

LA-UR- 09-01710

Approved for public release;  
distribution is unlimited.

<i>Title:</i>	On-Orbit Flight Results from the Reconfigurable Cibola Flight Experiment Satellite (CFESat)
<i>Author(s):</i>	Michael Caffrey <sup>1</sup> , Michael Wirthlin <sup>2</sup> , William Howes <sup>2</sup> , Daniel Richins <sup>2</sup> , Keith Morgan <sup>1</sup> , Diane Roussel-Dupre <sup>1</sup> , Scott Robinson <sup>1</sup> , Anthony Nelson <sup>1</sup> , and Anthony Salazar <sup>1</sup>
<i>Intended for:</i>	The Seventeenth Annual IEEE Symposium on Field-Programmable Custom Computing Machines  Napa, California, USA April 5-7 2009



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# On-Orbit Flight Results from the Reconfigurable Cibola Flight Experiment Satellite (CFESat)

Michael Caffrey<sup>1</sup>, Michael Wirthlin<sup>2</sup>, William Howes<sup>2</sup>, Daniel Richins<sup>2</sup>, Keith Morgan<sup>1</sup>, Diane Roussel-Dupre<sup>1</sup>, Scott Robinson<sup>1</sup>, Anthony Nelson<sup>1</sup>, and Anthony Salazar<sup>1</sup>

<sup>1</sup>Los Alamos National Laboratory  
Los Alamos, NM 87545, USA

<sup>2</sup>NSF Center for High-Performance  
Reconfigurable Computing (CHREC)  
Dept. of Electrical and Computer Engineering  
Brigham Young University, Provo, UT 84602, USA

**Abstract**—The Cibola Flight Experiment (CFE) is an experimental small satellite developed at the Los Alamos National Laboratory to demonstrate the feasibility of using FPGA-based reconfigurable computing for sensor processing in a space environment. The CFE satellite was launched on March 8, 2007 in low-earth orbit and has operated extremely well since its deployment. The nine Xilinx Virtex FPGAs used in the payload have been used for several high-throughput sensor processing applications and for single-event upset (SEU) monitoring and mitigation. This paper will describe the CFE system and summarize its operational results. In addition, this paper will describe the results from several SEU detection circuits that were performed on the spacecraft.

## I. INTRODUCTION

There is growing interest in using SRAM-based FPGAs within space systems due to low non-recurring engineering (NRE) costs, compressed life cycles and reduced costs (compared to ASICs), computational performance advantages, and reconfigurability. The ability to reconfigure SRAM-based FPGA devices after the spacecraft has launched allows them to be updated to accommodate evolving mission objectives, process data from multiple sensors, incorporate new scientific knowledge into the computational algorithms, or even to fix faults within the system. A variety of projects have demonstrated the benefits of using FPGAs in a spacecraft [1], [2]. Specific examples include the Mars rovers which use FPGAs for motor control and landing pyrotechnics [3] and the Australian satellite FedSat, which uses FPGAs as part of its high performance computing payload [4], [5].

While SRAM-based FPGAs offer a number of unique benefits for spacecraft electronics, they are susceptible to single event effects (SEE). SRAM-based FPGAs contain a large number of internal memory cells that can be upset by high energy particles. These include memory cells for user flip-flops, internal block memory, and for configuration memory. Single event upsets (SEU) within the configuration memory are especially challenging as these upsets may change the programming of the FPGA. Any system that incorporates

SRAM FPGAs within a high radiation environment such as space must provide a strategy for mitigating against such SEUs.

The Cibola Flight Experiment Satellite (CFESat), funded by the Department of Energy and developed by Los Alamos National Laboratory, tests the suitability of FPGAs as an on-board reconfigurable processors in a spacecraft [6]. The CFE instrument uses the reconfigurable logic to perform high-throughput RF sensor processing. The real-time processing demands of this system are immense and cannot be performed using multi-processing with traditional radiation hardened processor architectures. This platform also incorporates a number of techniques for detecting SEUs and mitigating the effects of SEUs within the FPGAs. This paper will describe the CFE system and the results of several SEU detection experiments operating on the satellite.

## II. CIBOLA FLIGHT EXPERIMENT (CFE)

The architecture of the processing payload of CFE is shown in Figure 1. As seen in this figure, the CFE payload includes an R6000 microprocessor, spacecraft communications interface, a digitally controlled radio tuner, a two channel, 12-bit, 100 MHz analog to digital converter, three reconfigurable computing processors using Xilinx Virtex FPGAs, and non-volatile memory to store program and FPGA configuration data.

The RAD6000 30 MHz microprocessor, a radiation hardened R6000 processor supplied by BAE, controls all of the payload digital modules and manages payload communications with the vehicle. The processor includes 8 Mbytes of radiation-hardened SRAM and executes the VxWorks operating system. Although this processor and its associated SRAM are radiation-hardened, the processor architecture is almost two decades old and does not have the computational power necessary to perform the on-board sensor processing.

One unique feature of the software architecture is the ability to dynamically link object code while in orbit. This unusual feature for software designs in space allows new software to be uploaded without reloading unchanged software components.

This work was supported by the Department of Energy (NNSA) NA-22. This work was also supported in part by the I/UCRC Program of the National Science Foundation under Grant No. 0801876.




Fig. 1. Payload Block Diagram

This allows the uplink communications channel to be used efficiently by eliminating the need to upload unchanged software. This dynamic object linking also allows the processor RAM to be conserved by only linking code for the currently running experiment rather than linking code from all experiments. The CFE software architecture also supports an unusual dynamic command dictionary. This dictionary has a command that inserts (and removes) additional commands. This allows operators and designers to conceive of new experiments, upload the object code, FPGA configurations, and insert new commands in the dictionary to execute the new experiment. These features enhance reliability by preventing bugs from migrating into the operational codes during whole system rebuilds.

The payload uses both EEPROM and flash memory for nonvolatile storage. Three banks of 1 Mbyte of EEPROM are available to store the operating system and binary user code objects for the microprocessor. Two banks of flash memory (24 Mbytes each) store compressed configuration bitstreams used to configure the Xilinx Virtex devices. More than 20 uncompressed FPGA bitstream configurations can be stored in each flash memory module. Error control coding (ECC) is incorporated to mitigate SEUs that occur during read or write operations in the nonvolatile memory.

The analog front end includes four RF channels, each connected to a distinct log-periodic antenna array, that can be “gang” tuned by microprocessor command between 100 and 500 MHz. This configuration is designed to make high fidelity interferometric measurements from a single lightning pulse. All four RF channels have an instantaneous bandwidth of 20 MHz. Two RF channels are combined into each of the 50 - 100 MHz intermediate frequency (IF) ADC inputs; this provides input from all four antennas simultaneously to the reconfigurable processors. The analog IF is sampled at 100 MHz with 12-bit resolution. The output of the payload ADCs is distributed across a network of point-to-point 200Mbyte/sec (32 bit x 50 MHz) LVTTL buses derived from the Front Panel Data Port (FPDP) specification. ADC data cascades through the three reconfigurable computers. Two reconfigurable computers each receive one channel of ADC data for preliminary processing, while the third RCC combines the two

intermediate results into a final measurement.

#### A. Reconfigurable Computer (RCC) Architecture

The processing payload was built around three reconfigurable computer (RCC) modules used to perform processing duties for a variety of experiments (see Figure 2). Each RCC module uses three Xilinx Virtex XQVR1000 CG560 FPGAs as the data processors. The FPGAs are organized in a ring and each has identical pinouts so they may share configuration files. This design strategy reduces the amount of nonvolatile memory needed for FPGA configurations, the required uplink bandwidth, and provides for greater reliability through redundancy. In addition, complex designs only need to be designed and verified once thus reducing design time on the ground. The nine FPGAs provide over 9 million system gates and over 1 Megabyte of block RAM memory.




Fig. 2. RCC Module Diagram

Each Virtex FPGA has 3 banks of independent memory; each bank is comprised of four Hyundai 64 Mbit SDRAMs organized as 8 M x 32-bit wide for a total of 288 Mbytes per module. Each RCC module also has microprocessor access through a radiation tolerant Actel RT54SX32S device that acts as a microprocessor interface and board controller. The Actel FPGA provides watchdog monitoring for the three Xilinx FPGAs as well as a configuration interface, which aids in CFE's FPGA SEU mitigation scheme.

While the use of Virtex FPGAs in this system may seem old when compared to FPGAs available today, the Virtex 1000 FPGA family was the most complex and dense FPGA available when the CFE system was first conceived. Since all satellite systems go through an extensive design, qualification, and testing procedure, the components used on orbiting satellites typically lag far behind the components available commercially. Furthermore, the Xilinx Virtex FPGA was the

first SRAM-based FPGA to go through extensive reliability qualification for radiation environments [7].

### B. SEU Mitigation

One of the most important issues that must be addressed when using SRAM-based FPGAs in a satellite is the presence of SEUs within the device. The CFE team investigated and developed techniques at the system level and application level for providing reliable operation of the SRAM-based FPGAs. First, the RCC boards used the QPro Virtex radiation hardened FPGAs [8]. These FPGAs use an epitaxial process that provides immunity against single event latch-up (SEL, an acute destructive failure) to an LET of 125 MeV/mg/cm<sup>2</sup>. The .25 micron process provides for approximately 100 K Rad(Si) total ionizing dose (a chronic and destructive radiation effect). These FPGAs are not, however, immune to SEUs within the user flip-flops, memories, and configuration data.

To detect and repair SEUs within the configuration memory, the CFE system employs a form of configuration scrubbing [9]. Configuration scrubbing is accomplished at the system level with the use of the radiation tolerant fused-based Actel FPGA. This device detects configuration SEUs by continuously reading the bitstream on each FPGA device through configuration readback. A cyclic-redundancy check (CRC) is calculated “on-the-fly” for each frame of a configuration bitstream. This calculated CRC is compared against the codebook CRCs that are precomputed on the ground. When an upset is detected by a CRC mismatch, a microprocessor interrupt is generated causing a reconfiguration of the upset frame from the onboard flash memory. When a SEU is detected, the exact bit that is upset and a timestamp is inserted into the ground telemetry. The use of configuration scrubbing prevents the accumulation of configuration upsets in order to significantly reduce the probability of having concurrent multiple configuration upsets.

In addition to configuration scrubbing, a variety of application-specific mitigation techniques have been developed for CFE. Specific techniques that have been applied include half-latch removal [10] and triple modular redundancy (TMR) [11], [12]. Half-latch removal involves the substitution of weak keeper circuits, which cannot be observed in the configuration bitstream, with logic structures for supply constant logic '1' and '0' values to the circuit. A tool for automatically applying TMR to a user design was also created [13]. This tool triplicates circuit resources and inserts majority voters to isolate any single upset caused by a configuration SEU. This tool provides the ability to “partially” mitigate user circuits when TMR proves too costly in hardware resources.

## III. CFE LAUNCH AND ORBIT RESULTS

CFE was launched into a circular low-earth orbit (560 km) on March 8, 2007 on a Lockheed Atlas-5 Medium rocket (STP1)<sup>1</sup>. Ground station connectivity was established quickly after the launch, and successful communication with

the satellite has been consistent over more than two years of flight time.

Since its launch, CFE has received configuration data from the ground more than three dozen times, both refining and increasing the portfolio of experiments within the reconfigurable payload. Over 20,000 experiments have been performed, where an experiment is the configuration of one or more Virtex devices and collection of data that is telemetered to the ground. This section will summarize the operation of the CFE in space including its power and thermal performance, SEU rates, and design applications.

### A. In-Orbit Temperature Control and Power Consumption

A major design concern of CFE was the satellite’s ability to adequately manage the heat generated by the FPGAs. While FPGAs are efficient in terms of watts/operation, absolute power consumption can be significant. Some CFE experiments require FPGA designs that consume more than 7 W in a single FPGA – this can be extreme for a space instrument in a vacuum. Further, the orbit environment includes wide temperature fluctuations due to solar exposure and Earth shading. The variation in power consumption by various reconfigurable experiments and the wide thermal dynamics create thermal cycles that must be carefully managed.

Managing the frequency and depth of these thermal cycles is critical for two reasons. First, timing performance of the FPGA design is governed to a maximum temperature limit. For a high reliability design like a space instrument, the manufacturer degrades the timing estimation in order to guarantee performance at high temperature (125 Celsius). By effectively designing the thermal system and managing operations to keep temperature below 70 Celsius, the devices can be used as though they are the fastest speed grade. Second, thermal cycling is a primary contributor to instrument failure due to wearout. Each cycle introduces stress on the mechanical assembly, due to mismatches in various material coefficients of expansion. The more extreme the cycle, the more severe the stress.

The most important mechanical interface for CFE is the printed circuit board to FPGA package interface. Due to the large package size (42x42mm), the stress on corner pins is the primary mechanical failure point. Two strategies were employed to minimize these thermal risks. First, an exotic printed circuit board material, thermount, was used that more closely matched the thermal properties of the FPGA package and hence reduce mismatch. Second, a system of heat pipes is used to transport heat from the FPGAs to the spacecraft radiators to limit maximum temperatures. The heat pipes are passive and particularly well suited for this application. Found in commodity personal computers, heat pipes use a working fluid, typically water, that goes through phase changes in order to transport large heat flows very efficiently. The heat pipes employed in CFE can transport more than 5W with less than 1 degree Celsius temperature drop across the pipe. The heat pipes use capillary action for fluid flow so that they operate in the absence of gravity. Also, when cold, the working fluid

<sup>1</sup>The current location of the CFE satellite can be tracked by visiting the following URL: <http://www.n2yo.com/?s=30777>

freezes and the pipe turns "off". The fluid is at low pressure and does not fill the pipe, so freezing does not cause damage.

Another important consideration on CFE is energy management. All power on the spacecraft is generated from solar panels and stored in Lithium Ion batteries. The power available to the payload is limited by the energy generation from these panels and that stored in the batteries. The FPGA configurations must be carefully scheduled to insure that the FPGA power consumption does not exceed the energy production for extended periods. These operational cycles result in discharge cycles on the spacecraft batteries and represents the primary wearout of the entire satellite. Similar to the thermal cycling mentioned above, the battery lifetime is determined by the number and depth of charge/discharge cycles, as well as the operating temperature. Ideally, the battery must be kept between 0 and 20 Celsius but preferably at 20.

Due to the configurability of the CFE payload, energy consumption is variable. Operational requirements demand that payload operations occur in the dark or that the payload perform experiments far in excess of the generation capacity of the panels. The dynamics of these operations can be appreciated by examining the current consumption of one of the FPGAs, shown in Figure 3. The maximum peaks in current consumption correspond to different experiments that are executed on the platform. These experiments, however, are not run for long periods of time as there is not sufficient energy to sustain them. These high-throughput, high-power applications are mixed with other low-power experiments to keep the average energy consumption within an acceptable range (note the relatively low mean current consumption). Power intensive designs run when more power is available, while lower power applications run when power constraints dictate. The highly dynamic power profile contributes to the thermal cycling and stresses the thermal management system.




Fig. 3. RCC1 Board Level +2.5V Current Consumption (Amps)

To mitigate against problems, payload operations are self-monitored and in the event of anomaly, the payload enters a safe state. A sensing circuit monitors a diode included on all FPGA die to measure their temperature: Figure 4 displays

the temperature history of one of the FPGAs to demonstrate the temperature fluctuations experienced by the reconfigurable computing boards. While the median temperature of the die is kept within a reasonable operating regime (between 5°C and 25°C), the minimum and maximum measured temperatures involve a wide range between -10°C and 35°C. Other parameters that are tracked include power consumption, operating voltages, and a variety of state variables. Those include the state of tasks within VxWorks, the status of data being sent (or not) to the spacecraft, and the processing of operation commands. A safety task monitors these variables against programmable upper and lower limits. In the event of an exceedance, the CFE instrument asserts a discrete "panic" signal to the satellite bus which responds according to one of three programmable possibilities: turn the payload off, reset the payload, or do nothing.




Fig. 4. RCC1 FPGA A Die Temperature History

#### B. In-Orbit SEU Rate

Several studies were performed to estimate the upset rate of the FPGAs in the 560 km low-earth orbit [14], [15]. These estimates were made using the CREME96 modeling environment and results from radiation testing on Xilinx Virtex devices [7]. The estimated SEU upset rates for the Virtex in this low-earth orbit are summarized in Figure 5. As seen in Figure 5, the estimated SEU rate for its orbit varies from 0.5 SEUs per device day (solar max, best case estimate) to 26 SEUs per device day (solar minimum, peak trapped protons).

All SEUs within the device have been logged during the CFE lifetime to measure the actual SEU rates of the system. Through configuration scrubbing, 759 SEUs have been detected over 2830.7 device days resulting in an average upset rate of .268 upsets per device day. The SEU upset rate is lower than the best case estimate and much lower than any worst-case conditions. With nine FPGAs in the payload, CFE averages 2.4 SEUs per day.

The SEUs do not occur uniformly as the spacecraft orbits around the earth. Figure 6 plots the location of the SEUs that have been detected in CFE FPGAs. The vast majority




Fig. 5. Anticipated FPGA Upset Rate

of SEUs, detected by configuration readback, occurred in the area known as the South Atlantic Anomaly (SAA). This is the region where the Van Allen radiation belt passes closest to the Earth's surface due to irregularities in the Earth's magnetic field.

### C. CFE Signal Processing Experiments

Several CFE signal processing payload experiments have been uploaded to reconfigurable platform and executed on the FPGAs. These signal processing experiments interface directly to the on-board ADC to process sampled data from the satellite antennae. Examples of this class of circuits that have been run on the satellite include several software defined radios (SDR), demodulators, decoders, and high-throughput FFT engines that exceeded a sustained computation rate of 10 Gops per second. The performance of the payload is two to three orders of magnitude better than what can be expected from radiation hardened microprocessors.




Fig. 7. 500-kbps QPSK Receiver Block Diagram

One of these signal processing experiment is a 500-KBPS QPSK receiver implemented on a single FPGA of the CFE reconfigurable computer (see Figure 7). The purpose of this experiment was to demonstrate a practical application using CFE's radio receiver and computing hardware. The experiment was successfully deployed and tested on-orbit in November 2008. A 500-kbps link was established between a transmitting ground station and CFE, with error free communication during the majority of the pass.

## IV. SEU DETECTION EXPERIMENTS

The second group of experiments tested on CFE are those used to measure SEU rates and test SEU mitigation strategies.

These experiments provide data collected in a real radiation environment that yields insight into the behavior of the FPGAs and appropriate mitigation approaches. One challenge faced by these SEU mitigation experiments is the low SEU rate compared to accelerator experiments or fault injection experiments. At a rate of one SEU upset/device day, it requires long experiment times to collect data and generate meaningful results to validate these techniques. As described in Section III-B, only 759 SEUs have been observed during the 741 days of flight time. This is several orders of magnitude less data than the data generated through artificial means.

An important class of applications that have executed on CFE are SEU detection experiments. These experiments were created to detect SEUs from *within* the FPGA fabric using additional logic and well known fault detection techniques. This in-circuit SEU detection operates in parallel with the scrubbing-based SEU detection scheme implemented in the Actel FPGA.

The SEU detection within the FPGAs on CFE employs a technique known as duplication with compare, or DWC (see Figure 8). To detect upsets in a circuit with DWC, two identical copies of the circuit run continuously while circuit outputs are compared at different points by comparator circuitry. Dual rail, self-checking comparator circuits are used so failures in the detectors can be identified. If the outputs of the two circuit copies disagree at any point, an error has occurred in one of the circuit copies, and the comparator at that point in the circuit outputs an error code. The outputs of all comparators in the circuit are merged to form a circuit-level error code which signifies the presence or absence of upsets in the FPGA [16].




Fig. 8. Duplication With Compare (DWC)

The in-circuit duplication experiments will only detect configuration SEUs for configuration bits that actively impact the circuit configured on the device. As described in [17], only a fraction of the upsets that occur will be detected because only a fraction of the configuration bits are used to configure any given circuit. Designs that utilize most of the FPGA logic resources typically use 10% or less of the configuration bits. Because of this, the DWC detection circuits will only detect a fraction of the SEUs that are detected by the readback mechanism.

One of the goals of these experiments is to detect as many upsets as possible. To do this, the design experiments must be as large as possible to provide sufficiently large "targets" for the high-energy particles. If the circuits are small, these experiments will not generate sufficient data. Further, these




Fig. 6. SEUs by Region

design experiments must consume as little power as possible. These experiments were scheduled onto CFE as low-power “background” tasks when other high-throughput experiments were not or could not be scheduled. The following subsections will summarize each experiment and the experimental results.

#### A. SEU1 - Configuration Upsets

The first SEU detection experiment, named *SEU1*, was designed as a low-power simple circuit that does *not* perform in-circuit detection. In this experiment, all 9 Xilinx FPGAs in the payload are configured with a simple circuit containing little logic and consuming minimal power. This simple experiment relied on satellite’s readback and configuration scrubbing technique for detecting and correcting SEUs in FPGA configuration memory. This experiment was created before the in-circuit detection techniques were available.

*SEU1* executed for 455.3 device days of operation (all 9 FPGAs operating for a total of 50.6 days). During this period, the readback process detected 216 SEUs indicating an upset rate of .47 upsets per day.

#### B. SEU2 - Online Detection

*SEU2* is a more sophisticated test than *SEU1* and the first experiment to implement the in-circuit detection using DWC. The base circuit of *SEU2* (i.e., the circuit before DWC is applied) includes a long 32-bit wide shift register driven by a gray code pattern generator (see Figure 9). A gray code was used to minimize the dynamic power. LUTs are inserted between each shift register with a pre-determined logic pattern. The use of LUTs between the registers provides more logic “area” for SEUs to hit. The output of each LUT drives the input of a flip-flop and the LUT inputs are driven by upstream flip-flops and the gray code counter.

In order to accommodate future detection designs, the base circuit of *SEU2* was designed to be entirely parameterizable in depth. Parameterization simplifies the process of creating a design that “fills” a device. This base circuit will be used in subsequent SEU test experiments. *SEU2* replaced *SEU1* on June 17, 2008 and operated for 101.6 device days. During this




Fig. 9. The Gray code generator and subsequent shift register.

time, 46 SEUs were detected with 4 of the 46 SEUs detected by the DWC circuitry (8.7% sensitivity).

#### C. SEU3 - BRAM

The *SEU3* experiment extended *SEU2* by detecting and reporting SEUs that occur within the block RAMs (BRAM) on the Virtex FPGAs. In the previous tests, there was no way to discriminate between upsets in any particular region of the device. For this test, a custom circuit was designed to detect BRAM SEUs by continuously scanning the entire BRAM memory, identifying SEUs, and reporting the total number to the processor. After receiving confirmation from the processor that the number has been recorded, the circuit proceeds to scrub (repair) the BRAM with predefined data.

Figure 10 shows the architecture of this BRAM scrubber and detector. This circuit includes a single gray code address generator (to reduce dynamic power) to drive all BRAM circuits, an error detector, and a scrubber. The scrubber inserts a 0xAAAA pattern into the BRAM and repairs this value whenever an upset is found. If any upsets were present when the BRAM has completed a scan, the processor is interrupted and the number of upsets is reported.

In addition to the BRAM scrubbing/detector circuit, *SEU3* includes the gray code shift register used in *SEU2* (Figure 9). The BRAM scrubber/detector is also sensitive to upsets. DWC is applied to this circuit as well to prevent erroneous data being reported due to SEUs in the detection circuitry. Should the DWC comparison circuitry detect an upset, an interrupt to the




Fig. 10. BRAM Scrubber and SEU Detector




Fig. 11. SDRAM scrubber

processor causes the FPGA to be reset and the configuration frame is fixed through conventional scrubbing.

*SEU3* replaced *SEU2* on July 14, 2008 and has been operational for 976.7 device days. During this time, 274 SEUs were detected with readback and 20 of these SEUs were detected by the DWC detection logic (7.3% sensitivity). One of these 20 SEUs occurred within the BRAM scrubbing logic. In addition, 23 BRAM upsets were detected by the scrubber. Twenty-two of these upsets involved single bit upsets while one of the 23 included two upsets within the BRAM.

#### D. SEU4 - SDRAM

The final detection test, *SEU4*, was designed to detect upsets within the SDRAM memory associated with each RCC FPGA (see Figure 2). A basic SDRAM controller was designed to run at the SDRAM clock rate of 52 MHz. As there are three 32 MB SDRAM banks for each FPGA, three SDRAM controllers were required in each FPGA (see Figure 11).

The SDRAM control circuit initializes the SDRAM memory by writing the pattern 0xAAAA in all 96 MB of SDRAM. After initialization, it continuously scans the SDRAM, one bank at a time, looking for deviations from this pattern. SDRAM refresh is integrated with the scanning/scrubbing process to ensure valid data. Upon completing a scan in which upsets are detected, the circuitry interrupts the processor to report the final upset count. When the processor has acknowledged

TABLE I  
CFE SEU APPLICATION RESULTS

	Device Days	Config SEUs	SEUs / D.D.	DWC SEUs	BRAM Upsets
SEU1	455.3	216	.47	N/A	N/A
SEU2	101.6	46	.45	4 (8.7%)	N/A
SEU3	976.7	274	.28	20 (7.3%)	23
SEU4	464.3	141	.30	9 (6.4%)	2

this interrupt, the SDRAM circuitry scrubs the entire array of SDRAM and resumes its scan.

The SDRAM controller and scrubber/scanner is triplicated using the BL-TMR tool in order to minimize the possibility of an upset occurring within the circuitry and causing erroneous data to be reported. The *SEU4* test merges the SDRAM circuitry with the existing *SEU3* detection test to provide detection for logic, BRAM, and SDRAM.

*SEU4* replaced *SEU3* on December 11, 2008 and has been operational for 464 device days. During this time, 141 SEUs were detected with readback and 9 of these SEUs were detected by the DWC detection logic (6.4% sensitivity). The lower percentage of SEUs detected by the DWC circuits is expected as fewer circuits are available for online detection (additional circuitry is needed for the SDRAM controllers). In addition to the configuration SEUs, two BRAM upsets were detected.

Unfortunately, no SEUs within the SDRAM have been detected by the *SEU4* experiment at the time of publication. A large number of SEUs within the SDRAM were expected and the lack of SDRAM SEU results suggests that there are problems with the *SEU4* SDRAM detection circuits. Future experiments are under development to resolve this problem and will be scheduled to operate at some future time.

#### E. Results

The four SEU detection experiments have been operational for a total of 1997.9 device days (see Table I). Within this time, 677 SEUs have been detected through readback, 33 SEUs have been detected with on-circuit error detection, and 25 SEUs within the BRAM have been detected with BRAM scrubbing. The detection experiments provide operational validation of the DWC approach and provide greater visibility into the impact of SEUs on the device. These and other SEU detection experiments will continue to be used on CFE.

#### V. CONCLUSIONS

The reconfigurable computing architecture within CFE has performed very well and continues to be used for a number of reconfigurable computing experiments. Future experiments include both real-time SEU mitigation tests and other signal processing tests.

The RCC modules within the CFE and the SEU mitigation approach used to protect them have proven successful. Several successful design techniques are worth mentioning. A symmetric layout was used for all FPGAs allowing the same

bitstream to be used for any FPGA. This symmetric layout was very helpful in simplifying the design and reuse of bitstreams across the platform. The dynamic command dictionary and on-orbit run-time linking was very convenient for allowing run-time scheduling and uploading of new FPGA bitstreams and the SEU scrubber design worked flawlessly.

While the Xilinx Virtex FPGAs have worked very well for this experiment, newer FPGA architectures will have a big impact on computational density and power. Specifically, the DSP48 primitives found in Virtex II and successor FPGAs would significantly reduce the size and power of the signal processing circuits used in this system. Also, the high speed serial I/O found on next generation FPGAs would significantly reduce the number of I/O pins needed for inter-FPGA communication. Even more importantly, a high speed serial network incorporating a runtime reconfigurable cross-point switch would allow the network topology to change for each application. This increased flexibility would allow the FPGAs to be used more efficiently than a hardwired network topology. This approach also increases the robustness of the system by allowing degraded or failed components to be gracefully removed from the system. A number of FPGA architectures succeeding the Virtex have been qualified for space operation and can be used on future missions. CFESat, as a technology pathfinder, has effectively demonstrated the importance of high-performance reconfigurable computing.

## REFERENCES

- [1] D. Weigand and M. Harlacher, "A radiation-tolerant low-power transceiver design for reconfigurable applications," in *Proceedings of the Earth Science Technology Conference (ESTC)*. ITT Industries Advanced Engineering & Sciences Division, 2002, p. Paper A1P2, [http://estc.nasa.gov/conferences/estc-2002/Papers/A1P2\(Weigand\).pdf](http://estc.nasa.gov/conferences/estc-2002/Papers/A1P2(Weigand).pdf).
- [2] K. Morris, "FPGAs in space," Tech Focus Media, Tech. Rep., August 2004, FPGA and Structured ASIC Journal.
- [3] D. Ratter, "FPGAs on Mars," Xilinx, Tech. Rep., August 2004, xCell Journal #50.
- [4] A. S. Dawood, S. J. Visser, and J. A. Williams, "Reconfigurable FPGAs for Real Time Image Processing in Space," in *14th International Conference on Digital Signal Processing (DSP 2002)*, vol. 2, 2002, pp. 711–717.
- [5] P. Bergsman, "Xilinx FPGA blasted into orbit," Xilinx, Tech. Rep., Summer 2003, xCell Journal #46.
- [6] M. Caffrey, "A space-based reconfigurable radio," in *Proceedings of the International Conference on Engineering of Reconfigurable Systems and Algorithms (ERSA)*, T. P. Plaks and P. M. Athanas, Eds. CSREA Press, June 2002, pp. 49–53.
- [7] E. Fuller, M. Caffrey, P. Blain, C. Carmichael, N. Khalsa, and A. Salazar, "Radiation test results of the Virtex FPGA and ZBT SRAM for space based reconfigurable computing," in *MAPLD Proceedings*, September 1999.
- [8] "QPro Virtex 2.5V radiation hardened FPGAs," Xilinx Corporation, Tech. Rep., November 5 2001, dS028, v1.2.
- [9] C. Carmichael, M. Caffrey, and A. Salazar, "Correcting single-event upsets through Virtex partial configuration," Xilinx Corporation, Tech. Rep., June 1 2000, xAPP216 (v1.0).
- [10] P. Graham, M. Caffrey, D. E. Johnson, N. Rollins, and M. Wirthlin, "SEU mitigation for half-latches in Xilinx virtex FPGAs," *IEEE Transactions on Nuclear Science*, vol. 50, no. 6, pp. 2139–2146, December 2003.
- [11] C. Carmichael, "Triple module redundancy design techniques for Virtex FPGAs," Xilinx Corporation, Tech. Rep., November 1, 2001, xAPP197 (v1.0).
- [12] F. Lima, C. Carmichael, J. Fabula, R. Padovani, and R. Reis, "A fault injection analysis of Virtex FPGA TMR design methodology," in *Proceedings of the 6th European Conference on Radiation and its Effects on Components and Systems (RADECS 2001)*, 2001.
- [13] B. Pratt, M. Caffrey, P. Graham, E. Johnson, K. Morgan, and M. Wirthlin, "Improving FPGA design robustness with partial TMR," in *Proceedings of the MAPLD Conference*, September 2005.
- [14] K. Morgan and M. J. Wirthlin, "Predicting on-orbit SEU rates," Brigham Young University, Tech. Rep., July 2005, <https://dspace.byu.edu/handle/1877/64>.
- [15] J. Engel, M. Wirthlin, K. Morgan, and P. Graham, "Predicting on-orbit static single event upset rates in Xilinx Virtex FPGA," Brigham Young University, Tech. Rep., November 2006, <https://dspace.byu.edu/handle/1877/431>.
- [16] D. L. McMurtrey, "Using duplication with compare for on-line error detection in FPGA-based designs," Master's thesis, Brigham Young University, Provo, UT, 2006.
- [17] M. J. Wirthlin, D. E. Johnson, N. H. Rollins, M. P. Caffrey, and P. S. Graham, "The reliability of FPGA circuit designs in the presence of radiation induced configuration upsets," in *Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines (FCCM '03)*, K. L. Pocek and J. M. Arnold, Eds. IEEE Computer Society Press, April 2003, pp. 133–142.