Neutron Soft Errors in Xilinx FPGAs at Lawrence Berkeley National Laboratory

Jeffrey S. George, Rocky Koga, Margaret A. McMahan

***Abstract--*The Lawrence Berkeley National Laboratory cyclotron offers broad-spectrum neutrons for single event effects testing. We discuss results from this beamline for neutron soft upsets in Xilinx Virtex-4 and -5 field-programmable-gate-array (FPGA) devices.**

I. INTRODUCTION

NEUTRON single event effects (SEE) in semiconductor devices are important for a wide variety of ground and space applications. Neutrons produced by cosmic ray bombardment of the atmosphere affect aircraft avionics. Reactors and particle accelerators can be intense sources of neutrons. Even spacecraft are not immune from neutrons generated in spallation processes in Earth's radiation belts and albedo effects around the moon and other bodies, as well as in the shielding of the spacecraft itself.

Ground-based neutron sources of neutrons in the relevant energy ranges often have limited access and availability for single event effects testing. The Lawrence Berkeley National Laboratory (LBNL) 88" cyclotron currently has tunable broad spectrum neutrons for SEE and plans to add two mono-energetic beams during 2008 [1]. These beams offer experimenters additional resources for testing.

We measured soft neutron upsets in Xilinx Virtex-4 and Virtex-5 field-programmable-gate-array (FPGA) devices using broad-spectrum neutrons at LBNL during October 2007 and again in February 2008. Our main purpose was to explore the capabilities of the new LBNL beam line and our intent was to use the Xilinx parts as beam monitors for other testing. These devices are convenient for this purpose because they contain a large array of SRAM-based storage cells and they have a convenient method for loading and checking the cells.

The Xilinx devices are, however, interesting in their own right to the industrial community. They are used extensively in applications where neutron upsets are a known concern, especially avionics. There is also significant interest in flying them in space. The manufacturer has performed extensive testing of its own to show that neutron soft error rates are

low. Our measurements are not directly comparable to these prior data because of differences in test methods, knowledge of the device structure, and in the characteristics of the neutron beams used for testing. Nevertheless, the experiment illustrates the opportunities for neutron testing at LBNL as well as the unique challenges of measuring neutral particles compared to the charged beams.

II. THE LBNL NEUTRON BEAMLINE

The LBNL broad spectrum neutrons are produced by bombarding a thick beryllium target with energetic deuterons from the cyclotron. This is a highly-efficient and well-known method employed by light-ion accelerators for many years. The method is nicely described for a comparable facility. [2]. A dedicated beam line and irradiation station was added to the cyclotron and commissioned during 2007. It is now available on request to outside users.

The basic production mechanism is straight-forward. The deuteron (consisting of a single proton and neutron) is weakly bound. When a deuteron beam impacts a light ion target such as beryllium (Be), the deuteron falls apart into its separate constituents, each carrying roughly half the energy of the original particle (see Fig. 1). The neutrons easily escape the target and emerge with a Gaussian angular profile with a typical spread of 10-15 degrees from the beam axis due to the kinematics of the breakup reaction. The target thickness is sized to ensure that the remaining deuterons and the proton products are fully stopped in the target. The cross sections for charge-exchange reactions are low enough at these energies that the protons do not make any significant contribution to the forward-peaked neutrons. The beam power dumped into the target by the stopping deuterons and protons is removed by a water cooling system.

Neutron dosimetry is a challenge but well established standards exist for reactor environments and for mono-energetic neutron generators [3, 4] which can be adapted to the accelerator environment. For SEE applications, one would like to have a direct measurement of the neutron fluence on target. Neutral particles, however, do not ionize materials and can have very long mean free paths. The usual beam monitoring devices used for charged particle beams simply do not apply. In practice, the neutron beam intensity is inferred by more indirect means, and the dosimetry for SEE testing at the 88-Inch Cyclotron is still under evaluation and development. Our fluence measurements relied on

J. S. George and R. Koga are with The Aerospace Corporation, El Segundo, CA 90245 USA (telephone: 310-336-2089, email: Jeffrey.S.George@aero.org)
M. A. McMahan is with the Lawrence Berkeley National Laboratory, MS88R1092, 1 Cyclotron Rd, Berkeley, CA 94720 USA.

various combinations of neutron detectors (NE213 scintillator or fission chamber), neutron dose monitors, and activation foils.

Organic or inorganic scintillators such as stilbene or NE213™ have traditionally been used as neutron detectors [5]. They indirectly detect the neutrons through measuring the light produced by the scattered protons in the detector medium. They are sensitive to both gamma rays and neutrons; thus, pulse shape discrimination techniques must be employed to separate the two. Scintillators have very poor energy resolution. At the Cyclotron, the NE213 detector was used to measure the flight time and thus the energy of the neutrons, relative to the RF pulse of the Cyclotron. In addition, during the October 2007 run, it was used to monitor the beam current as a function of time.

A fission chamber is essentially a gas ionization chamber with a thin coating of fissile material, in this case ^{238}U, on an inner foil. A neutron interacts with the coating causing a nuclear fission reaction. The reaction products are charged and produce a large ionization pulse in the gas that can easily be counted. The thickness of the fissile layer is selected to maximize the neutron capture efficiency while avoiding saturation of the detector count rate. ^{238}U has an energy threshold of 1.8 MeV for neutron-induced fission, thus is insensitive to low energy neutrons. The LBNL fission chamber – a commercial unit made by LND, Inc. [6] - was not available during the October experiment but was used for some runs in the February 2008 exposure.

Neutron counters detect gamma rays from neutron-capture reactions. The gamma-ray detector is typically surrounded by a block of paraffin to slow the neutrons and greatly increase the capture efficiency [7]. These detectors are part of the standard radiation safety monitoring program at the laboratory and are permanently installed at numerous locations around the facility, including the "roof" directly above the irradiation cave outside the concrete shielding. Portable units are also available for monitoring at any location. The neutron counters are sensitive to all neutron energies. In the February 2008 run, the counter directly above the target (through 2 feet of concrete) was used to get a relative measure of the beam intensity over time.

Finally, activation foils are the only means of obtaining a direct measure of absolute neutron fluence and are a long-established technique [8]. Energetic neutrons can induce small amounts of radioactive nuclei in metal foils. The unstable isotopes decay by emitting gamma rays at characteristic energies with half-lives of minutes or hours. By measuring the intensity of the gamma lines at known intervals after the irradiation it is possible to extrapolate the abundance of the radioactive parent at the time of exposure. The cross sections for neutron activation of aluminum and other foils are well known [9]. This makes it possible to determine the total neutron fluence required to produce enough initial activation to match the radioactive decay measurements. The nuclear reactions leading to a particular activation product may have an energy threshold, which can

also give some information on the relative number of, e.g, thermal to fast neutrons. Several activation products with different half-lives serve as a self-check for the measurement. The advantage of this procedure is that the foils can be placed directly in line with the beam and can be used with very high intensity neutrons. In fact, the decay counting becomes easier with higher activation levels. The down side is that the foil packs must be prepared ahead of time and are single use. After exposure, each must be measured repeatedly at known times in a low-background counting facility. This makes them impractical for measuring individual test runs. They can, however, measure the total neutron fluence for a group of similar runs while using room counters to assign a relative share of the total to each exposure.
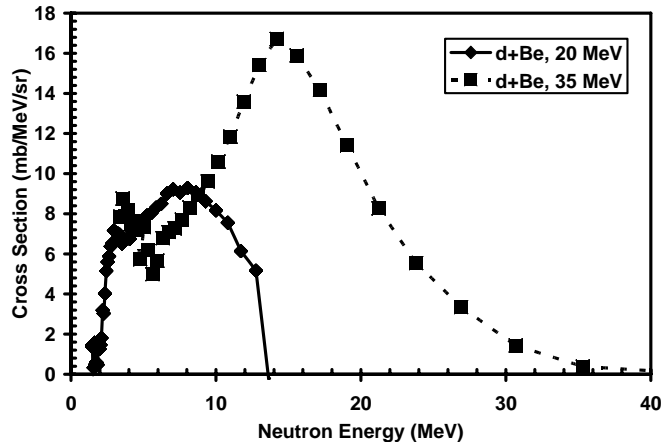

Fig. 1: Neutron production cross sections for deuterium on a beryllium target. The mean neutron energy is peaked at 14.2 MeV for 35 MeV incident deuterons. The 20 MeV beam produces neutrons peaked at about 7.5 MeV. The energy was calculated from the flight time of the neutrons relative to the radiofrequency pulse of the cyclotron [10].

III. TEST SETUP

We irradiated one Xilinx Virtex-4 (4VSX55) and one Virtex-5 (V5LX50T) field-programmable-gate-array (FPGA) device on commercial development boards. We loaded a known design into the part using a Xilinx Parallel IV configuration cable. After irradiation, the design verified against the original bitstream to identify upsets. The Virtex-4 bitstreams were read back and saved for later analysis. We were not able to save the Virtex-5 bitstreams because of changes in the Xilinx development software that support newer generation devices. In this case we could only record the total number of errors reported during verification by the Xilinx IMPACT download tool. Neither part was delidded as neutrons easily penetrate the device package.

A. Test designs

The Virtex-4 test used a 4VSX55 device socketed in a Xilinx AFX development board. The test design was nearly blank with only one route from a fixed internal '1' value to and external pin leading to an LED on the test board. The illuminated LED gave a quick way to verify that the design had been loaded. The device contains 22.7M bits available for upset, nearly all in the zero state.

The Virtex-5 test used a V5LX50T device soldered on a Xilinx ML-505 evaluation platform. For this test we loaded a Microblaze soft-processor design that we normally use as a testbed for single event effects testing of external memories. The only impact on this simple static upset test was that a larger fraction of the 14M available bits began in the '1' state.
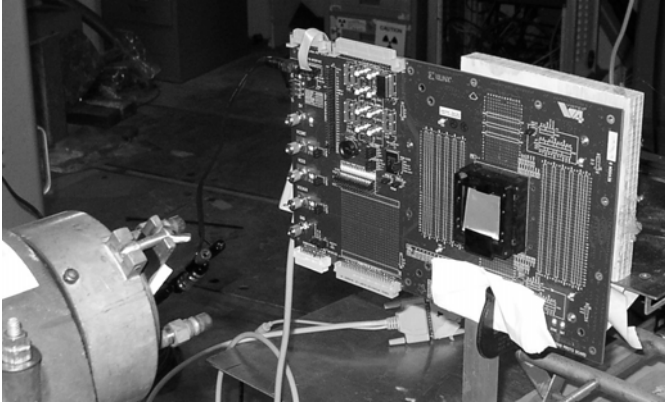

Fig. 2: Xilinx AFX development board with a Virtex-4 FPGA installed in front of the neutron beam. The beryllium target at the end of the cyclotron beamline can be seen at the left.

The parts were mounted 30 cm in front of the target during the October run and 8 cm in front during February. In both cases, we used deuteron beams of 20 and 35 MeV, resulting in neutron beams peaked around 8 and 14 MeV.

B. Beam monitoring

A wrapped packet of activation foils was placed directly behind the part on the back side of the board. After each of several groups of exposures the foil pack was removed and taken to a low-background counter where line strengths and decay rates were measured for several isotopes to establish the activation exposure level. This provided a direct measure of the total neutron fluence at the device for each group of exposures.

Room neutron counters were monitored at several locations inside and outside the irradiation cave. The one inside the room was placed off the beam axis and at some distance from the device to avoid saturating the count rate. A facility neutron counter located on the roof of the cave outside the concrete shielding gave a very nice record of the relative beam intensity as well as the absolute beam start and stop times for the activation analysis. This information was logged by the facility at one minute intervals. We used these logs to apportion the total neutron fluence from the activation foil among the several exposures in the group.

During the February runs we also had the ^{238}U fission chamber available. It had been newly reinstalled and a few initial glitches prevented us from obtaining readings for all the exposures but we did get some values for comparison. This should become a valuable standard fixture for future neutron testing as it can provide a real-time measurement directly behind the device. The fission detector has a

manufacturer's quoted efficiency of 0.001, which can be confirmed through comparisons with the activation foil results.

IV. ANALYSIS

The neutron soft error cross section σ (on a per-bit basis) are obtained by dividing the number of observed errors N_{err} by the neutron fluence F (n's/cm2) and the total available bits N_{bits} as follows:

$$\sigma = \frac{N_{err}}{F \cdot N_{bits}}$$

This straight-forward formulation immediately begs the question: what are the appropriate values for N_{err}, F, and N_{bits}? Our reliance on commercial software to verify the part and our indirect methods for monitoring the beam lead us to adopt certain choices that deserve explanation.

A. Number of errors

The number of upset bits after irradiation was taken to be simply the number of errors reported during design verification with the Xilinx IMPACT programming tool. The total number of errors per run ranged from a few hundred to a maximum of nearly twenty thousand. These numbers are large enough for reasonable statistical uncertainties while still corresponding to a tiny fraction of the total available bits. It is very unlikely that any single bit would have upset more than once. We had no way to determine whether a single neutron could upset multiple bits.

The IMPACT tool only reports the total number of verification errors, it does not indicate the location of those errors in the device. We were able to read back and save the Virtex-4 bitstreams with errors. Based on our previous experience with these devices, we were able to categorize the errors by functional blocks. For these devices, the cross section for an upset in the embedded memory contents portion of the bitstream is roughly twice that for an upset in the configuration logic block section. However, while we do not know whether an upset bit within a particular section is actually connected to a physical structure inside the device.

The manufacturer does know what the individual bits do. We submitted some of our bitstreams to Xilinx for further analysis and received some more detailed information on the CLB and BRAM errors that correspond to physical structures. However, it is not always practical to send designs to Xilinx for analysis. For the present purpose, we simply report the total number of errors reported by IMPACT, a measure available to anyone. The manufacturer's test results, however, may include information not available to the general user.

B. Number of available bits

The total number of target bits is a similar issue. The obvious number is simply the number of bits in the bitstream. One assumes that when the bitstream is loaded into the device, each bit programs something and may be upset by the

beam. In reality, the IMPACT verification process does not examine every bit in the bitstream. Some parts of the bitstream may be changed by the application, for example, shift register latches, lookup table storage elements, and the contents of embedded memory. For this reason, the verification step requires the existence of a 'mask file' that overlays the bitstream. This tells IMPACT which bits to compare for errors. If the number of errors is taken to be IMPACT verification errors then the total number of bits should be that reported by IMPACT. This should also correspond to the number of '0' bits in the mask file. For our designs, we used 22.7M total bits for the Virtex-4 and 14.0Mbits for the Virtex-5.

C. Neutron fluence

The neutron fluence for each exposure depends on the beam dosimetry. As described above, we used a combination of activation foils, neutron counters, and a fission chamber to determine the neutron fluence. It should be carefully noted that the neutrons used in this experiment are not mono-energetic but are rather sampled from a broad distribution peaked at a known energy. When we speak of the total number of beam particles, we are referring to an integration of the beam spectrum above some threshold response energy imposed by the counting method. This is yet another reason why it is very difficult to compare these results directly with data taken with a simulated atmospheric neutron spectrum at LANSCE or with true atmospheric neutrons as done with Xilinx's Rosetta project.

1) Activation foils

The following figure illustrates the analysis of activation data for a single ^{27}Al foil. Neutron bombardment of the foil generates two unstable isotopes through the production reactions ^{27}Al(n,α)^{24}Na and ^{27}Al(n,p)^{27}Mg. The ^{24}Na decays to ^{24}Mg with a 15 hour half-life by emitting an electron (β⁻). The excited nuclear state relaxes by emitting two gamma rays (1.368 MeV, 2.754 MeV) with known branching ratios that are counted in a low-background counting facility at known times after the exposure. The ^{27}Mg decays back to ^{27}Al with a 9.5 minute half-life, also emitting two gamma rays (843 keV, 1.014 MeV) to reach a stable configuration.

The gamma counts over time indicate the remaining abundance of the radioisotope. The solid lines in Fig. 3 extrapolate back to the abundance at the end of the exposure, using the known decay half-lives and gamma ray branching ratios. The efficiency for detecting the gammas in the counter must also be included. It is an energy-dependent number and is typically less than 0.1%. The two gammas from each radioisotope provide a consistency check.

The final step is to calculate what neutron fluence is required to generate the necessary amounts of radioisotopes at the end of the exposure using known nuclear activation cross sections. When the exposure time is long compared to the decay half-life of the products, the competing rates for production and decay must be taken into account. Note that the generated amounts of ^{24}Na and ^{27}Mg need not be the same

as they have different production rates. The two isotopes also have different energy thresholds for production (3.25 MeV for ^{24}Na and 1.896 MeV for ^{27}Mg) so the calculated neutron fluence calculated from each case may also differ depending on the spectrum of the neutron beam.
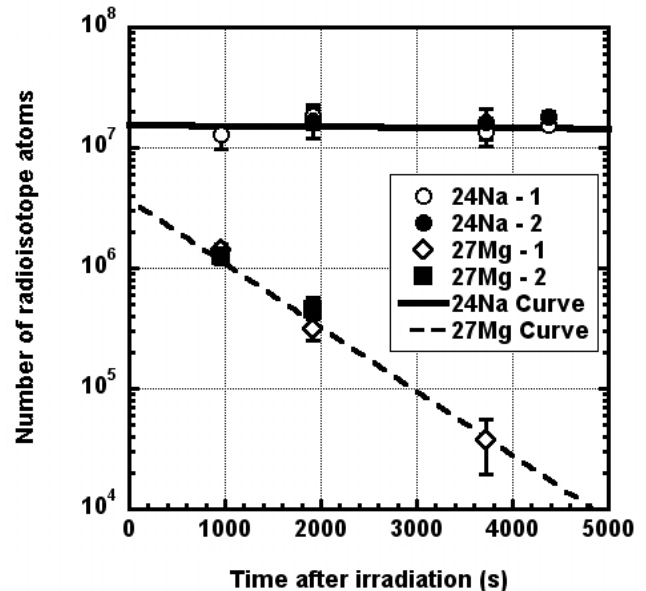

Fig. 3: Radioisotope abundance measured in an aluminum activation foil after exposure. The abundances inferred from the two gamma lines for each isotope are consistent, showing that the detector efficiencies and branching ratios have been properly accounted for. The ^{27}Mg has a lower energy threshold for production and a shorter half-life for decay.

2) Neutron counters

The facility neutron counters are logged at one minute intervals. They do not give the absolute flux of neutrons in the beam path but they can give a very good relative indication of the beam intensity. It is assumed that the spatial characteristics and spectrum of the beam do not change over the exposures of interest. Retuning of the beam optics after changing energies will affect the beam spread but once tuned, it should remain stable.
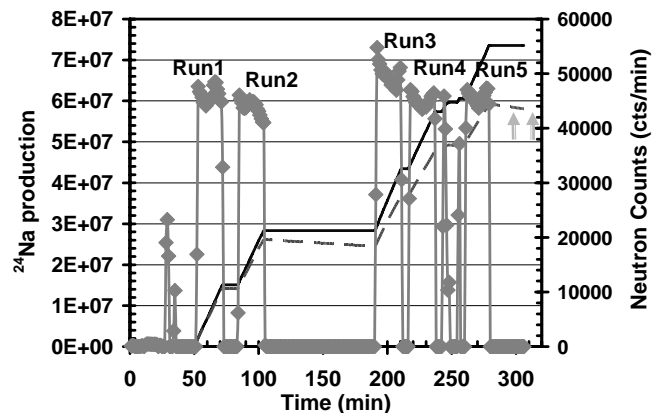

Fig. 4: Neutron counter log from roof monitor during February 2008 exposures. The increasing solid line indicates ^{24}Na production in the activation foil. The dashed line includes the effects of decay during the activation period. Small arrows at the end indicate when the activation foils were counted.

Error! Reference source not found. shows the logged neutron counter data for a portion of the February experiment. Each main run lasted 20 minutes with a few short tests in between the main ones. The increasing solid line on the plot shows the accumulation of ^{24}Na in the activation foil. The dashed line shows the same ^{24}Na abundance with decay over the time of exposure taken into account. The solid and dashed lines were normalized to give the correct total abundance when the foil was counted (times noted with arrows). This normalization establishes the total neutron fluence seen by the foil. The relative data from the neutron counter permits that fluence to be assigned on a run-by-run basis.

3) Fission chamber

The ^{238}U fission chamber (manufactured by LND) had just been reinstalled after refurbishment when our runs began. The electronics chain (amplifier, discriminator, scalar) had not yet been adjusted and so we were only able to obtain reliable fission chamber data for four of our runs. The detection efficiency is 10^{-3} based on the manufacturer's specification but no energy dependence is given.

We found that the fission chamber gave a consistently lower estimate of the neutron fluence than we inferred from the activation analysis. This may be a simple distance effect. The fission chamber was mounted separately in the beam axis at more than one meter from the Be production target. The results given below for the February exposures use the fission chamber results so an upward correction to the neutron fluence will only drive the soft error cross sections downward.

V. RESULTS

The results of the neutron upset testing are shown in Fig. 5 for both tested devices as a function of the neutron energy. The upset cross section per available bit is rather low and falls rapidly with lower beam energy. Also, as expected, the Virtex-5 had consistently fewer upsets than the Virtex-4 on a per-bit basis. Data from the two test trips were nicely consistent with the exception of the Virtex-5 low energy point. This difference is not understood but likely reflects a systematic error in our fluence calculation.

The statistical error in Fig. 5 is smaller than the plotted points. The uncertainty in the neutron fluence is not shown but may be as large as a factor of five. This uncertainty comes mainly from accounting for competing production and decay in activation foils that were left in the beam for times that were long compared to the radioisotope half-life. Using more foils for shorter exposures would have minimized the issue and simplified the analysis. In addition, the neutrons are not truly mono-energetic. The horizontal bars on some points are meant to indicate the full width of the energy spread from the production target.

The results shown here cannot be compared directly to Xilinx reported results of 1.55×10^{-14} cm^2/bit for Virtex-4 configuration bits and 0.67×10^{-14} cm^2/bit for the Virtex-5

from the LANSCE spectrum (>10MeV) [11]. (BlockRAM cross sections are given as $3-4\times10^{-14}$ cm^2/bit). There are significant differences in method and large uncertainties in the current measurements. Our peaked spectrum is unlike the LANSCE Hess spectrum commonly used to approximate the atmospheric neutron spectrum. Since we assume that any upset in the bitstream matters, even our error counting may differ. Given all these issues, we are encouraged that our results are roughly as expected but we caution against reading too much into either the absolute cross section or the energy dependence at this time.
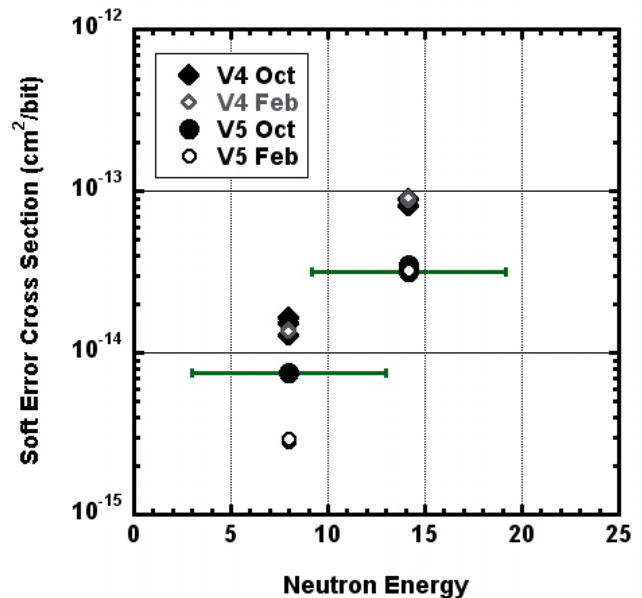

Fig. 5: Neutron single event upsets in Xilinx Virtex-4 and Virtex-5 FPGAs. The colors distinguish data taken on two different tests. The statistical errors are small compared to the points. Uncertainty in the dosimetry is not shown. Horizontal bars on some points indicate the neutron energy spread.

VI. FUTURE WORK

The LBNL 88" Cyclotron is now providing broad-spectrum neutrons for experimenters and will soon be adding mono-energetic beams. This offers new options for neutron testing. We have shown the results of neutron soft upsets in Xilinx Virtex-series FPGAs as an example of what can be done with the broad spectrum beam as well as to illustrate some of the differences from the usual proton and heavy ion beam experiments.

Many of the issues we encountered, particularly with dosimetry, were directly related to the developmental nature of the neutron beams at the time. LBNL is moving these beams into a regular production mode to provide a smoother experience for outside users. The fission chamber is being calibrated independently at the laboratory to verify the detection efficiency over the range of available neutron energies. With this calibration, the fission chamber will become the default beam monitoring device, giving a direct measurement of the neutron beam at the location of the tested part. In addition, two recently acquired mono-energetic neutron generators (2.5 MeV and 14 MeV) will be installed

during 2008. These will offer additional capability for neutron testing at the Laboratory.

VII. ACKNOWLEDGMENTS

VIII. REFERENCES

[1] M.A. McMahan et al., "Neutron beams from deuteron breakup at the 88-Inch cyclotron at Lawrence Berkeley National Laboratory" in Proceedings of the International Conference on Nuclear Data for Science and Technology, April 22-27, 2007, Nice, France, editors O.Bersillon, F.Gunsing, E.Bauge, R.Jacqmin, and S.Leray, EDP Sciences, pp 411-414, 2008.

[2] J. Collot et al., "A neutron irradiation facility featuring cryogenic temperatures and dedicated to Large Hadron Collider detector design," *Nucl. Inst. Meth. A*, 350, pp. 525-529, 1994.

[3] ASTM Standard #1854-07, "Standard Practice for Ensuring Test Consistency in Neutron-Induced Displacement Damage of Electronic Parts", published by ASTM International, www.astm.org

[4] ICRU Report #26, "Neutron Dosimetry for Biology and Medicine", published by ICRU Publications, Bethseda, MD USA, 1976

[5] G. F. Knoll, "Radiation Detection and Measurement", John Wiley and Sons, New York, 1979, pp. 570-586

[6] TM, LND Inc., Oceanside, NY USA, see www.lndinc.com

[7] See Ref. 5, pp. 552-558

[8] See Ref 5, pp 765-771

[9] W.S. Lyon, Jr., et al, "Guide to Activation Analysis", D. van Nostrand Company, Inc., Princeton, NJ, 1964

[10] D. Bleuel et al., "Characterization of a Tunable Quasi-Monoenergetic Neutron Beam from Deuteron Breakup**,** in Proceedings of the CAARI 2006, Fort Worth, TX, pp.

[11] A. Lesea, "Continuing experiments of atmospheric neutron effects on deep sub-micron integrated circuits," Xilinx White Paper WP286, http://www.xilinx.com/support, 2008.