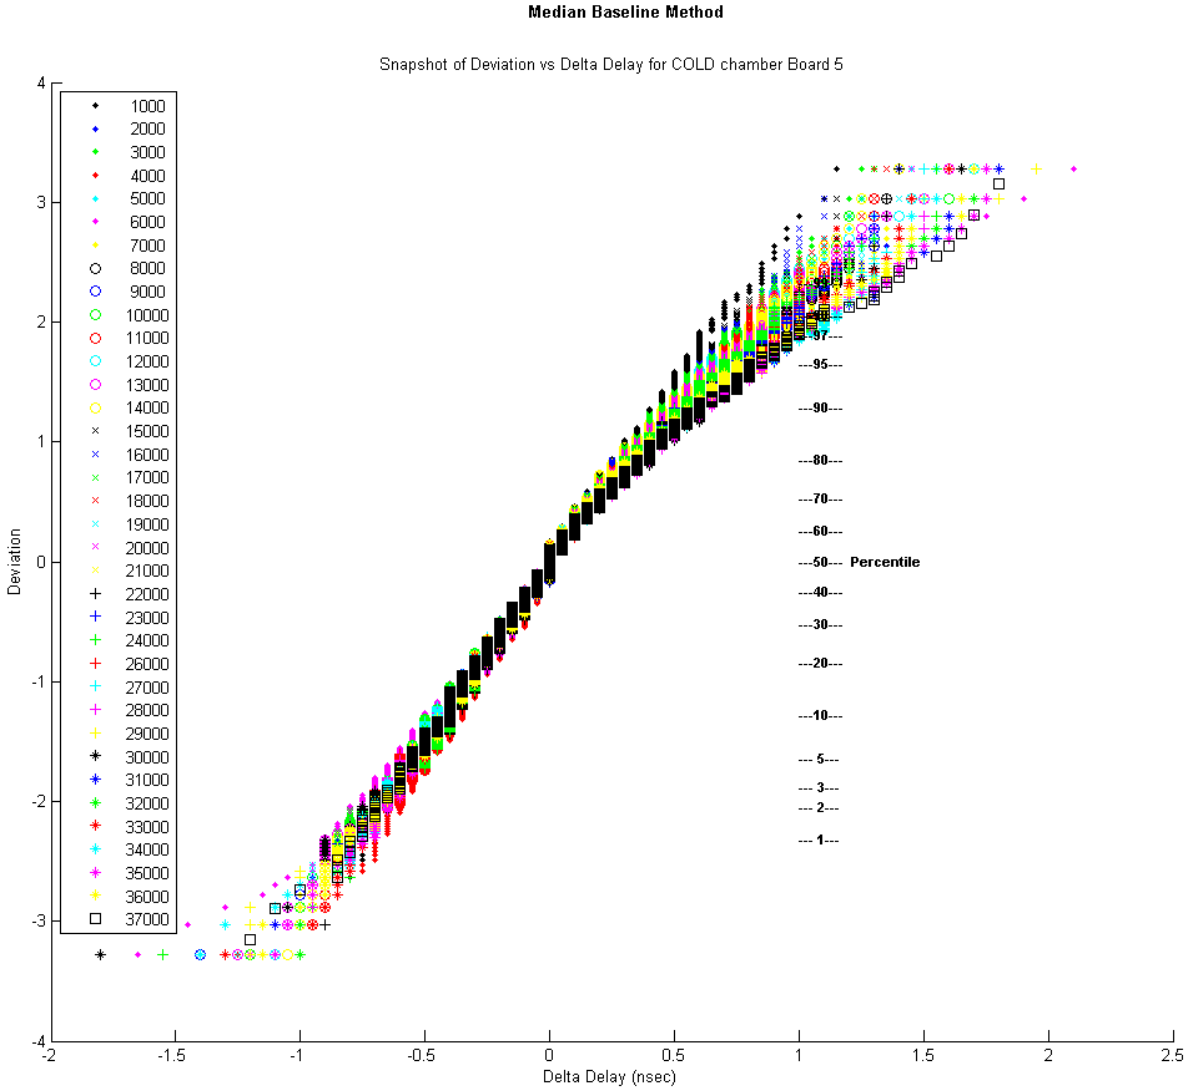

Figure 27. Data from board 5C in the cold chamber with all parts plotted except for 6, 11, and 16. The single point that exceeds the +2 ns bound is from part 4 at 4000 hours. Because the deviation was not persistent, there is no indication of antifuse failure.

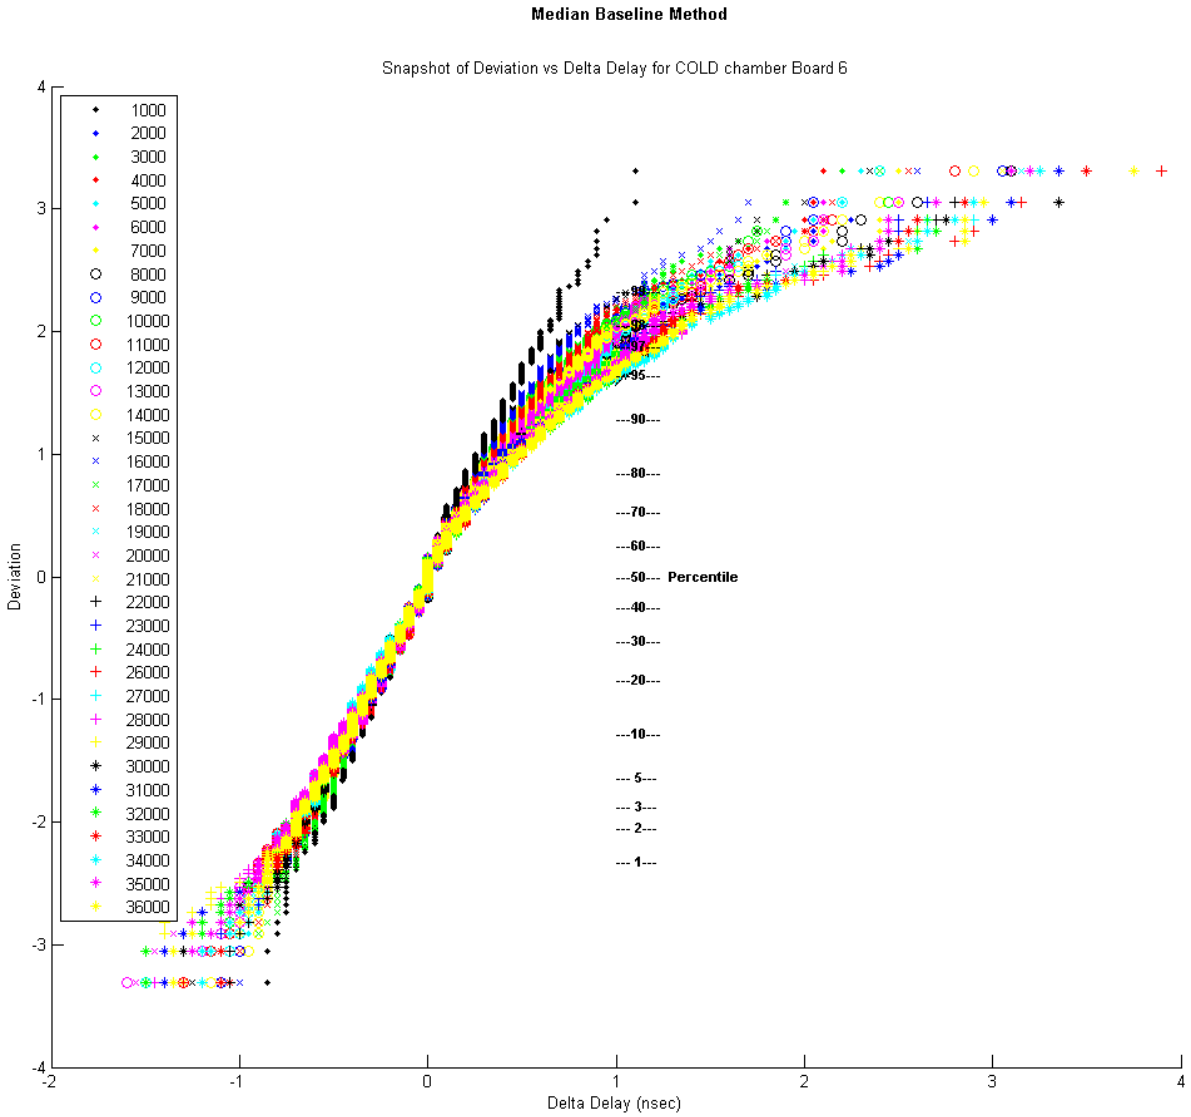

Figure 28. Data from board 6C in the cold chamber. The shift in the data for the parts locally warmed on this board is shown in the spread of the distributions. The data outside the +2 ns bound is all from part 16 as can be seen in the next figure that plots the probability distribution for all parts except for 16.

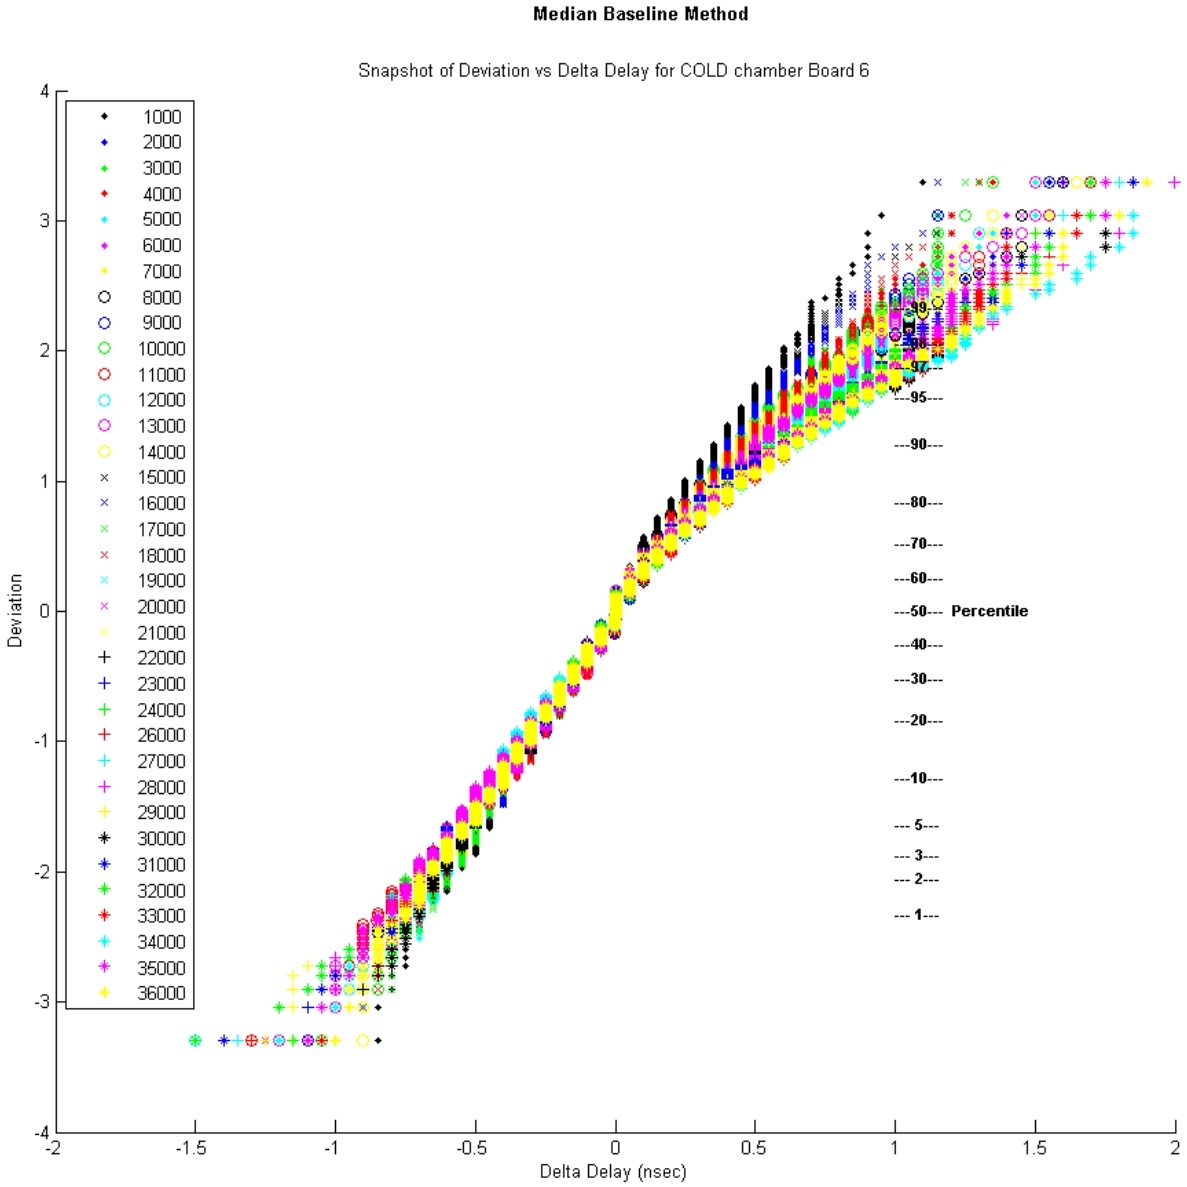

Figure 29. Data from board 6C in the cold chamber with all parts plotted except part 16.

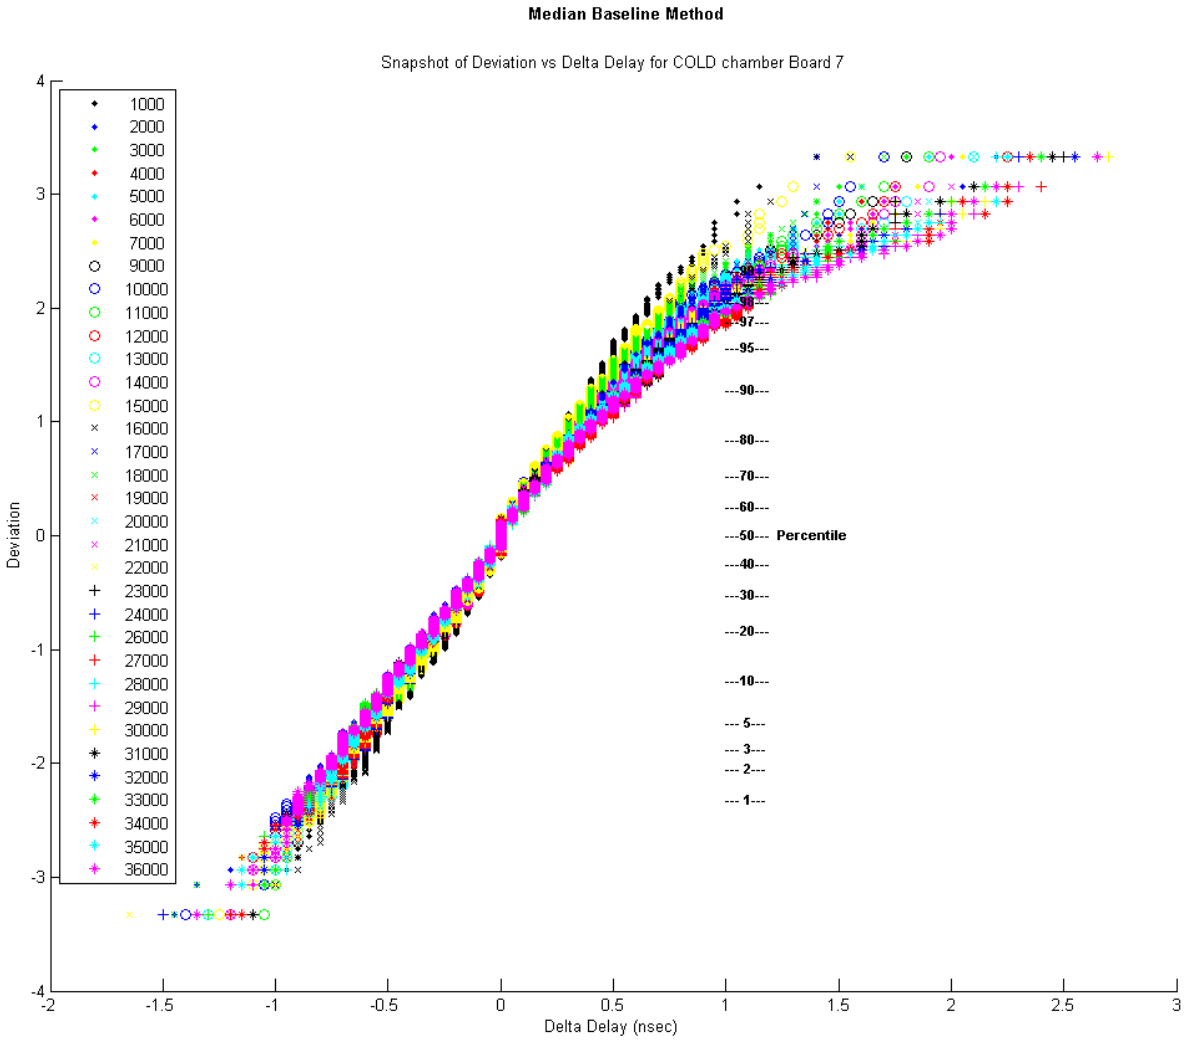

Figure 30. Data from board 7C in the cold chamber. The data that is outside the +2 ns bound is from parts 11 and 16. This can be seen in Figure 31 when part 16 data are removed from the data distribution plot.

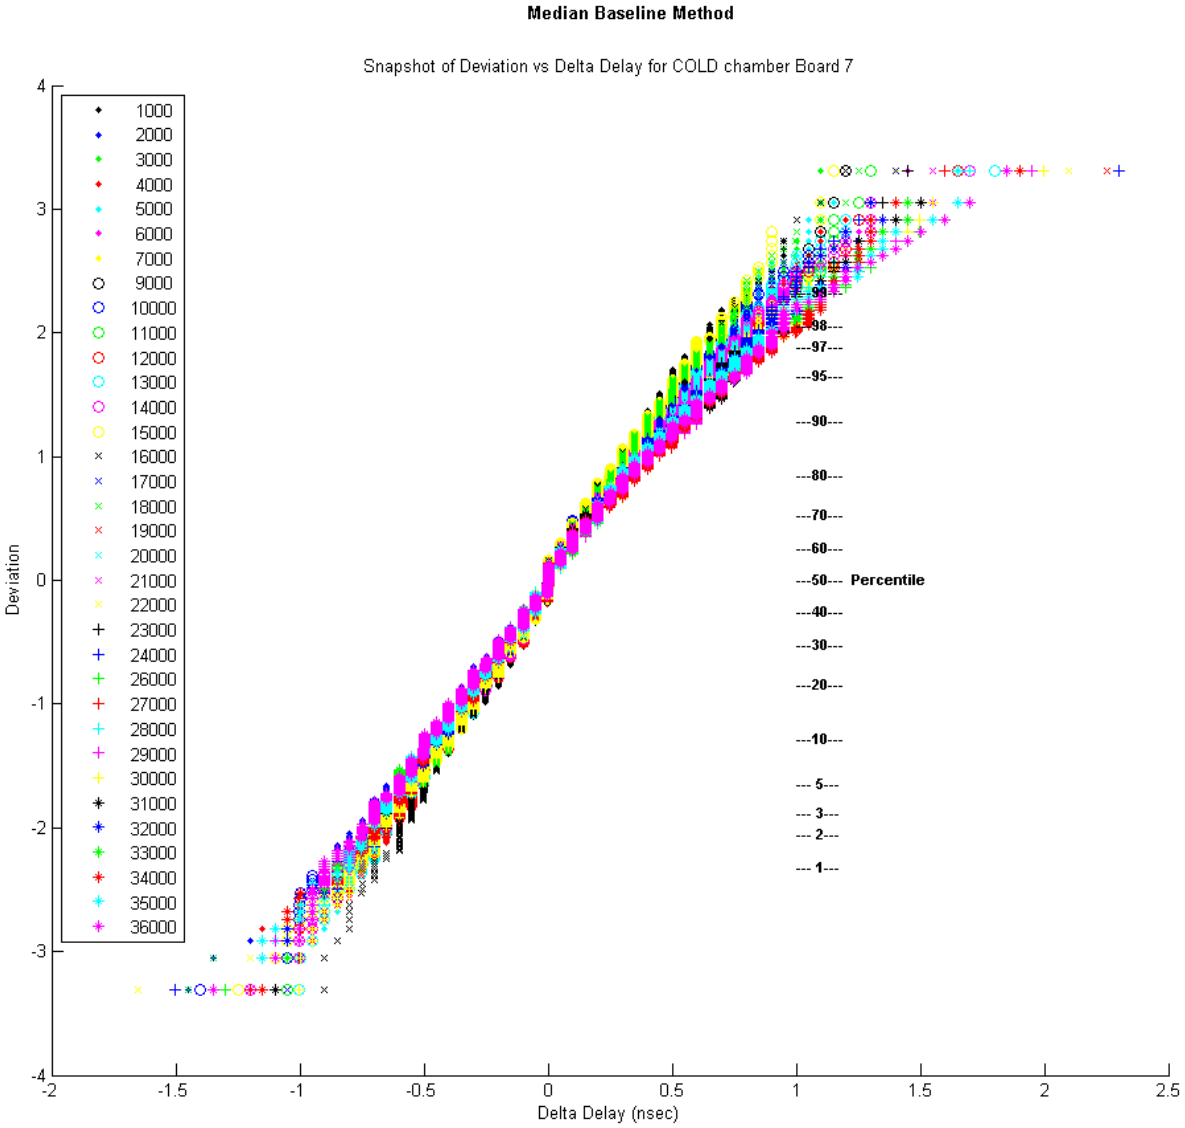

Figure 31. Data from board 7C in the cold chamber with all parts plotted except for part 16. The 3 data points still outside the positive bound of 2 ns are from part 11 at 7000 hours, 21000 hours, and 24000 hours. Because these deviations are not persistent, there is no evidence of antifuse failure.

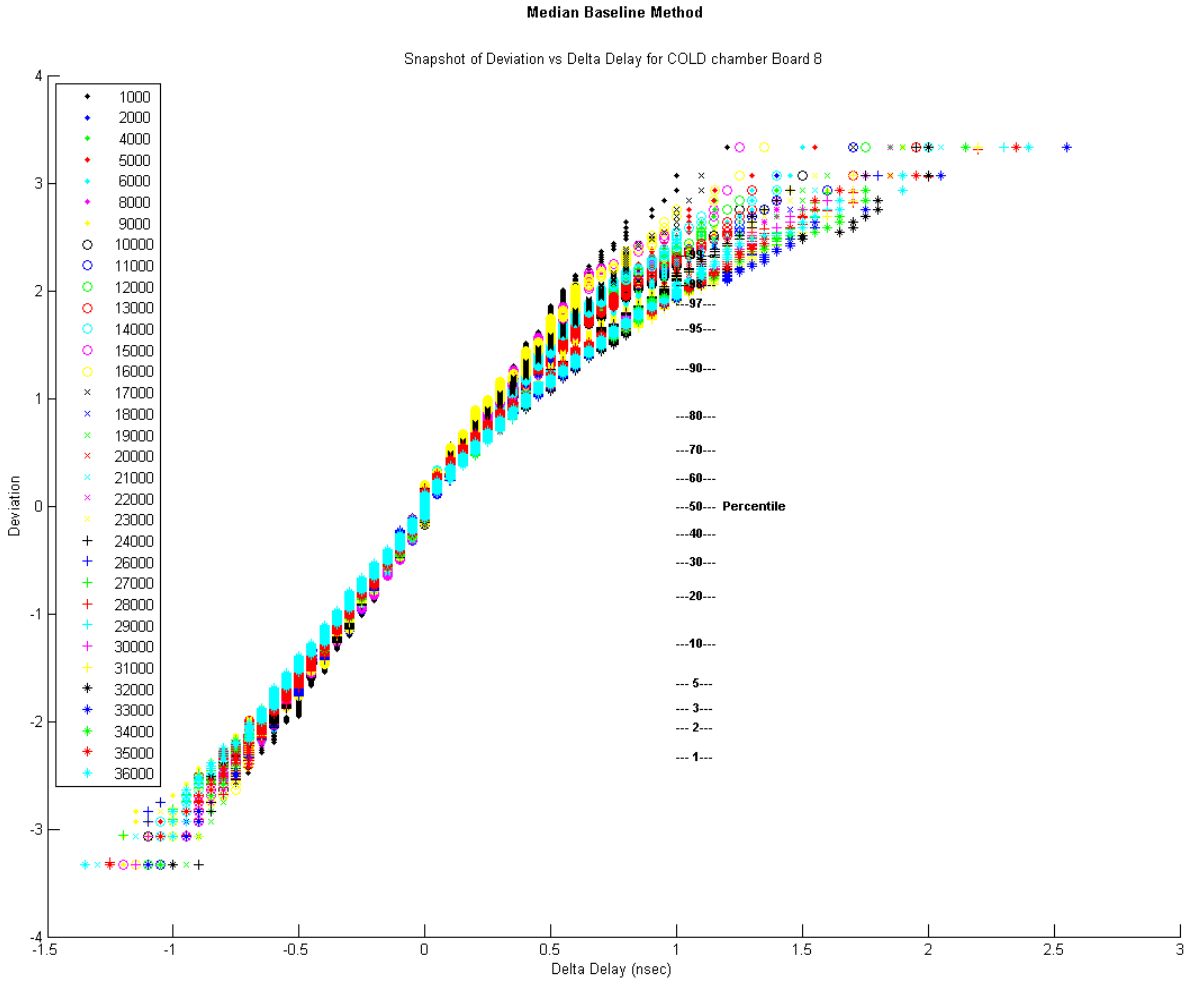

Figure 32. Data from board 8C in the cold chamber. The data that is outside the positive bound of 2 ns are all from part 16.

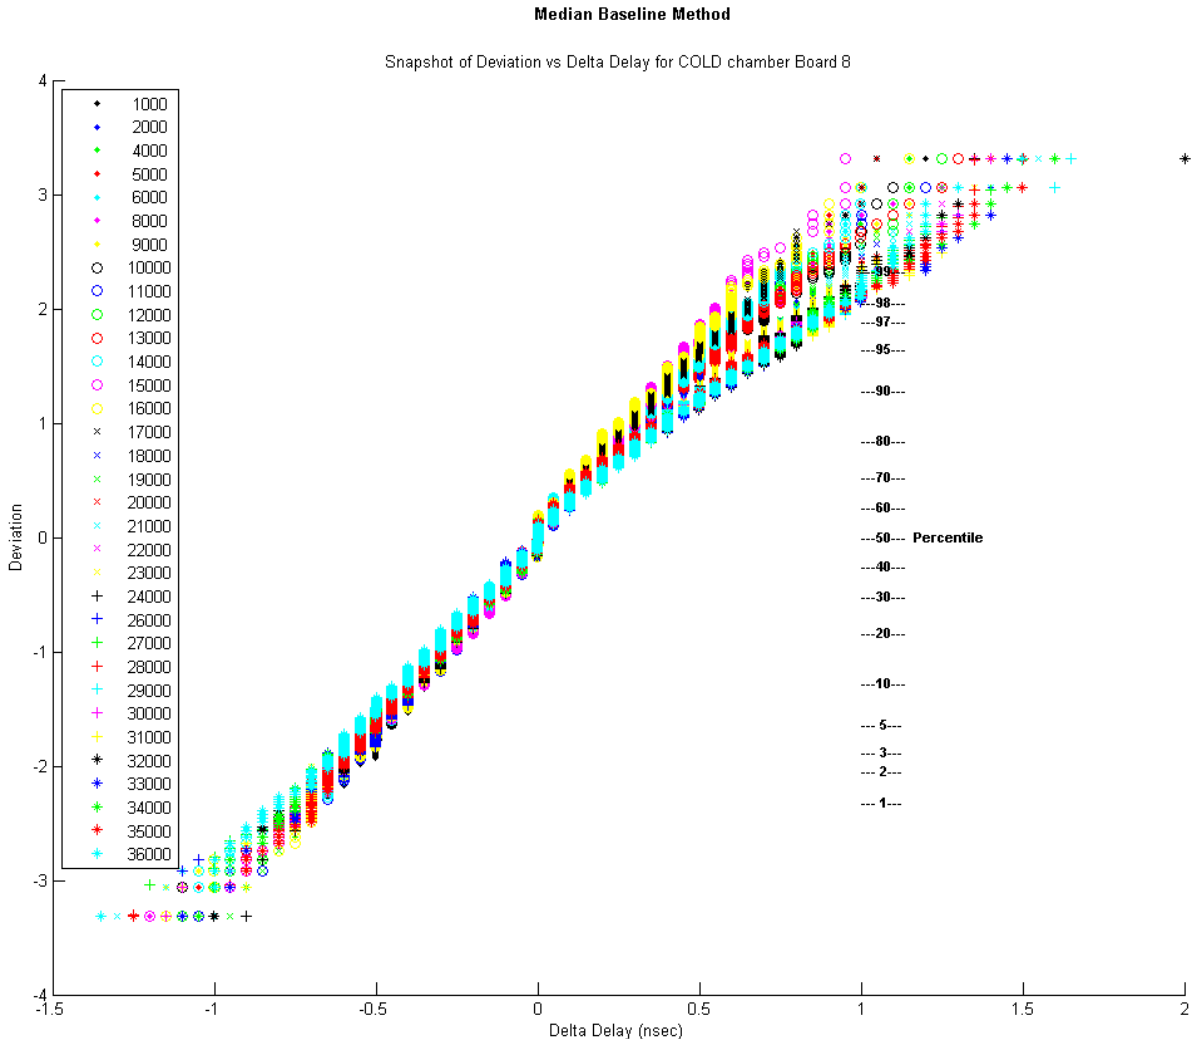

Figure 33. Data from board 8C in the cold chamber with all parts plotted except for part 16.

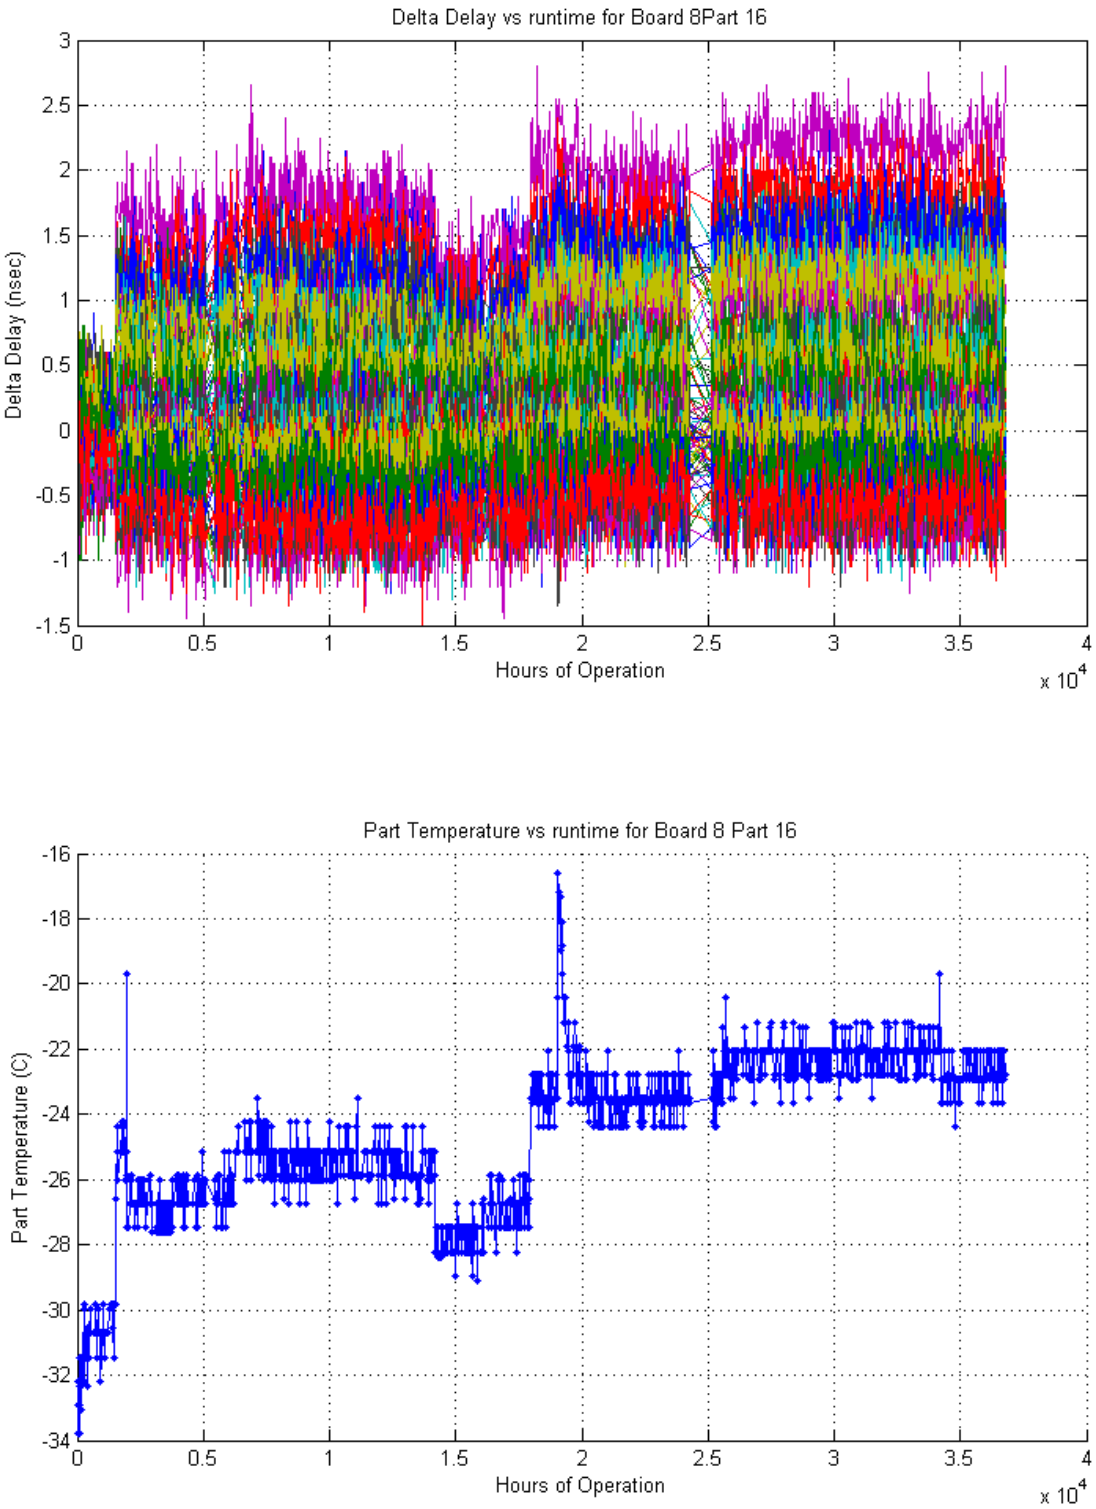

Figure 34. Data from part 16 on board 8C in the cold chamber. The chart at top shows the delta delay from the baseline measurement for every track on part 16. The spread in the data is correlated to the ambient temperature for this part in the cold chamber, shown in the bottom chart.

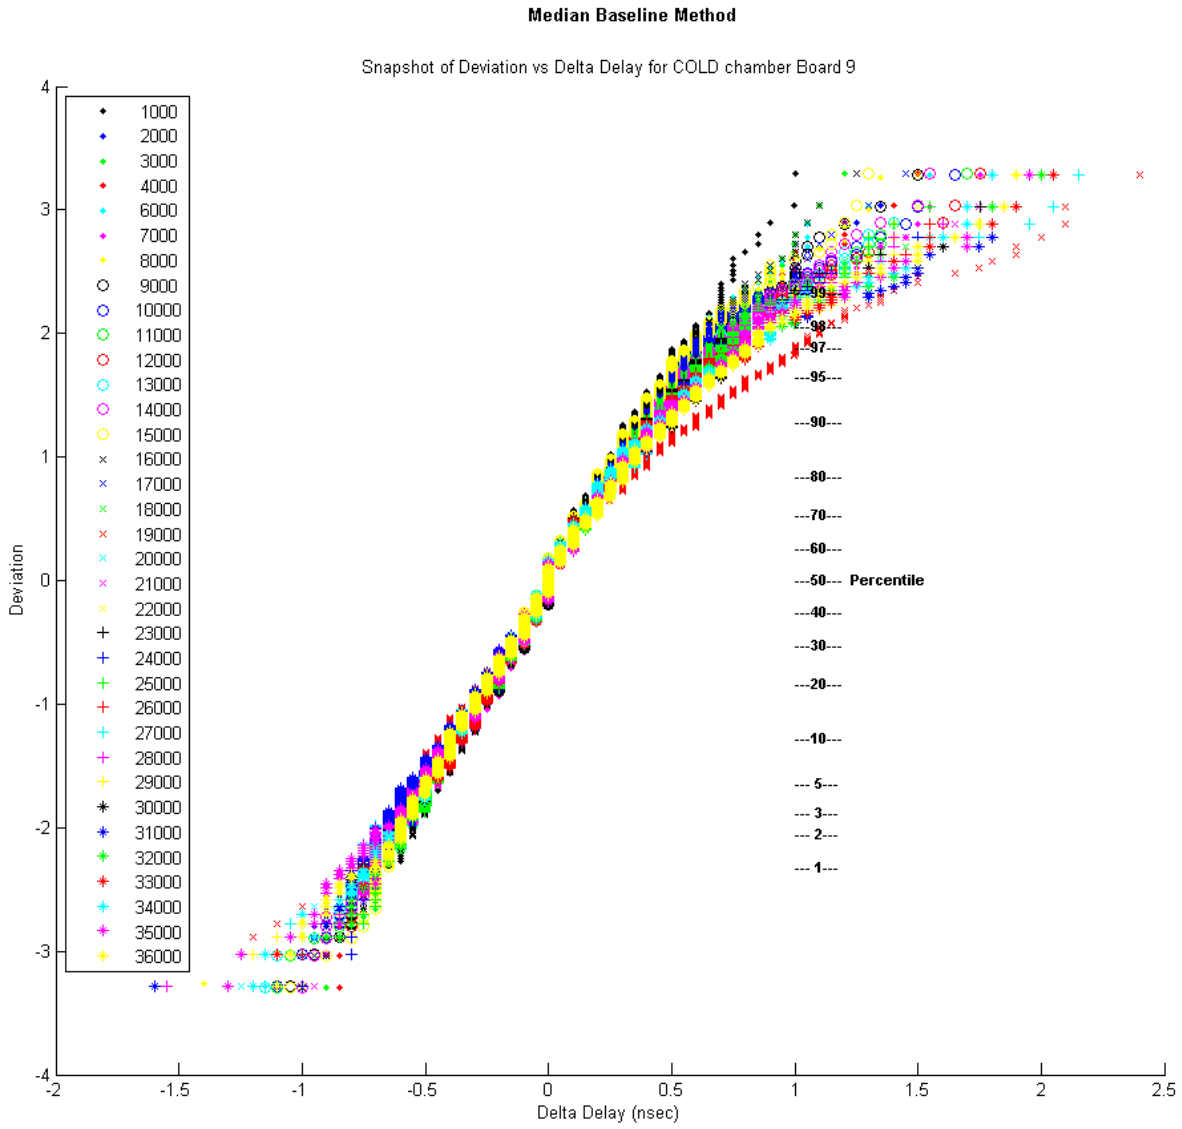

Figure 35. Data from board 9C in the cold chamber. There are 6 data points outside the +2 ns bound that are all from part 16. Three occurred at 19000 hours, 2 at 27000 hours, and one at 33000 hours. The deviations were not persistent and so do not indicate an antifuse failure.

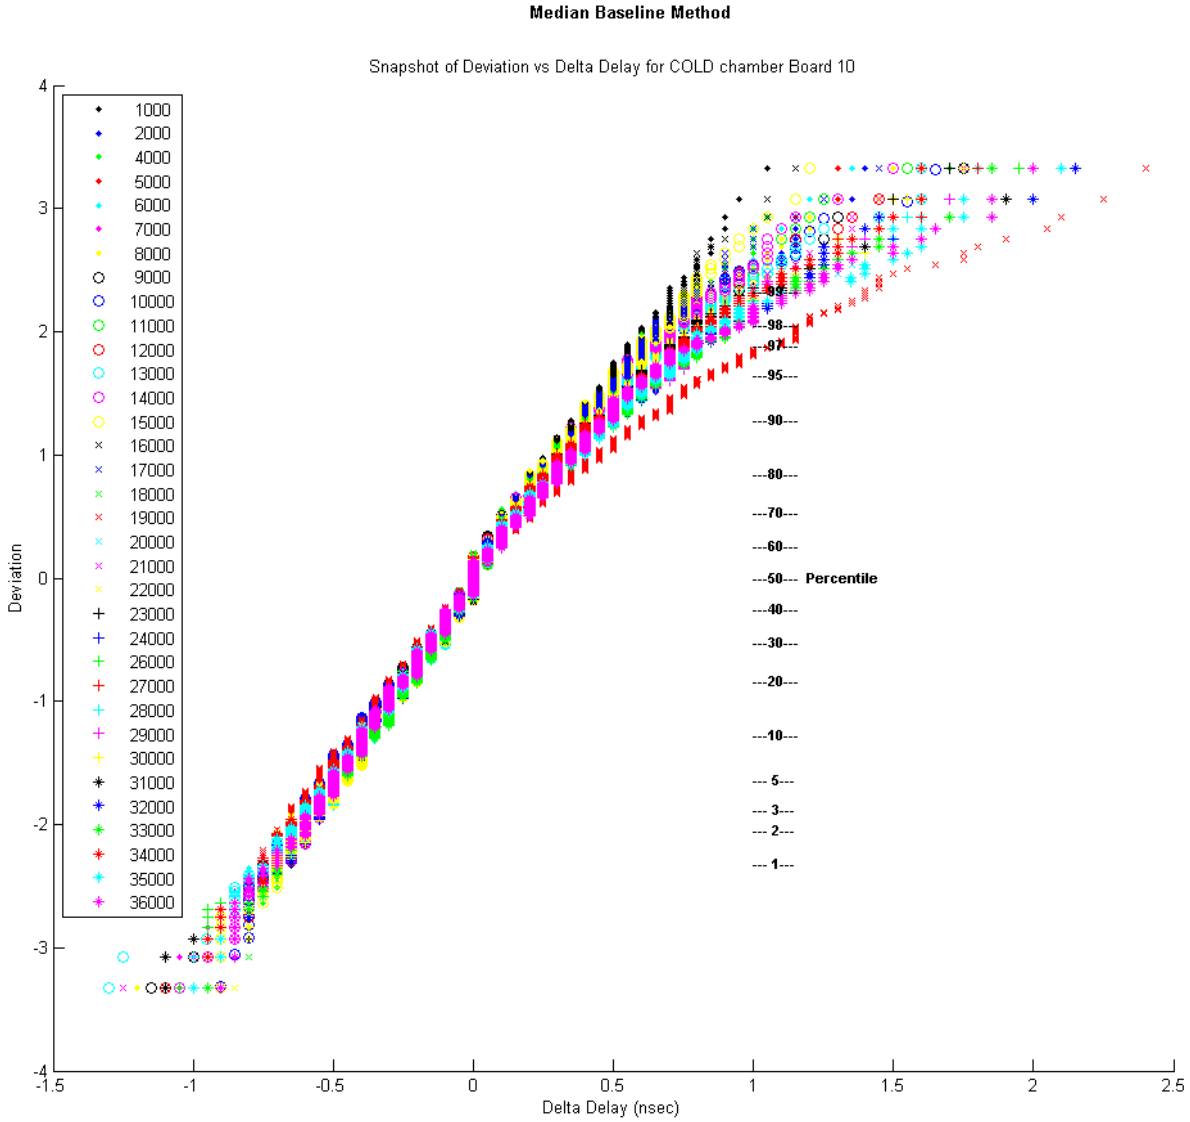

Figure 36. Data from board 10C in the cold chamber. The six data points that exceed the +2 ns bound are from parts 11 and 16. One point for part 11 and 3 for part 16 occurred at 19000 hours. The other 2 points were from part 16 at 32000 and 35000 hours. None of the deviations were persistent so no antifuse failure is suspected.

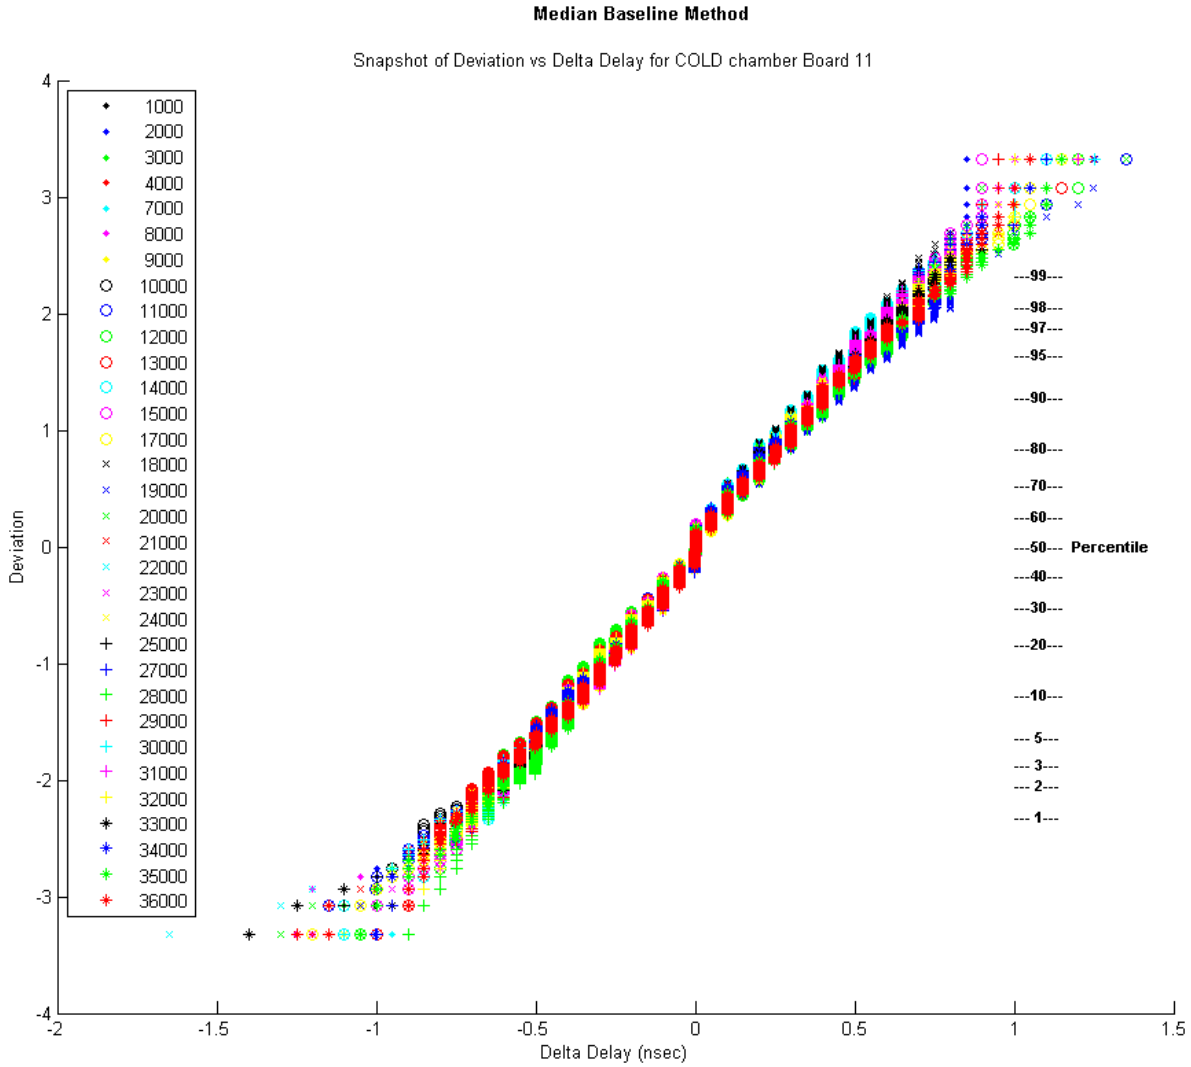

Figure 37. Data from board 11C in the cold chamber in a distribution plot of the deviations from the baseline delay measurements.

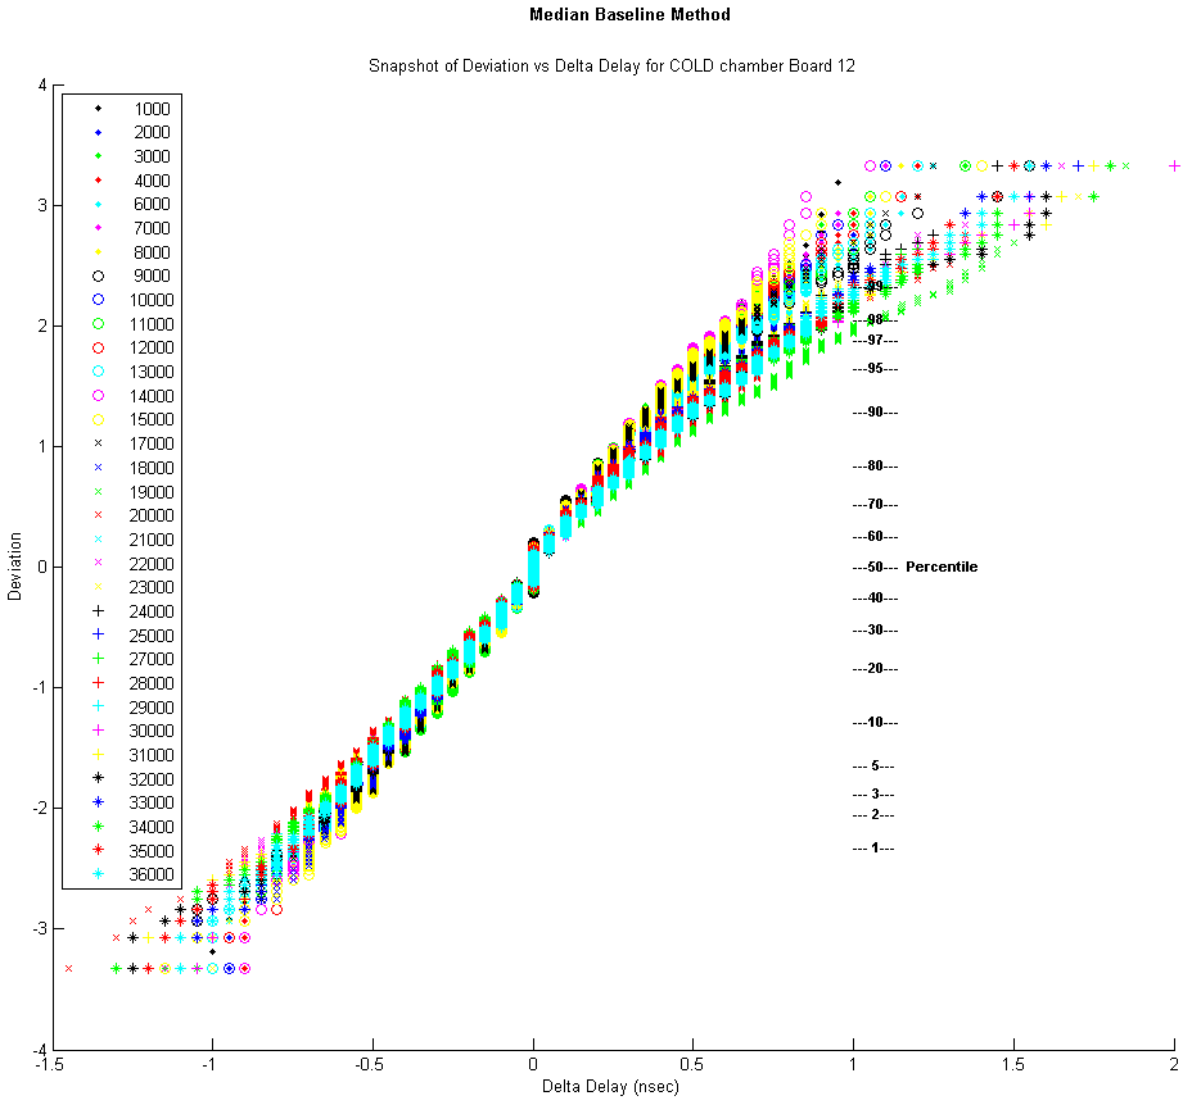

Figure 38. Data from board 12C in the cold chamber.

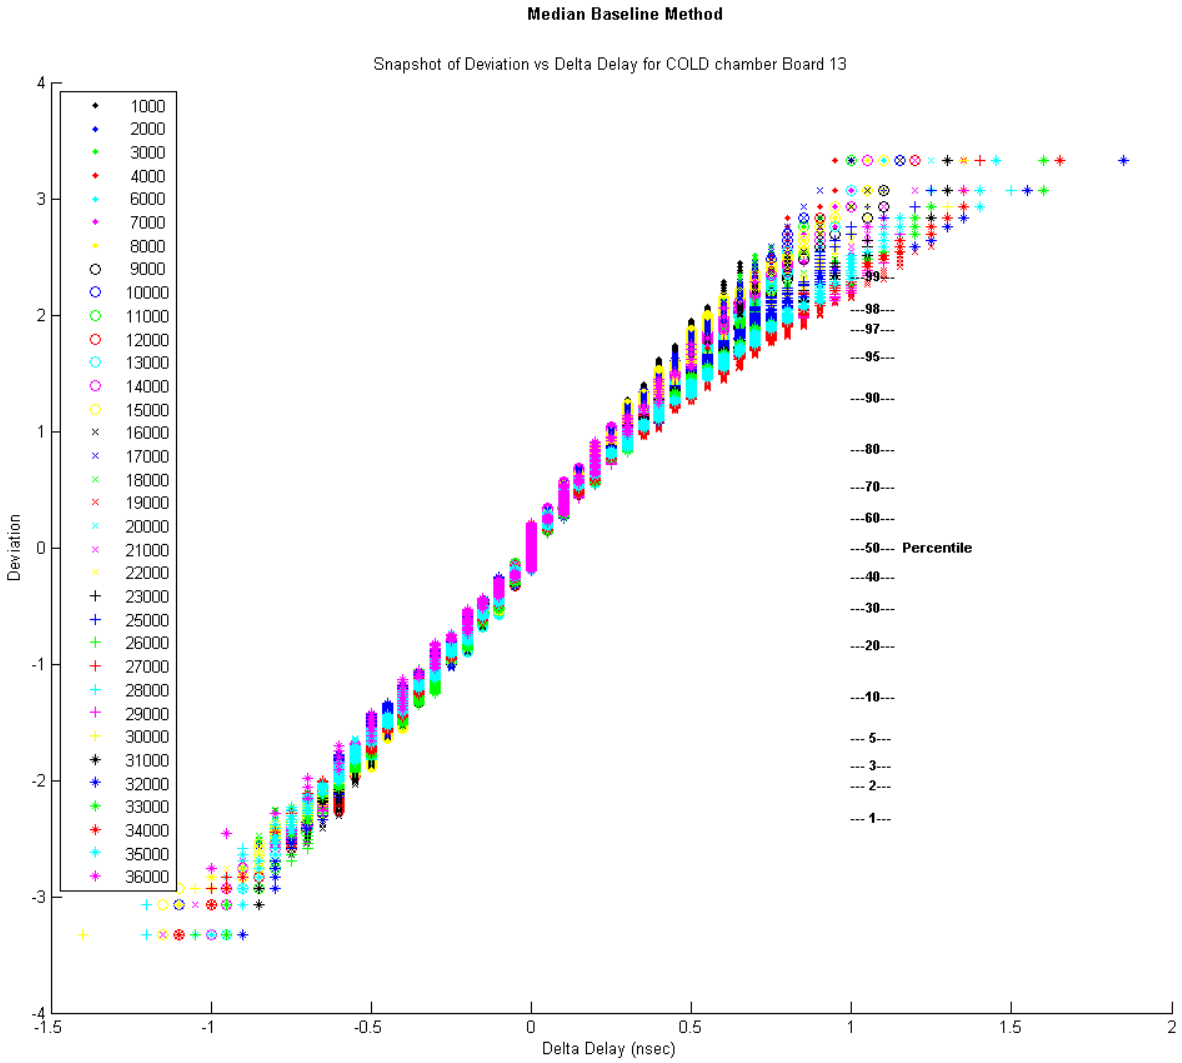

Figure 39. Data from board 13C in the cold chamber.

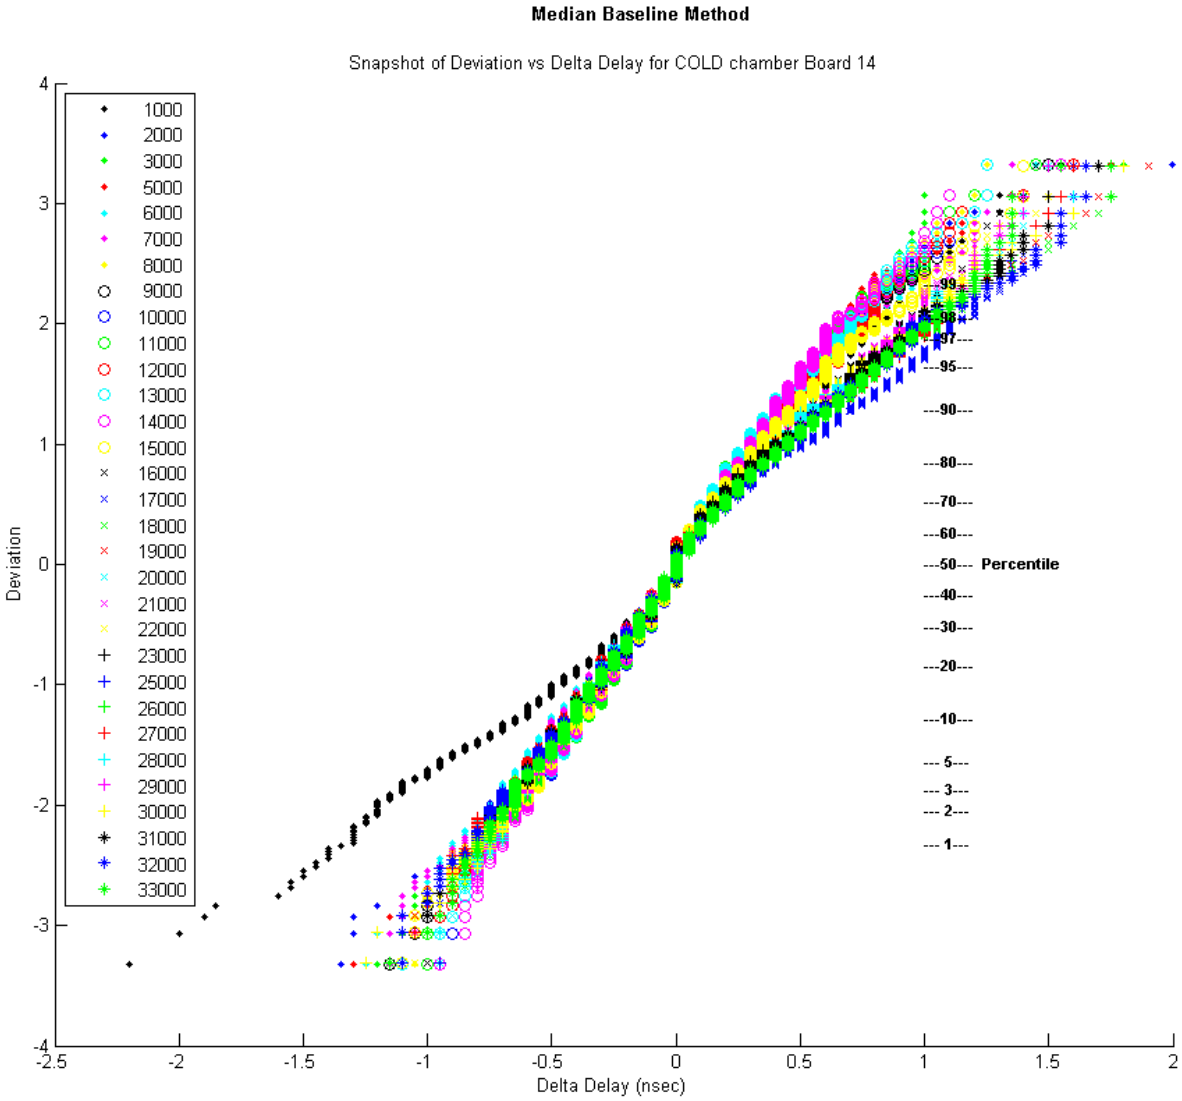

Figure 40. Data from board 14C in the cold chamber. The one data point that exceeds the −2 ns bound is for part 17 at 1000 hours when all the data experienced a shift because of a change in the chamber ambient temperature. The shift was not permanent so no antifuse failure had occurred.

Figures 41−54 show the data for the boards in the hot chamber. There were no deviations in the data other than the temperature dependencies discussed previously. The data here is presented in its entirety, including the data for parts 1, 6, 11, 16, 17, and 18 for every board.

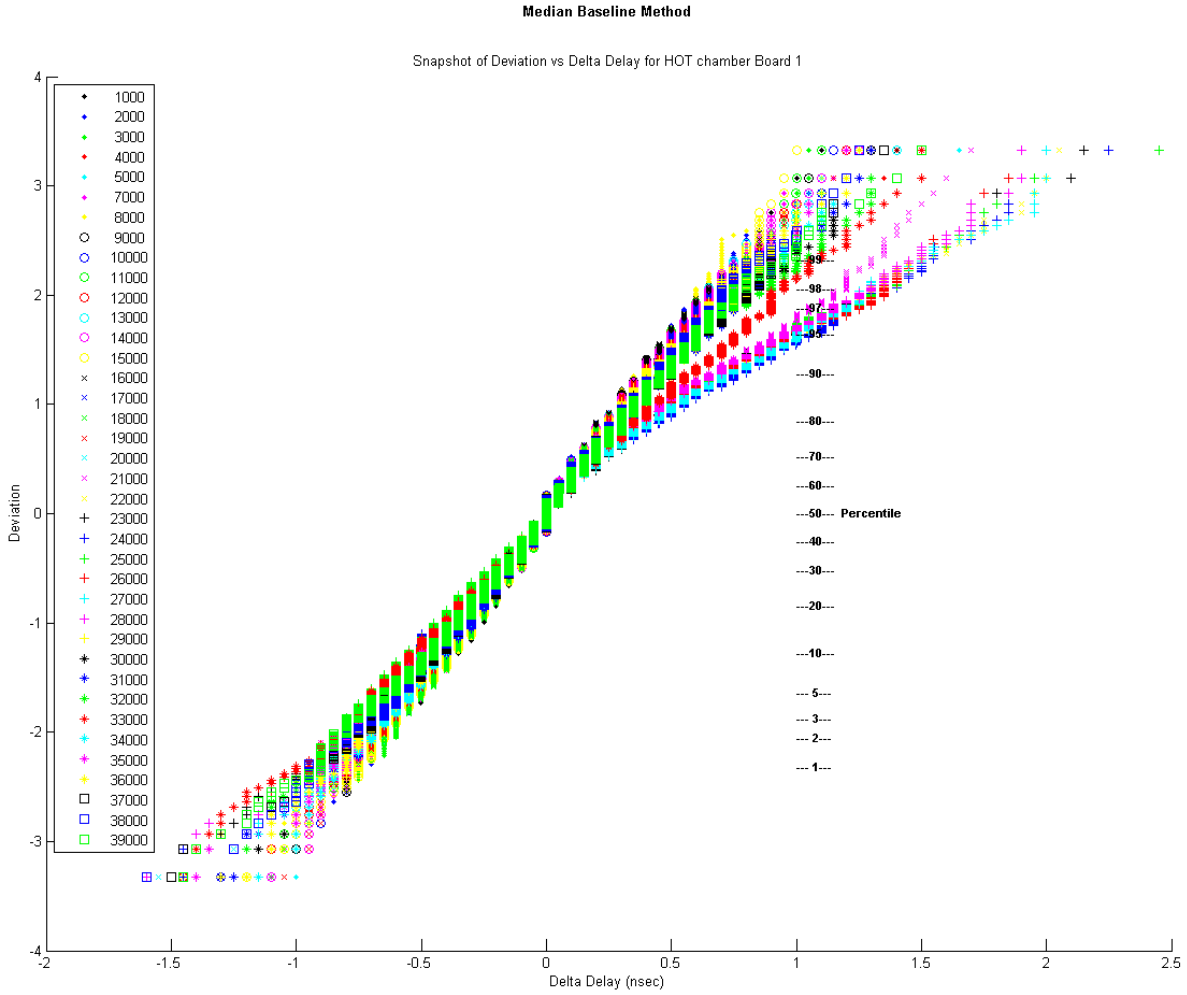

Figure 41. Data from board 1H. In order to ensure the devices were operated no warmer than the specified 85°C, the hot chamber was held at about 73°C ambient once all the boards were installed and running. During the testing, from about 22000 to 28000 hours, the temperature in the chamber increased slightly and the parts along the top of the board ran warmer. The 4 points that exceed the +2 ns bound are from parts 1 and 6. The deviations were not permanent and so no antifuse failure had occurred.

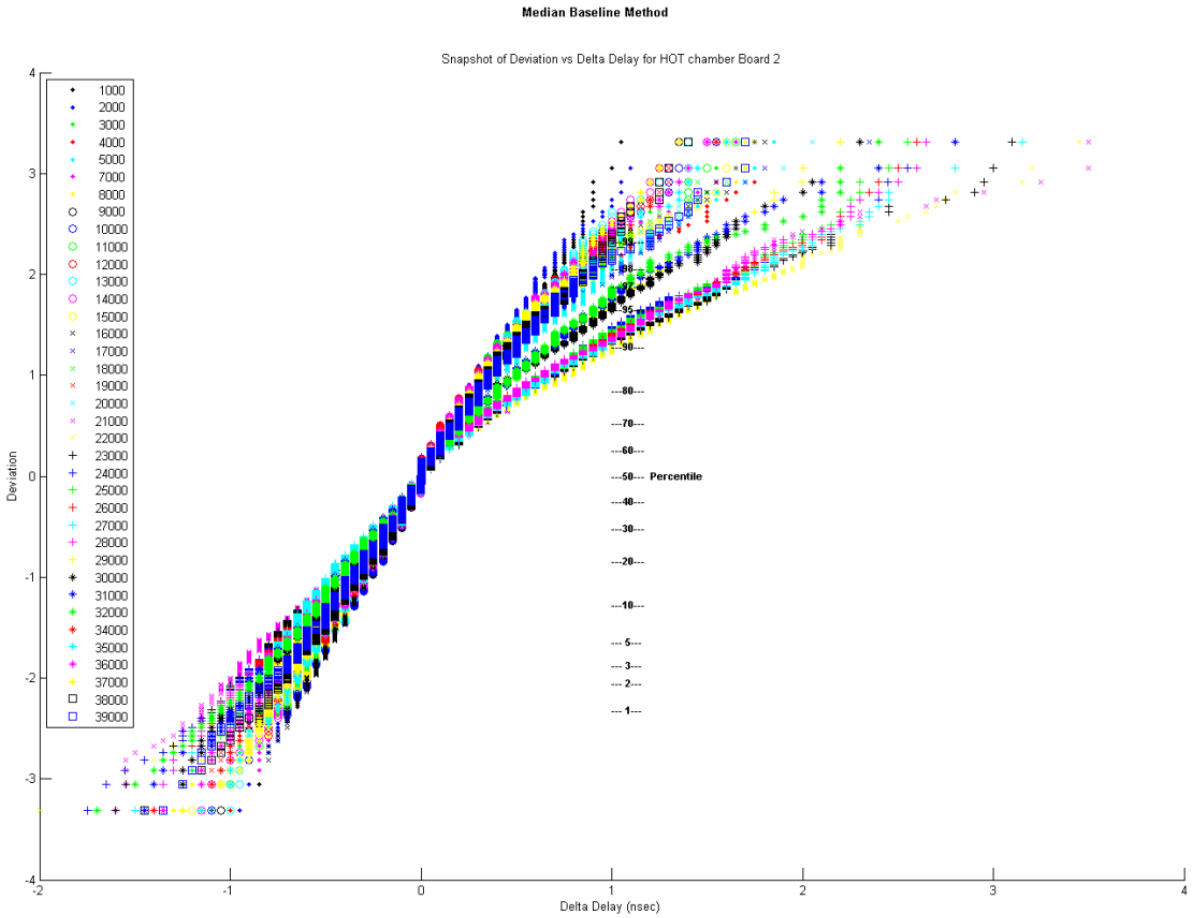

Figure 42. Data from board 2H in the hot chamber. This data shows a similar response to the increased temperature on the boards that came from moving things inside the temperature chamber so that the local heat on the boards increased. This board ran warmer from about 21000 hours to 31000 hours. The deviations outside the +2 ns bound are primarily from parts 1, 6, 11, 16, and 17. Since the data later in the testing shows a return to baseline values, there is no evidence of antifuse failure.

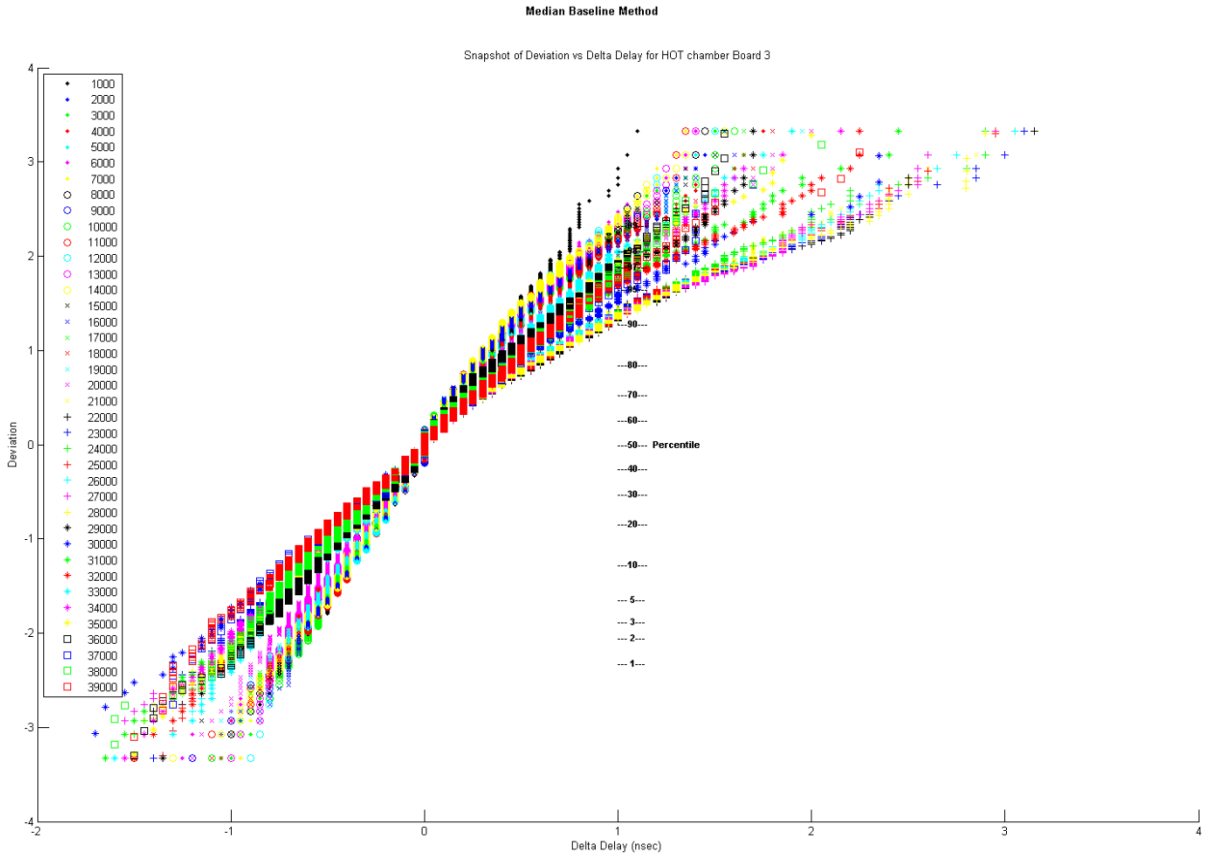

Figure 43. Data from board 3H in the hot chamber. The data outside the 2 ns bound returns to within the range after about 28000 hours. There is no evidence of permanent damage that would indicate an antifuse failure.

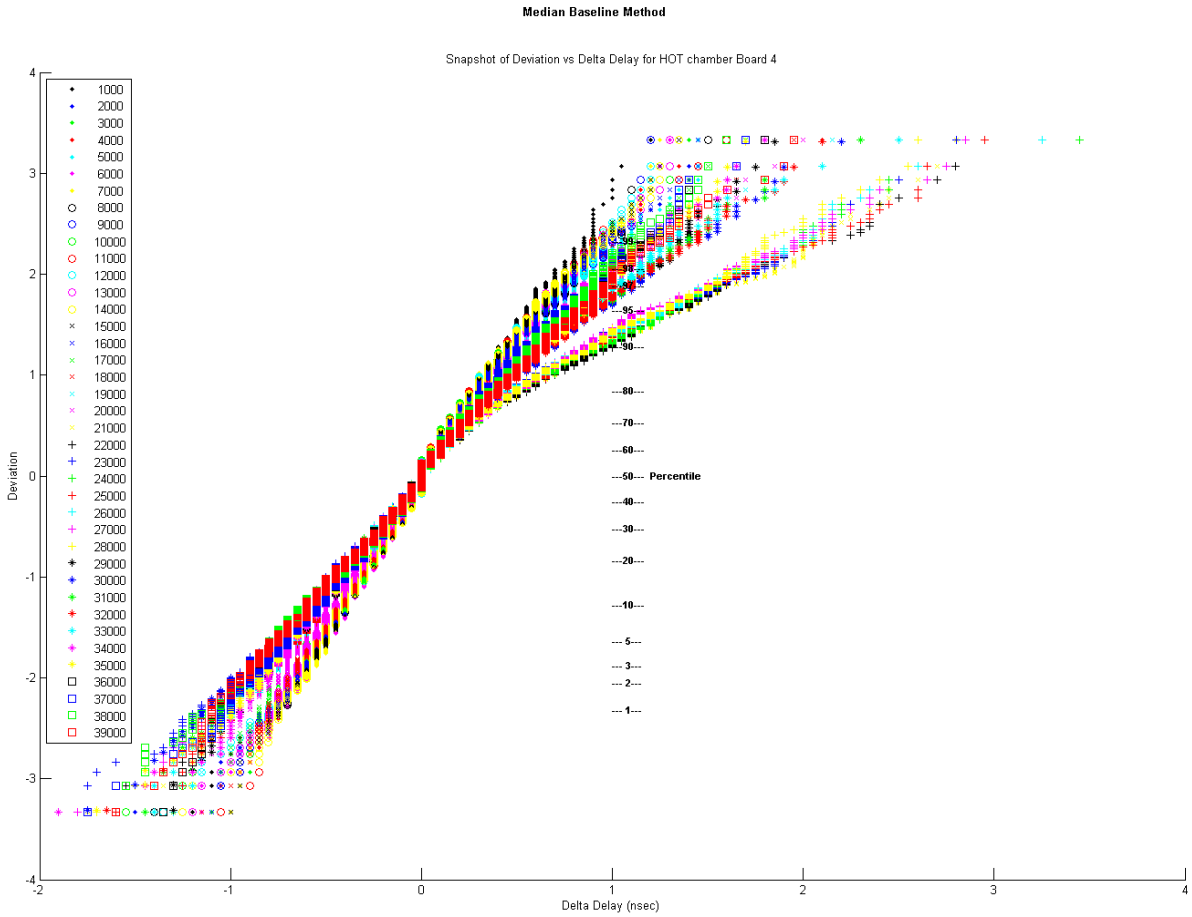

Figure 44. Data from board 4H in the hot chamber. While the data shows that parts 1, 6, 11, 16, 17, and 18, (all parts at the edge of the board) had deviations greater than the +2 ns limit for a time during testing, the delta delays returned back to within range. This indicates there was no antifuse failure.

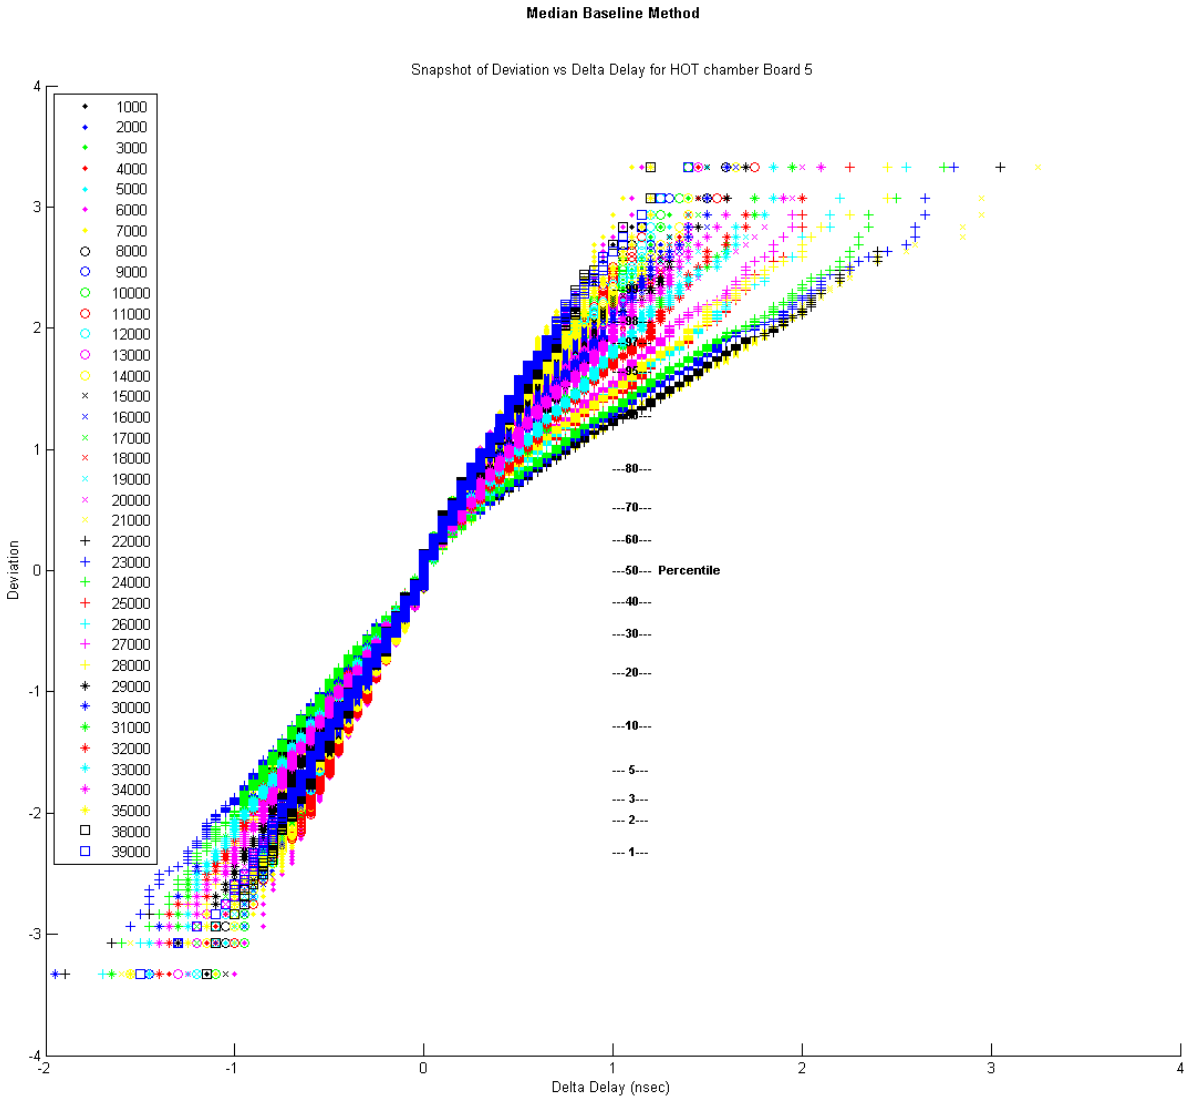

Figure 45. Data from board 5H in the hot chamber. The distribution of the deviations from baseline for these parts follows the same pattern as the other boards in this chamber with a response to the elevated ambient temperature that recovers as the temperature does. There is no indication of antifuse failure.

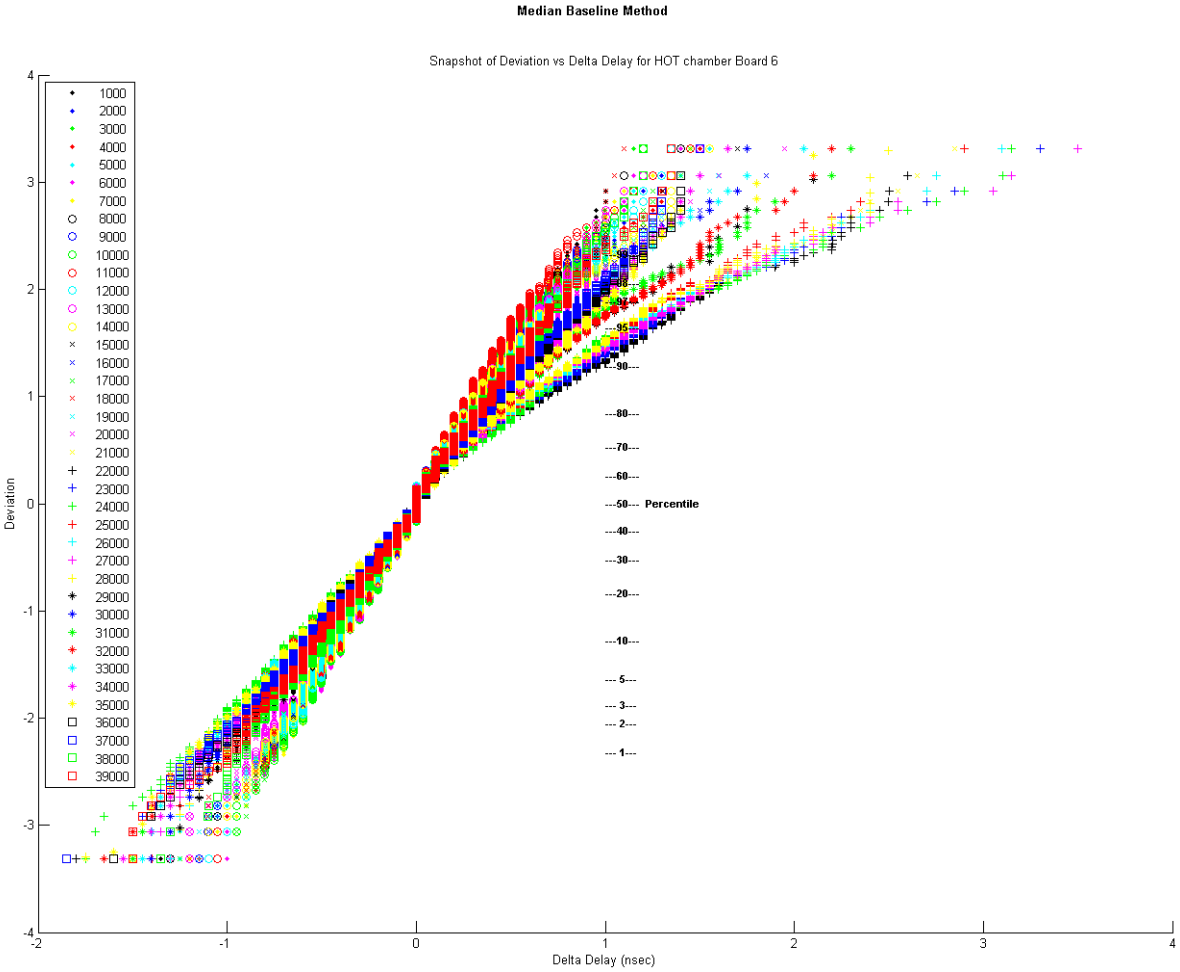

Figure 46. Data from board 6H in the hot chamber. The data at the end of the test is closely following that for the beginning of the test indicating no antifuse rupture. The increase in the deviation from the baseline for the data from 21000 hours to 25000 hours is due to an increase in the ambient temperature in the chamber.

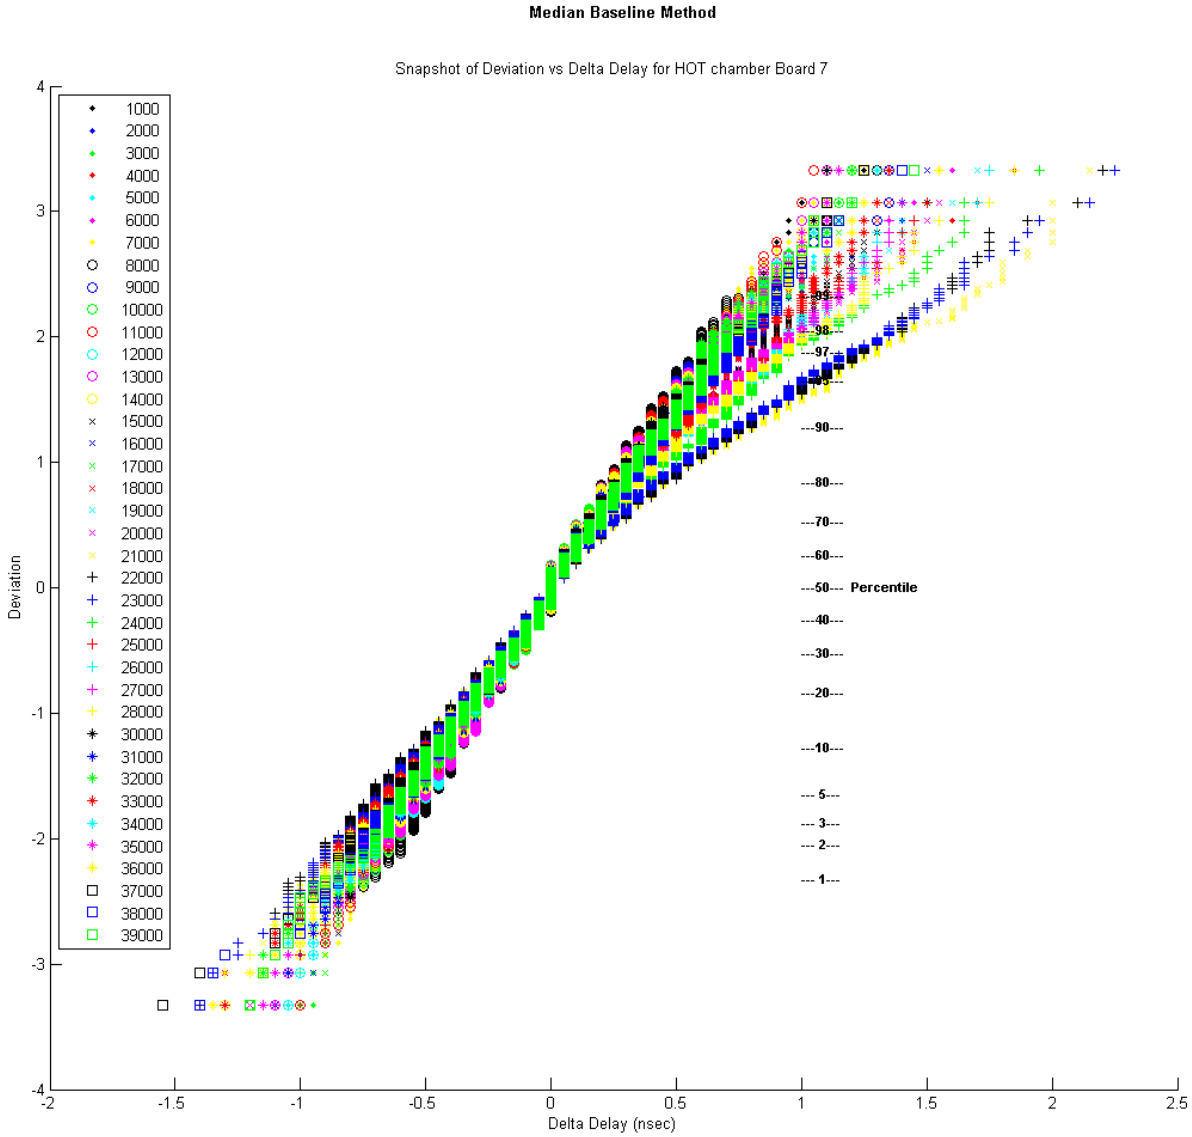

Figure 47. Data from board 7H in the hot chamber. These parts follow a pattern similar to the other boards in this chamber. The 5 data points outside the +2 ns bound are from parts 6 and 16. The three points from part 6 are 1 at 21000 hours and 2 at 22000 hours. The remaining 2 points are from part 16 at 23000 hours. There is no indication of antifuse failure.

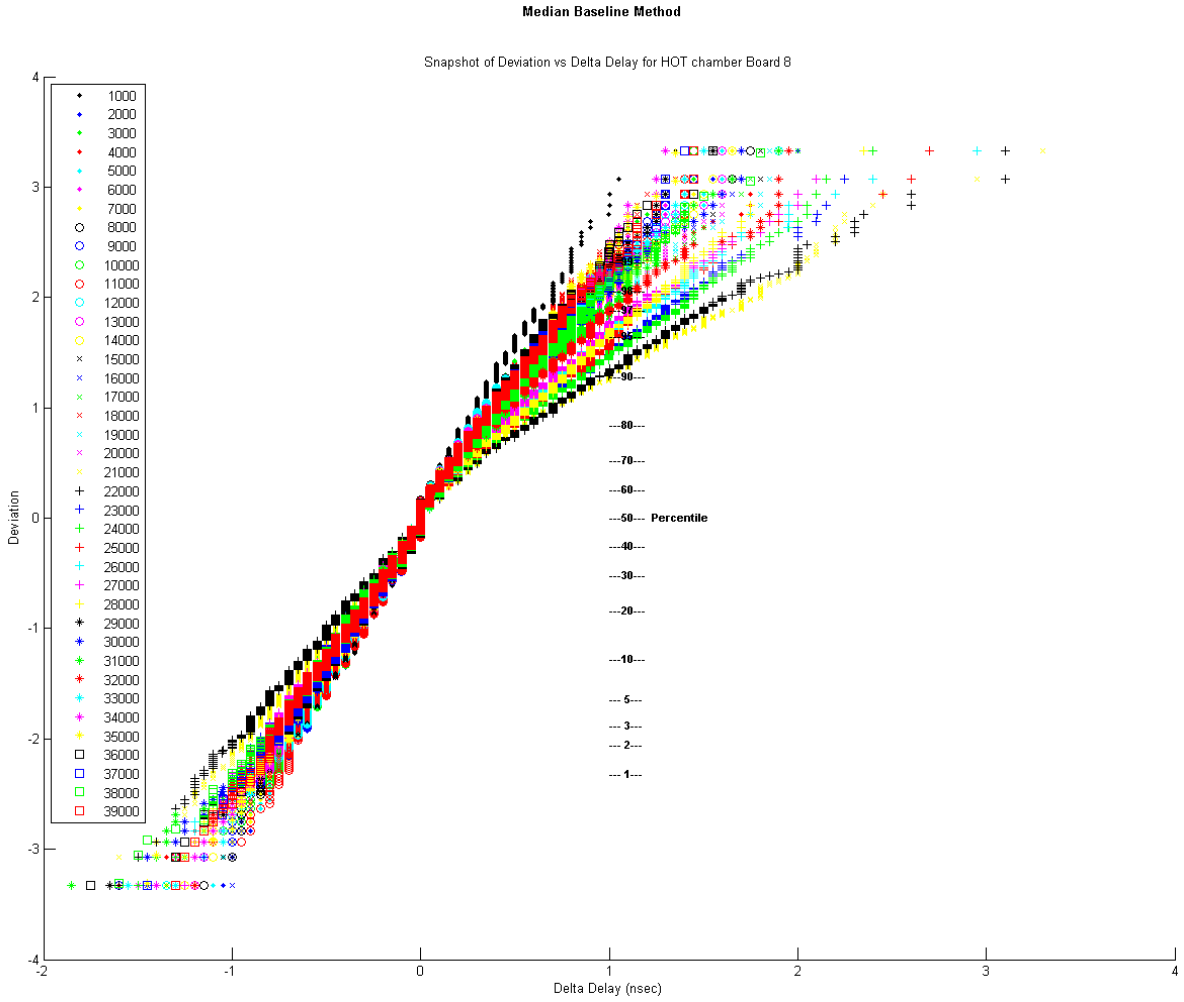

Figure 48. Data from board 8H in the hot chamber. The data follows the pattern where while the chamber ran warner from 22000 hours to 28000 hours, the data deviated more than the +2 ns from the baseline. However with a return to baseline temperatures, the deviations came back within bounds. There is no evidence of antifuse failure.

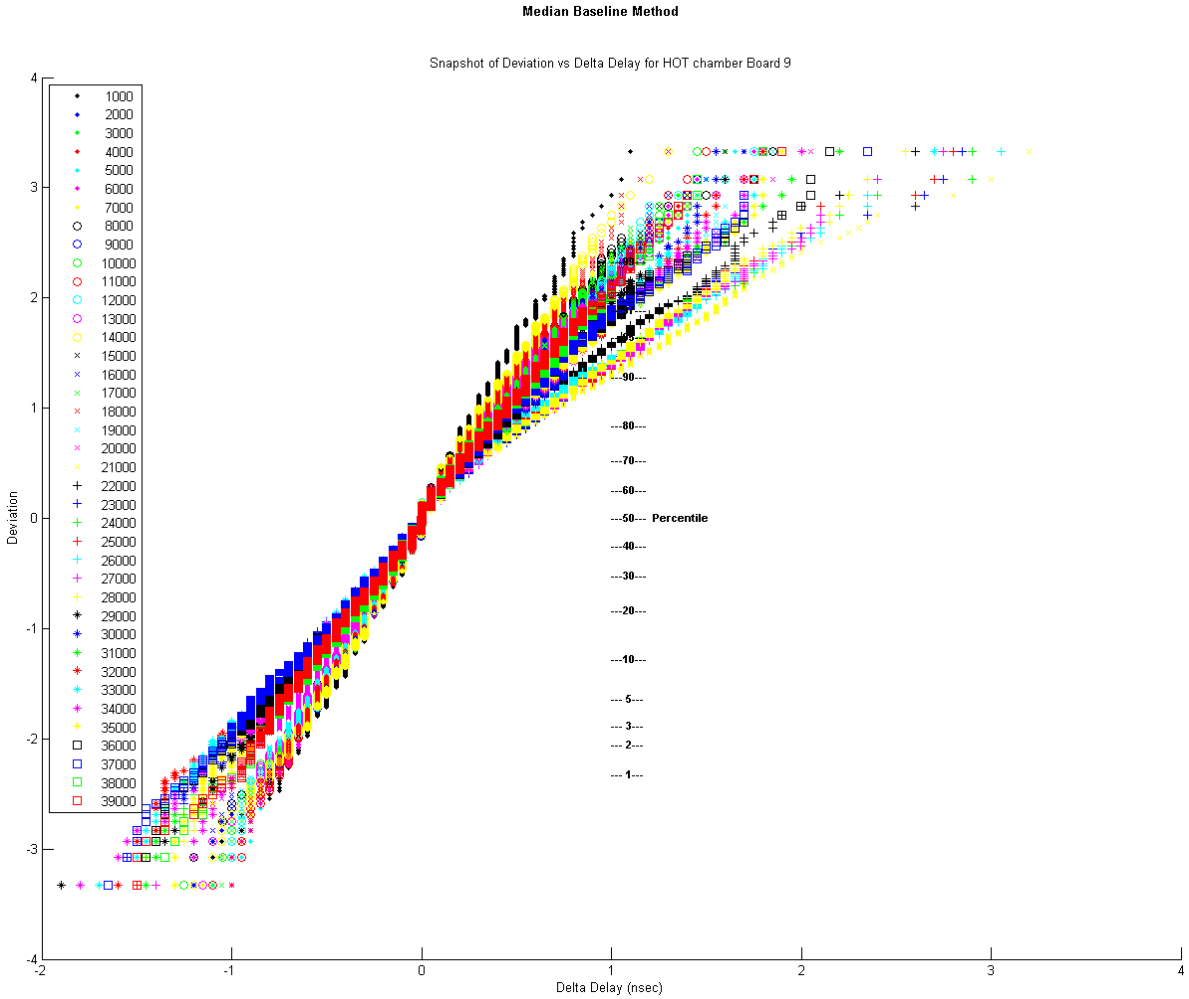

Figure 49. Data from board 9H in the hot chamber. The deviations outside of the +2 ns bounds happened between 20000 and 37000 hours. There is no indication of permanent change so no antifuse failure is indicated.

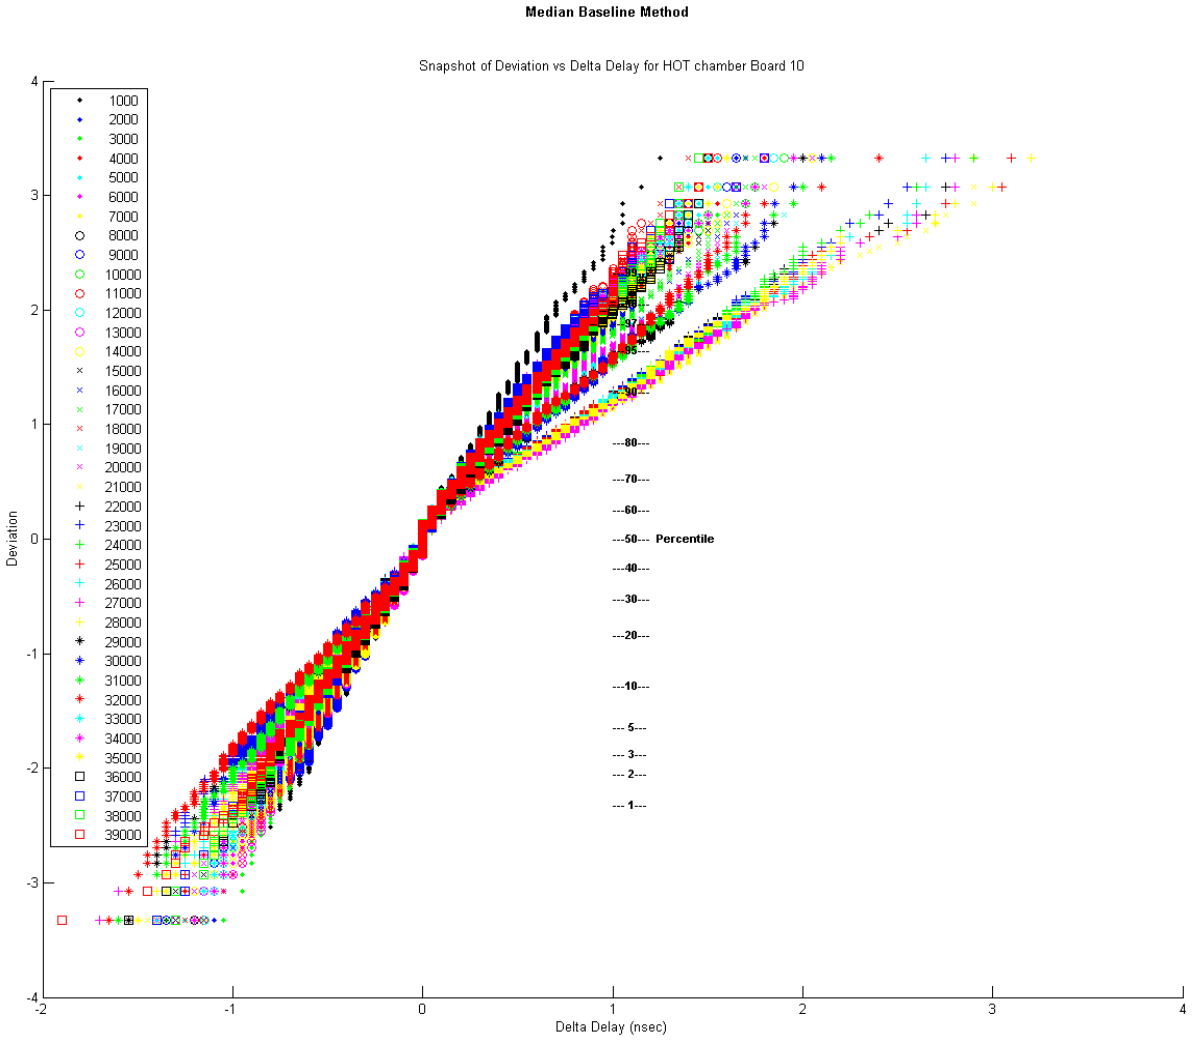

Figure 50. Data from board 10H in the hot chamber. The data for this part indicates an increase in deviations between 20000 and 30000 hours. The shift is not permanent so no antifuse failure is indicated.

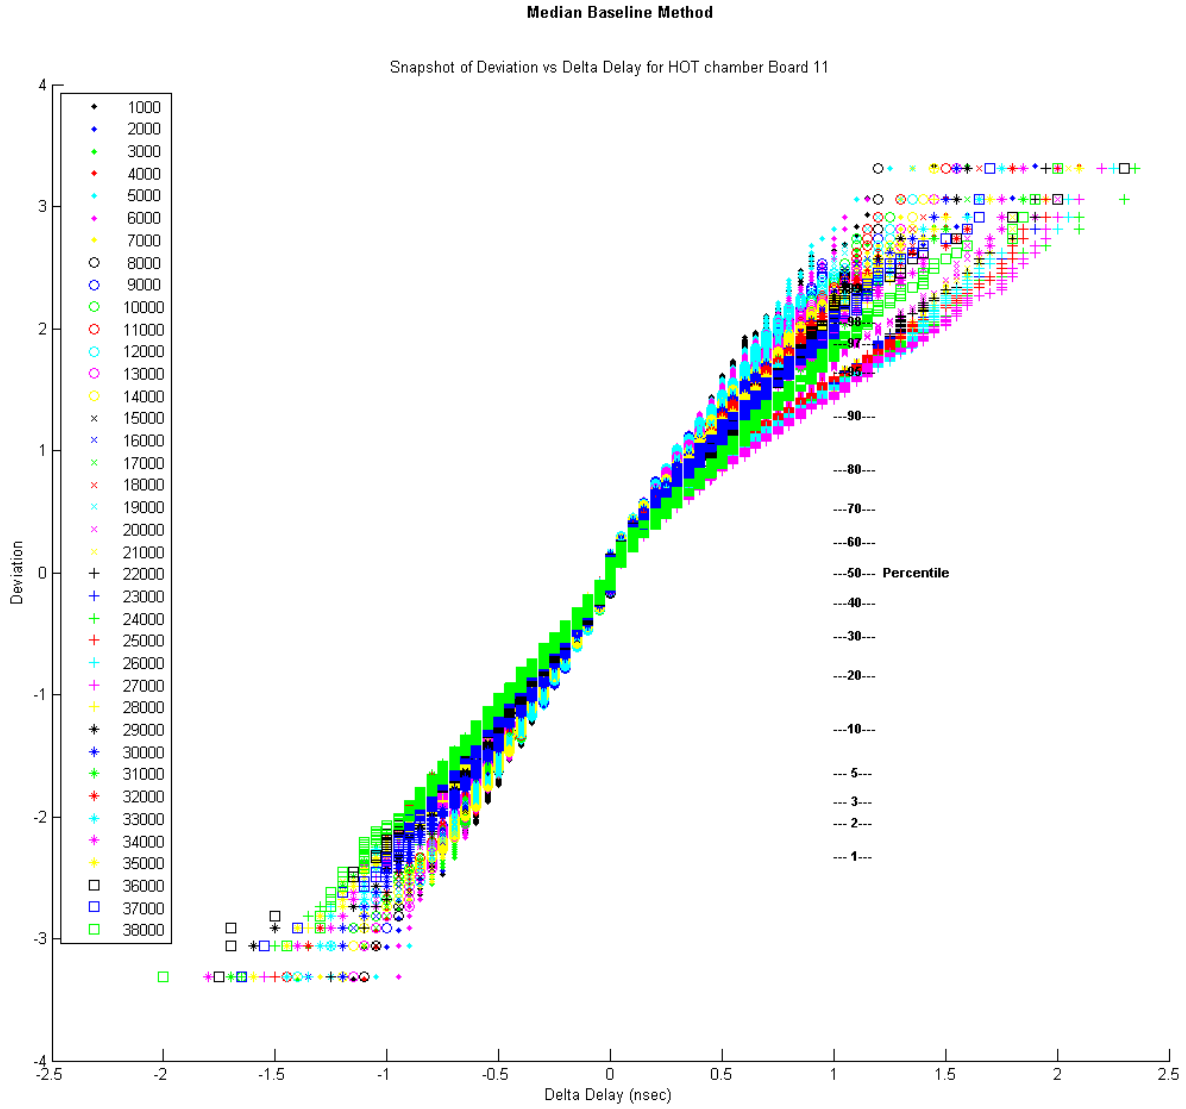

Figure 51. Data from board 11H in the hot chamber. The 13 data points outside the +2 ns bound are from the parts at the top of the board and along the back side. The tracks are all the earliest read out following the pattern seen for local heating. The deviations are not permanent so no antifuse failure is indicated.

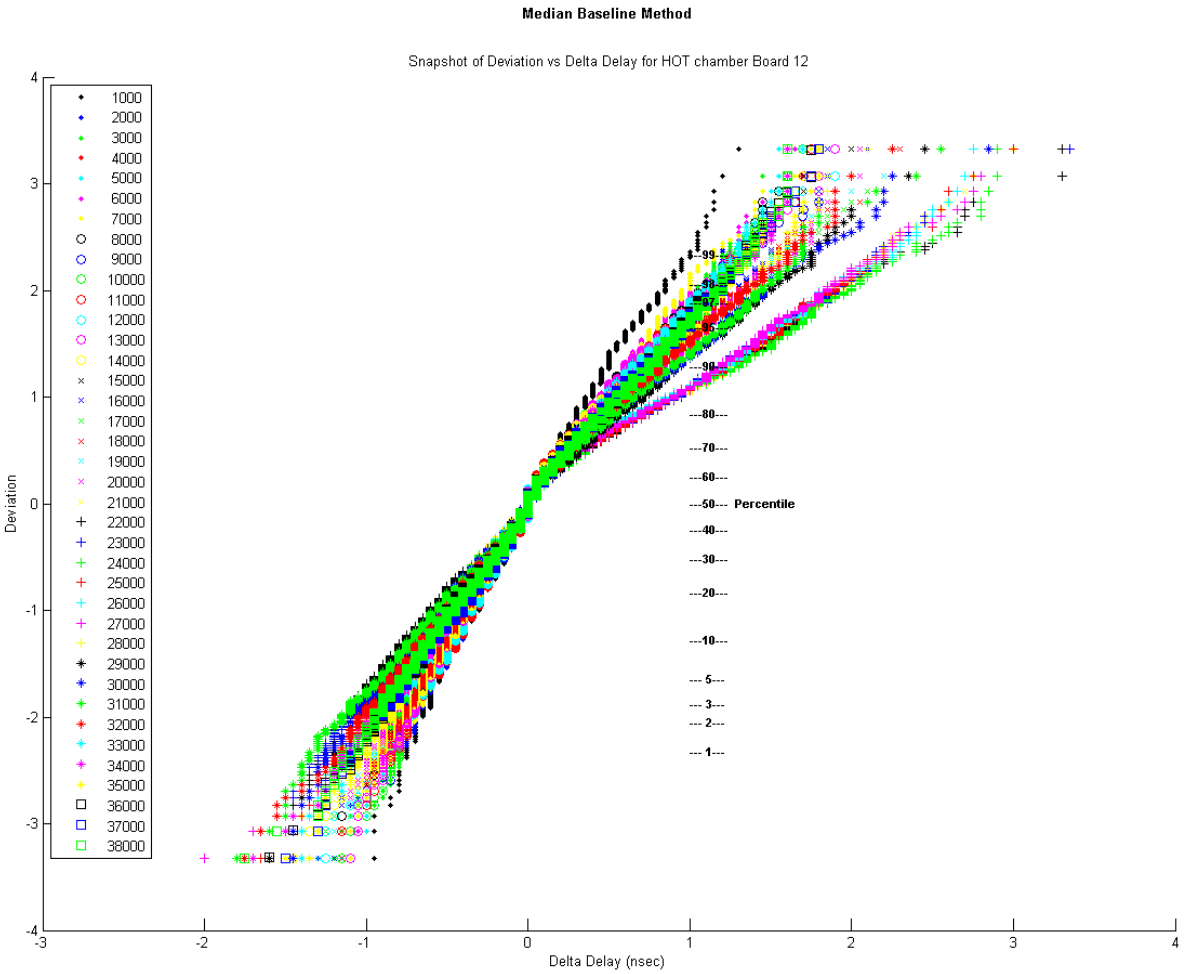

Figure 52. Data from board 12H in the hot chamber. This board was in the warmest part of the chamber. Between 20000 and 30000 hours when the parts along the top and back of the board were running warmest, these parts exhibited the temperature-dependent behavior that accounts for all of the data outside the +2 ns bound. Once the temperature was brought back down to baseline, the data came back in to range. There is no evidence of antifuse failure here.

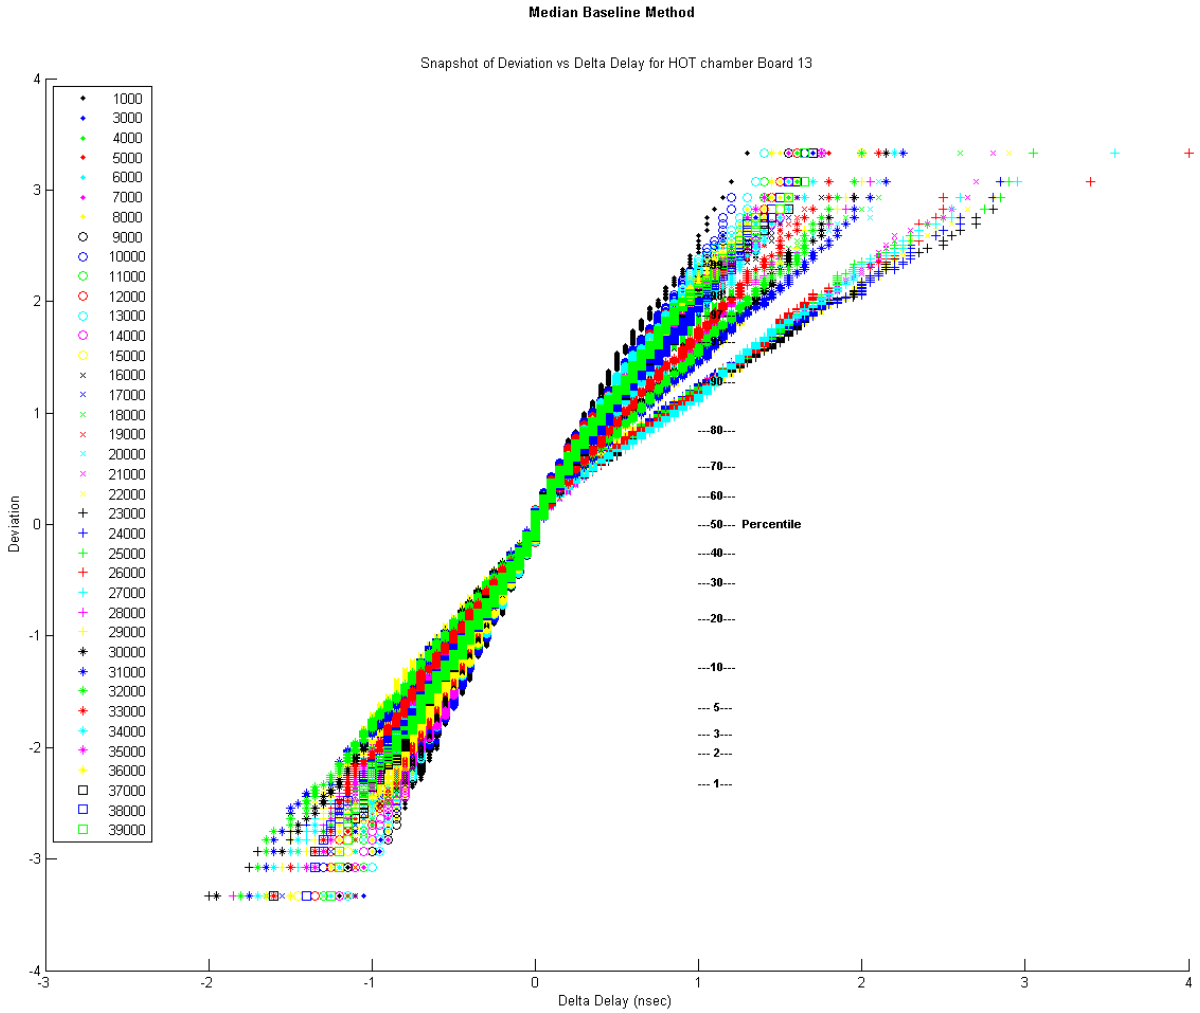

Figure 53. Data from board 13H in the hot chamber. The shift in the data seen between 20000 and 32000 is not permanent so no antifuse failure is indicated.

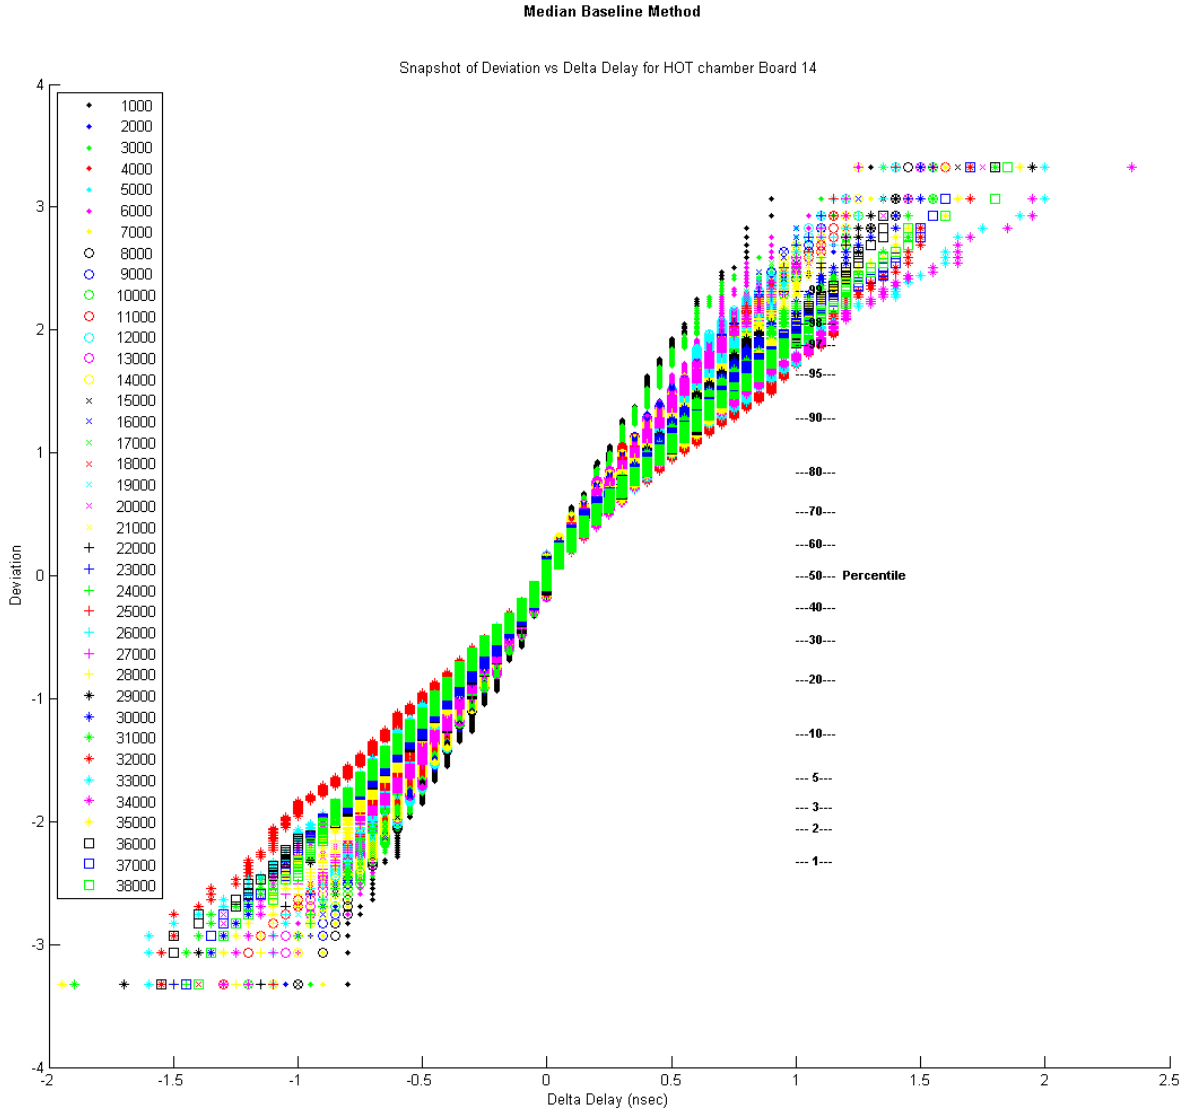

Figure 54. Data from board 14H in the hot chamber. The single data point outside the +2 ns bound is from part 9. The deviation is not permanent so it is not evidence of an antifuse failure.

Beyond the deviations from the measured baseline that are attributed to temperature, there were no timing exceedances measured on any of the tracks of cells. No permanent changes were observed indicating no antifuse failure in any device in this lifetest.

THIS PAGE INTENTIONALLY LEFT BLANK

9. Failures in Time (FIT) Calculation

The results of this test were incorporated into a failure rate calculator (see Appendix B). The calculator uses only the data from the Aerospace testing*. The total number of hours was 26.5 million device-hours, equating to 6.55×10^{12} antifuse hours, with no antifuse failures. The calculator does not include any acceleration factors.

Without any actual antifuse failures, it is impossible to model the failure distribution; however, we can estimate the upper bound of the failure rate. This is done by assuming that the failure times of the antifuses will be completely random and follow an exponential or random distribution (as opposed to a wear-out distribution or an infant mortality distribution). Confidence bounds are calculated using the Chi Square distribution.

$$100\cdot(1-\alpha)\%\ \text{Confidence Bound on FIT Rate for Single Antifuse} = \left(\frac{\chi^2(\alpha,2)}{2*(\text{Total Time on Test})}\right)\times 10^9$$

The failure rate for any particular device is ultimately dependent on the specific device and operating conditions. It is estimated that for an average part with 185,000 dynamic antifuses and 100% duty cycle that a 60% confidence bound on the failure rate due to the antifuse portion of an RTAX FPGA is 25.88 FITs. To obtain a total RTAX FPGA failure rate, this FIT rate needs to be combined with the failure rate for the CMOS portion of the FPGA. This was not estimated.

* Additional tests were conducted by Microsemi; incorporation of this data reduced the failure rate by 1 FIT.

THIS PAGE INTENTIONALLY LEFT BLANK

10. Description of Raw Data Products

The raw data is contained in thousands of text files that are resident on Aerolink under Physical Sciences Laboratories, an Organization in the Aerospace's Engineering and Technology Group. The folder that contains all the data and the program files used to look at the data are in a folder named AX2000_Antifuse_Lifetest. The MatLab code used to examine the data is resident here in the directory AX2000_Antifuse_Lifetest\MatLab. It contains all the program files as well as subdirectories containing the data for each chamber in separate directories names Hot, Cold, and TC.

A data file was made for each DUT every day that it was tested. The board-level power voltage and current readings are recorded in a separate data file for each board on each day as well. Finally, every chamber system recorded a log file for every day of test in that chamber. The data files have been concatenated into 760 device files containing all the data measured on each device and 38 board files containing the power housekeeping data taken for each board over the entire lifetest. These concatenated files of all the data for each part is named as BxxPyy.dat where xx is the board number and yy is the part number on that board. The log files are as originally written and are named for the date (e.g., A19Oct08.log) preceded by the A for Aerospace.

The other directory in AX2000_Antifuse_Lifetest contains the files for the reliability calculator. The Microsoft Excel spreadsheet for the reliability calculator and a Word document describing it is saved at AX2000_Antifuse_Lifetest\Reliability Calculator\Reliability Calculator.xlsx.

The original raw data files made for each part on each day are stored on the corporate server called Ampere. Anyone wanting access to these files should contact this author, Susan Crain, for information.

THIS PAGE INTENTIONALLY LEFT BLANK

References

1. Benard, A. and Bosi-Levenbach, E. C., “The Plotting of Observations on Probability Paper,” *Statistica Neerlandica*, **7**, 163−173 (1953).

2. Ross, Sheldon M., *Introduction to probability and statistics for engineers and scientists* (4th ed.). Associated Press (2009).

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A—Calibration of Instrument Used in Data Collection

Technical Reports Addendum Asset Summary

AEROSPACE

TRAAS ID #: 20121023125722172l5

Report Name: Default Addendum Name 20121023125722172l5

Aerospace Report Number: TOR-2013(1459)-2

Start Date of Test: 2007-05-07 Created By: 17215 Crain, Susan H JO: 1459-30

End Date of Test: 2012-09-30 First Aerospace Author / PI: 17215 Crain, Susan H

Program:

Description: Lifetest of AX2000 Field Programmable Gate Arrays

Keywords: AX2000 RTAX Axcelerator Actel Microsemi FPGA Lifetest FIT antifuse

Asset:	Manufacturer:	Model:	Usage Start Date:	Usage End Date:	Asset Comment:
AAL223	AGILENT	53132A	2012-09-25	2012-09-28	

Date:	Calibration Due Date:	Comment:	Certificate Number:
2012-05-15	2013-03-10	TMT-NORMAL	ba1a40c0c985554c9cc6ff54bf840b19

Asset:	Manufacturer:	Model:	Usage Start Date:	Usage End Date:	Asset Comment:
CT198	HEWLETT-PACKARD CORP.	34401A	2012-09-25	2012-09-28	

Date:	Calibration Due Date:	Comment:	Certificate Number:
2012-02-09	2014-01-05	TMT-NORMAL	db533f4c55cfec48aa2937ff32c04a89

*Support and Auxillary Equipment are not calibrated.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B—Method of Actel FPGA Failure Rate Assessment, and FIT Rate Calculator for RTAX Family

The failure rates described here are applicable only to the Actel Axcelerator FPGA products, particularly the RTAX250S, RTAX1000S, RTAX2000S, and RTAX4000S. In the ideal case, all the testing would have been performed on these spaceflight-qualified devices alone. However, because of their high cost and lower availability, life tests have been performed on the commercial device type AX2000S, which employs the identical antifuse structures, programming methods, and usage as the radiation-tolerant versions. Early in the life test experiment development, the antifuse structures were independently examined by Aerospace using Scanning Electron Microscopy (SEM) and Focused Ion Beam (FIB) imaging on cross-sections of both RTAX and AX devices. Due to Actel proprietary sensitivities, the SEM and FIB images are not included here, but Aerospace concluded that the commercial part would serve as an accurate test of the antifuse reliability. Using these commercial parts in this lifetest, no antifuse failures have been observed.

The failure rates described here are applicable only to the Actel FPGA products RTAX250S, RTAX1000S, RTAX2000S, and RTAX4000S. In the ideal case, all the testing would have been performed on these devices alone. However, because of their high cost and lower availability, life tests have been performed on the commercial device types AX2000. No antifuse failures have been observed.

Excel Spreadsheet Calculator

An Excel spreadsheet calculator has been included that provides predicted failure rates for a particular user FPGA design and mission parameters. Instructions for using this spreadsheet are as follows.

1. **Define the user FPGA Design.** Run the Actel-provided executable program that provides antifuse link counts by type. The total number of antifuses should be input into cell B15 of the tab entitled “1. Total FIT Rate”.
2. **Environmental and Operating Parameters.**
   Input the duty cycle in cell B17. The duty cycle is the percentage of calendar time that the system in which the component is operating is in an operational state.
3. **Results.**
   The 60% upper confidence bound for the antifuse portion of the FPGA is provided in cell B23 (gray shading) in the first tab “1. Total FIT Rate”.
4. **Confidence Bounds**
   Upper 60% and upper 90% confidence bound results are provided for the antifuse failure rate in the tab titled “3. Antifuse FIT Rate”.
5. **Data Cells**
   Do not change the data cells or intermediate calculations in any of the other tabs.
6. **Actual Data**
   The actual data used for all antifuse calculations is shown in tabs 4 and 6.

Enclosure

Excel File: RTAX Calculator.xls (found at aero-org\Ampere\shared\Actel_Burnin\ AX2000Test\AX2000TestData\Reliability Calculator.xlsx)

Tab 1. Total Fit Rate	basic user inputs and outputs
Tab 2. Antifuse FIT Rate	provides 60% and 90% confidence bounds for antifuse FIT rate
Tab 4. Antifuse Data	contains data used
Tab 6. Antifuse Types	contains antifuse breakdowns

Appendix C—Wafer Lot Numbers for Devices Used in Test

The devices used in this life test were from three Wafers with the same Lot Date Code (LDC) of 0638. Sixty devices were used from wafer D2GF51, with 20 of those going into the cold chamber and 40 in the hot chamber. Another 140 were from wafer D2A5A1 with 65 of these in the cold chamber, 41 in the hot, and 34 in the cycling chamber. The remaining 560 were from wafer D203Y1 with 195 in the cold chamber, 199 in the hot, and 166 in the cycling chamber. The distribution of the parts is shown in Tables C1−C3 for the three chambers.

Table C1. Wafer Lot Codes for the Hot Chamber Devices Where Lot 1 is D2GF51, Lot 2 is D2A5A1, and Lot 3 is D203Y1. Devices that were removed from the test set are shaded gray. Part 11H16 was removed for SRAM failure analysis.

Part	1H	2H	3H	4H	5H	6H	7H	8H	9H	10H	11H	12H	13H	14H
1	3	3	3	3	3	3	3	3	3	2	3	2	1	1
2	3	3	3	3	3	3	3	3	3	2	3	2	1	1
3	3	3	3	2	3	3	3	3	3	2	3	2	1	1
4	3	3	3	2	3	3	3	3	3	2	3	3	1	1
5	3	3	3	2	3	3	3	3	3	2	3	3	1	1
6	3	3	3	3	3	3	3	3	3	2	3	2	1	1
7	3	3	3	3	3	3	3	3	3	2	3	2	1	1
8	3	3	3	2	3	3	3	3	3	2	3	2	1	1
9	3	3	3	2	3	3	3	3	3	2	3	3	1	1
10	3	3	3	2	3	3	3	3	3	2	3	3	1	1
11	3	3	3	3	3	3	3	3	3	2	3	2	1	1
12	3	3	3	3	3	3	3	3	3	3	3	2	1	1
13	3	3	3	3	3	3	3	3	3	2	3	2	1	1
14	3	3	3	2	3	3	3	3	3	2	3	2	1	1
15	3	3	3	2	3	3	3	3	3	2	3	3	1	1
16	3	3	3	3	3	3	3	3	3	2	3	2	1	1
17	3	3	3	3	3	3	3	3	3	3	3	2	1	1
18	3	3	3	3	3	3	3	3	3	2	3	2	1	1
19	3	3	3	2	3	3	3	3	3	2	3	3	1	1
20	3	3	3	2	3	3	3	3	3	3	3	2	1	1

Table C2. Wafer Lot Codes for the Cold Chamber Devices Where Lot 1 is D2GF51, Lot 2 is D2A5A1, and Lot 3 is D203Y1. Devices that were removed from the test set are shaded gray.

Part	1C	2C	3C	4C	5C	6C	7C	8C	9C	10C	11C	12C	13C	14C
1	3	3	3	3	3	3	3	2	2	3	2	3	3	1
2	3	3	3	3	3	3	3	2	2	3	2	3	3	1
3	3	3	3	3	3	3	3	2	2	2	2	3	3	1
4	3	3	3	3	3	3	3	2	2	3	2	3	3	1
5	3	3	3	3	3	3	3	2	2	3	2	3	3	1
6	3	3	3	3	3	3	3	2	2	3	2	3	3	1
7	3	3	3	3	3	3	3	2	2	3	2	3	3	1
8	3	3	3	3	3	3	3	2	2	2	2	3	3	1
9	3	3	3	3	3	3	3	2	2	3	2	3	3	1
10	3	3	3	3	3	3	3	2	2	3	2	3	3	1
11	3	3	3	3	3	3	3	2	2	3	2	3	3	1
12	3	3	3	3	3	3	3	2	2	3	2	3	3	1
13	3	3	3	3	3	3	3	2	2	2	2	3	3	1
14	3	3	3	3	3	3	3	2	2	3	2	3	3	1
15	3	3	3	3	3	3	3	2	2	3	2	3	3	1
16	3	3	3	3	3	3	3	2	2	3	2	3	3	1
17	3	3	3	3	3	3	3	2	2	3	2	3	3	1
18	3	3	3	3	3	3	3	2	2	2	2	3	3	1
19	3	3	3	3	3	3	3	2	2	3	2	3	3	1
20	3	3	3	3	3	3	3	2	2	3	2	2	3	1

Table C3. Wafer Lot Codes for the Temperature Cycling Chamber Devices Where Lot 1 is D2GF51, Lot 2 is D2A5A1, and Lot 3 is D203Y1. Devices that were removed from the test set are shaded gray. Part 7TC-3 was removed for SRAM failure analysis.

Part	1TC	2TC	3TC	4TC	5TC	6TC	7TC	8TC	9TC	10TC
1	3	3	2	3	3	3	3	3	3	2
2	3	3	3	3	3	3	3	3	3	2
3	2	3	3	3	3	3	3	3	3	2
4	2	3	3	3	3	3	3	3	3	2
5	3	3	2	3	3	3	3	3	3	2
6	3	3	2	3	3	3	3	3	3	2
7	3	3	3	3	3	3	3	3	3	2
8	2	3	3	3	3	3	3	3	3	2
9	2	3	3	3	3	3	3	3	3	2
10	3	3	2	3	3	3	3	3	3	2
11	3	3	2	3	3	3	3	3	3	2
12	3	3	3	3	3	3	3	3	3	2
13	3	3	3	3	3	3	3	3	3	2
14	2	3	3	3	3	3	3	3	3	2
15	3	3	2	3	3	3	3	3	3	2
16	3	3	2	3	3	3	3	3	3	2
17	3	3	3	3	3	3	3	3	3	2
18	3	3	3	3	3	3	3	3	3	2
19	2	3	3	3	3	3	3	3	3	2
20	3	3	2	3	3	3	3	3	3	2

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix D—Actel Programming Report

Table D1 shows the final usage report generated by Actel's FPGA programming software.

Table D1. – Usage report for the DUT

```
Post-Combiner device utilization:
    SEQUENTIAL (R-cells)         Used: 10391  Total: 10752   (96.64%)
    COMB (C-cells)               Used: 21480  Total: 21504   (99.89%)
    LOGIC (R+C cells)            Used: 31871  Total: 32256   (98.81%)
    RAM/FIFO                     Used:   64  Total:   64
    IO w/Clocks                  Used:   63  Total:  586
    CLOCK (Routed)               Used:    4  Total:    4
    HCLOCK (Hardwired)           Used:    1  Total:    4
    PLL                          Used:    0  Total:    8

    Input I/O Register  :    0
    Output I/O Register :    0
    DDR Register        :    0
    Comb-Comb (CC)      :    0
    Carry Chain         :    1

I/O Information:
    Input Pads                 :    29
    Output Pads                :    34
    Bidirectional Pads         :    0
    Differential Input Pairs   :    0
    Differential Output Pairs  :    0

Net Information:

Following nets are assigned to HCLOCK (Hardwired) resources:
        Fanout    Name
         66       HCLK

Following nets are assigned to CLOCK (Routed) resources:
        Fanout    Name
         9856     KNETCLK
         402      $5I386/SYSCLK
         237      $5I386/RSTN
         32       CLKSR

Remaining high fanout nets:
        Fanout    Name
         15       $5I386.i_ram.RD_SEL_0[2]
         14       $5I386.i_ram.RD_SEL[0]
         13       $5I386/i_ram/RD_SEL_1[2]
         12       $5I386.i_ram.RD_SEL_0[1]
         12       $5I386/wd_q_1[4]
         11       $5I386/cw1_q[3]
         11       $5I386/i_ram/RD_SEL_2[2]
         10       $5I386.i_ram.RD_SEL_0[4]
         10       $5I386.i_ram.RD_SEL_0[5]
         10       \\CTRL_DATA[0]\\

Connectivity Information:
        58.02 percent Connectivity Fanout <=3
        86.93 percent Connectivity Fanout <=6
        99.99 percent Connectivity Fanout <=12
        100.00 percent Connectivity Fanout <=28
There were 0 error(s) and 1729 warning(s) in this design.
```

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix E—Software

After being powered on, each DUT was configured such that the mode for all the tracks of muxes were set to zero, the SSO switching was started, and the clock was set to the nominal 20-MHz setting and then enabled. The SRAM test was cleared and then enabled, and finally the auto-toggling of the mode for the muxed tracks was initiated.

From this initial state, the data collection started. For each DUT, the supply voltage and current for both the I/O and the core was read. The temperature at the case of the DUT also was recorded. A flag indicating whether there was an error in the SRAM and a flag for an overvoltage on either supply since the last read were read and the oscillator frequency was measured. Each of the sequential tracks was then exercised, and a measure of the pulse period and duty cycle made. The clock was stopped, and a readout of the combinational tracks followed. For the tracks made up of inverter cells, the signal propagation time was measured as a single pulse width. For the tracks made up of mux cells, the same measurement was repeated for each of the four muxed modes.

Finally, a current measurement was made on the core in this quiescent state. Then the oscillator was restarted and the part was reconfigured back to the operational state. The tester cycled through all the parts in this manner. Figure E1 shows the flow for (a) this configuration/reconfiguration and (b) the data collection.

```
for (boardNum = 1; boardNum <= MaxNumBoards; boardNum++) {
        disable all board outputs except boardNum
        if board_present {
                for (partNum = 1; partNum <= 20; partNum++) {
                        write the mux control for the K-fuse nets = 0
                        disable the ring oscillator to make sure it is really
stopped
                        reset all K-fuse test networks, then start them running
                        disable operation of SSO switching
                        load the SSO initial pattern; configure to clock from
CLK/4
                        select ring oscillator frequency and start oscillator
                        clear the SRAM test error status
                        enable the SRAM test for this part
                        enable auto-toggling of the MUX4DELAY control bits
                        clear the over-voltage monitor
                } // partnum
        } // board_present
} // boardNum
```

(a)

```
for (boardNum = 1; boardNum <= MaxNumBoards; boardNum++) {
        disable all board outputs except boardNum
        if board_present {
                disable auto-toggling of the MUX4SET control bits
                for (partNum = 1; partNum <= 20; partNum++) {
                if part_enabled {
                        read/check VCCI(V), VCCA(V), VCCI(I), VCCA(I), Temperature
                        read/check ring oscillator frequency
                        read/check SRAM tester status
                        read/check over-voltage latch
                        for (trackNum = 0; trackNum < 32; trackNum++)
                                read/check SeqNetwork pulse period and duty cycle
                        disable ring oscillator
                        for (trackNum = 0; trackNum < 32; trackNum++) {
                                ground all CombNetwork clock inputs for other tracks
                                if (MUX4_type_CombNetwork) {
                                        for (muxState = 0; muxState < 4; muxState++) {
                                                set MUX4NET control = muxState
                                                read/check CombNetwork pulse duration
                                        } // muxState
                                } // MUX4_type_CombNetwork
                                else
                                        read/check CombNetwork pulse duration
                        } // trackNum
                        read VCCA current while ring oscillator stopped (=IddQ)
                        re-configure part for normal operation (see Software (re)Configuration
                } // part_enabled
                enable auto-toggling of the MUX4DELAY control bits
        } // board_present
} // boardNum
```

(b)

Figure E1. (a) Software flow for configuring and reconfiguring the devices under test and (b) Data Collection flow for each Actel.