

On-Orbit Results for the Xilinx Virtex-4 FPGA

Heather Quinn, Paul Graham, Keith Morgan, Zachary Baker, Michael Caffrey, Dave Smith, Randy Bell

Abstract—This paper provides information regarding the use of the Xilinx Virtex-4 field-programmable gate array in a spacecraft deployed to low-earth orbit. The results are compared to pre-deployment accelerated and fault-injection testing.

I. INTRODUCTION

For the past several years the Xilinx Virtex-4 FPGA has been studied by a number of organizations for use in space applications. Several organizations have also launched the Virtex-4 components into low-earth orbit (LEO), including Sandia National Laboratories and the National Aeronautics and Space Administration (NASA). Recently, Los Alamos National Laboratory (LANL) launched an experimental payload, called the Mission Response Module (MRM), on a Department of Defense satellite. The MRM system includes two units, called Los Alamos Experimental Units (LEUs). Each LEU includes two Xilinx Virtex-4 FPGAs: one XQR4VLX200 and one XQR4VSX55. The intention of the payload is to experiment with commercial technologies to determine whether these components would be useful to high-availability, high-throughput missions, such as the Department of Energy (DOE) program on space-based nuclear detonation detection (SNDD) that needs to withstand both long duration space missions and man-made radiation environments.

The LEO orbit for the MRM payload is a particularly harsh orbit for commercial components as it transits

This work has been authored by an employee of the Los Alamos National Security, LLC (LANS), operator of the Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396 with the U.S. Department of Energy. The U.S. Government has rights to use, reproduce, and distribute this information. The public may copy and use this information without charge, provided that this Notice and any statement of authorship are reproduced on all copies. Neither the Government nor LANS makes any warranty, express or implied, or assumes any liability or responsibility for the use of this information. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness. This paper is published under LA-UR-12-22759.

H. Quinn, P. Graham, K. Morgan, Z. Baker, M. Caffrey and D. Smith are with Los Alamos National Laboratory, Los Alamos, NM, 87545, USA (e-mail: hquinn@lanl.gov).

R. Bell is with the Department of Energy, Washington, DC 20585, USA.

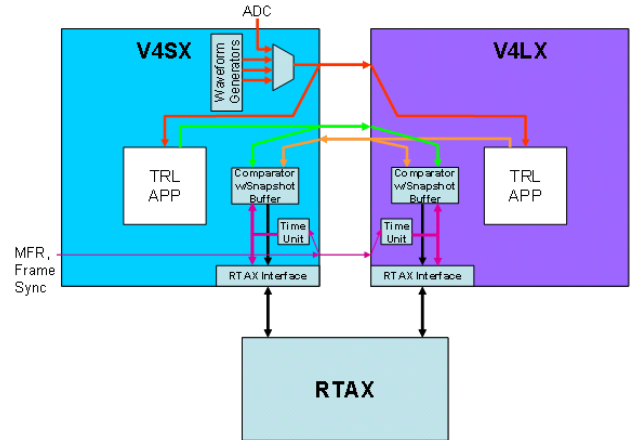


Fig. 1. A block diagram of the TRL application architecture shows how the two FPGAs compare results.

the peak of trapped proton belt. For commercial components with proton sensitivities, such as the Virtex-4, the predicted error rates can be quite high, which can make fault-tolerant computing difficult. One of the applications is a technology-readiness-level (TRL) application that allows the system to monitor the commercial components for computational errors. This application provides information about SEUs that are scrubbed from the FPGAs' programming data, computational failures in the FPGAs, and SEUs in the auxiliary memories that could affect the FPGAs' computation. This information is time stamped and returned to the ground station so that the errors can be correlated to where they happen on orbit. The system has also been instrumented to indicate the location on the FPGA that has been affected, so that SEUs can be correlated to accelerated radiation and fault-injection testing completed before launch.

In this paper, we present data for the MRM system from October 2011 to June 2012. Section II provides information about the MRM system and the TRL application. Section III provides information regarding accelerated radiation test data collected before launch that is compared to on-orbit results in Section IV.

II. THE TRL APPLICATION AND THE MRM SYSTEM

Fig. 1 illustrates the general user circuit architecture for the TRL application. The TRL application is a digital signal processing (DSP) application that has been

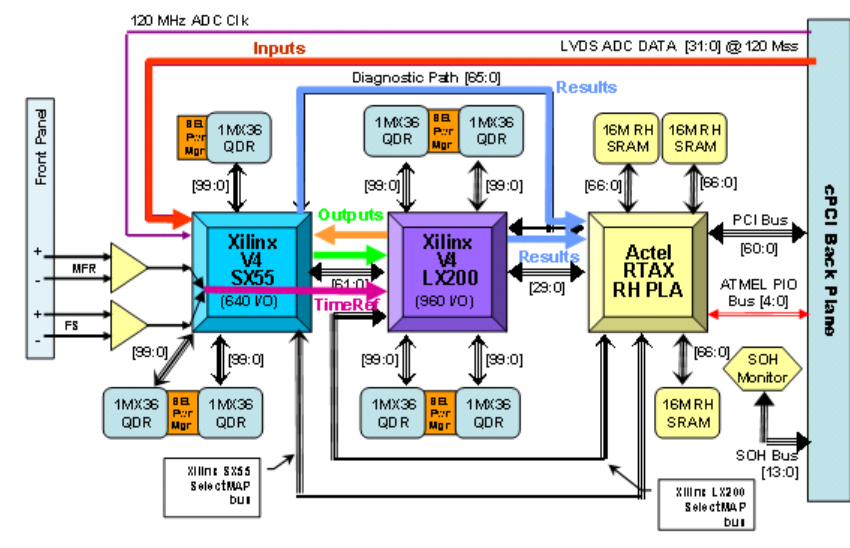


Fig. 2. A block diagram of the board-level shows the data flow for TRL application.

implemented so that the same algorithm can run in lockstep on both FPGAs, despite the size differences in the FPGAs. The TRL application has been mitigated for SEUs using the BL-TMR tool for automated implementation of triple-modular redundancy (TMR) [1]. The user circuit includes intellectual property (IP) processing cores that were generated in the Xilinx CoreGen tool. Due to CoreGen's use of proprietary design formats, the BL-TMR tool triplicated these cores at the core level instead of at the gate level. The rest of the circuitry was triplicated at the gate level.

Fig. 2 provides a similar board-level view of the data flow for the overall application. This figure also shows that the two FPGAs are scrubbed using an Actel RTAX FPGA. The scrubber algorithm has a frame-based readback and repair methodology. Readback-based scrubbers continuously read FPGA configuration memory frames and only repair single frames of data as SEUs are detected. The scrubber uses radiation-hardened static random-access memories (SRAMs) from Aeroflex to store the "golden" checksums. Because SEUs can affect the golden checksums, the scrubber for this system can also detect and scrub SEUs in the golden checksums.

As the two FPGAs in each LEU have the same algorithm and the same input data, it is then possible to use one FPGA to monitor the other FPGA for observable output errors. This process is done by sending the output of the TRL application code to both a local comparator and a comparator on the other Xilinx FPGA. The comparators themselves are triplicated for added reliability. Additionally, the comparator redundancy between the two FPGAs provides additional confidence in the reported errors. If the two comparators report the




Fig. 3. The percentage of MBUs for the Virtex-4 found during heavy ion testing.

same errors and the same captured data streams, there is a significant probability that the errors are from the TRL application and not the comparators themselves. If only one reports an error, then the likelihood is that a comparator error occurred.

When an error occurs, each comparator circuit notifies the RTAX and samples of the output streams are sent to the ground for further processing. Furthermore, information about SEUs and SEFIs are sent to the ground for processing. The system provides an approximate time for all of these different types of events to aid in of correlating output errors with SEFI or SEU effects.

III. PRE-LAUNCH ACCELERATED AND FAULT INJECTION TESTING

Several organizations have performed radiation testing of the Virtex-4 and the results are summarized in [2]–[6]. In [2], the saturated heavy ion cross-section for

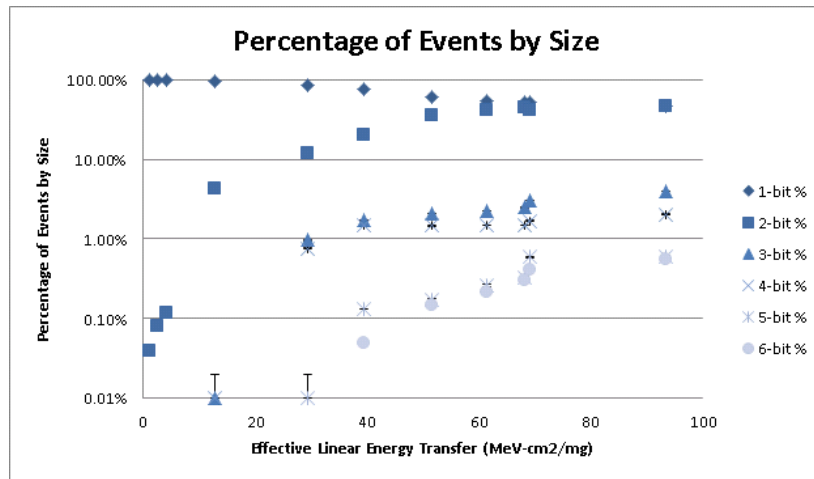


Fig. 4. The percentage of events by size found during heavy ion testing.

configuration memory is stated as $2.6 \times 10^{-7} \frac{cm^2}{bit}$ and for BlockRAM memory is $3.5 \times 10^{-8} \frac{cm^2}{bit}$. The proton cross-sections for configuration memory and BlockRAM memory are the same at $4.5 \times 10^{-14} \frac{cm^2}{bit}$ [2]. The data in report [2] were used as the inputs to CREME96. LANL also collected its own heavy ion and proton test results from 2004 to 2006, which correlate to previous test results. The LANL testing also included a validation of the scrubber using 63.3 MeV protons to ensure proper scrubbing before launch.

As the application was to be triplicated, studies on multiple-bit upsets (MBUs) were completed, as MBUs have been shown to be a limitation to TMR [7]. The MBU analysis for the Virtex-4 was published in [6]. These results indicated that $3.25\% \pm 0.05\%$ of all SEUs caused by 63.3 MeV protons would be MBUs, although the configuration memory could have up to $6.75\% \pm 0.10\%$ MBUs. In heavy ions, the MBU rates could be much higher, as shown in Fig. 4. For the LET ranges that cover the naturally occurring space particles, single-bit upsets (SBUs) occur between 86.18 – 99.96% of the time. Of the MBUs that occur, most of the upsets are 2-bit upsets, as shown in Fig. 4. Only at $30 \frac{MeV-cm^2}{mg}$ and above do 3-bit and 4-bit events occur at 1% or higher rates and 5-bit and larger events are exceptionally rare unless the charged particle strikes the component at a large angle of incidence.

Fault injection testing of both FPGAs was completed prior to launch using a similar methodology to the one discussed in [8]. Fault injection is a useful process for finding *sensitive bits* in the application. Sensitive bits are SRAM bits in the FPGA user circuit that can cause an observable output error if flipped. These sensitive bits are sometimes caused by problems with the application of TMR, especially in cases where TMR has not been

completely applied to the user circuit. We have found from previous experiments that the number of sensitive bits in an FPGA user circuit can be used with the static cross-section information to define an observable output error rate and that this rate is often much smaller than the SEU rate for the components, due to utilization and logical masking. The fault injection tests showed that there were some single points of failure remaining in the application, but no persistent cross-section [1].

Fault injection was particularly challenging for this application, because the application requires one second to properly initialize and synchronize. As the fault injection methodology that LANL follows requires resetting the system after the fault injection process is completed for one bit, it would have taken two years to complete fault injection. For this application, the system was reset only after injecting a fault that caused an observable output error. Because it is possible that corrupted state created the output error and not the fault injection into a particular bit, we completed the standard LANL fault injection process for a window of bits around the location that caused the output error.

After running multiple tests using the faster fault injection process, we found that there was a subset of bits found from fault injection that caused output errors for every fault injection pass, which we call high frequency sensitive bits. When we completed the detailed fault injection process for the windows, we found that high frequency sensitive bits were very likely to occur in the detailed fault injection processes. While we were able to perform one pass of fast fault injection on either FPGA in about four hours, it took months to complete the detailed fault injections on the windows.

For this orbit, CREME96 predicted an SEU rate of 68-89 SEUs/device-day for each LEU, depending on solar

conditions. When we deployed we had no concept of what the MBU rate would be, as there is no tool to generate such a calculation, due to chip design dependencies. Fault injection results for the application indicate that 0.4% of all SEUs in the configuration memory of the XQR4V5X55 would cause observable output errors. Likewise, 0.1% of all SEUs in the configuration memory of the XQR4VLX200 would cause observable output errors. When using the worst-case predicted SEU rates, each Virtex-4 should have an observable output error approximately every 15 to 25 days. These results also indicate the application should be able to withstand over 99% of the SEUs that occur in the component without manifesting an observable output error.

CREME96 was also used to determine the predicted SEU rate for the radiation-hardened SRAM that store the golden checksums. In this case, the SEU rate is 12 SEUs/device-month, depending on solar conditions. As the golden checksums do not completely fill the component, we expect the observed error rate to be lower than predicted for the entire device, as we do not observe SEUs in unused portions of the memory components. Therefore, we are expecting error rates to be about 1.2 SEUs/device-month.

It should be noted that incomplete information about the satellite could affect the SEU rates. In particular, we only have partial knowledge of how much metal is surrounding the payload. The SEU rates were estimated using a standard 100 mils of aluminum in absence of better information.

IV. ON-ORBIT RESULTS

In this section we compare the accelerated and fault injection test results with the on-orbit behavior of the system by looking at several months of operation. This discussion focuses on where the SEUs occur in the orbit, how often MBUs occur, where MBUs occur in the orbit, and validation of the CREME96 predicted SEU rates. These results use extensions of the analysis tools described in [6] and [8] that were modified to take the on-orbit data as input. All of the results presented in this paper are for the TRL application only.

A. SEUs

During this time period the two LEUs have seen a total of 3,170 SEUs between the four FPGAs. Fig. 5 shows a breakdown of SEUs/device-day for each of the four FPGAs. The TRL application runs exclusively on LEU1 and rarely on LEU2, so the data for LEU2 is not as complete as LEU1. On average the LX FPGA has 14.4 SEUs/device-day \pm 0.09 SEUs/device-day and the SX




Fig. 5. SEUs/device-day is broken down by month and FPGA.

FPGA has 5.2 SEUs/device-day \pm 0.09 SEUs/device-day. The ratio of device-days between the LX and SX is consistent with the variation in the FPGA sizes.

Figs. 6 and 7 show a projection of the sub-satellite locations for SEUs for both of LEU1's FPGAs. These plots are done in a heat map style that shows the density of SEUs in 3° squared regions. This plot shows the majority of the SEUs occur in and around the South Atlantic Anomaly (SAA). The handful of non-SAA SEUs occur predominantly around the northern polar cap. As yet, the payload has not had much interaction with the southern polar cap.

For the same time period, CREME96 predicted that the number of SEUs for the two LEUs should have been 14,830 SEUs based on operational usage. This value indicates that the predicted SEU rate is approximately five times higher than the actual SEU rate. A breakdown of the per month ratio of actual to predicted SEUs is shown in Fig. 8. This graph shows not only very little month-to-month variation but also very little SX-to-LX variation. While we find this result to be consistent with our experience with flying the Cibola Flight Experiment (CFE), which has nine Virtex-1000 components, and other organizations' experience with their deployed Xilinx Virtex family components [9], we have examined the results in more detail. To this end, we studied the uptime in the SAA and the effect that shielding could possibly have on these results.

We noticed anomalous days during the turn-on phase of the payload where the ratio of the actual to predicted SEUs was very high and occasionally above 100%. We were able to determine that these anomalous days occurred when the payload was very lightly scheduled and coincidentally scheduled to be on almost exclusively in the SAA. We also noticed in March 2012 that there was regular per-orbit maintenance work occurring in the southern part of the orbit. This situation caused the payload to be off during a large portion of the SAA. As the SAA is where nearly all of the SEUs are predicted to occur, we changed the maintenance work to minimize the




Fig. 6. The density of SEUs on LEU1's LX shown by sub-satellite location.




Fig. 7. The density of SEUs on LEU1's SX show by sub-satellite location.

amount of time the payload is off. While this scheduling change caused an approximate 15% increase in uptime on an average day, it did not increase SEUs by 15%, as not all of the increased time is in the SAA. In Fig. 9 we show the amount of time each month that the payload is up in the SAA. We find this graph explains some of the dips in the SEU/device-day plot, such as in February. We took the data from duty cycle in the SAA and used that to predict SEU rates. This prediction is shown in Fig. 10. From this analysis we have found that the ratio of actual to predicted SEU rates is on average 52%.

Our other concern is the amount of surrounding metal in the satellite. While one side of the payload is on an outer edge of the satellite, the other side is surrounded by several inches of aluminum. In discussion with other researchers 10 cm of aluminum would shield the protons below 200 MeV, oxygen below 400 MeV and iron below 800 MeV [10]. Therefore, it is possible that extra shielding on the interior wall affects the SEU rate. Unfortunately, without a mechanical model of the

spacecraft, which we are unable to get, we are unable to quantify fully the effect of shielding on the SEU rates.

Because the actual SEU rate is much lower than predicted the time between observable output errors should also increase. The original estimate is from 15-25 days based on 68-89 SEUs/device-day, but should be amended to 75-125 days to reflect 19 SEUs/LEU-day. There have been two output errors since October 2011. One was triggered on in October 2011 by coincident SBUs in the LX and SX within seconds. Both SEUs were not in the list of known bit locations that trigger output differences found through fault injection. A second output difference was triggered in May 2012. Not much is known about this incident as the satellite was already having problems that affected messages being sent properly. The current rate for output errors is once every 136 days. Given the very low rate of occurrence the error bars put the output error rate at between 38 and 1360 days per output error. As expected, the payload is immune to over 99% of the SEUs that occur in the FPGAs. Furthermore, the LEUs

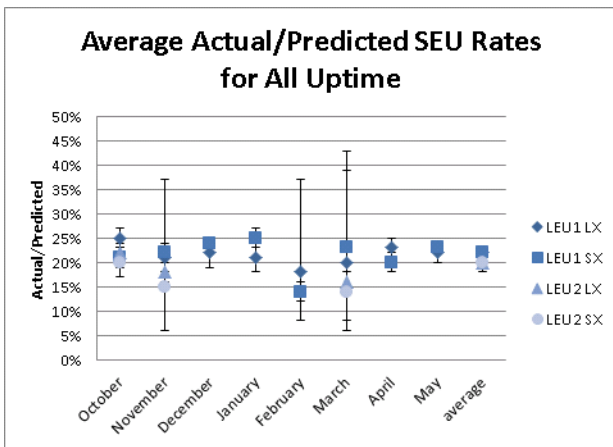


Fig. 8. The ratio of actual to predicted SEU rates based on duty cycle.

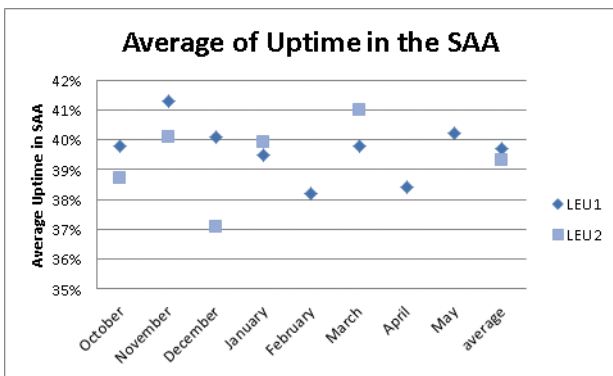


Fig. 9. The average duty cycle is given for the SAA region.

have been available for 0.999999 of all uptime.

Since October 2011 there have been five SEUs in the RH SRAM of LEU1: two SEUs in November 2011, one SEU in December 2011, one SEU in January 2012, and one SEU in February 2012. The CREME96 prediction for this part is 3 SEUs/device-week with full device utilization. For the given utilization of LEU1 and a 30 day month, LEU1 has been 100% utilized for the equivalent of 4.66 months and should have seen 56 SEUs in the RH SRAM for the scrubber since October 2011. We believe that the CRCs are using approximately 10% of one RH SRAM, which indicates the expected number of SEUs should be around 5.6, matching on-orbit experience well.

While there have been a number of flares since October 2011, only one flare caused a change in the SEU rate. In January 2012 there was an M8 solar flare with a full-halo coronal mass ejection and an X1 solar flare with an asymmetrical-halo coronal mass ejection. These flares created a significant disturbance to the proton fluxes as measured by the GOES satellites. During this time, the SEU rate increased by 25-50% without triggering observable output errors.

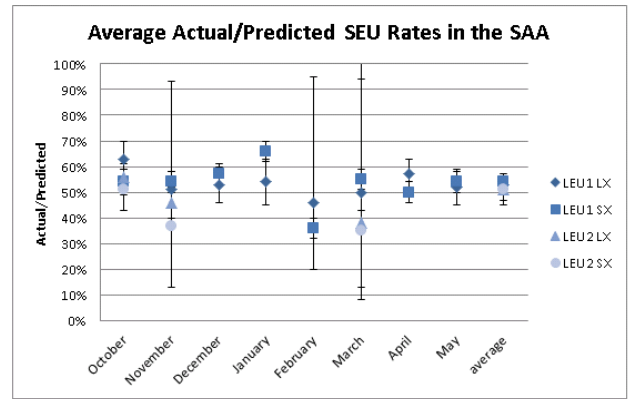


Fig. 10. The ratio of actual to predicted SEU rates based on the SAA duty cycle.

B. SEFIs

To date, one single-event functional interrupt (SEFI) has occurred. While the worst-case predictions for these components indicate SEFIs occur on the time frame of several years to decades, these error rates do not predict arrival times. On March 15, 2012, LEU1 was observed to have a SelectMAP SEFI. The SelectMap is an interface for programming the FPGA. When the SEFI occurs it might not be possible to read or write to the FPGA's programming data. In this case, the scrubber was unable to write to the component. The scrubber is designed such that each frame that is written to the component is checked for accuracy. If the same frame is written to the component multiple times and the SEU still exists, most likely a SelectMAP SEFI has occurred. In that event, the scrubber completes a full reconfiguration of the component to recover from the SEFI. When this SEFI occurred, the scrubber attempted to fix the affected frame the set number of times and then did a full reconfiguration of the component. Once the SEFI was cleared, the component returned to fault-free operation.

We have also seen a handful of events that are clearly not typical SEUs. In these cases, many bits in a single configuration frame were corrupted. While these types of events did occur occasionally in accelerated radiation testing, the event's occurrence was sporadic and we were not able to predict an on-orbit error rate. We are also not certain what the mechanism is that causes these events to manifest. The scrubber was able to correct these events without intervention and without triggering an observable output error.

C. MBUs

Figs. 11 and 12 show where SBUs and MBUs occur in sub-satellite plots. As with Figs. 6 and 7 these plots are heat maps that show the density of SBUs and

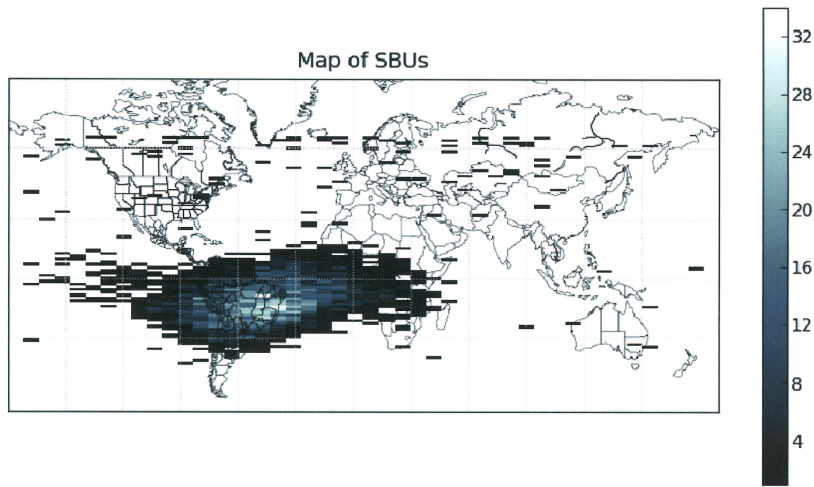


Fig. 11. The density of SBUs on both LEUs given by sub-satellite location.

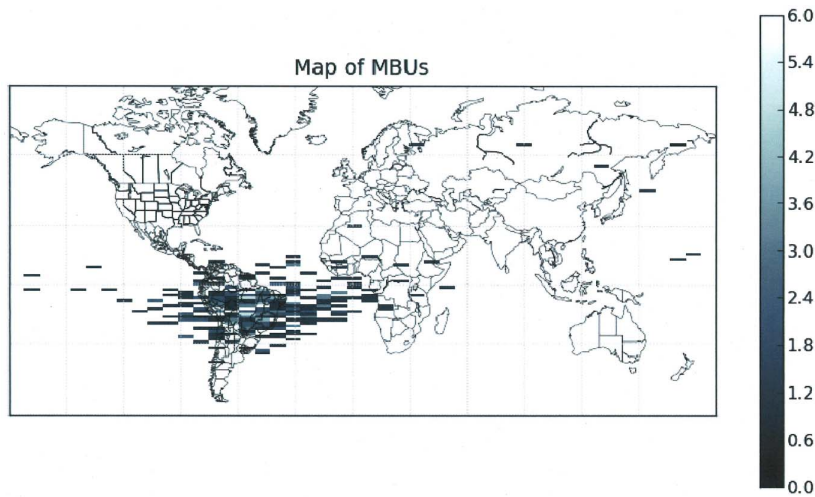


Fig. 12. The density of MBUs on both LEUs given by sub-satellite location.

MBUs in 3 squared regions. These plots show that the majority of MBUs occur in the SAA. We determined that $8.42\% \pm 0.39\%$ of the on-orbit SEU events are MBUs. None of the output errors have been tied to MBUs. Table I shows the percentage of MBUs out of all events for each FPGA. The data from the LX FPGAs are consistent with normal incidence 63.3 MeV proton tests, but the MBU rates for the SX FPGAs are twice as high as expected from proton testing. Table II shows the breakdown of SEUs by size. Given the results of accelerated testing for the component, large MBUs were not expected. While 6-bit events did occur in proton and heavy ion tests, 6-bit events only occurred 0.06% in 63.3 MeV protons and between 0-0.056% in heavy ions. Further analysis of these events determined that these events were exclusively occurring in the South Atlantic Anomaly (SAA), randomly distributed in time and location on the SX FPGAs.

Of the 54 4-bit and larger MBU events that have

TABLE I
RATIO OF MBUs OUT OF ALL EVENTS FOR EACH FPGA

Component	MBU
LEU1 LX	$6.82\% \pm 0.50\%$
LEU1 SX	$10.56\% \pm 0.88\%$
LEU2 LX	$6.52\% \pm 0.88\%$
LEU2 SX	$12.92\% \pm 1.46\%$

TABLE II
BREAKDOWN OF EVENTS BY SIZE

Size of Event	Percentage of All Events
1-bit	$91.58\% \pm 1.72\%$
2-bit	$6.59\% \pm 4.78\%$
3-bit	$0.92\% \pm 0.13\%$
4-bit	0.71% (0.20%, 1.66%)
5-bit	0.08% (0.06%, 0.28%)
6-bit	0.12% (0.07%, 0.36%)

occurred, only four of the large MBUs have been in the LX FPGAs. These four MBUs were all 4-bit MBUs. All of the 5-bit and 6-bit MBU events have occurred on the SX FPGA. MBU testing was exclusively completed on a Virtex-4 LX FPGA, which means that it is possible that the Virtex-4 SX FPGA has a tendency toward larger MBUs that was not accounted for by only testing the Virtex-4 LX FPGAs. These abnormal MBU events are also part of the reason why the SX MBU rates are higher than the LX MBU rates.

D. Memory Types

While the FPGAs' memory is made of standard SRAM cells, its memory cells are laid out in three feature sizes. From testing, we have found that these different memory cell layouts have subtly different sensitivities to radiation. Approximately $78\% \pm 1.53\%$ of the programming data upsets occurred in configuration logic blocks (CLBs), which make up the majority of the reconfigurable fabric, and $15\% \pm 0.54\%$ of SEUs affected the BRAM interconnect that connects the BlockRAM to the CLBs. In accelerated radiation testing the BRAM interconnect was particularly sensitive to radiation despite being a small portion of the SRAM cells. The BlockRAM interconnect is comprised entirely of routing logic, which tends to be the most sensitive type of memory on the Virtex-4. Of the remaining SEUs $6\% \pm 0.32\%$ of the SEUs occurred in the input/output blocks. Data on the BlockRAM is not available as it is not being monitored for SEUs.

When the data are analyzed in terms of memory cells used for routing or logic, most of the SEUs affect routing resources. SEUs in the programmable routing used to connect the programmable logic comprise $76\% \pm 5\%$ of the SEUs and the programmable routing used to interface to the hard cores, such as the digital signal processing units and the BRAM, comprises $13\% \pm 1\%$ of the SEUs. SEUs in the programmable logic, which consists of lookup tables (LUTs), comprise $6\% \pm 1\%$ of the SEUs. Previous testing has shown us that as a rule of thumb about two thirds of the observable output errors come from SEUs in the programmable routing, whereas a third of the observable output errors come from the programmable logic. While TMR protects both the programmable routing and the programmable logic, designers should remain aware that SEUs manifest most commonly in the most sensitive memory cells.

V. CONCLUSIONS

In summary, we have presented on-orbit results from a deployed Virtex-4. These early results show that the

CREME96 predicted SEU rate is four to five times higher than the actual SEU rate. We believe that the SEU rate is artificially low due to the duty cycle in the SAA and the amount of metal to one side of the payload. We found that $8.42\% \pm 0.39\%$ of the SEUs are MBUs and that there is potentially an anomalously higher rate of 6-bit MBUs than expected in SX FPGAs. We also determined that $78\% \pm 1.53\%$ of the SEUs are occurring in the CLBs. Finally, we had two incidents of observable output errors, which proves that the payload is immune to over 99% of all SEUs.

REFERENCES

- [1] K. Morgan, M. Caffrey, P. Graham, E. Johnson, B. Pratt, and M. Wirthlin, "SEU-Induced Persistent Error Propagation in FPGAs," *IEEE Transactions on Nuclear Science*, Vol. 52, No. 6, pp. 2438 – 45, 2005.
- [2] G. Allen, G. Swift, and C. Carmichael, "Virtex-4VQ static SEU characterization summary," Xilinx Radiation Test Consortium, Tech. Rep. 1, 2008.
- [3] M. Berg, C. Poivey, D. Petrick, D. Espinosa, A. Lesea, K. LaBel, M. Friendlich, H. Kim, and A. Phan, "Effectiveness of internal versus external seU scrubbing mitigation strategies in a xilinx fpga: Design, test, and analysis," *Nuclear Science, IEEE Transactions on*, Vol. 55, No. 4, pp. 2259 –2266, aug. 2008.
- [4] J. George, R. Koga, G. Swift, G. Allen, C. Carmichael, and W. Tseng, "Single event upsets in xilinx virtex-4 FPGA devices," in *Radiation Data Workshop of the Nuclear and Space Radiation Effects Conference*, 2006, pp. 109–113.
- [5] D. M. Hiemstra, F. Chayab, and Z. Mohammed, "Single event upset characterization of the virtex-4 field programmable gate array using proton irradiation," in *Radiation Effects Data Workshop, 2006 IEEE*, July 2006, pp. 105 –108.
- [6] H. Quinn, P. Graham, J. Krone, M. Caffrey, and S. Rezgui, "Radiation-induced multi-bit upsets in SRAM-based FPGAs," *IEEE Transactions on Nuclear Science*, Vol. 52, No. 6, pp. 2455 – 2461, December 2005.
- [7] H. Quinn, K. Morgan, P. Graham, J. Krone, M. Caffrey, and K. Lundgreen, "Domain Crossing Errors: Limitations on Single Device Triple-Modular Redundancy Circuits in Xilinx FPGAs," *IEEE Transactions on Nuclear Science*, Vol. 54, No. 6, pp. 2037 – 43, 2007.
- [8] H. Quinn, P. Graham, M. Wirthlin, B. Pratt, K. Morgan, M. Caffrey, and J. Krone, "A test methodology for determining space readiness of xilinx sram-based fpga devices and designs," *Instrumentation and Measurement, IEEE Transactions on*, Vol. 58, No. 10, pp. 3380 –3395, oct. 2009.
- [9] M. Caffrey, K. Morgan, D. Roussel-Dupre, S. Robinson, A. Nelson, A. Salazar, M. Wirthlin, W. Howes, and D. Richins, "On-orbit flight results from the reconfigurable cibola flight experiment satellite (cfesat)," in *Field Programmable Custom Computing Machines, 2009. FCCM '09. 17th IEEE Symposium on*, April 2009, pp. 3 –10.
- [10] J. George, Personal Communication.