

National Aeronautics and Space Administration



FPGA NEPP FY09 Summary Report

Douglas Sheldon
Jet Propulsion Laboratory
Pasadena, California

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL Publication 10-1 01/10



FPGA NEPP FY09 Summary Report

NASA Electronic Parts and Packaging (NEPP) Program
Office of Safety and Mission Assurance

Douglas Sheldon
Jet Propulsion Laboratory
Pasadena, California

NASA WBS: 724297.40.43
JPL Project Number: 103982
Task Number: 03.02.01

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

<http://nepp.nasa.gov>

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration Electronic Parts and Packaging (NEPP) Program.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

Copyright 2010. California Institute of Technology. Government sponsorship acknowledged.

TABLE OF CONTENTS

1	Overview.....	1
2	Device Physics-Based Design Modifications	2
3	Contention Study—Xilinx Virtex-4.....	6
	3.1 Contention Experiment	7
	3.2 LVDS Summary.....	11
	3.3 LVTTTL and LVCMOS3 Data.....	13
4	IDDQ Study in Actel A54SX FPGAs.....	15
	4.1 IDDQ Experimental Design.....	16
	4.2 FIR Design	17
	4.3 FPGA Testing	19
	4.4 FPGA Testing Results.....	20
5	Conclusions.....	33
	References.....	34
	Appendix A. LVCMOS and LVTTTL Contention Statistics.....	35
	Appendix B. Control Board Statistics for Virtex-4 Contention Test.....	36

1 OVERVIEW

This report documents the activities and results of the fiscal year 2009 (FY09) funding for the NASA Electronic Parts and Packaging (NEPP) program for re-programmable field programmable gate arrays (FPGAs).

The FY09 task was divided into three efforts:

1. Device physics-based design modifications
2. Contention study of Xilinx Virtex-4
3. IDDQ study of Actel A54SX

The FY09 NEPP FPGA study was organized into these three sections in recognition of the continued and increasing importance of FPGAs to NASA. FPGAs represent the state of the art in electronic components with millions and millions of transistors integrated into a single device. Modern NASA spacecraft design has dozens and dozens of FPGAs implemented onboard. FPGAs are used in critical command and data handling, instrument control and monitoring, and communications protocols, to name just a few. The reliability of FPGAs is fundamental to mission success for NASA.

Ensuring reliability means having both a detailed theoretical understanding of FPGA aging and wear-out mechanisms as well as an empirical quantification of possible degradation phenomena. This FY09 NEPP task was organized to provide a firm foundation in terms of understanding of device physics-based reliability issues as well as exploring practical details of contention and IDDQ testing.

2 DEVICE PHYSICS-BASED DESIGN MODIFICATIONS

Device physics-based design modifications means developing precise physics-based models of the transistor structures and materials that go into making a modern FPGA. These models deal with environmental effects such as temperature and current density over time. Once these models exist, the programmable nature of FPGAs allows many options to be implemented to help minimize and mitigate their effects on overall device performance. The objective of work in FY09 was to develop research contacts in this field and propose research collaborations. The results of these discussions and findings are summarized below.

Device physics models exist for defects due to both the manufacturing process (extrinsic) as well as the fundamental material/device interaction (intrinsic). Intrinsic failures are caused by aging and wear-out of devices due to specific operating conditions over a period of time, leading to device degradation or complete failure. Intrinsic failures can be due to such effects as:

- Electro-migration (EM) in wires and metal traces
- Time-dependent dielectric breakdown (TDDB)
- Negative bias temperature instability (NBTI)
- Channel hot carrier (CHC) induction
- Stress migration, thermal cycling

These various processes have been well characterized in terms of dependence on temperature, electric field, current density, etc. [Sheldon 2009]. With the ever-increasing density of FPGAs due to constant device scaling from manufacturers, these effects have manifested themselves into three practical areas of concern:

- Power density
- Junction temperature
- Leakage currents

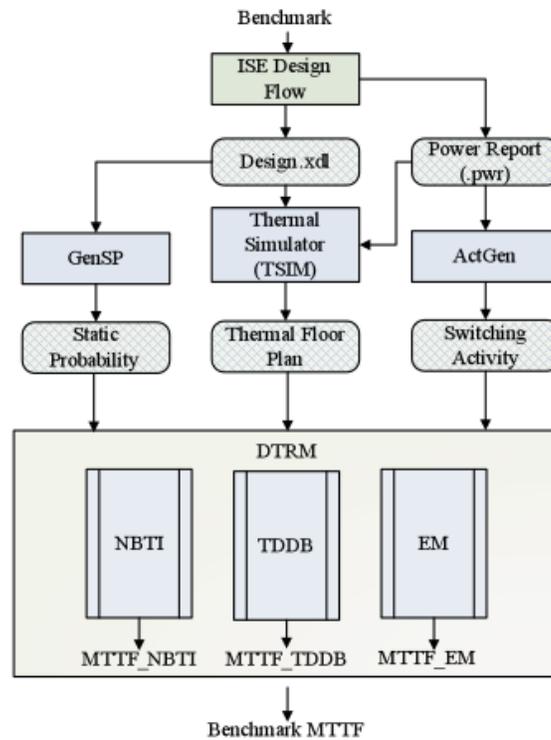


Figure 1. Design for reliability framework of FPGAs

Opportunities for mitigation of these three areas of concern are possible in the FPGA design environment and design process. Such a FPGA design process is shown in Figure 1.

In Figure 1, the initial ISE Design Flow is the standard integrated system development environment provided by Xilinx for their FPGAs. The ISE environment is a complete design suite providing simulation, timing and analysis, input/output (I/O) development, and synthesis. Figure 1 describes a modified flow where output .xdl files are used as input for a variety of simulators. Each of these simulators is designed to address power density (switching activity), junction temperature (thermal floor plan), and leakage current (static probability). These outputs go into the device physics-based models. These include the NBTI, TDDB, and EM blocks shown. Once these models have been established, their effect on device behavior is then integrated into the overall design and development flow of the device. Iteration and optimization can then take place in this environment. Techniques can be developed to minimize degradations and wear-out effects. Some of the techniques that have been developed as a result of the iterative FPGA design environment include [Mangalagiri et al. 2007]:

- Gate leakage optimization
- Region-constrained placement for reliability
- Selective alternative routing technique
- Bit relaxation

Gate leakage optimization, for example, includes optimizing input look up table (LUT) vectors that ensure conditions of minimum gate leakage. Using this technique, FPGA mean time to failure (MTTFs) have been increased by 24%, as shown in Figure 2.

This type of work is well established in literature and exists at a high level of sophistication. Several large research universities work on the development of these device physics models with strong support from FPGA vendors Xilinx and Altera. Often the results of the work end up in commercial FPGA design software as improvements and updates to algorithms that are transparent to the user.

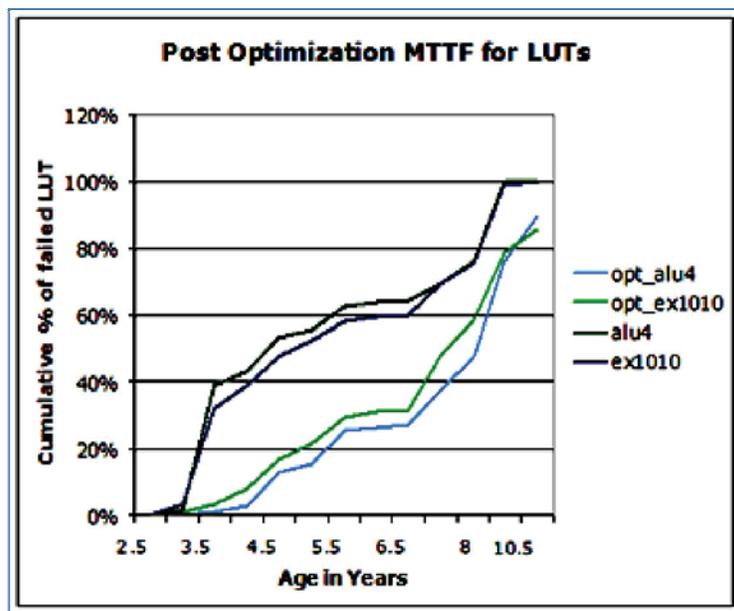


Figure 2. Optimized MTTF (blue-green lines) for various FPGA LUTs

One goal of the FY09 FPGA NEPP task was to establish contacts in the FPGA device physics research community. Such contacts would provide a baseline for future collaborations. Meetings were held in FY09 with Professor Vijay Narayanan of Penn State. Professor Narayanan has published more than 300 refereed articles in journals and conferences in the areas of power-aware and reliable systems, embedded systems, reconfigurable architectures, nano-architectures and computer architecture. As a result of this collaboration, JPL submitted a specific device-physics-focused FPGA proposal for FY10.

The goal is to integrate the techniques and results that Professor Narayanan has developed into NASA's FPGA design flow. This would allow for the maximum possible impact on FPGA reliability to occur at the earliest point (and most critical point) in the development flow, the design step. JPL plans to continue work with Professor Narayanan in FY10 and hopes to expand the collaborative level.

3 CONTENTION STUDY—XILINX VIRTEX-4

At the opposite end of the FPGA reliability spectrum from device physics models is the area of user-induced reliability faults. Independent of how robustly constructed and designed an FPGA may in fact be, FPGA users can inadvertently compromise FPGA reliability by making mistakes in the board manufacture, system integration, and final test. These user-induced reliability faults and concerns are often the result of some electric overstress condition being applied to the FPGA for a relatively short period of time, usually several hours to several days. Determining whether or not these situations actually degrade FPGA performance is a major risk mitigation concern for NASA projects.

FPGA device installation and in-system testing can produce (usually unintentionally) stressful circumstances to the devices represented by timeframes in between normal long-term (thousands of hours) accelerated life test conditions and short-term (tenths of seconds) ESD/EOS stress. Extrapolating either one of these classical models to fit a unique engineering circumstance often produces questionable and hard to justify results. The subsequently described contention experiment was designed to provide accurate results for short- and long-term contention conditions and to quantify possible device degradations as a result.

Contention is a practical concern for all FPGA designs. Contention can be defined as simultaneously driving two (or more) different electrical signals across a single electrical conductor. This is literally connecting a “1” to a “0,” for example. The result of such a connection can range from a short-term output error to a full-scale device malfunction.

For example, bus contention occurs when multiple devices simultaneously attempt to control a bidirectional bus during power-up, which can affect I/O reliability. Similar undesirable conditions where a “sneak” path from the rising input rail(s) to ground is temporarily created by transistors in unknown states can cause a digital device to pull large inrush currents, which may cause immediate damage or cumulative long-term reliability concerns. Recently, FPGA and Digital Signal Processor (DSP) manufacturers

have improved protection circuits to reduce the risk of latch-up, bus contention, and similar undesirable states [Falin 2006].

Contention can result from a hardware mistake but can also result from a software error. This is of particular concern for FPGA usage. While designs may operate fine independently, once integrated into the board, unknown contentions between FPGAs and other devices can easily occur due to the complexity of programming the I/Os of FPGAs. Configuration functions can be disrupted by signal contention between configuration inputs and the FPGA user outputs that become active at the end of configuration.

3.1 Contention Experiment

The contention experiment was done in collaboration with Xilinx Design Services [Muse 2009]. A characterization study was developed to focus on the effects of I/O contention on a Virtex-4 family device. This test consisted of tying two I/O drivers of a Virtex-4 FPGA together. One I/O was configured to output logic 1, the other to logic 0. This was performed for the following I/O standards, each configured with the maximum drive strength:

1. LVTTTL
2. LVCMOS 3.3V
3. LVDS

The study was organized in the following sequence of events:

1. A baseline data sample is taken.
2. A design is loaded into the FPGA to create the I/O contention.
3. After a predetermined amount of time, the contention is removed from the I/O and a new design is loaded into the FPGA for purposes of collecting a data sample.
4. Repeat steps 2 and 3 for varying contention times.
5. Collect data and analyze.

The contention study was designed to measure changes in the output I/O parameters by creating a simple design to output clock pulse signals of varying frequency on the pins used to create the contention and measuring changes in the following properties at ambient temperature:

- Signal amplitude
 - A 100 MHz clock signal is driven out of the device and measured across the appropriate termination header test pin using an oscilloscope. Maximum and minimum voltages are recorded over 10 K cycles.
- Signal jitter measurement
 - A 100 MHz clock signal is driven out of the device and measured at the test pin using an oscilloscope. The cycle-to-cycle jitter is recorded at the 50% voltage level as a maximum over 10 K cycles.
- Signal drive strength
 - This is a measurement of the voltage and current the FPGA is able to drive over a specified load. In the single ended cases (LVTTTL and LVCMOS), it is calculated by measuring the voltage drop across a resistor to ground as the FPGA output drives a logic “1” for source, and across a resistor to 3.3 V as the FPGA drives “0” for sink. In the differential LVDS case, the voltage drop over a resistor between the P and N sides of the driver is measured.
- Signal rise/fall times
 - A 100 MHz clock signal is driven out of the device and measured across the appropriate termination at the header test pin using the oscilloscope. Rise and fall are defined as 10% to 90% of the full voltage reached as an average of 10 K cycles.
- FPGA power consumption
 - Power consumption is calculated for the three voltage rails that power the FPGA: 1.2 V, 2.5 V and 3.3 V. Board power consumption for each rail can be calculated by measuring the voltage drop across a low ohm series resistor on each power supply. All power measurements are made with the multimeter, and while the I/O is driving a 100 MHz clock

LEDs indicate when the heat circuit is active. An external fan and heatsink are used to reduce the device temperature automatically if it exceeds 64°C. The minimum and maximum temperature measurements are stored for display, but can be reset at any time with the center button of the button array. The present, minimum, and maximum temperature readings are printed to the onboard LCD and are also printed to the serial UART interface where they can be read using a terminal emulation program on a PC.

In the LVTTTL and LVCMOS33 cases, the outputs are configured to drive opposite logic levels, and externally the pins are jumpered together to create the desired contention condition. In the LVDS case, one differential driver is configured to drive logic “1,” while a second is set to drive “0” and externally the pins of both pairs are tied together.

The contention time on the I/O of the FPGA was organized for (30 min, 60 min, 300 min, 24 hour, and then 24 hour increments) up to 900 hours. The testing was performed on three separate Virtex-4 XC4VFX20-FF672 devices. These devices were tested on the Xilinx ML405 development board with multiple banks setup for the I/O contention testing. A picture of the ML405 development board is shown in Figure 4.

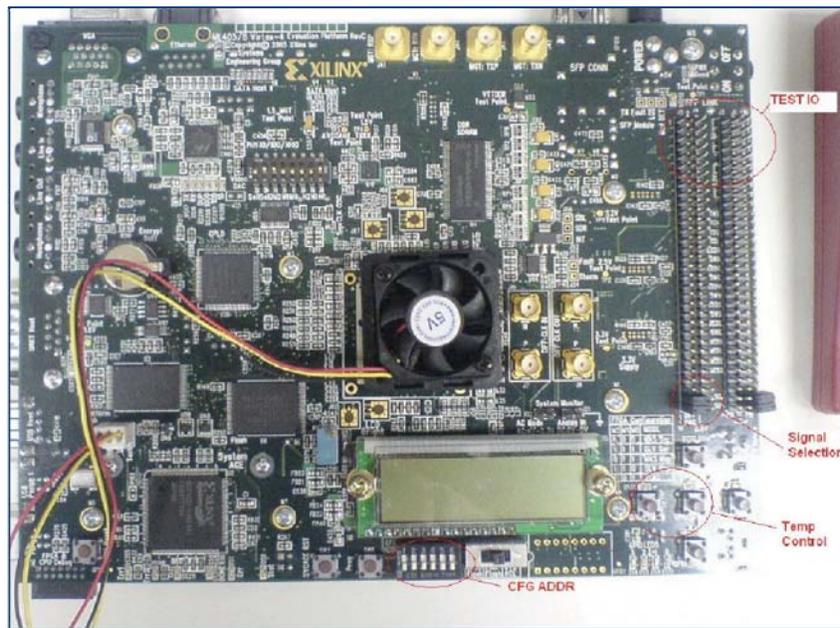


Figure 4. Xilinx ML405 test board

3.2 LVDS Summary

In each of the measurement categories, data did not vary considerably from the baselines in a way that would suggest the drivers were significantly damaged through their exposure to contention with other LVDS drivers. In each case, drivers did not appear to be negatively affected by prolonged periods of contention, nor did accumulated contention time show a clear sign of degrading the driver in such a way that would cause a loss of signal integrity. LVDS pair 1 data are summarized in Table 1.

Table 1 shows the rise/fall times first decreasing by approximately 9% then increasing to values approximately 4% higher than baseline. As the testing continued, the rise/fall times then started to decrease again. These data are graphed in Figure 5. This trend was not observed in the other measurements.

This trend in rise and fall times was seen on both LVDS boards. This occurred in measurements of both pairs. One point of view is that it is unlikely that this drift was caused by contention, as this would mean that the driver became stronger, then weaker, then stronger again as contention time accumulated.

Table 1. Pair 1 LVDS data

Time (hrs)	Rise (ps)	Fall (ps)	max (mV)	min (mV)	dper - (ps)	dper + (ps)
0	345.69	360.38	424.43	-428.07	-159.60	164.20
0.5	343.52	356.30	416.12	-429.57	-155.50	162.00
1	358.98	374.91	418.11	-427.58	-154.50	155.00
5	317.02	326.26	427.95	-435.66	-154.00	156.50
24	316.76	326.45	423.51	-435.23	-159.00	165.50
48	319.75	329.36	424.03	-434.58	-161.50	166.00
72	316.19	326.12	419.97	-437.10	-160.50	162.00
96	340.14	338.00	422.78	-435.61	-159.00	164.00
120	352.42	367.84	420.54	-427.57	-157.50	158.00
192	374.90	392.96	421.51	-425.75	-157.50	157.50
216	360.24	375.92	421.63	-425.77	-153.50	162.50
288	370.44	386.34	421.87	-425.08	-163.50	165.00
312	377.62	395.98	421.30	-426.36	-160.50	159.00
336	369.10	385.86	419.50	-424.55	-158.50	165.00
360	383.98	403.53	422.49	-426.38	-153.00	158.00
384	364.29	380.57	422.09	-423.60	-162.50	161.50
456	366.28	383.01	418.68	-423.50	-152.00	158.50
552	361.26	379.23	419.66	-429.12	-154.50	157.00
648	371.20	385.50	419.23	-430.64	-153.50	156.00
768	342.26	353.65	423.36	-429.69	-151.50	159.00
888	341.48	354.25	424.98	-432.76	-156.50	159.00

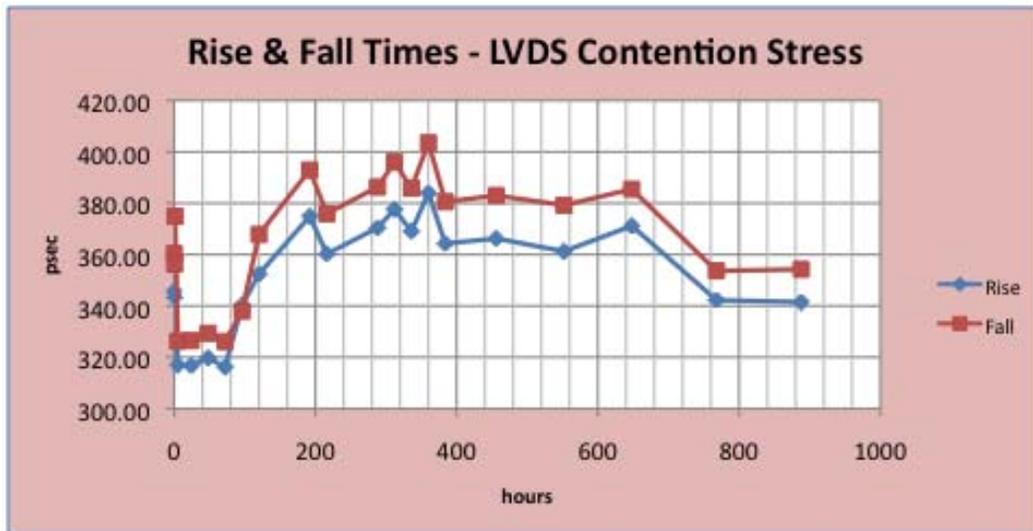


Figure 5. LVDS rise and fall times vs. contention time

The drift could then be explained by some external environmental variation such as ambient temperature or probe placement. Further investigation has shown that small variations in probe placement do cause like variations in the measured rise and fall times. Additionally, and perhaps more importantly, the range of change seen (again on the order of 30–40 ps) is beyond the bandwidth capabilities of the equipment (6 GHz scope and 7.5 GHz differential probe), and therefore could simply be attributed to sampling error.

Maximum and minimum differential voltages varied within very small ranges from the baseline average. Cycle to cycle jitter measurements varied by a very small amount from the baseline average as well. Static drive voltages measured at static high and low logic levels varied by less than 1 mV over the course of the experiment.

Power consumption varied within a small range. The current was calculated by measuring the voltage drop across the low ohm (0.003 ohm) resistor on each supply. The voltage across the 0.003 ohm resistor for the 1.2 V supply varied by +0.04 mV/-0.01 mV, the 2.5 V measurement varied by +0.009 mV/-0.004 mV, and the 3.3 V measurement varied up to 0.04m V. These measurements tended to be slightly higher than the baseline. See Figure 6 for the currents graphed against contention time.

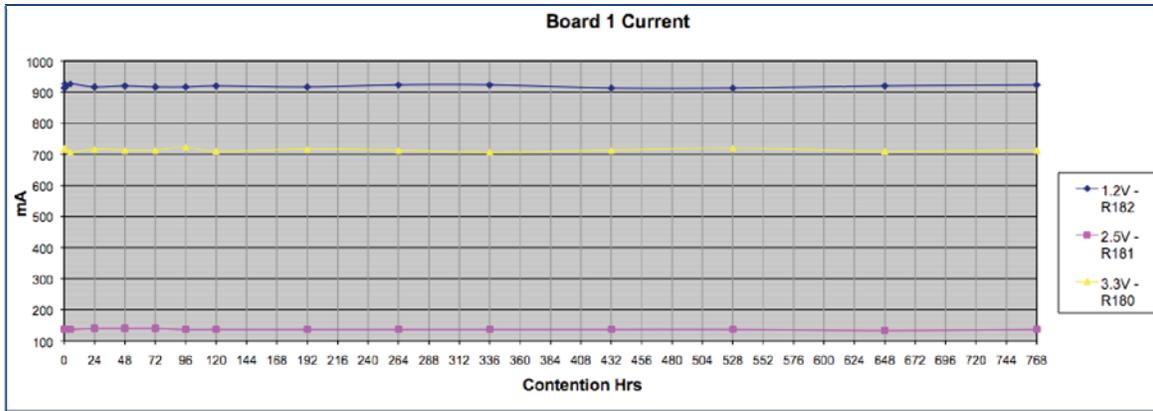


Figure 6. Current versus total contention time

In each of the measurement categories, data did not vary considerably from the baselines in a way that would suggest the drivers were significantly damaged through their exposure to contention with other LVDS drivers. In each case, drivers did not appear to be negatively affected by the periods of contention experienced, nor did accumulated contention time show a clear sign of degrading the driver in such a way that would cause a loss of signal integrity.

3.3 LVTTTL and LVC MOS3 Data

Table 2 provides sample data for LVC MOS. Five percent changes in rise and fall time were noted. Additional statistical analysis was performed on the LVC MOS and LVTTTL data to determine if the observed data are statistically significant. This test is a hypothesis testing result. Data are summarized in Appendix A.

Table 2. LVC MOS data

Hrs	Rise (ns)	Fall (ns)	max (mV)	min (mV)	dper - (ps)	dper + (ps)	Ave Max (mV)	Ave Min(mV)
0	1.58	1.70	2,850.24	511.66	-302.80	315.20	2,794.70	559.78
0.5	1.62	1.75	2,856.20	522.65	-348.50	332.50	2,791.30	565.55
1	1.62	1.76	2,858.50	525.45	-351.00	315.50	2,798.00	573.75
5	1.62	1.76	2,843.00	502.70	-291.00	295.50	2,793.90	559.30
24	1.63	1.75	2,843.55	507.30	-277.50	284.00	2,790.65	565.95
48	1.61	1.75	2,842.40	521.30	-318.00	308.50	2,789.00	565.90
72	1.62	1.74	2,850.80	517.80	-275.50	310.50	2,792.15	564.90
96	1.63	1.76	2,848.10	521.25	-280.50	289.00	2,794.45	565.15
120	1.62	1.75	2,858.40	507.65	-294.50	296.50	2,793.60	563.65
192	1.55	1.68	2,854.85	507.70	-273.00	258.00	2,796.05	560.95
264	1.57	1.70	2,851.45	515.80	-332.00	343.00	2,794.60	560.35
336	1.61	1.73	2,827.20	525.80	-301.00	298.00	2,776.95	573.70
432	1.54	1.62	2,855.55	505.30	-329.50	333.00	2,806.65	552.45
528	1.61	1.73	2,849.35	509.55	-300.50	308.00	2,792.15	560.65
648	1.65	1.77	2,837.60	527.85	-295.50	319.50	2,779.05	572.05
768	1.57	1.69	2,857.20	517.20	-343.00	326.50	2,793.90	567.10

When the two tailed probability test is $<5\%$, this means there is a greater than 95% likelihood that the two distributions are different. For this contention study, this means that if this condition occurs, then the average values of rise and fall times are different than the baseline values to a high degree of statistical significance. The results from Appendix A show that 100% of the LVTTL rise/fall times were different between baseline and average contention stress while only 25% of the LVCMOS rise/fall times were statistically significant.

As a verification of this result, a control board was also used during the test. No contention conditions occurred on this board. The control board was subjected to the same measurement procedures as the test boards. Sample results are shown in Appendix B for the control board. The statistical analysis in Appendix B shows that end of test data was not statistically different than the beginning experiment data, as expected. This means the changes in the LVCMOS and LVTTL should be considered real.

The existence of these changes in rise and fall times of I/O performance does not mean that the I/Os failed data sheet conditions. To the contrary, the I/Os continue to perform as expected and continue to meet manufacturers specifications. Work is continuing to provide a device physics level explanation of these measured changes. From a practical FPGA implementation point of view, maintaining signal integrity is a first order design concern. Due to the reconfigurable nature of the Xilinx Virtex-4 device, if the rise/fall times were to reach an unacceptable level, they could always be “sped up” by either increasing the drive strength, changing the slew rate to “FAST,” or both. As long as the board was designed initially to the lowest settings possible, there is considerable margin available.

4 IDDQ STUDY IN ACTEL A54SX FPGAS

Antifuse-based FPGAs remain the vast majority of FPGA technologies used by NASA. The inherent rad hard nature of the antifuse is the main reason for choosing these devices. Passing a high current through a titanium-based metal electrode into a silicon oxide/carbide dielectric layer to another titanium-based metal electrode forms the antifuses. This high current produces a local (~10 nm) high temperature that can exceed the melting point of silicon (>1400°C). This very high temperature melts the titanium and thermo-migrates it through the dielectric to form a thin filament of connecting material between the two electrodes. This filament can be used to connect drive signals into FPGA resources or to provide interconnection paths for signal propagation across a chip.

Once the antifuse structure is formed, the FPGA is considered completely “programmed” and ready for operation. Historical generations of antifuse FPGAs have shown reliability failures as a result of the programming process [Sakaide et al. 2004]. These programming failures resulted in NASA and other spacecraft providers having to enact multilevel risk mitigation steps [Sheldon 2009]. Many of the risk mitigation steps are trying to address the quality and reproducibility of the programming process. The FPGA programming machine does not provide quantitative information to the user, only qualitative. The device either “passes” or “fails.” The details of how the device passed, whether or not it was near functional current limits, etc. are not available to the user.

This experiment was designed to address this issue of programming repeatability. By quantifying programming repeatability, JPL hopes to gain insight into the variation of the antifuse resistance values. The possible presence of a large variation in antifuse resistance values for the same design would translate at least qualitatively into an increased reliability risk. Being able to measure antifuse resistance at least indirectly will set a stage for an evolution of risk mitigation schemes.

Historically, FPGA risk mitigation involved a tri-temperature test of the design once programming was completed. This tri-temp test is complicated in the fact that the programmed FPGA is now a custom-designed ASIC and requires custom/design-specific

test vectors to adequately implement. In most, if not all practical situations, tri-temp testing was waived as being too costly and providing too little “actionable” information.

This experiment, however, performed tri-temp testing to determine usefulness and applicability to overall FPGA risk mitigation. Along with the functional tri-temperature test, an IDDQ test campaign was also performed. IDDQ testing has been shown to be a very accurate predictor of reliability degradation.

4.1 IDDQ Experimental Design

For this experiment, the Actel SX-A family of devices was chosen. Specifically, the commercial AS54SX72A device was used exclusively throughout the testing. The AS54SX72A device is a 220 nm complementary metal-oxide semiconductor (CMOS)-based device. The device has a maximum of 108,000 system gates with 6,036 logic modules.

A single design was chosen for this test. This design was programmed into several different FPGAs and the resulting IDDQ current was measured. Differences in the IDDQ current can be initially assumed to be due to the variations in the antifuse programming. The design chosen was a finite impulse response (FIR) design. This design can be scaled in terms of the number of logical resources that are implemented. Scaling the design is an important variable in helping to determine variation in antifuse programming.

The basic experimental design is shown below in Table 3. There are three different variables—design resources used, temperature, and voltage. Each of these three variables has three different “levels.” This makes for a total of 27 different combinations of testing/device conditions. A total of 15 separate AS54SX72 FPGAs were used. Five FPGAs were used for each design resource level. Then each of these five FPGAs was measured at three different voltages and three different temperatures. This is a total of 135 separate experimental design inputs.

Table 3. 3³ experimental design for IDDQ testing

Design Resources	Voltage	Temperature
10%	Nominal – 10%	-40°C
50%	Nominal	25°C
80%	Nominal + 10%	85°C

4.2 FIR Design

An FIR filter is a type of a digital filter. FIR filters have been used for many years in DSP applications, including signal conditioning, anti-aliasing, and convolution [Knapp 1988]. The impulse response, the filter’s response to a Kronecker delta input, is finite because it settles to zero in a finite number of sample intervals. There is no feedback in the FIR filter. This is in contrast to infinite impulse response (IIR) filters, which have internal feedback and may continue to respond indefinitely. The impulse response of an Nth-order FIR filter lasts for N+1 samples, and then dies to zero. An example 8-bit FIR filter is shown in Figure 7.

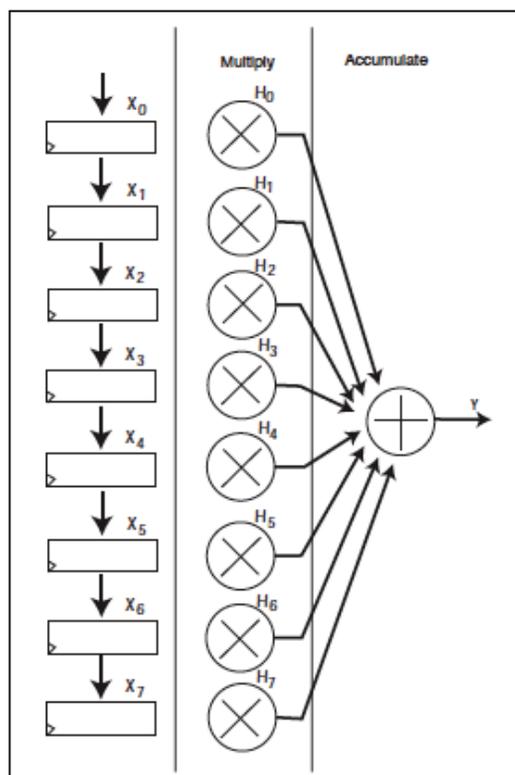


Figure 7. 8-bit 8-tap FIR filter

The FIR filter in Figure 7 receives a data sample X_0 every data clock cycle. All the other values, X_{1-7} , are time delayed values of X_0 . The filter multiplies each data sample X_n by its corresponding constant filter coefficient, H_n . The filter then sums the products from all the multiplications to produce the final result Y . Mathematically, this operation can be expressed as:

$$Y = \sum_{n=1}^N X_n H_n \quad (1)$$

N is the filter order. An N th-order filter has $(N + 1)$ terms on the right-hand side; these are commonly referred to as taps. The number of FIR taps is an indication of the amount of memory required to implement the filter, the number of calculations required, and the amount of “filtering” the filter can do. More taps means more stopband attenuation, less ripple, and narrower filters. Sample FIR filter output is shown in Figure 8 and Figure 9. These figures are simulations of FIR filters made with MATLAB’s discrete-time filter object *dfilt*. The filters simulated here were for a low pass filter with a 9.6 KHz frequency and a stop frequency of 12 KHz. Doubling the amount of taps/order from 12 to 24 taps substantially improves the filters response.

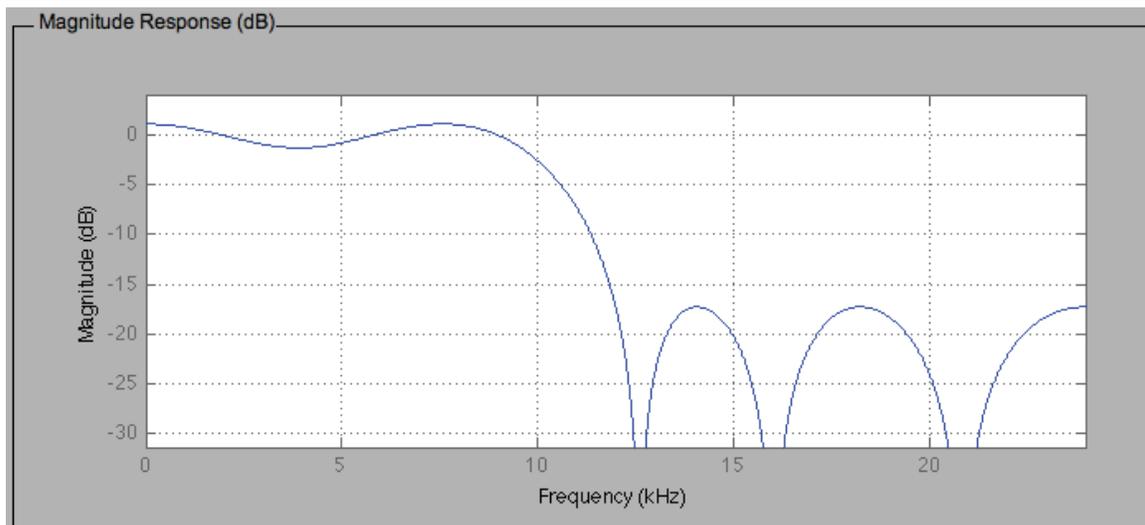


Figure 8. 12 order FIR filter (low pass/ $F_{PASS}=9.6$ kHz, $F_{STOP}=12$ kHz)

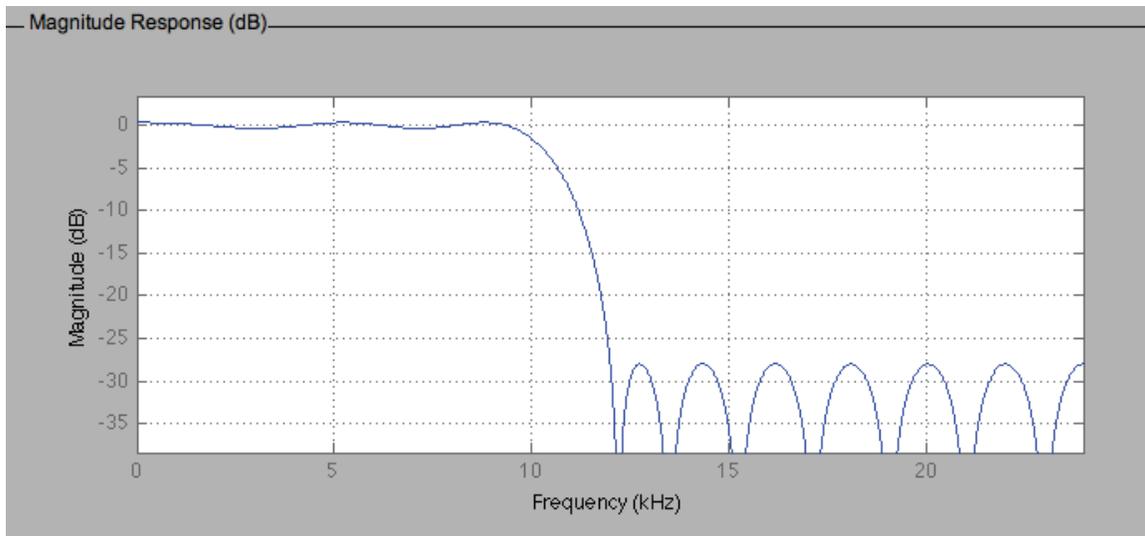


Figure 9. 24 order FIR filter (low pass/ $F_{PASS}=9.6$ kHz, $F_{STOP}=12$ kHz)

For the FIR filters used in this experiment, Actel’s CoreFIR filter generator were used. The CoreFIR is an Actel-supplied and optimized RTL generator that produces a FIR with an intellectual property (IP) core generator optimized for use with Actel’s FPGAs. Using the core generator, designers can input the number of taps, bit width and coefficient bit width, and the ratio between system clock frequency and data sampling rate for a variety of configurations, including signed or unsigned inputs, fixed or variable coefficients, and embedded RAM. In addition, the core generator outputs RTL code and a test bench.

The number of combinatorial and sequential cells can be varied for the CoreFIR design. This reflects changes to the number of input bits, coefficient size, and number of taps. For this experiment, these values were varied so that the total amount of resources consumed could be varied. The result was three different designs that used 10%, 50%, and 80% respectively of the overall AS54SX72 resources.

4.3 FPGA Testing

Each FPGA was programmed as a 10%, 50% or 80% resource utilization device. Then, the FPGAs were tested at -40°C , 25°C , and 85°C . At each different temperature, the voltage of the device was varied from nominal to (nominal -10% or low) and finally to (nominal $+10\%$ or high). Each device was subjected to 10,000 different test vectors at each voltage and temperature combination. These 10,000 test vectors were random

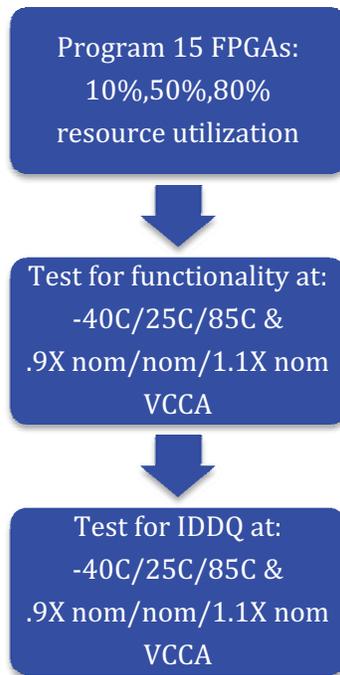


Figure 10. Testing flow for FIR IDDQ experiment

numbers generated as part of the verification test bench. After each unique vector was sent in, the resulting output was then compared and confirmed to be correct. IDDQ was measured at both the periphery as well as the array (VCCI and VCCA) between each unique application of these random vectors. The amount of time between a vector application/confirmation and the completion of an IDDQ measurement was 12 msec. The current from VCCI and VCCA was measured to ground. Figure 10 summarizes this test flow.

4.4 FPGA Testing Results

Both IDDQ and functional performance were obtained for each FPGA DUT for each combination of voltages and temperatures. Functional voltage and timing information is plotted in a two-dimensional fashion to obtain a schmo plot [Schmoo]. Examples of this type of analysis are shown in Figures 11 through 13. The schmo plot shows functionally passing and failing areas, marked in green and red, respectively. The FIR filter's output is compared to a known correct output generated during FPGA synthesis. For a specific set of input values to the FIR filter, the output can be exactly calculated. The FIR filter is then operated over the voltage and temperature range previously mentioned.

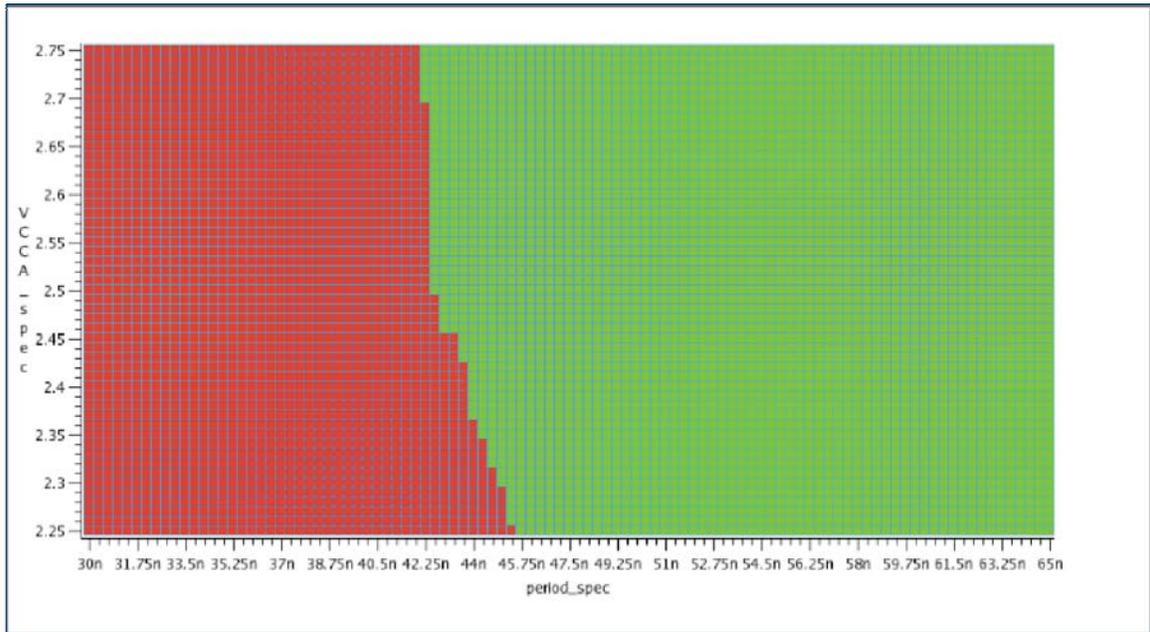


Figure 11. Schmoo plot of frequency versus VCCA – 80% design @ -40°C

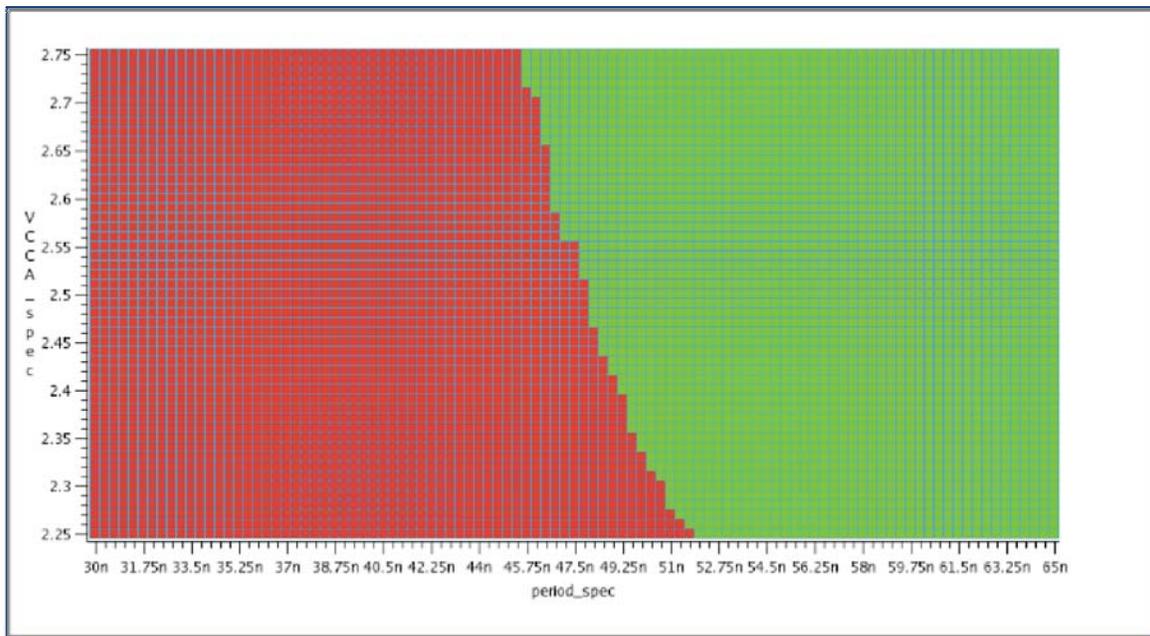


Figure 12. Schmoo plot of frequency versus VCCA – 80% design @ 25°C

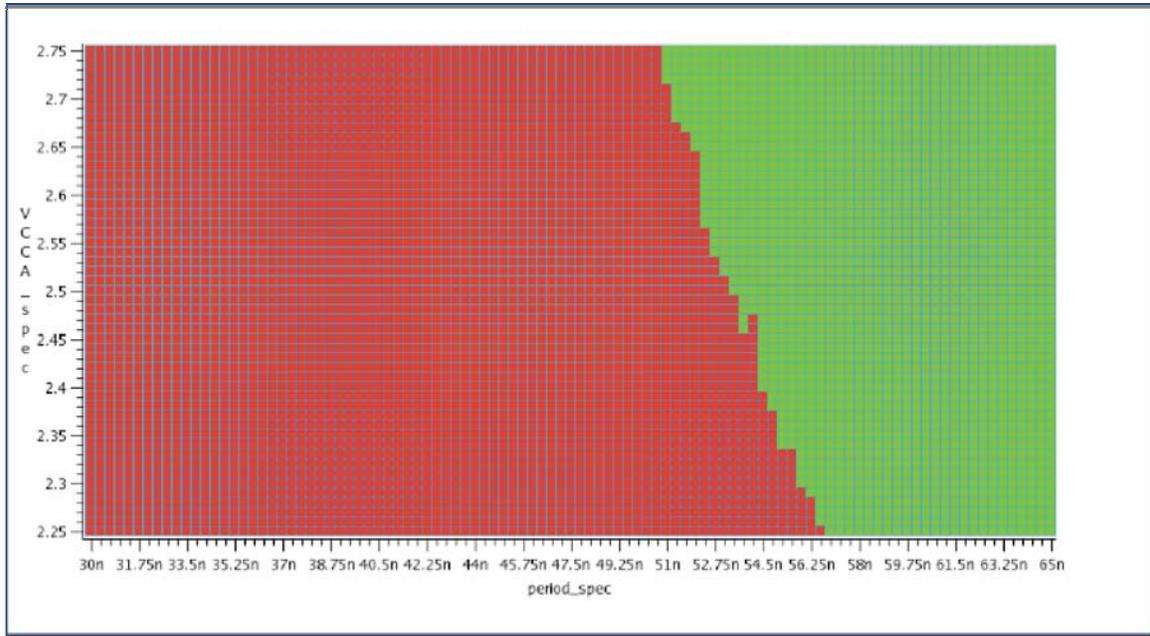


Figure 13. Schmoo plot frequency versus VCCA – 80% design @ 85°C

The 80% design usage DUTs show a strong dependence on temperature for minimum passing time. There is also a well-defined negative slope for timing dependence on VCCA voltage. Increasing the voltage decreases the passing period (or increases the passing frequency). Figure 14 summarizes the schmoo plot information.

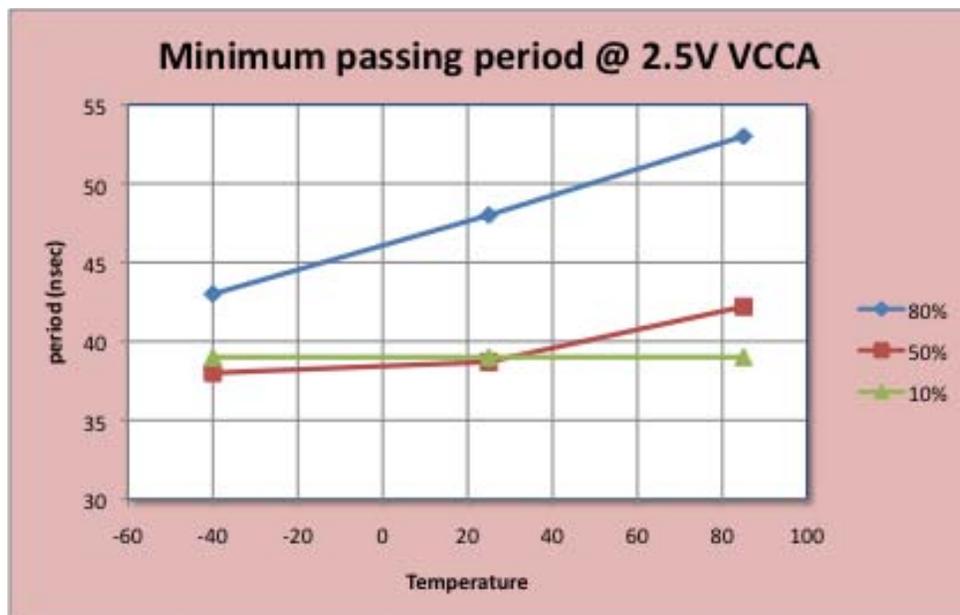


Figure 14. Minimum passing period from schmoo plots

Figure 14 shows that the 80% resource design has a strong dependence on temperature, while the 10% and 50% design resource splits show almost no dependence on temperature. This temperature dependence may be due to local die heating effects caused by the large resource utilization. Lowering the temperature increases transistor switching speed and decreases leakage. Increased contribution of interconnect signal paths may also be contributing to this temperature dependence for the 80% resource utilization design.

Figure 15 is a cumulative box plot of all three design resource utilizations. The data are taken from the 25°C and VCCA=nominal split. Box plots show both mean and variation in terms of inter-quartile ranges (IQR). Each box in Figure 15 represents nearly 5,000 data points. Figure 15 shows between 35% to 75% variation for the mean values for each design resource level. This amount of variation is directly indicative of the antifuse variation from FPGA to FPGA. This important conclusion is that independent of the FPGA design, the amount of variation due to the programming step and the underlying antifuse process is between these 35% to 75% values. This is the inherent variation of this particular FPGA device and process.

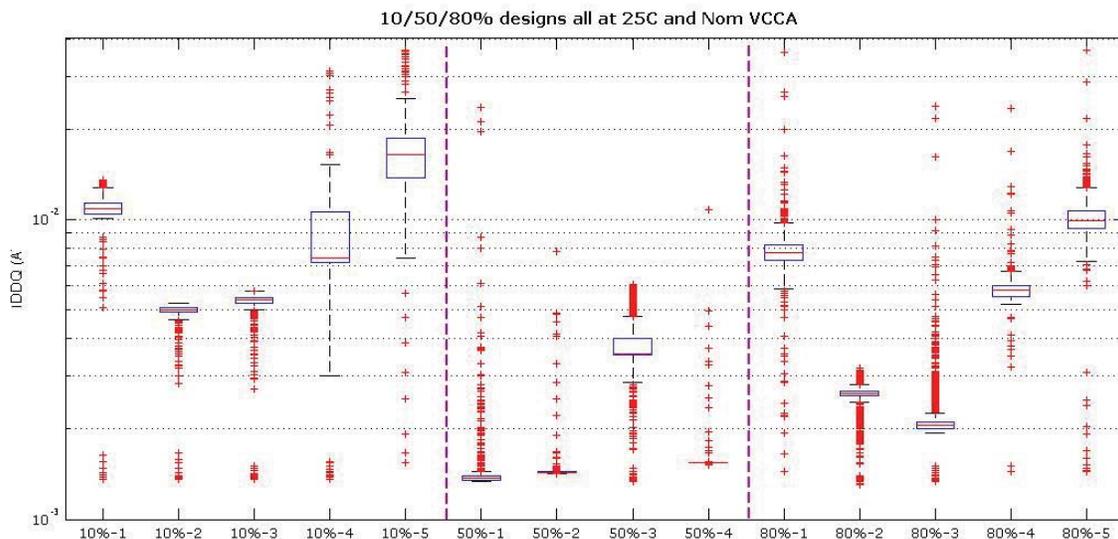


Figure 15. Cumulative box plot of IDDQ, all design levels for 25°C/VCCA=nominal

This result provides positive re-enforcement for the use of the tri-temperature test as a risk mitigation step for NASA and JPL FPGAs. If the proposed design consumes a large amount of current, somewhere with 50% of the maximum values listed by the manufacturer, then there is a strong likelihood that several attempts at programming FPGAs will result in devices that produce current values in excess of manufacturers' recommendations.

The variation in Figure 15 shown by the IQR values is not evenly distributed about the median value. This is indicative of systematic bias in the raw data. Examples of the raw IDDQ data are shown in Figures 16 and 17.

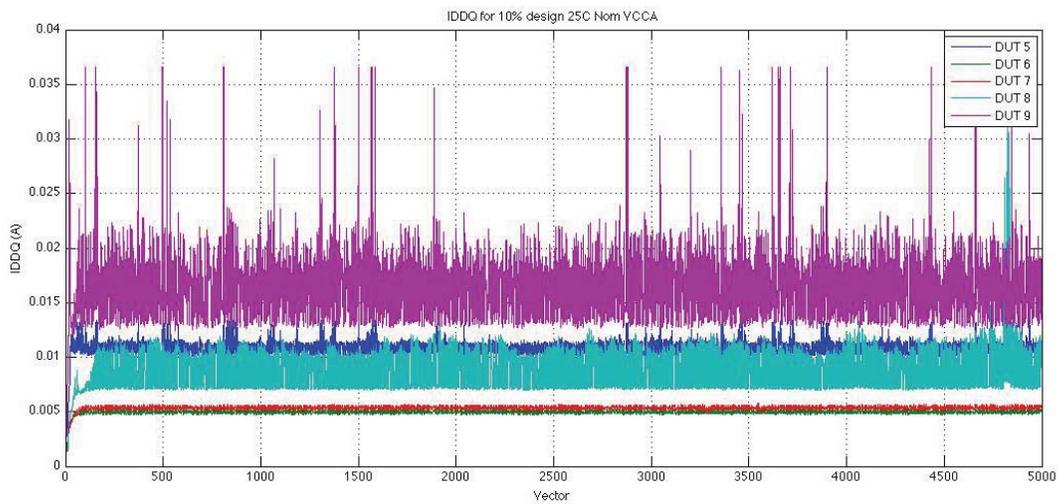


Figure 16. IDDQ vs. vector for 10% designs at 25°C and nominal VCCA

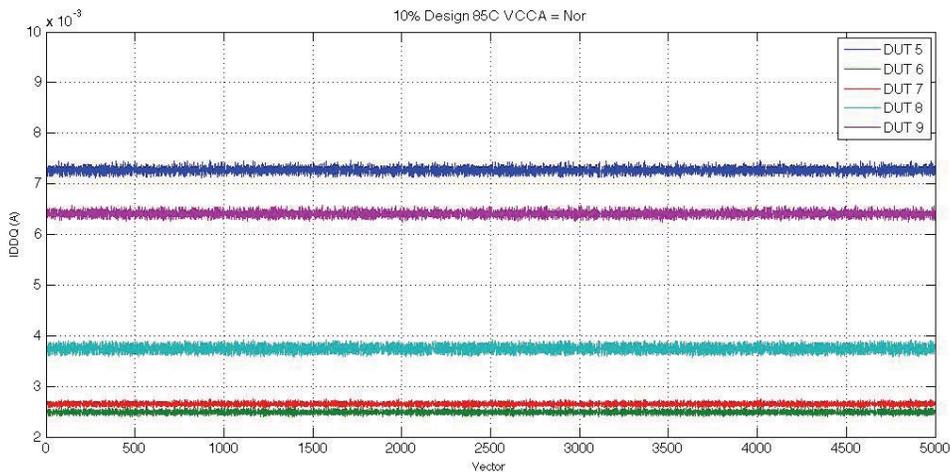


Figure 17. IDDQ vs. vector for 10% designs at 85°C and nominal VCCA

The IDDQ data at 85°C shows well-defined values and small amounts of variation. The nearly 65% variation from 2.5 mA to 7.2 mA is clearly illustrated. The data at 25°C in Figure 16 shows well-defined IDDQ results from two FPGAs (#6 and #7), while the other three FPGAs show a well-defined floor result with excursions from 5 mA up to 20 mA. These excursions above a floor value explain the box plot data in Figure 15. The IDDQ data for -40°C has a large amount of non-linear structure in it. An example of this is shown in Figure 18.

All devices tested at -40°C show at least one large spike in IDDQ of usually 1.5 orders of magnitude, from ~1.5 mA up to 35 mA. There is often more than one spike. The instrumentation was current-limited to 35 mA. This explains why the current spikes show a clipped result. A rough graphics estimate suggests that the actual maximum current spike value is between 60 and 70 mA.

Once 10% to 40% of the test vectors have been applied, the IDDQ current begins to rapidly decrease to a stable value. This cold (-40°C) IDDQ result was seen independent of design resource, voltage, or test temperature. Antifuse FPGAs have been shown to work down to -120°C without any current surges while SRAM-based FPGAs can produce a large start up current surge at the core supply [Burke et al. 2004]. The authors found that start-up current surges increased with decreasing temperature.

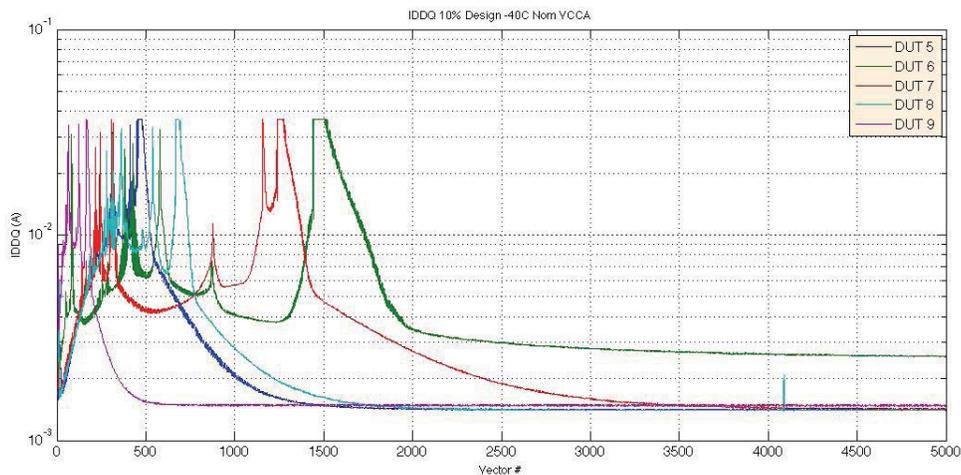


Figure 18. IDDQ vs. vector for 10% designs at -40°C and nominal VCCA

This low temperature behavior could be a system reliability risk. If such a FIR filter was implemented and operated at low temperatures, but only for a very minimal operational duty cycle, the system would experience these significant current spikes

A “well-behaved” IDDQ curve is defined as a stable IDDQ value independent of test vector with no “ramp up” or “decay” period of time. These well-behaved curves only occurred for the 25°C and 85°C test conditions. The summary of all well-behaved IDDQ curves is shown in Figure 19.

It can be seen from Figure 19 that there is an increase in well-behaved IDDQ curves as temperature is increased. There is no evident correlation with voltage.

At this time, the true nature of the variation in the IDDQ curves is unknown. Authors have reported non-linear results in IDDQ measurements at 0.18 μm size devices and below [Okuda and Furukawa 2003]. JPL is continuing to evaluate the possibility of such non-linear effects occurring for this particular measurement setup and experiment. The present working assumption is that the measurements are correct and will continue to be investigated. Actel has been contacted for further detailed design and verification help. With this stated assumption, the following discussion regarding the IDDQ results is presented.

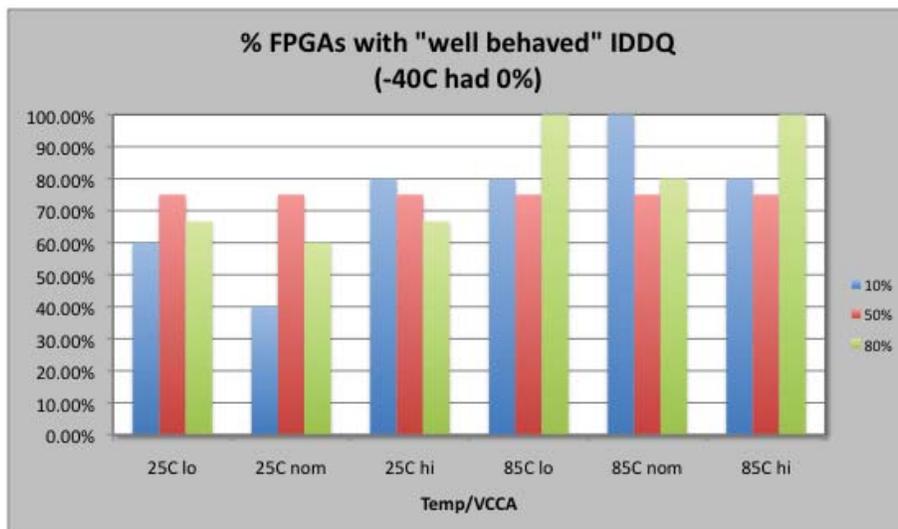


Figure 19. Percentage of well-behaved IDDQ curves for different test conditions

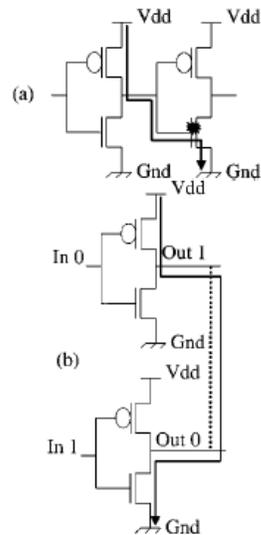


Figure 20. CMOS inverter with defects detected during IDDQ testing [Sabade]

IDDQ has been an important part of CMOS device testing since the mid-1980s. IDDQ stands for quiescent IDD, or quiescent power-supply current. In a steady state situation, when all switching transients are settled-down, a CMOS circuit dissipates almost zero static current. The leakage current in a defect-free CMOS circuit is negligible (on the order of few nanoamperes). However, in case of a defect such as gate-oxide short or short between two metal lines, a conduction path from power-supply (Vdd) to ground (Gnd) is formed and, subsequently, the circuit dissipates significantly higher current. This faulty current is a few orders of magnitude higher than the fault-free leakage current. Thus, by monitoring the power-supply current, one may distinguish between faulty and fault-free circuits. This has been the main reason for doing IDDQ testing. An example of a defective CMOS inverter is shown in Figure 20.

For this particular experiment, the FPGAs are all considered “good” because they have passed the programming step. Therefore, it was not expected to have significant IDDQ current variations that would indicate FPGAs with CMOS faults due to manufacturing. One concern is the existence of undefined high current states in the circuit during this particular FPGA testing. Such high current states are not faults. Some examples of high current states might be:

- Not having all flip flops (registers) in a known, initialized state. Usually, this type of initialization is done by a set/reset signal or through scan operation.

- The existence of some form of static current dissipating logic that has not been switched off. Usually, this includes memory sense-amps, dynamic logic, asynchronous logic, pull-up/pull-down resistors, special I/O buffers, and analog circuitry.
- The circuit not being stable at the strobe point. The stability is defined as there being no pending events.
- Improper termination of all inputs and bi-directional pins. These should be either at 0 or at 1.

IDDQ increases with vector number have been reported for various CMOS circuits [Sabade and Walker 2005]. Their results are reproduced in Figure 21.

The data in Figure 21 shows two CMOS die, A and B. Die A has no manufacturing defects while die B has several. These types of defects are “active” defects that respond by having a higher than normal leakage. The defects in die B stay active once they are subjected to the appropriate test vector.

The use of IDDQ in this experiment was as a characterization tool and not as a means to identify manufacturing defects. All the FPGAs used had, of course, passed manufacturers test and were considered good. The fact that time- (vector) and temperature-sensitive IDDQ results have been demonstrated may indicate the existence of a quasi-stable device condition that is, in fact, related to the programming operation. Further testing will be required to explore this possibility.

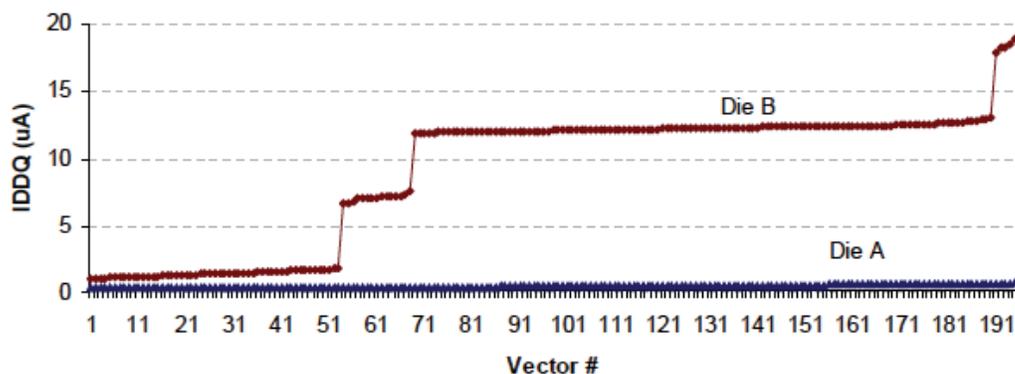


Figure 21. Reproduction of IDDQ increase with applied vectors [Sabade]

Another possible explanation of the results obtained is the possible “hysteresis” effects in the IDDQ measurements. Recent literature results demonstrate test pattern and device speed dependence on IDDQ that results in a hysteresis phenomenon [Okuda and Furukawa 2003]. The term hysteresis is used here to define oscillations in IDDQ as function of increasing test vector. This is somewhat similar to the results obtained in this task.

Okuda provides a plot of the difference in IDDQ vectors between the 100th vector and the 1st vector. This delta condition is then plotted against the IDDQ value of the 1st vector. The results are shown in Figure 22.

Figure 22 shows the hysteresis effect. As the 100th IDDQ vector increases, so does the difference between the 100th IDDQ value and the original value. The results shown in Figure 22 are for 423 different die. This correlation indicates that the values are not random, but must depend upon a physical condition. This could be either a systematic defect or an intrinsic IDDQ defect. A similar calculation was performed on the 80% resource utilization data from this experiment. These results are shown in Figure 23.

These data show a similar positive correlation. This is for a reduced sample size of three out of the five samples in the 80% resource utilization, nominal VCCA, and 25°C split. The authors define a charge trapping mechanism to explain these results. Part of that derivation is reproduced here for further discussion.

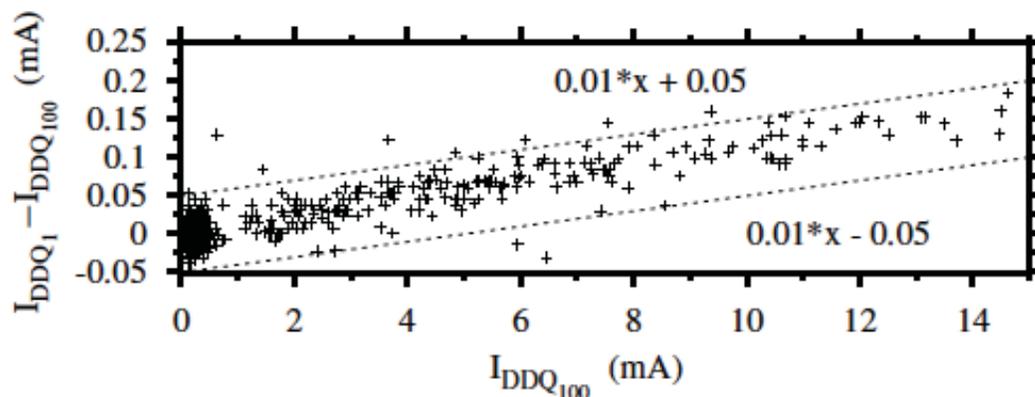


Figure 22. Delta IDDQ versus 100th IDDQ vector [Okuda]

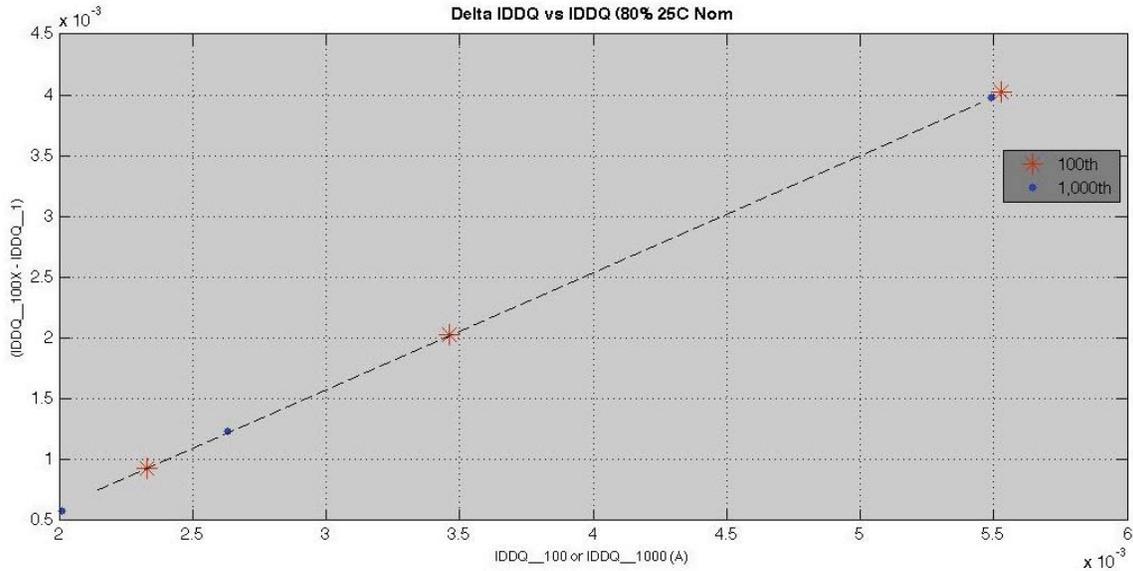


Figure 23. [IDDQ 100 (or IDDQ 1000) – IDDQ 1] versus IDDQ 100 (or IDDQ 1000).

Starting with the definition that IDDQ is dominated by subthreshold transistor leakage, I_{sub} , we can derive an expression to relate changes in IDDQ with accumulated charge. This accumulation of trapped charge is an aging mechanism and indicator of future long-term reliability performance. Subthreshold current I_{sub} has temperature, voltage, and device dimension dependence and can be expressed as:

$$I_{ddq} = g(T)f(V,d) \quad (2)$$

When repeated IDDQ vectors are applied to the circuit, Equation 2 becomes:

$$I_{ddq_i} = \sum_i g_i(T)f_i(V,d) \quad (3)$$

Now IDDQ is the result of a combination of the vectors and the timing of the vectors applied to the device. Repeated application of IDDQ vectors will cause hot carrier and high electric field conditions to exist, even if briefly, in the transistors of the FPGA being tested. This repeated application would result in charge trapping to occur for a variety of the previously mentioned mechanism, even if the contribution per test cycle is small. Charge trapping in MOS devices is logarithmically dependent on time and the amount of charge applied:

$$Q_{total} = Q_0 \log(t) \quad (4)$$

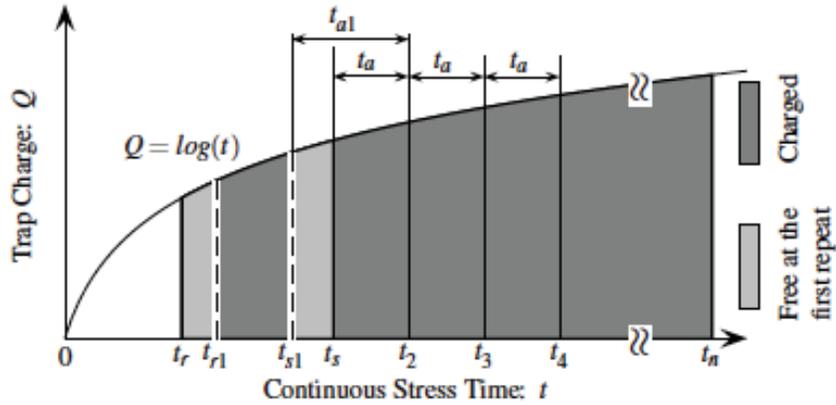


Figure 24. Charge trapping versus IDDQ stress [Okuda]

Consider the diagram in Figure 24 for modeling the charge trapping and stressing mechanism during repeated IDDQ measurements.

At the first functional vector scan, the traps will charge from $t = 0$ to t_s , where t_s is the corresponding continuous stress time of the scan-in time; total charge is then expressed as $Q_t = \log(t_s)$. At the first quiescent mode, discharging occurs from the traps from time t_s to time t_r . This means the total amount of charge is:

$$Q_{tot} = \log(t_s) - \log(t_r) \quad (5)$$

At the next functional test mode, re-charging occurs to the traps. If this new charging time interval is t_a , this next charging and discharging cycle will result in a total charge of

$$Q_{tot} = \log(t_s + t_a) - \log(t_r) \quad (6)$$

This means that over a large number of repeated applications of functional stress vectors followed by quiescent IDDQ tests, the resulting IDDQ current value can be expressed as

$$I_{ddq_i} \approx I_{ddq_0} (1 - \alpha (\log(\frac{t_s}{t_r}) + \log(i))) \quad (7)$$

The expression in Equation 7 shows a semi-qualitative explanation for non-linear IDDQ current that can be dependent on both frequency of applying functional stress vectors and the number of repeating IDDQ measurements that are made. Test results confirming the functional form of Equation 7 are shown in Figure 25.

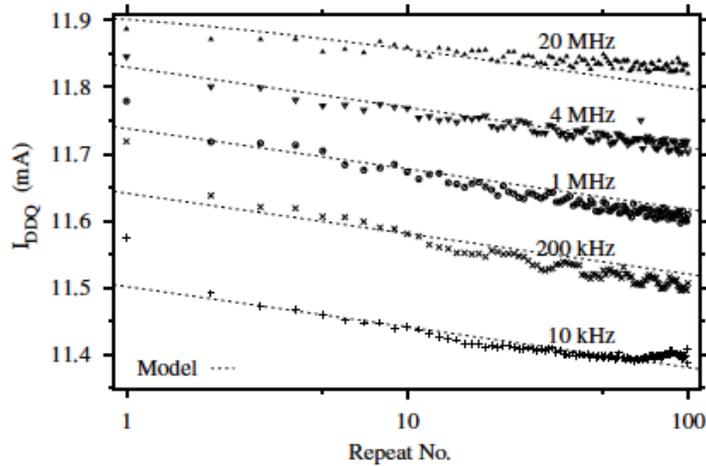


Figure 25. IDDQ measurements showing non-linear dependence [Okuda]

The experimental space was not designed to match Okuda's and, therefore, it cannot be completely verified that Okuda's results explain this study's results. However, the initial strong positive correlation obtained using his analysis provides a fruitful path for further exploration.

Further investigation on IDDQ testing in FPGAs is recommended as a way to understand quantitative changes in device behavior as a function of operational conditions. These changes can be integrated into both short-term burn-in screens and long-term reliability analysis to ensure NASA missions have the most effective FPGA risk mitigation schemes possible.

5 CONCLUSIONS

This report documents the work done in FY09 under the NASA NEPP program for FPGA reliability. Efforts focused on experimental analysis of practical problems experienced by both types of FPGAs NASA uses, reprogrammable and one-time programmable antifuse devices.

Long-term contention studies on reprogrammable 90 nm FPGAs showed statistically significant changes in rise and fall times for a variety of I/O standards. Some I/O standards appear to be more sensitive to this effect than do others. The measured changes will not affect FPGA signal integrity, however.

Quiescent current measurements were made on 220 nm antifuse FPGAs as a function of voltage, design resource, and device temperature. Significant variations of almost an order of magnitude in quiescent current were detected. This amount of variation and dependence upon experimental condition has not been completely explained. Variations due to antifuse programming are convoluted with possible testing-induced charge trapping mechanisms. Further testing is planned in FY10 to resolve these results.

Both experimental results have shown that modern FPGAs are complex devices that require significant levels of testing precision to accurately define possible reliability related degradations. Continued experimental and theoretical evaluations of these results are recommended as part of NASA's overall FPGA risk mitigation process.

REFERENCES

- Burke, Gary R.; Cozy, Scott; Lacayo, Veronica; Bakhshi, Alireza; Stern, Ryan; Mojarradi, Mohammad; Johnson, Travis; Kolawa, Elizabeth; Bolotin, Gary; Gregoire, Tim; and Ramesham, Rajeshuni, "Operation of FPGAs at Extremely Low Temperatures," MAPLD 2004.
- Falin, Jeff, "Powering today's multi-rail FPGAs and DSPs, Part 1," *Texas Instruments Analog Applications Journal*, 1Q 2006.
- Knapp, Steve, "FPGAs furnish fast, furious FIR filters," *Personal Engineering*, December 1988.
- Mangalagiri, P., S. Bae, R. Krishnan, Y. Xie, and V. Narayanan, "Thermal Aware Reliability Analysis for Platform FPGAs," *IC-CAD*, 2007.
- Muse, Mike, "Test Report for JPL FPGA Contention Characterization Study," XDS Project number: 1125-0002, 2009.
- Okuda, Yukio, and Nobuyuki Furukawa, "Hysteresis of Intrinsic IDDQ Currents," ITC International Test Conference 2003, Paper 22.2.
- Sabade, Sagar, and Duncan Walker, "IDDQ Test: Will It Survive the DSM Challenge?" Texas A&M University, 2005.
- Sakaide, Yasuo, K. Kariu, K. Numata, A. Tsukino, M. Urano, K. Chiba, K. Sugimoto, Y. Tsuchiya, and T. Arimitsu, "Effective Post-Programming Screening of Antifuse FPGAs for Space Applications," MAPLD 2004.
- Schmoo plot, http://en.wikipedia.org/wiki/Schmoo_plot.
- Sheldon, Douglas, "Reliability Considerations," NSREC 2009 Short Course.

APPENDIX A. LVCMOS AND LVTTTL CONTENTION STATISTICS

P1 LVCMOS		rise	fall	max	min	dper-	dper+
	Before	1.46	1.59	2862	500	-305	287
	After	1.5	1.64	2853	506	-318	311
	Sigma	0.02	0.03	12.43	7.9	24.3	22.5
	Z score	2.000	1.667	-0.724	0.759	-0.535	1.067
	Z-statistic	0.977	0.952	0.235	0.776	0.296	0.857
	2 tailed p	4.550%	9.558%	153.097%	44.756%	140.734%	28.612%
P2 LVCMOS		rise	fall	max	min	dper-	dper+
	Before	1.49	1.63	2860	499	-284	280
	After	1.52	1.67	2849	507	-308	296
	Sigma	0.02	0.08	42.15	32.3	67	63
	Z score	1.500	0.500	-0.261	0.248	-0.358	0.254
	Z-statistic	0.933	0.691	0.397	0.598	0.360	0.600
	2 tailed p	13.361%	61.708%	120.589%	80.438%	127.981%	79.952%
P3 LV CMOS		rise	fall	max	min	dper-	dper+
	Before	1.48	1.56	2854	505	-269	258
	After	1.53	1.6	2841	517	-296	279
	Sigma	0.03	0.11	38.6	23.5	152	112
	Z score	1.667	0.364	-0.337	0.511	-0.178	0.188
	Z-statistic	0.952	0.642	0.368	0.695	0.430	0.574
	2 tailed p	9.558%	71.613%	126.372%	60.960%	114.099%	85.127%
P4 LVCMOS		rise	fall	max	min	dper-	dper+
	Before	1.58	1.69	2850	511	-302	315
	After	1.6	1.72	2849	515	-307	308
	Sigma	0.03	0.04	8.58	8.1	26.5	21.4
	Z score	0.667	0.750	-0.117	0.494	-0.189	-0.327
	Z-statistic	0.748	0.773	0.454	0.689	0.425	0.372
	2 tailed p	50.499%	45.325%	109.278%	62.143%	114.966%	125.641%
P1 LVTTTL		rise	fall	max	min	dper-	dper+
	Before	1.49	1.64	2887	523	-343	300
	After	1.57	1.71	2863	529	-374	358
	Sigma	0.03	0.03	9.84	7.04	31	33
	Z score	2.667	2.333	-2.439	0.852	-1.000	1.758
	Z-statistic	0.996	0.990	0.007	0.803	0.159	0.961
	2 tailed p	0.766%	1.963%	198.527%	39.406%	168.269%	7.882%
P2 LVTTTL		rise	fall	max	min	dper-	dper+
	Before	1.51	1.66	2885	524	-318	288
	After	1.58	1.73	2859	527	-344	350
	Sigma	0.03	0.03	8.8	6.7	20.1	114
	Z score	2.333	2.333	-2.955	0.448	-1.294	0.544
	Z-statistic	0.990	0.990	0.002	0.673	0.098	0.707
	2 tailed p	1.963%	1.963%	199.687%	65.433%	180.417%	58.654%
P3 LVTTTL		rise	fall	max	min	dper-	dper+
	Before	1.50	1.60	2877	523	-271	251
	After	1.58	1.70	2851	533	-343	313
	Sigma	0.03	0.03	10.8	6.6	37	38
	Z score	2.667	3.333	-2.407	1.515	-1.946	1.632
	Z-statistic	0.996	1.000	0.008	0.935	0.026	0.949
	2 tailed p	0.766%	0.086%	198.393%	12.973%	194.834%	10.277%
P4 LVTTTL		rise	fall	max	min	dper-	dper+
	Before	1.57	1.71	2883	513	-295	283
	After	1.68	1.79	2860	531	-327	327
	Sigma	0.0466	0.038	8.3	6.14	19.6	22.9
	Z score	2.361	2.105	-2.771	2.932	-1.633	1.921
	Z-statistic	0.991	0.982	0.003	0.998	0.051	0.973
	2 tailed p	1.825%	3.527%	199.441%	0.337%	189.746%	5.468%

**APPENDIX B. CONTROL BOARD STATISTICS FOR VIRTEX-4
CONTENTION TEST**

	Rise (ns)	Fall (ns)	max (mV)	min (mV)	dper - (ps)	dper + (ps)	Ave Max (r)	Ave Min (n)
Baseline Ave	1.571	1.816	2,848.900	467.600	-369.000	385.800	2,776.433	549.100
Average	1.573	1.804	2,865.320	469.823	-378.933	378.000	2,786.777	550.903
Median	1.575	1.811	2,862.975	468.200	-366.500	380.250	2,785.225	550.700
Range	0.090	0.119	41.100	40.450	118.500	62.000	25.700	21.750
STDEV	0.025	0.034	13.043	12.375	34.458	19.272	7.696	6.613
Average Dev	0.019	0.029	10.574	10.070	28.164	16.175	6.489	5.678
Max	1.607	1.872	2,887.200	487.000	-341.000	408.500	2,801.450	560.250
Min	1.517	1.753	2,846.100	446.550	-459.500	346.500	2,775.750	538.500
Max (+) from Base	0.036	0.056	38.300	19.400	28.000	22.700	25.017	11.150
Max (-) from Base	0.054	0.063	2.800	21.050	90.500	39.300	0.683	10.600
Z-score	0.046	-0.346	1.259	0.180	-0.288	-0.405	1.344	0.273
Z-statistic	0.518	0.365	0.896	0.571	0.387	0.343	0.911	0.607
2 tailed p	96.33%	127.04%	20.81%	85.74%	122.69%	131.43%	17.90%	78.51%