The Rosetta Experiment: Atmospheric Soft Error Rate Testing in Differing Technology FPGAs

Austin Lesea, Saar Drimer, Joe Fabula, Carl Carmichael, and Peter Alfke

***Abstract*—Results are presented from real-time experiments that evaluated large FPGAs fabricated in different CMOS technologies (0.15µm, 0.13µm and 90nm) for their sensitivity to radiation-induced single-event upsets. These results are compared to circuit simulation (Qcrit) studies, as well as to LANSCE neutron beam results and Crocker Nuclear Laboratory (University of California, Davis) cyclotron proton beam results.**

***Index Terms*—soft error, single event upset, cosmic rays, SRAM, FPGA, NSEU**

I. Introduction

"Rosetta refers to the crucial breakthrough in the research regarding Egyptian hieroglyphs. It especially represents the "translation" of "silent" symbols into a living language, which is necessary in order to make the whole content of information of these symbols accessible." [1]

Just as the Rosetta Stone enabled researchers to decode the unsolvable and mysterious Egyptian hieroglyphs by comparing them to the same text written in a known language, Xilinx' Rosetta experiments link two prior known and well documented techniques of estimating atmospheric neutron single event upsets with the hitherto undocumented real effects of atmospheric neutrons on integrated circuits. The known techniques are accelerated testing in a neutron or proton beam and software simulation of the circuit to determine the critical charge a particular node or latch can handle before it changes state. In our experiments, the actual upset rate of Xilinx' FPGAs due to cosmic ray induced atmospheric neutron cascade is determined. Having a good understanding of the real effect of these neutrons enables us to validate the design and technology choices that can be used to mitigate the effects of such single event upsets.

A good primer on the effects of cosmic ray induced particle showers on integrated circuits is by Dr. Keith E. Holbert [2].

In this paper we concern ourselves solely with single error events (SEE), or soft errors, or single event upsets (SEU). Single event latchup (SEL) is not considered here, as we are able to design and have fabricated devices that are immune to destruction by parasitic bipolar structure latchup. All Xilinx commercial and industrial, as well as military FPGAs must be qualified to not exhibit a latchup mechanism before they are production qualified. Extensive testing in accelerated neutron beams (at LANSCE) has shown that this has been accomplished.

Single event transients (SET) are temporary electrical events (glitches) where a logic signal momentarily changes state and then returns to its original value. Such events are caused by coupling of the circuit to the charge trail that arises from the impact of the particle. We do not consider single event transients in this paper although Xilinx is investigating mitigation methods against this effect in other studies. The FPGA structure by its nature is highly resistant to SETs due to the large capacitive loading of signal paths. This loading is many times that of the loading in an application specific integrated circuit (ASIC) or application specific standard circuit (ASSP). Our testing, when compared with projections of the failure rate of ASICs and ASSPs, indicates that we are at least at parity

Manuscript received October 31, 2004. This work has benefited from the use of the Los Alamos Neutron Science Center at the Los Alamos National Laboratory. This facility is funded by the US Department of Energy under Contract W-7405-ENG-36. The authors would also like to thank the California Institute of Technology Submillimeter Observatory at Mauna Kea, Hawaii and the White Mountain Research Station at White Mountain, California operated by the University of California Office of Research in San Diego, California.

Austin Lesea is a Principal Engineer in the IC design group of the Advanced Products Group at Xilinx, San Jose California. (phone: 408-559-7778; fax: 408-559-7114; e-mail: austin@ xilinx.com).

Saar Drimer is an IC Design Engineer in the IC design group of the Advanced Products Group at Xilinx, San Jose California. (phone: 408-559-7778; fax: 408-559-7114; e-mail: saar.drimer@ xilinx.com).

Joe Fabula is a High Reliability and Radiation Effects Engineer with the Military/Aerospace/Automotive Group at Xilinx, San Jose. (phone: 408-559-7778; fax: 408-559-7114; e-mail: joe.fabula@ xilinx.com).

Carl Carmichael is the Manager of Aerospace & Defense Product Applications Engineering at Xilinx, San Jose. (phone: 408-559-7778; fax: 408-559-7114; e-mail: carl.carmicahel@ xilinx.com).

Peter Alfke is a Distinguished Engineer in the IC design group of the Advanced Products Group at Xilinx, San Jose California. (phone: 408-559-7778; fax: 408-559-7114; e-mail: peter.alfke@ xilinx.com).

and perhaps as much as ten times less sensitive to SETs and SEUs on a basis of logic blocks.

Predicting atmospheric neutron flux is not an exact science. In his paper [3], J.F. Ziegler makes this clear (see Figure 13, which depicts the many studies of terrestrial neutron flux). If the particle beam experiments are intended to predict actual failure rates, one could be in error by a factor of as much as 10 to one as there is a wide variation in energy vs. flux as reported by the various studies. This means that anyone can create and defend almost any result.

Additional work by J.F. Ziegler [4] indicates that many of the atmospheric flux estimates have used detectors with very narrow angles of incidence, so that the total flux from all directions is again the result of estimation. The FPGA, as a detector of energetic neutrons and protons, has a variation of only 2:1 in cross section (sensitivity) over all angles [11], making it an accurate means of measuring neutron flux for all angles of incidence, albeit one of the least efficient in terms of collecting particle events. To improve on this efficiency of collection we use large numbers of the largest devices available.

II. Experiment

Each Rosetta experiment consists of multiple sets of 100 of Xilinx' largest FPGAs of differing technologies and is located at four different altitudes. All tested components were fabricated by UMC in a 300mm fabrication line using standard logic CMOS processes.

A picture of one set of 100 experimental setup is shown in Figure 1.

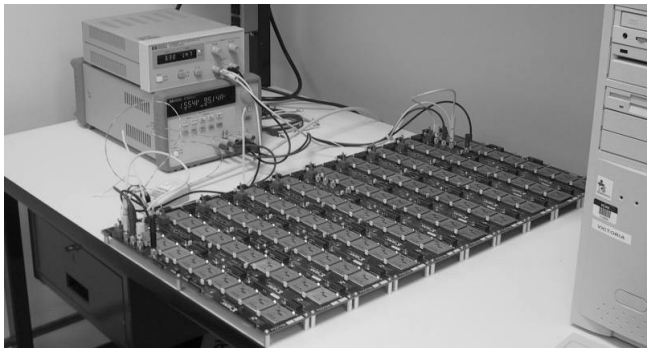

Figure 1. One set of 100 Devices

See Table 1 for locations, and Table 2 for the number of devices, and device hours at each location.

Facility	Location	Alt. (ft)	Physical Placement
Xilinx	San Jose, CA	0	Floor 1 of 2 in concrete slab "tilt-up" style
Xilinx	Albuquerque, NM	5,100	Floor 1 of 2 in concrete slab "tilt-up" style
WMRS	White Mountain, CA	12,470	Top story of a Quonset hut building
CSO	Mauna Kea, HI	13,200	Bottom floor of observatory

Table 1. Description of the location at which Rosetta experiments are present. [5,6]

Node (nm)	Die	Location	# DUTs	Device hours
150	2V6000	San Jose, CA	100	1060000
		New Mexico, NM	100	1670000
		White Mountain, CA	100	856800
		Mauna Kea, HI	100	535000
130	2VP50	San Jose, CA	200	1191000
		New Mexico, NM	200	709000
		White Mountain, CA	200	655000
130	2VP40	San Jose, CA	100	
90	S31500	San Jose, CA	500	256000
90	V4LX25	San Jose, CA	100	20000

Table 2. Devices under test.

In the integrated circuit design of the Xilinx FPGAs, the individual memory cells (actually static latches) that are used for configuration, look-up tables (LUT), and Block RAM (BRAM) were all simulated for their sensitivity to single event upsets.

Simulations of the memory cell (latch) used in the FPGA showed that the design of this latch is more robust than that of a standard SRAM cell fabricated in the same process technology by the same foundry. Xilinx has always designed its static latch configuration cells to different criteria than that those used to design a commercial SRAM cell, where the dimensions are pushed to the lithographic limits in order to achieve the fastest read and write performance.

Since the look-up tables are also part of the configuration memory space and can suffer from single event upsets; they are also designed with the same rules as the configuration latch cells.

The BRAM is designed with similar small size cells similar to a commercial SRAM. Upset performance of the BRAM is not as robust as the configuration cells, however.

The overall goal of a FPGA IC designer is to keep the largest device's FIT rate lower than one functional failure per 200 years per user design. Earlier this goal was considered to be at risk in the design of the Virtex-II family. Beam testing and simulations provided us with information that was incomplete, sometimes misleading

and sometimes overly optimistic. The Rosetta experiments results allowed us to predict the reliably that the user would see in actual operation.

Each group of 100 FPGA devices is composed of ten separate boards, with ten devices on each board. All 100 devices are wired into a serial JTAG test chain, so that the configuration memory (which includes the look-up table RAM, and the BRAM) can be written, read back, and verified.

Each group of 100 devices is powered by multiple output power supplies, which have remote sense lines to provide for accurate voltages at the devices. A computer is connected to the Xilinx intranet, to the power supplies, and to a standard PCIV™ programming cable. Software on the computer notifies us by email when an upset occurs. An event is not considered an upset until it is verified by being read back one more time, and then reprogrammed to be sure that the flip is not a faulty memory cell. If an event is recorded, and on a consecutive readback is not present, it is considered a signal integrity issue, error in the computer, or an error on the serial link being used. The event is logged only after verification, and an email generated with the date, time, and information about the file locating the event. This process is detailed in Figure 2.

All 100 devices are scanned for events continuously, 24 hours a day, seven days a week. Due to the large number of devices (over 1.9 gigabits of configuration memory) and the nature of the scan, scanning takes about two hours per setup, limiting the time resolution of detection of an upset. This limit is based on the speed of the computer being used for the verification process, and the speed of the JTAG readback, which is 5 Mb/s in this experiment. Events are located physically by recording the bitstream address of the event and using Xilinx proprietary software to locate the exact function and physical location of the flipped bit. In this way we are able to identify the loading, layout, circuit, and electrical parameters for each upset.

In order to detect alpha contamination in packaging and assembly, we rotate the experimental groups through the three altitudes. If there were any evidence of a constant upset rate due to alpha particles, it would be observed as a non-altitude dependent factor in the resulting upsets. The solder bumps used in the flip-chip assembly process are certified ultra-low alpha lead solder. If there were any observable constant upset rate, these upsets would also cluster about the solder bumps.

Alpha counts per hour for the solder bumps was measured at less than 0.01 (the limit of the instrumentation). Additional measurements of the underfill material had less than 0.001 counts per hour. The underfill material will shield alpha particles in less than 10 microns of travel.

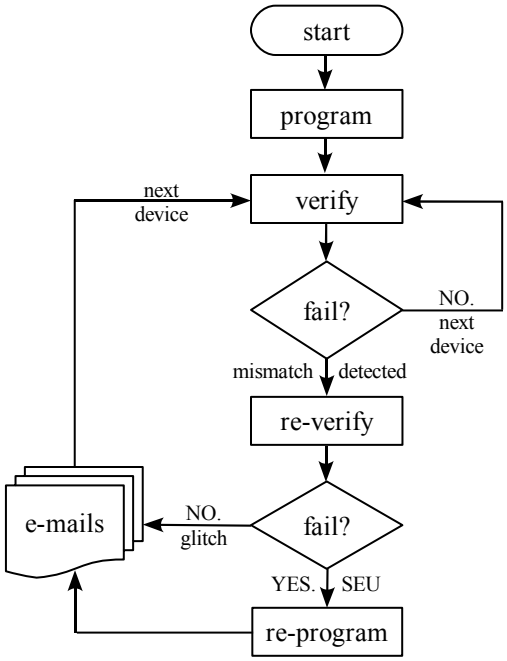

Figure 2. Flowchart of the Rosetta monitoring software program

Figure 3 is an example of a map of four atmospheric upsets by location on a 2V4000 die.

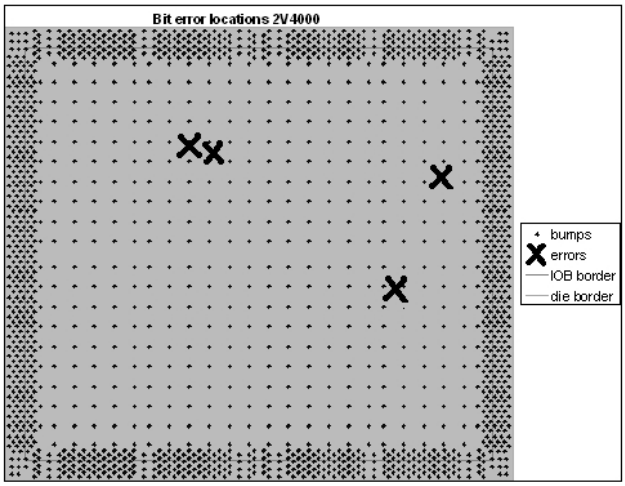

Figure 3. Location of Bit Events on a 2V4000

By assembling 100 devices with regular solder, we were able to calibrate the contribution of alpha emitting solder in terms of counts per hour and mean time to an upset. From this we were able to conclude that whatever alpha contributions there are in the ultra low alpha solder bump-packaged devices, they are much less than one upset per 200 years per device. This alpha particle experiment also gave us a calibration point for low alpha, and ultra low alpha packaging materials and compounds that are required to reduce the number of upsets to an acceptable level.

The power supplies run at the nominal voltage level for the core: 1.5 volts for 0.15µm and 0.13µm, and 1.2 volts for 90nm. The devices that are flip chip packaged are mounted with the active side of the die facing down (towards the board), and the backside of the die facing the sky (up). The devices that are wire bond packaged have their active side facing up, and backside of the die facing down.

We have measured only a 2:1 variation in the upset rate with beam incidence angle at the Los Alamos facility. Normal to either the front side or backside of the die yields the largest cross-section, with the smallest cross section being with the beam edge on. Spallation products from the neutron beam interacting with the flip chip solder balls add to the beam flux, making radiation through the flip chip solder balls the worst case.

San Jose	N37.25	W121.94
Albuquerque	N35.17	W106.59
White Mountain	N37.63	W118.25
Mauna Kea	N19.80	W155.50

Table 3. Latitude and longitude of sites

Table 4 shows the magnetic field information for all components, as well as the force vector are all similar in first three locations: [7]

Location (nT*)	X	Y	Z	H**	Total
San Jose	23119	6184	43118	23932	49314
New Mexico	23178	4254	44880	23565	50909
White Mountain	22598	5823	44829	23336	50270
Mauna Kea	27730	4763	20680	28136	34919

* nano Torr (magnetic field strength)
** H = horizontal intensity
Table 4. Geomagnetic properties of the sites

Using the above data to correct for effects of magnetic latitude indicates that the expected difference in flux for the first three locations should be less than ±5% (ignoring altitude). The effect of magnetic latitude on Mauna Kea is more pronounced, with as much as a 2:1 reduction in neutron flux expected.

The atmospheric acceleration rates for these locations can be found by assuming a 30% increase for every 1,000 feet in altitude (a commonly used best fit to the measured flux vs. altitude below 40,000 feet), or by the methods outlined in the JESD89 standard (ignoring magnetic rigidity). These are noted in Table 5, along with the experimental altitude factors derived from 0.15µm data set.

	30% / 1000ft.	JESD89	Actual 0.15µm
San Jose	1	1	1
New Mexico	3.812	3.75	4.6
White Mountain	26.356	16	23.4
Mauna Kea	33.639	12.3	17.1

Table 5. Altitude acceleration factors compared with 0.15µm data

The one-sigma estimate of the error in the actual data from the count of errors above is: 40%, 13%, 7%, and 11%, respectively.

	30% / 1000ft.	JESD89	Actual 0.13µm
San Jose	1	1	1
New Mexico	3.812	3.75	3.4
White Mountain	26.356	16	10.6

Table 6. Altitude acceleration factors compared with 0.13µm data

Due to the similarity in the magnetic field data for the first three locations, we have used the 30%/1000 feet altitude factors for all of our calculations. A 50% reduction for Mauna Kea due to the magnetic field data for that location is applied as well as the altitude estimate which also results in an excellent fit to the measured factor (33.639/2 = 16.8 vs. 17.1).

This method yields factors that are a best fit to the experimental data sets, excepting White Mountain 0.13µm data. The one-sigma estimate of the actual data from the count of errors for 0.13µm is: 20%, 12%, and 7%, respectively.

The one-sigma uncertainty in the San Jose actual results could account for the lack of agreement.

Another explanation for this difference is that the sensitivity to upset is also based on energy and on the mix of neutrons and protons. The flux by itself says nothing about the spectral density of the energy. If one calculates the flux using the LANSCE beam cross section and the actual atmospheric upsets the of 0.15µm technology node, it would imply that the flux is 8.4 neutrons/centimeter squared hour. The 0.13µm technology node beam tests with actual atmospheric upsets calculates to 13.3 neutrons/centimeter squared hour. LANSCE beam measured cross section is 3.18E-14 and 2.98E-14 respectively for 0.15µm and 0.13µm configuration memory technologies. Only more time will allow resolution of this issue.

If the Qcrit for the two technologies were different, then this could be explained; however, the Qcrit from simulation, as well as the LANSCE beam tests for the 0.15 µm and 0.13µm technologies are almost identical, so this apparent difference in flux causing upsets is still a mystery. The energy distribution of the neutron spectrum is known to vary with altitude, so it should be expected that the

sensitivity of the static latches would also show this variance. The mix of protons vs. neutrons is significantly higher at higher altitudes, and this may also be a factor [4].

The spectrum of the LANSCE beam is a reasonable match for the atmospheric energy distribution at 40,000 feet for neutron energies from 1.5MeV to 600MeV, but does not include the proton distribution expected at that altitude. As you come down in altitude, the proton contingent of the flux decreases to near zero while the neutron contingent shifts towards lower energies. This may explain why a calculated "Rosetta Factor" changes with altitude and is not addressed particularly well in JESD89 [13].

A. Atmospheric Test Results for 0.15µm

At the 0.15µm technology node, the resulting data for the configuration cell is 295 FIT/Megabit, where 1 FIT is one bit flip in a billion hours. The BRAM is 265 FIT/Megabit.

This is based on 350 upsets and 7,050 device years of data. The 95% confidence interval is ±11%.

A single bit upset in the configuration memory is unlikely to cause a user functional failure as fewer than 10% of the configuration memory cells are ever used in the customer's FPGA design. Research has shown that the ratio of flipped bits to actual functional failures ranges from as low as 6, to as high as 100, depending on the actual user design [11, 12]. We refer to this effect as the Single Event Upset Probability Impact, or SEUPI derating factor.

In the iRoC Technologies paper [15], events or upsets are referred to as SEFI, or Single Event Functional Interrupts. We have chosen not to use this terminology, as we had previously defined a SEFI to be a non-recoverable error which requires re-programming the entire device, or a power off and on cycle. We prefer to use the concept of a factor on the basic event rate (single event upset rate divided by the SEUPI derating factor) to distinguish the much more severe (and rare) case of a unrecoverable functional error or SEFI from a single bit upset or SEU which results in corrupted user data. A SEU by our definition is correctable by re-programming just the one offending bit. Re-programming may be done while operating by using the Internal Configuration Access Port (ICAP) feature of the devices.

There have been no SEFI events in any of our atmospheric testing to date. In beam testing we can cause a complete functional failure requiring a power cycle or reset and reprogram, but the cross section is so small, that it is unlikely to ever happen in atmospheric testing. Work is ongoing to remove all SEFI causes by hardening those elements of the design that might be responsible.

The error rate is either stated in Failures in Time per billion hours (FIT), or in Mean Time Between Events (MTBE) in hours, days, or years. A functional failure of the user data due to the single event upset rate then becomes Mean Time Between Functional Failure (MTBFF) in hours, days, or years.

The BRAM upset is likely to cause a functional failure of a user's design, so we have an extra bit per byte for use as a parity bit by the customer. Error check and correct (ECC) logic is easily programmed into the FPGA to mitigate the effects of upsets in the 0.15µm and 0.13µm nodes, with the ECC logic being designed into the 90nm node to simplify usage.

B. Atmospheric Test Results for 0.13µm

At 0.13µm, the configuration cell is at 290 FIT/Megabit, and the BRAM is at 530 FIT/Megabit. The BRAM cells used at the 0.13µm technology node were more aggressive (smaller) in their layout than those used at the 0.15µm node. The configuration latches are roughly the same design as previous technologies.

These results are based on 317 upsets and 5,900 device years of data. The 95% confidence interval is ±12%.

C. Atmospheric Test Results for 90 Nanometers

At the 90nm technology node, the first FPGA product produced was the Spartan-3. At the present time we have 200 units under test at the San Jose location. The next product is the Virtex-4. Virtex-4 utilizes 0.13µm "mid" oxide transistors for the configuration memory and 90nm transistors for the block RAM. There are 500 units under test at the San Jose location. We plan to have 600 units of the larger device by June, 2005. Consult the Xilinx website for periodic updates on the Rosetta experiment.

The Spartan-3 data at this point is 520 FIT/MB with a 95% confidence interval from 50 to 1,920 FIT/Mb. Virtex-4 data at this point is 130 FIT/Mb with a 95% confidence interval from 0 to 483 FIT/Mb. Obviously, more time is needed to have a good estimate of the FIT rate.

III. Effects of Solar Activity on Experimental Results

Solar activity is reported daily by various sources (spaceweather.com, NOAA National Geophysical Data Center, Oulu Cosmic Ray Station) and can be compared with the upset time, date and location.

Comparing observed experimental data to solar activity, the only correlation we see is the one associated with the sunspot cycle. Before December, 2003 we were seeing the beginning of the decline from the solar sunspot

maxima. After January, 2004 we had entered a solar sunspot minimum. The change in cosmic ray flux as recorded by the neutron monitors (for example, the Oulu Cosmic Ray Station in Finland) is an increase of 3 to 6% from the sunspot maxima to sunspot minima.

A 2.1% increase is noted in the experimental data for the 0.15µm technology devices at White Mountain when compared between the period before January, 2004 to the period after January, 2004. Given the ±12.9% uncertainty in the data (confidence interval), this may be a coincidence.

IV. Comparison to Commercial SRAM

The upset rate for both the configuration memory and the BRAM is much less than that of commercial SRAM as detailed by Normand [8].

The CMOS configuration cell used by Xilinx is showing an improvement of a factor up to eight over the Boeing study (which did not include FPGAs). This does not take the SEUPI derating factor in account. If one also considers that it takes more than 10 upsets on average to create a single functional failure (a conservative assumption as we have seen it taking as many as 100 upsets to create a functional failure), the factor improvement is then at least 80 times over that of a commercial SRAM. The improvement can be as high as 100 times, or 800 times as great as the data in the Boeing paper, depending on the user pattern programmed into the device.

V. Results of Simulations of Qcrit on a per Structure Basis

The Xilinx IC design group uses models and methods learned from our fabrication partners to determine the sensitivity of the memory cells to upsets. These models and methods have been used in their production of standard products, and ASICs. Their prediction of Qcrit to atmospheric upsets is used to compare with our observations in the Rosetta experiment.

The typical injected charge profile is modeled by a time varying current pulse that has a rise time of 5 picoseconds, and a fall time of 30ps. Total integrated charge delivered to the node in question is stepped from below 1 femtocoulomb, to more than 20 femtocoulombs to find where the circuit is affected. Many hundreds of thousands of simulations are run on the extracted layout netlist nodes of the cells in question. All results are collected, and displayed as histograms of Qcrit sensitivity. Since the basic functional logic and interconnect block of the FPGA (the CLB or configurable logic block) has many hundreds of memory cells controlling many features, the Qcrit varies by as much as a factor of 12:1 due to the loading on the memory cells.

The FPGA architecture uses memory cells to control every programmable function and feature. It is therefore infeasible to add triple redundancy, error check and correct, or some other means to every memory cell to prevent upsets. Adding the required circuitry would cause the area to explode and make the FPGA commercially not viable.

The circuit design of a radiation hard FPGA is a subject of present work. A device that is radiation hardened may not be profitable to manufacture. One approach used today is to use triple modular redundancy (TMR) in the design of the user logic pattern [9] on a standard device.

Xilinx FPGA memory cells are designed to be as robust as possible, so that subsequent generations do not experience higher failure rates. Known design techniques, such as maximizing capacitive load where speed is not critical, as well as proprietary design techniques (patents pending), are being used to keep the Qcrit high as the technology nodes get smaller. Block RAM cells are allowed to follow the technology curve of decreasing Qcrit and increasing FIT rate, as their upsets are correctable through the use of standard error check and correction (ECC) techniques, which are now designed into hard logic in the Virtex-4 FPGA family. Previous generations may make use of programmable ECC cores if needed by the customer. The predicted FIT/Mb rate from Qcrit simulation results for the memory cells are detailed in Table 7.

Qcrit/Node	Configuration	BRAM
0.15µm	101 FIT/Mb	105 FIT/Mb
0.13µm	106 FIT/Mb	114 FIT/Mb
90nm (S3)	111 FIT/Mb	222 FIT/Mb
90nm (V4)	61 FIT/Mb	222 FIT/Mb

Table 7. FIT/Mb based on Qcrit vs. technology node.

Note that the simulation results are predicting roughly half the results of the actual atmospheric tests. However, they do not show a large variation in performance with technology, which is in agreement with the atmospheric test results for configuration cells. For the BRAM we are seeing a larger failure rate in 0.13µm than we would have expected from the Qcrit simulations. LANSCE beam testing also indicates that the BRAM in 0.13µm has a larger cross section. Rosetta results agree with LANSCE BRAM test results for 0.13µm. Why the Qcrit simulations for 0.13µm BRAM are not accurate is being studied.

VI. Accelerator Tests

The best resource currently available to simulate atmospheric neutrons is the high-energy Neutron Testing

Facility at the Los Alamos Neutron Science Center (LANSCE). At LANSCE, high-energy neutrons are produced by spallation. A linear accelerator produces an 800MeV pulsed proton beam that strikes a water cooled tungsten target. The impact produces a spectrum of neutrons whose energy distribution and intensity is precisely measured. The spectrum has energies in the range of 1 MeV up to approximately 600 MeV and is very similar in shape to the atmospheric (Hess) spectrum. A neutron flight path with a station (ICE House) at 20 meters from the neutron production target has been designed for accelerated neutron testing of semiconductor devices. The flight path consists of a small building for the irradiation which also encloses the testing equipment, isolated from the beam by a substantial concrete barrier. The devices to be tested are placed in the neutron beam line (in air) in the irradiation building. The experimenters control the neutron beam by opening and closing a shutter external to the irradiation building, and the number of neutrons on the sample is continuously monitored and recorded. Corrections to flux for the $1/R^2$ distance from the source must be included.

The integrated neutron flux (energies above 1 MeV) for the LANSCE beam is approximately 10^6 n/cm^2 with a beam spot size of approximately 10cm. Uniformity across the beam is better than 5%. Since it was expected that between 10^{10} n/cm^2 and 10^{11} n/cm^2 would be required to provide a statistically significant number of upsets, long beam run times of l hour or more were usually needed for each data point.

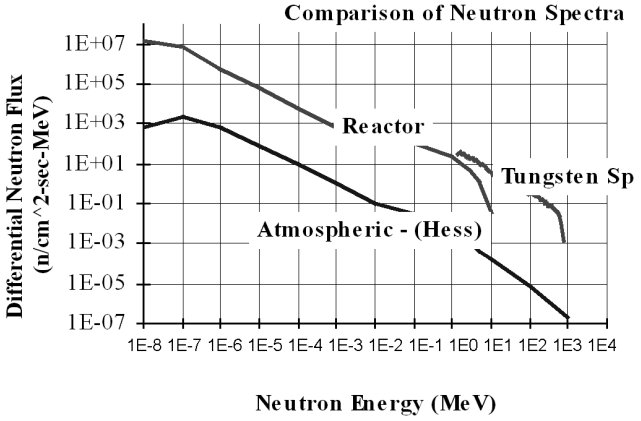

Figure 4. Neutron energy flux distribution comparison.

The test configurations used often allowed several parts to be tested during each run. This is important in order to maximize the amount of data that could be collected in the long runs required for these tests. Because high-energy neutrons penetrate most materials easily, it was not necessary to decap (remove the top metal heat spreader) the parts for exposure. We used an internal test system and software tool that utilized the JTAG test capability of the Spartan-3 , Virtex-II Pro and Virtex-II architectures. Each device can be individually tested via the JTAG port. For the Virtex-II test we also utilized two additional parts that were not in the beam as controls. The Virtex-II test set-up was an adaptation of the Rosetta fixture that normally required 10 samples on each board to complete the JTAG scan. On each of two PC boards, 2 devices would "stand-up" in the beam line while the third part was well outside the beam penumbra. The seven sockets on each board between the two parts under test and the single control at the end of the board were jumpered. For Spartan-3 and for Virtex-II Pro commercial performance boards, available from Xilinx, were used for the test. Each board was driven by a laptop PC in the control area over a 40-foot communications cable. The power supplies were also remotely located in the control room, with remote sense lines used avoid voltage drops at the parts should high currents flow. The supplies were current limited to detect any possible latch. A set of voltage monitoring wires was used to make sure that the desired voltage levels were applied at each part. The test boards are shown in Figures 5,6,7 and 8.

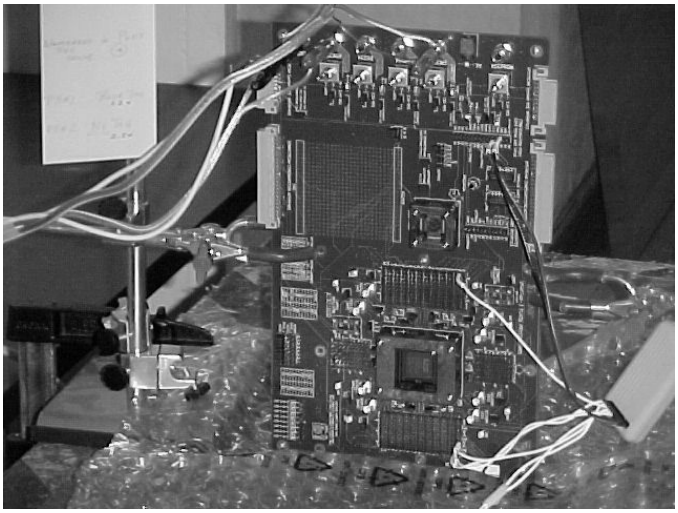

Figure 5. Spartan-3 test fixture.

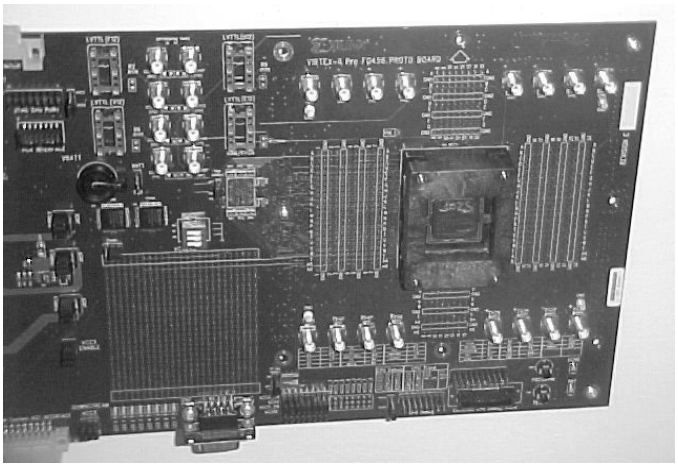

Figure 6. Virtex-II Pro test fixture.

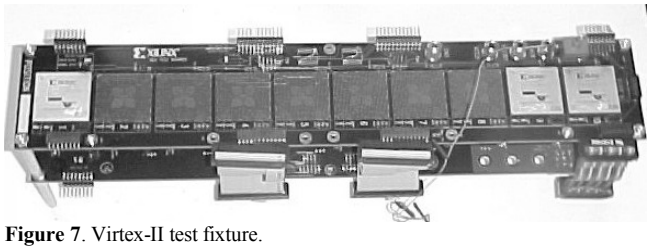

Figure 7. Virtex-II test fixture.

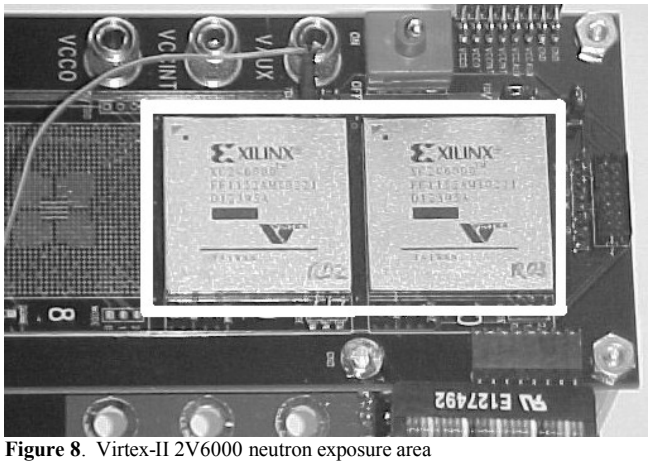

Figure 8. Virtex-II 2V6000 neutron exposure area

Many of these technologies had been previously tested in proton, neutron and/or heavy ion test environments, and both static and dynamic testing had been carried out to consider the issues of static mode vs. dynamic mode upset. Details of this testing have been published elsewhere and some of these publications can be found on the Xilinx web site at "Xilinx.com". Accordingly, the accelerated NSEU testing at LANSCE measured only the static mode upset sensitivity due to high-energy neutrons. Through the JTAG port of each device, the configuration bit stream was be written and read. Test software (FIVIT) was developed that would write to the configuration bit stream of each part and parse the readback results. Because an arbitrary pattern could result in contention or violation of illegal conditions in an FPGA, previously documented patterns were selected for the static test. The test algorithm was implemented as follows

1. Select the device to be programmed
2. Write configuration bit stream with the data pattern
3. Verify correct configuration with read back
4. Repeat for all samples
5. Pause while the neutron beam is applied (could be for a number of hours)
6. Verify post radiation configuration with readback
7. Compare data before and after radiation.
8. Record bit upsets including frame and location data
9. Reprogram to the original pattern to assure no stuck-bits
10. Run test long enough to accumulate up to 100 bit upsets per part for statistical validity.

Initial testing was done at room temperature and nominal power supply as appropriate. In the specific case of Virtex-II, the core voltage was varied from above the max operating voltage of the array (1.60 volts versus a maximum specification of 1.575 volts) to a voltage well below the minimum specification voltage of the array (1.20 volts versus a minimum specification of 1.425 volts) to investigate the robustness of the configuration latches. The test samples used included the XC3S1000, the XC2VP4, the XC2VP7 and the XC2V6000. Icc was monitored throughout the testing in order to look for possible neutron induced latch-up. Neutron induced latch-up has not seen for any of the devices in over 120 hours of beam exposure (equivalent to over 13,000 years of atmospheric exposure). Xilinx was allotted three slots at LANSCE of 48 hours each during the past year for this research.

D. Results for Virtex-II

Virtex-II testing has been performed at LANSCE on three different occasions: August 2002, December 2002 and September 2003. Each test utilized the 2V6000 devices. The following summarizes the test results from each date.

In Table 7, please note that experimental configuration B has one set of the Virtex-II samples (S/N 5 and 6) in the front of the beam line with the second set of Virtex-II samples behind them. Experimental configuration A has the position reversed (i.e. S/N 2 and 3 are in front of the beam line) Experimental configuration C moved the location of the Virtex-II Pro board (results reported later) while position D rotated position A at a 45 degree angle to the beam in an attempt to investigate the effect of the neutron beam at a non-orthogonal incidence.

Table 7 summarizes the measurements that were made on all of the Virtex-II parts. The upset averages of the first pairs of parts in the neutron beam were utilized to calculate the bit cross section. The difference between the first parts in the beam and the second parts averaged a ratio of 1.23, and the magnitude of this difference was not expected and has been shown to be the result of enhanced dose from the spallation products resulting from the flip chip solder mounting balls. More about this effect appears later in this paper in the discussion section. Therefore for this calculation, only the data from the front units was used in the calculation of the cross section.

E. Results for Virtex-II Pro

Virtex-II Pro testing has been performed at LANSCE on two different occasions: August 2002 and September 2003. The August 2002 test utilized a 2VP4 while the September 2003 test was on a 2VP7.

Only one device was available for the Virtex-II Pro experiment. The fixture used to expose the part is shown in Figure 6 and the results of the test runs are summarized in Tables 9 and 10. The part was similarly programmed and interrogated (as was done with the Virtex-II parts). In run 3 the 2VP4 part experienced a SEFI and control of the bit stream was lost. This run was excluded from the calculation of the bit cross section, as the error count was fallacious. However this one SEFI did allow a crude estimate of the SEFI Cross Section for the 2VP4. The calculated cross section for this SEFI was about 2.0E-18 and the SEFI cross section, based on this single data point, is estimated to be somewhere between 10E-17 and 10E-19.

Calculation of the bit cross section was made using the neutron fluence estimates for the eight successful runs with an energy >1.5 MeV and again for an energy >10.0 MeV. The actual data taken on the Virtex-II Pro (2VP4) FPGA along with the calculation of the critical cross-section is shown in Tables 9 and 10.

The average CLB Cross-Section for the Virtex-II Pro, based on the data above, is 2.98E-14 for neutrons >10MeV. This figure (for a 0.13µm technology) is lower than the cross section calculated for Virtex-II (a 0.15µm technology). Due to the number of bit upsets that went into these calculations, the value arrived at has a high confidence of being correct. This implies that, while the Virtex-II Pro technology is fabricated on a tighter technology than the Virtex-II and hence has a lower Qcrit, the geometry factor predominated and resulted in a lower cross section calculation. The effect is believed to be real and a consequence of the lower probability of a neutron striking a critical area in the Virtex-II Pro technology.

F. Results for Spartan-3

Spartan-3 testing has been performed at LANSCE on two different occasions: September 2003 and November 2003. All tests utilized the 3S1000.

The fixture used to expose the part was shown previously in Figure 5 and the results of the test runs are summarized in Table 11. The part was similarly programmed and interrogated (as was done with the Virtex-II and Virtex-II Pro parts). In this case, custom software (FIVIT) was used to interrogate the parts and to calculate average cross sections.

Run #9 was not used in the calculations of cross sections because of a operator error which resulted in the loss of data. Note that there was negligible current increase during the dosing and absolutely no indication of single event latch in any of the testing. The cross sections were calculated based on the results and are consistent with earlier data taken on these same devices. The effects of process migration and improvements in the design of the configuration latches (to enhance latch stability) are detailed later in this report.

The average CLB Cross-Section (>10MeV) for the Spartan-3 based on all runs is 2.87e-14.

Papers have been presented at the Military Applications of Programmable Logic Devices (MAPLD) conference showing the results of these beam tests [11,12].

Identical devices from different fabrication facilities (UMC and IBM), as well as different fabrication processes (low-K dielectric and standard Fluorinated Silicon Glass – FSG, and copper and aluminum copper interconnect) have been tested side by side and shown to be similar. It should be noted that at, at 220nm and below, ^{10}B (boron-10) has been effectively removed from all Xilinx devices, rendering them insensitive to low energy thermal neutrons. This was accomplished by the elimination of BPSG reflow glasses, replacing this leveling technique with CMP (Chemical Mechanical Polishing) and by utilizing purified ^{11}B (boron–11) source gas for all Boron implantations at our foundries.

G. Accelerated Testing Conclusions

There are several observations that can be made from the data that has been obtained.

First (and simplest), the upset rates were relatively linear with fluence. This indicated that the statistics of upset were relatively random in nature, and that the effects were not synergistic.

Second, since the total number of upsets was kept to a small percentage of the configuration file size, there was no need to correct the data for the possible effect of flipping a disturbed bit back to its initial value. This is illustrated in the curves below (Figure 9) for data taken on the Spartan 3 samples.

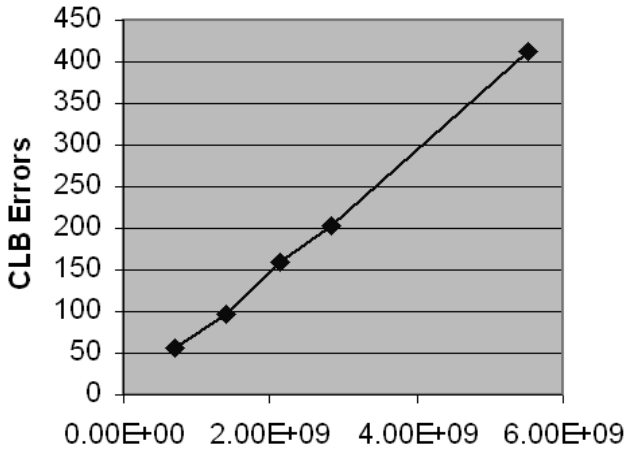

Figure 9. CLB errors vs. > 10 MeV fluence upset linearity

Third, care must be taken in the directionality of the beam, especially for flip chip packaged devices. The reason for this is illustrated below. A Virtex-II part was rotated in the beam in 15 degree increments and a measure of the cross section was made as a function of beam angle. This was originally proposed to look for multiple bit errors from a single particle at low angles (evidence of multiple bit errors was not seen), which is consistent with the distributed nature of the configuration latches used to program Xilinx FPGA. However, there is a very definite effect of flux enhancement by the flip chip solder balls used to connect the die to the substrate. This effect is not seen in wire bonded devices. While the same effect is occurring from radiation from the back of the die, the range of the spallation products produced from the backside is inadequate to affect the device operation. However, when the beam is incident from the top of a flip chip die, the spallation products are within range of the active area of the device and do affect the measured cross sections. The fortunate outcome of the effect is that the local dose enhancement is only a factor of two or less. See the following illustrations (Figures 10 and 11) for details on this effect.

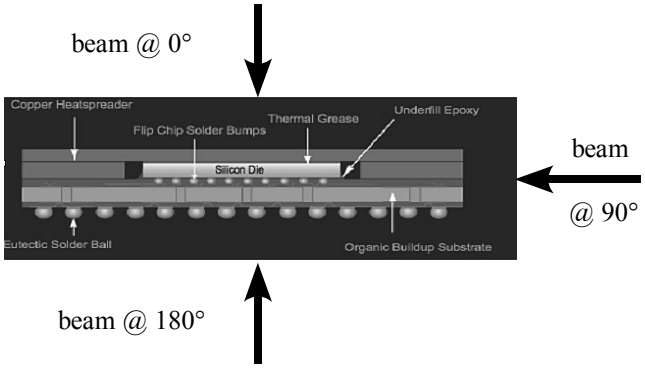

Figure 10. Flip chip geometry.

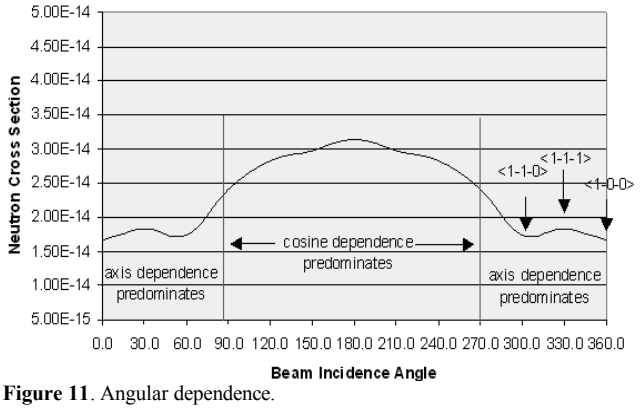

Figure 11. Angular dependence.

Fourth, the static latches that are used to configure an FPGA are not the same as the flip flops used in a commercial SRAM. The flip flops in a commercial SRAM are designed with three criteria, namely size, ability to quickly write, and ability to quickly read. Hence commercial SRAM cells have been shown to be relatively sensitive to NSEU effects. This is illustrated by the sensitivity of the upset rate of a commercial SRAM to operating voltage. This is not the case for the static latches used to program an FPGA. The design criteria for these static latches have historically been and continue to be stability, stability and stability. The size of an FPGA is dominated by interconnect, not by the configuration latches. Also, these latches are typically written to once and hold that data, they do not have to be particularly fast or easy to write to. What they have to be and have been historically designed to be is stable. This is well illustrated by the plot of cross section versus operation voltage for a 150nm latch used in Virtex-II. The cross section was measured over the range of 1.20 volts to 1.60 volts, well outside the published operating voltage range of the device (1.50 ±5%) and the results are illustrated below in Figure 12. This is one proof of the stability of these static latches.

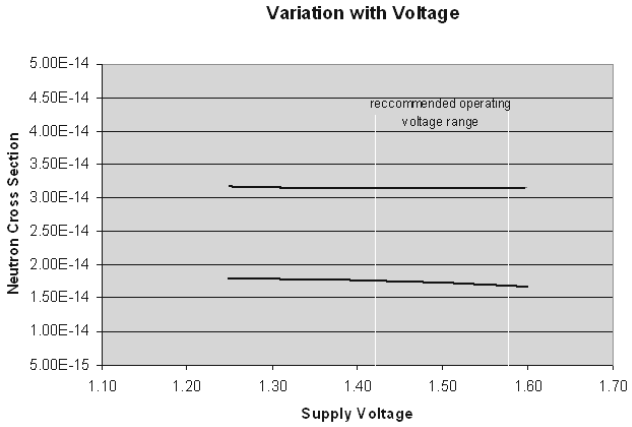

Figure 12. Voltage dependence.

Fifth, the Xilinx Rosetta experiments (performed on large quantities of FPGAs at four different altitudes and multiple technologies) have demonstrated that the correct energy model to use for the calculation of FPGA upset rates from LANSCE data (the >10.0 MeV model), and is different from the extrapolations of >1.5 MeV commercial SRAM data that others have used to calculate these effects. The clear need to do Rosetta-like experiments to calibrate results from accelerated NSEU testing was illustrated by Zeigler in recent presentations and is illustrated below in Figure 13. The choice of a flux and energy distribution to use in accelerated upset calculations is not possible based on the published literature over the past forty years, and must be determined for each specific technology by Rosetta-like experiments.

Sixth, A failure in a configuration memory cell statistically does not cause a functional failure. Research by BYU and LANL using their bit corrupter on a V1000 device has shown the multiplier to be between 25 and 100. SEUPI analysis of specific customer applications has shown SEUPI derating factors from 10 to 80 with a mean of 42.

The conclusion is that a multiplier of 10 can be used conservatively.

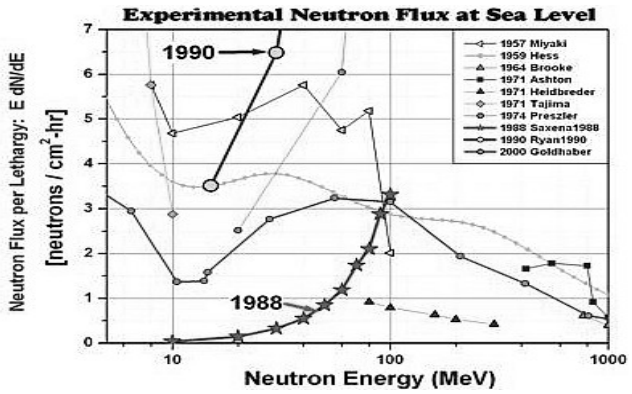

Figure 13. Experimental neutron flux measurements [3].

Seventh**,** and finally, the projections of ever worsening neutron cross section and upset rates as technology has shrunken through the 220nm, the 180nm, the 150nm and now the 90nm range have not been shown to be correct. Rather the decreased size of the shrinking latches combined with robust design efforts have stabilized the critical neutron cross sections of state of the art static latches and reversed this trend. Indeed the MTBF that can be calculated for a million logic gate FPGA fabricated with 90nm technology is better than that of the MTBF of a similar complexity array fabricated in 150 or 130nm processes. This is illustrated by the Table 12 below which illustrates the fact that the neutron cross section (and hence the upset rate) for 90nm product is actually lower than the neutron cross section of similar 150nm product.

Technology node	Cross section >1.5MeV	Cross section >10Mev
150nm	1.56E-14	3.18E-14
130nm	1.54E-14	2.98E-14
90nm (S3)	1.49E-14	2.87E-14

Table 12. Measured LANCE cross sections.

To derive the FIT/Mb rate from the cross section, one multiplies the cross section by 14 (for the neutron flux per square centimeter per hour at sea level for New York – JESD89), and again by 1E6 (for one megabit) and again by 1E9 (the billion hours in the FIT rate definition). We get then for FIT rates from Table 12: 445, 417, and 402 FIT/Mb.

Recent tests of Virtex-4 90nm devices at LANSCE have shown the cross section to be 2.15E-14 (>10MeV). In that same beam test, a Virtex-II device was also present, and measured 2.83E-14. If we use the ratio of the LANSCE results alone (and not their absolute value, which is known to vary from test to test), we can conclude that Virtex-4 is 0.76 the failure rate of Virtex-II (or that Virtex-4 has 0.76 the cross section of Virtex-II). This is in agreement with the Qcrit FIT rate prediction, which implies an improvement of 0.61. So far, the Rosetta results for Virtex-4 indicate a similar improvement, presently at 0.50. The improvement was by design, with the capacitive loading increased on every memory cell, in order to improve the FIT rate.

VII. Conclusion

Xilinx has presented data from the Rosetta Experiment, which correlates actual atmospheric upsets to simulated results, and to beam tests. LANSCE testing utilizing the 10 MeV model appears to be the most accurate predictor of atmospheric upset rate. Sensitivity to atmospheric upsets is shown to be at least eight times less than that predicted by the Boeing summary of experiments for SRAM devices. Upsets at different altitudes are following the accepted predictions. JESD89 predictions are too low at the higher altitudes. Cosmic ray flux variation with sunspot cycle is suspected of being observed, but it amounts to a variation of only 3 to 6% at present, which is less than the 95% confidence interval of 11 to ±12%. The actual upset rate experienced by users is much less than that of the rate of configuration memory cell upsets. This derating factor (SEUPI) varies depending on the actual user design and ranges from 6 to 100. Qcrit studies have aided the Xilinx IC designers to improve the field failure rate experienced by customers. Improving the Qcrit has involved both traditional as well as proprietary design techniques. Beam testing when used as a comparative measurement, and subsequently calibrated by the atmospheric tests, yields the best predictor of a new technology before the atmospheric test results are available.

Run #	Part Position	Count	Fluence >10 MeV	Control* S/N 1	S/N 2	S/N 3	Control* S/N 4	S/N 5	S/N 6	Average Errors	Cross-Section >10MeV
1	B	156,697	2.62E+09	1	1644	1708	0	1377	1358	1368	3.15E-14
2	B	466,852	7.82E+09	2	4936	4875	1	4005	3981	3993	3.08E-14
3	B	1,293,428	2.17E+10	9	13643	13779	2	10886	10915	10901	3.03E-14
4	A	562,097	9.41E+09	1	5044	5045	4	6745	6424	5045	3.23E-14
5	A	1,953,857	3.27E+10	1	17604	17413	1	22618	22615	17509	3.23E-14
6	C	311,031	5.21E+09	0	2436	2446	1	3286	3203	2441	2.83E-14
7	C	996,496	1.67E+10	1	7934	7701	1	10328	9974	7818	2.82E-14
8	D	468,256	7.84E+09	1	3873	3733	1	1579	2474		
9	D	1,799,424	3.01E+10	5	14402	14121	5	6685	9633		

* Control devices are not in the beam

Table 7. Neutron bit cross section data for Virtex-II (8/02)

Run	Vdd (V)	Count	Fluence >10 MeV	Fluence >1.5 MeV	CLB cross section	
					Errors	>10 MeV
Neutron SEU test Sept. 1st, 2003 (2V6000)						
14	1.50	74975	1.33E+09	2.55E+09	650	2.95E-14
Neutron SEU test, Dec. 2002 (2V6000)						
1	1.50	126013	2.14E+09	4.08E+09	1213	3.42E-14
	1.50	126013	2.14E+09	4.08E+09	1137	3.20E-14
5	1.50	98091	1.67E+09	3.17E+09	894	3.24E-14
	1.50	98091	1.67E+09	3.17E+09	858	3.11E-14
6 @ 60° angle	1.50	90780	1.54E+09	2.94E+09	906	3.54E-14
	1.50	90780	1.54E+09	2.94E+09	880	3.44E-14
7 @ 30° angle	1.50	95821	1.63E+09	3.10E+09	913	3.38E-14
	1.50	95821	1.63E+09	3.10E+09	873	3.24E-14

Table 8. Neutron bit cross section data for Virtex-II (12/02 and 9/03)

Run #	Epermnt. Config.	Pulse count	Counts/min	Run time	Fluence >10 MeV	Fluence >1.5MeV	Cross section Errors	Cross section >10 MeV	Cross section >1.5 MeV
1	A	156,697	3,561	44	2.62E+09	5.13E+09	142	2.19E-14	1.12E-14
2	A	466,852	3,112	150	7.82E+09	1.53E+10	367	1.90E-14	9.71E-15
3	A	1,293,428	3,804	340	2.17E+10	4.23E+10	174,960		
4	A	562,097	3,386	166	9.41E+09	1.84E+10	515	2.21E-14	1.13E-14
5	A	1,953,857	3,428	570	3.27E+10	6.39E+10	1,840	2.27E-14	1.16E-14
6	B	311,031	2,592	120	5.21E+09	1.02E+10	331	2.57E-14	1.31E-14
7	B	996,496	3,163	315	1.67E+10	3.26E+10	978	2.37E-14	1.21E-14
8	B	468,256	3,345	140	7.84E+09	1.53E+10	502	2.59E-14	1.32E-14
9	B	1,799,424	3,395	530	3.01E+10	5.89E+10	1844	2.47E-14	1.27E-14

A= Part behind 2V6000's

B= Part in front of beam

Table 9. Neutron bit cross section data for the Virtex-II Pro (8/02.)

Run #	Vdd (V)	Counts/min	Run time	Fluence >10 MeV	Fluence > 1.5MeV	Errors	Cross section >10 MeV	Cross section >1.5 MeV
15	1.50	72,826	30	1.29E+09	2.48E+09	150	4.06E-14	2.11E-14
16	1.50	354,730	120	6.28E+09	1.21E+10	737	4.10E-14	2.13E-14
18b	1.50	49,929	15	8.84E+08	1.70E+09	110	4.34E-14	2.26E-14
25	1.50	374,224	105	6.62E+09	1.28E+10	715	3.77E-14	1.96E-14
26	1.50	176,636	60	3.13E+09	6.02E+09	347	3.87E-14	2.01E-14

Table 10. Neutron bit cross section data for the Virtex-II Pro (9/03.)

Run	Vdd (V)	Count	Time (min)	Fluence >10 MeV	Fluence > 1.5 MeV	CLB Cross section Errors	CLB Cross section >10 MeV	CLB Cross section > 1.5 MeV	BRAM Cross section Errors	BRAM Cross section >10 MeV	BRAM Cross section > 1.5 MeV
Spartan-3 3S1000 testing (Sep. 1st, 2003)											
1	1.2	73,206	25	1.30E+09	2.49E+09	106	3.08E-14	1.60E-14	28	4.89E-14	2.54E-14
2	1.2	39,755	15	7.04E+08	1.35E+09	57	3.05E-14	1.58E-14	18	5.78E-14	3.00E-14
3	1.2	78,584	30	1.39E+09	2.68E+09	96	2.60E-14	1.35E-14	35	5.69E-14	2.95E-14
4	1.2	120,305	45	2.13E+09	4.10E+09	160	2.83E-14	1.47E-14	47	4.99E-14	2.59E-14
5	1.2	160,103	60	2.83E+09	5.46E+09	202	2.68E-14	1.39E-14	65	5.19E-14	2.69E-14
6	1.2	311,946	120	5.52E+09	1.06E+10	412	2.81E-14	1.46E-14	126	5.16E-14	2.68E-14
7	1.2	300,332	120	5.52E+09	1.02E+10	417	2.95E-14	1.53E-14	111	4.72E-14	2.45E-14
8	1.2	160,020	85	2.83E+09	5.45E+09	213	2.83E-14	1.47E-14	54	4.31E-14	2.24E-14
9	1.2	161,337	60								
10	1.2	30,743	14	5.44E+08	1.05E+09	31	2.15E-14	1.11E-14	10	4.15E-14	2.16E-14
11	1.2	211,320	90	3.74E+09	7.20E+09	254	2.56E-14	1.33E-14	87	5.26E-14	2.94E-14
12	1.2	87,320	30	1.55E+09	2.98E+09	112	2.73E-14	1.42E-14	40	5.85E-14	3.04E-14
13	1.2	233,173	90	4.13E+09	7.94E+09	356	3.25E-14	1.69E-14	100	5.48E-14	2.85E-14
27	1.2	201,972	100	3.57E+09	6.88E+09	235	2.48E-14	1.29E-14	75	4.74E-14	2.46E-14
28	1.2	132,054	50	2.34E+09	4.50E+09	173	2.79E-14	1.45E-14	43	4.16E-14	2.16E-14
29	1.2	97,026	25	1.72E+09	3.31E+09	111	2.43E-14	1.26E-14	43	5.66E-14	2.94E-14
Spartan-3 3S1000 testing (Nov. 11th, 2003)											
10A	1.2	157,176	60	2.78E+09	5.36E+09	244	3.30E-14	1.72E-14	59	4.79E-14	2.49E-14
11A	1.2	157,473	60	2.79E+09	5.37E+09	253	3.42E-14	1.78E-14	76	6.16E-14	3.20E-14
17A	1.2	2,031,744	725	3.60E+10	6.92E+10	3537	3.70E-14	1.92E-14	539	3.39E-14	1.76E-14

Table 11. Spartan-3 results.

References

[1]Ulrich Schade, Richard Wäsch, *The Stone of Rosette*, 1996. http://www.ba.dlr.de/ne/pe/virtis/stone1.htm

[2] Keith E. Holbert, *Single Event Effects.* http://www.eas.asu.edu/~holbert/eee460/see.html

[3] J. F. Ziegler, *Trends in Electronic Reliability -Effects of Terrestrial Cosmic Rays.* http://www.srim.org/SER/SERTrends.htm

[4] J. F. Ziegler, *Terrestrial cosmic ray intensities*. IBM Journal of Research and Development, Volume 42, Number 1, 1998 GMR, oscillatory coupling, and related studies. http://www.research.ibm.com/journal/rd/421/ziegler.html

[5] White Mountain Research Station, University of California, Office of Research, San Diego, Ca. http://www.wmrs.edu/facilities/BAR/default.htm

[6] Caltech Submillimeter Observatory Hilo, Hawaii. http://www.cso.caltech.edu/

[7] NOAA, *Estimated Values of Magnetic Field Properties*. http://www.ngdc.noaa.gov/seg/geomag/jsp/struts/calcPointIGRF

[8] Eugene Normand, *Single Event Upset at Ground Level.* http://www.boeing.com/assocproducts/radiationlab/publications/

[9] Xilinx, *TMRTool.* http://www.xilinx.com/products/milaero/tmr/index.htm

[10] Candice Yui, Gary Swift and Carl Carmichael, *Single Event Upset Susceptibility Testing of the Xilinx Virtex II FPGA.* http://parts.jpl.nasa.gov/docs/P29_yui.pdf

[11] Joseph J Fabula, Austin Lesea, Carl Carmichael, Saar Drimer, *The NSEU Response of Static Latch Based FPGAs.* Presented at the MAPLD Conference, September 2003.

[12] Joseph Fabula, Jason Moore, Austin Lesea, Saar Drimer. *The NSEU Sensitivity of Static Latch Based FPGAs and Flash Storage CPLDs.* Presented at the MAPLD Conference, September 2004.

[13] JEDEC STANDARD, *Measurement and Reporting of Alpha Particles and Terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices.* JESD89, August 2001, JEDEC Solid State Technology Association.

[14] IEC TS62396 , *Standard for the Accommodation of Atmospheric Radiation Effects via Single Event Effects within Avionics Electronic Equipment*. p (draft AR-9).

[15] Overview of iRoC Technologies' Report "Radiation Results of the SER Test of Actel, Xilinx and Altera FPGA Instances" http://www.wmrs.edu/facilities/BAR/default.htm

Austin Lesea (M'75) Austin received his Bachelor of Science in Electrical Engineering and Computer Sciences, University of California at Berkeley in 1975. In 1976 he received his Masters of Science in Electrical Engineering and Computer Sciences, University of California at Berkeley in optimization, control, and communications theory

He has spent the first year of his career teaching, and writing the book, *Microprocessor Interfacing Techniques (1976,* Sybex Publishing, Berkeley, California). He spent the next 13 years designing telecommunications transmission (copper, fiber optic, digital microwave), switching (PBX and ACD systems), and synchronization systems (building integrated timing systems). During that time, he spent eight years on the ATIS/ANSI T1 Committees helping write the SONET and TI CPE standards.

Mr. Lesea has spent the last six years at Xilinx where as a Principal Engineer in the IC design group where he has been part of the design efforts for five product families. Recent interest in radiation effects has led to the Rosetta study work, however he is also actively involved in the configuration design group with special emphasis on encryption and security. He holds 23 patents.

Saar Drimer received a Bachelors of Science degree in Computer Engineering from the University of California at Santa Cruz.

Upon graduation he joined the Advanced Product Division at Xilinx as a Design Engineer. He developed the Rosetta experiments and participates in the characterization and verification of new FPGA families. Currently, he is focusing on security, cryptography and random number generation from an FPGA and hardware perspective. He is generally interested in the broad application space of FPGAs, and finding solutions and enhancements to accommodate industry needs.

Mr. Drimer currently has six patents pending.

Joseph J. Fabula earned a Bachelor of Engineering (with Honors) degree, Stevens Institute of Technology, 1965 and a Masters of Business Administration (in Management), Fairleigh Dickenson University, 1981.

He started his career at RCA, holding various engineering and later management positions spanning 18 years there. He was a key contributor to the establishment of RCA's radiation tolerant and radiation hardened standard cell product lines, establishing production capabilities for radiation hardened bulk silicon and silicon on sapphire CMOS technologies. He is the co-inventor of and has several basic patents issued (to RCA Corporation) in the field of radiation hardening of channel oxides and in edge-defined guard banding for the control of radiation induced leakage currents. He was a recipient of the RCA Sarnoff Laboratories Achievement Award for his contributions in the field of radiation hardening. He also served as the Program Manager for several Air Force Avionics Laboratory and Air Force Materials Laboratory funded Manufacturing Technology Programs which established RCA's manufacturing capabilities for radiation hardened CMOS processes. Subsequently he joined the ITT Corporation for a period of 8 years, first serving as the Director of Semiconductor Process Engineering at ITT Semiconductors in Shelton CT. and later as the Director of Quality Assurance for ITT Power Systems in Tucson, AZ.

Mr. Fabula joined the Xilinx Corporation in 1991 as the Director of Quality Assurance, a position he held for eight years. He is currently supporting Xilinx terrestrial and deep space radiation testing efforts and providing customer applications and quality support for the Aerospace, Defense and Automotive products businesses at Xilinx Corporation.

Carl Carmichael holds a BSEE degree from the California State University, Fresno CA.

He has performed much of the radiation and SEU characterization for the Virtex, Virtex-E, and Virtex-II families of Xilinx FPGAs. He developed and verified mitigation methods for these technologies and aided several Xilinx aerospace customers in their development of scientific, defense, and commercial telecommunication re-configurable systems for use in orbital applications. He has been an invited speaker at various venues and has published numerous papers at the Nuclear Space and Radiation Effects Conference (NSREC) and the Military and Aerospace Applications for Programmable Logic Devices Conference (MAPLD).

Mr. Carmichael was awarded three United States patents on embedded system design, with four new patent applications pending for novel approaches to addressing the effects of Single Event Upsets in advanced commercial CMOS devices. He is well known within the United States and European Military and Aerospace industries, as well as the international radiation effects community. He led the establishment of a consortium of United States based aerospace companies and organizations (including the Aerospace Corporation, Sandia National Laboratories, Jet Propulsion Laboratory, Los Alamos National Laboratory, SEAKR Engineering, Ball Aerospace, General Dynamics, ITT, Boeing Satellite Systems, and Northrup Grumman ST), combining the unique talents of these government research facilities and industry leaders to further the advancement of leading programmable technologies for the aerospace industry.

Peter Alfke came to the US in 1966, with a German MSEE degree and nine years experience in digital systems and circuit design at LM Ericsson and Litton Industries in Sweden.

He has been manager, later director of applications engineering for over 35 years, at Fairchild, Zilog, AMD, and, since 1988, at Xilinx.

Mr. Alfke holds twenty patents, has written many Application Notes, presented at numerous design conferences, and has given many applications-oriented seminars in the US and in Europe. He is an active participant in the best newsgroup for FPGA users, comp.arch.fpga.