Necessity of Fault Tolerance Techniques in Xilinx Kintex 7 FPGA Devices for Space Missions: A Case Study

Louis van Harten[†], Roel Jordans^{†‡}, and Hamid Pourshaghagh^{†‡}

[†] Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

[‡] Radboud RadioLab, Department of Astrophysics/IMAPP, Radboud University, Nijmegen, The Netherlands
l.d.v.harten@student.tue.nl, r.jordans@tue.nl, h.r.pourshaghagh@tue.nl

Abstract—In FPGA applications in space, implementations are generally protected using radiation-error mitigation techniques such as triple modular redundancy. For high-performance systems, such fault tolerance techniques can prove problematic due to large power overhead. This paper presents a case study on the Digital Receiver System (DRS) in the Netherlands-China Low-frequency Explorer (NCLE), which is implemented using a Xilinx Kintex 7 SRAM FPGA. Estimates for the critical cross-section of the system are presented, as well as estimated fault rates for a five-year mission to the second Earth-Moon Lagrange point. This includes simulations on the expected radiation environment, an analysis on the applicability of the used Xilinx Kintex 7 FPGA in these conditions and an analysis on the feasibility of implementing the DRS with minimal mitigation techniques for this mission. The steps performed during the analysis are described in detail, as to provide a guideline for replicating such an analysis for different space missions.

Index Terms—FPGA, fault tolerance, space missions, radiation error mitigation, Kintex 7

I. INTRODUCTION

The Netherlands-China Low-frequency Explorer (NCLE) is a low-frequency radio instrument payload aboard the Chinese Chang'e 4 relay satellite. This satellite is scheduled to be launched to the second Earth-Moon Lagrange point (L2) in 2018. Its goal is to collect astrophysical data at radio frequencies below 80MHz, in order to form a low frequency sky map and to attempt astrophysical study of the cosmic dawn, along with several other science objectives. The NCLE will be the first international payload on a Chinese space mission, and the first Dutch instrument to be sent to (and beyond) the Moon.

The NCLE Digital Receiver System (DRS), implemented using a Xilinx Kintex 7 SRAM FPGA, is tasked with processing and storing the data obtained from three monopole antennas. This processing entails large Fourier transforms, as well as filtering operations. There is an average power budget of approximately 3 Watts for the FPGA. The reason for this is not a shortage of power; the limiting factor is the amount of heat that can be dissipated, as the capsule in which the payload resides is not pressurized. Neither convective nor liquid cooling is available, putting a severe limit on the amount of power that can be dissipated without the system heating up to dangerous temperatures.

Because of the small power budget and high performance requirements, an efficient implementation of the system is

imperative, in order to allow continuous measurements to be performed. The system could operate on a sub-100% duty cycle to allow the components to cool down after a time of intensive processing, but this could interfere with a number of the science goals and should preferably be avoided.

The Xilinx Kintex 7 FPGA is known to be susceptible to radiation-induced upsets [1], which may introduce faults in the computed results. Traditionally, fault tolerance techniques such as triple modular redundancy (TMR) are used [2] to mitigate or eliminate these faults. However, these techniques incur significant overhead, dramatically increasing both the required FPGA area and the required power. For example, TMR increases both the required area and power by approximately a factor of three.

Overhead in this order of magnitude would likely undermine a portion of the desired science goals, as smaller-than-preferred Fourier transformations would have to be used, reducing the quality of the gathered science data.

This paper describes a case study on whether it would be a viable approach to implement the system with only limited application of fault tolerance techniques in the design. Section II shows a more elaborate overview of the mission and its goals, and Section III gives an overview of the logical lay-out and physical situation of the NCLE payload in the Chang'e 4 relay satellite. Section IV discusses simulation results of the radiation environment at the mission destination, in order to provide an estimate for the expected radiation sustained by the system over the mission duration.

After that, Section V gives an overview of the susceptibility of the system to various types of radiation errors and degenerative effects, as well as their impact on the system. Where applicable, vulnerable cross-sections are determined.

Section VI provides an attempt at quantifying two definitions of the critical cross-section of the FPGA implementation: faults that propagate to the observed data, as well as the subset of those faults which have a significant and possibly catastrophic effect on the science data, together with their respective estimated incident rates.

II. MISSION OVERVIEW

The NCLE is expected to operate in the second Earth-Moon Lagrange point for a mission duration of between three and five years. It is set to collect astrophysical science data in the

Fig. 1. The prototype of the NCLE Digital Receiver System, in development. The FPGA board is mounted on a backplane together with the ADC board, which will be connected to the antenna system.

range below 80 MHz, observing several phenomena in specific sub-sets of this band.

The different science cases require different types of data being produced. In this paper, the term “science mode” refers to an application running on the FPGA, collecting data for a specific science case. The various science modes are mostly run on a sequential schedule, although data collection for at least one of the burst phenomena (Jovian S-bursts) is planned to be implemented as a triggered mode.

For observations on static emission sources such as the galactic background or the cosmic dawn signal, there is no requirement on the time resolution. This means these science modes can average large amounts of measurements into single data points, improving the effective sensitivity.

Other science cases, such as measuring specific planetary emissions, do have set temporal resolution requirements. Typically this is in the range of one second (Earth Radio-Frequency Interference) to ten seconds (Auroral Kilometric Radiation), meaning significant amounts of averaging is still performed for each data point.

Exceptions are measurements on Jovian S-bursts, giant pulses and pulsar emissions, which have temporal resolution requirements in the range of milliseconds. Note that due to the limited downlink bandwidth to Earth, the observation time of these science modes is limited to an order of seconds. The significance of these requirements with respect to the radiation fault impact on the system is clarified in Section VII.

III. SYSTEM SET-UP

The prototype of the NCLE DRS is pictured in Fig. 1. In this picture, the longer board sticking out from the backplane is the studied FPGA board containing a Xilinx XC7K160T. This board contains the data flash memory as well. The shorter board contains three ADC channels, servicing the three antennas on the NCLE.

On the other side of the backplane, pins are broken out to allow the command and data handling system (CDHS) access to the configuration flash, making it possible to track housekeeping data and push updates to the FPGA design.

Fig. 2. Five-year ionizing dose in near Earth interplanetary orbit at 1.0 AU from the sun, starting from March 2018. Results obtained from SHIELDOSE simulations via Spenvis.

CDHS is tasked with communicating gathered science data back to Earth, and it also acts as a monitor for the DRS FPGA: it can restart and reconfigure the FPGA if this is deemed necessary.

The DRS board will operate in a vacuum environment which, as mentioned in the introduction, inhibits cooling. A passive heat-coupling system is currently being designed, which is expected to allow the FPGA to operate at safe temperatures with an average power usage of between 3 and 4 Watts.

At the time of writing, the shielding thickness is a to be decided quantity. If no further justification is available, the ESA standard recommends assuming 1 g/cm² of shielding [3], which is equivalent to approximately 3.7mm of aluminium. However, as there was uncertainty within the project on whether this amount of shielding would be present, a more conservative estimate of 2mm of aluminium shielding was used for the analysis.

IV. RADIATION ENVIRONMENT AT MISSION DESTINATION

Information about the radiation environment near the second Earth-Moon Lagrange point is needed to provide a relevant error model. Ideally, both the radiation dose and the spectrum shape of particle energies should be known, along with the type of expected particles. The latter two determine the effective LET (linear energy transfer) of the particles to the device. The effective energy transfer of a particle impact is (approximately linearly) related to upset rate [4, Fig. 1].

By combining the spectrum of expected LETs with the effective cross-section of the device (faults per amount of flux), the expected number of faults can be obtained. A more thorough explanation of this is given in Section V.

A. Total dose estimates

Spenvis (the SPace ENVironment Information System by ESA) [5] was used to find total dose estimates for a five-year flight. As Spenvis does not contain a model for the L2

Fig. 3. Simulated five-year fluence plotted as LET spectrum, 2mm of Al shielding.

orbit, it was approximated by an arbitrary point in space at a distance of 1 AU from the sun. The simulation ignores the intermittent shielding effect of the Moon and the Earth, as well as the effects of intermittently passing through the Earth's magnetotail, but it should still provide a reasonable estimate.

From Spensiv, the SHIELDOSE-2Q [6] simulation was run for a spherical aluminium shield around a silicon target. The total dose results from this simulation can be found in Fig. 2. These results show an estimated total dose of approximately 8 krad(Si) for 2mm of shielding.

B. Particle fluence estimates

Also using Spensiv, a prediction for the long-term LET spectrum was obtained, with total fluence as a function of particle energy transfer. These results can be found in Fig. 3. Note that the unit for integral fluence in this figure is $m^{-2}sr^{-1}$: particles per area, per steradian.

The simulation calculates flux through a spherical shield, hence the steradian in the unit. To convert these results to particles hitting a flat surface of a certain area, it is necessary to calculate the projection of the flat chip area to a sphere. Integrating over both the azimuth and elevation, this boils down to

$$\int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} A |\cos(\theta) \sin(\theta)| d\phi d\theta = 2\pi A, \quad (1)$$

i.e., the results should be multiplied by 2π in order to get the amount of particles per square meter per second. This results in a total integral fluence of $1.056 \cdot 10^{14} m^{-2}$ over the mission lifetime, equivalent to an average flux of $66.9 cm^{-2}s^{-1}$.

C. Worst case flux

Due to the nature of varying solar conditions, average flux can differ wildly from the worst case scenario. In additional simulations, results for various worst case scenarios were obtained. These results are shown in table I.

TABLE I
EXPECTED PARTICLE FLUX

Situation	Flux ($cm^{-2}s^{-1}$)
average	$6.69 \cdot 10^1$
worst week	$6.37 \cdot 10^3$
worst day	$2.93 \cdot 10^4$
worst 5 minutes	$1.09 \cdot 10^5$

An important observation is that during the worst five minutes, the expected flux is approximately 1600 times higher than in the average case. In addition, this implies that in the median situation, the flux is likely significantly lower than in the average case.

V. OVERVIEW OF RADIATION ERRORS

Radiation errors in FPGAs come in several categories: Configurable Logic Block (CLB) errors, BRAM upsets, configuration (SRAM) upsets, destructive latch-ups and total ionization dose failure. The following sections will go over each of these, listing the effects, expected incident rate, and (if needed) propose defensive strategies.

For single event effects, the effective cross-section of the system is discussed. This is defined as the amount of events per unit of fluence. It combines the probability that a particle impact will cause an upset with the amount of particles passing through the chip. The quantity of fluence is the inverse of area, which makes area per bit (or area per system) the quantity for the cross-section. A meaningful interpretation of this concept is the critical area which a particle has to strike in order to cause an event.

A. Destructive latch-ups

When a heavy ion strikes a silicon CMOS microcircuit, there is a chance for a latch-up to occur: a self-sustaining parasitic short-circuit, which draws high current and may break the circuit due to high temperatures. Surprisingly, the Kintex 7 does not seem to suffer from the same destructive latch-ups found in many other FPGAs [4] [7].

Heavy ion testing in the TAMU K500 Cyclotron facility [4] has shown that latch-ups in the Kintex 7 only draw approximately 125 mA from the VCCAUX line (1.8V, meant for auxiliary circuits such as clock managers and dedicated configuration pins), which is not enough to cause any lasting damage to the circuit. The exact cause for the draw is not clear.

In the initial tests, the FPGA is operated above its normal operating voltage in order to trigger and study the latch-up behavior. Additional testing in the same facility at nominal voltages has shown that the event only occurs for very high energy particles; the lowest effective LET at which this phenomenon was observed is $1.5 \cdot 10^4 MeV cm^2/g$, at which the estimated cross-section was determined as approximately $5 \cdot 10^{-7} cm^2$.

These results were confirmed by heavy-ion tests at the Cyclotron Resource Center in Louvain [7], where similar

behaviour was observed with a threshold SEL of $1.56 \cdot 10^4 \text{ MeVcm}^2/\text{g}$.

The expected five-year integral fluence in L2 for events of at least this LET is only approximately $1.0 \cdot 10^5 \text{ cm}^{-2}$, meaning the chance that a single event of this type occurs during the five-year mission is approximately 5%. Results for additional test are available [8], in which a Kintex 7 device is irradiated with $1.9 \cdot 10^{11} \text{ cm}^{-2}$ fluence of 105 MeV protons. In this test, not a single latch-up was detected.

An additional note from the Cyclotron Louvain tests is that the latch-ups in the Kintex-7 do not seem to cause any loss of part functionality. Power cycling the device removed all symptoms.

These results imply that for Kintex 7 applications in space, latch-ups do not pose a threat in the form of lasting damage to the FPGA. It can be concluded that because of the extremely low incident rate and low impact on the system, the effects of latch-ups in the Kintex 7 can effectively be ignored.

B. Total Ionizing Dose effects

The total ionizing dose (TID) effects mostly consist of the transistors in the FPGA slowly breaking down by particles hitting the doped silicon and slowly weakening the doping. This results in the transistors slowing down (and eventually breaking), resulting in a longer delay in the critical path.

Few elaborate tests have been published researching the TID effects in the Kintex 7 specifically. However, these test have been performed on other FPGA devices, such as on the Lattice ECP3 [9, Fig. 4]. It is likely reasonable to assume the results for the Kintex 7 would follow a similar pattern: negligible slowdown up to approximately halfway the device failure point, after which the slowdown gradually increases up to the point of total device failure.

The Lattice ECP3, of which test results were mentioned above, is an FPGA which uses 65nm technology [10]. The Kintex 7 used in the NCLE project uses 28nm technology. Perhaps unintuitively so, smaller nanometer technologies are generally less prone to TID transistor slowdown effects due to their smaller oxide thickness [11]. This means that the Kintex 7 is unlikely to perform worse than the LFE3-35EA under similar large ionizing dose conditions.

There is some data available on Kintex 7 TID failure points and all data points seem to confirm the above mentioned assumption. One paper presented a Kintex 7 FPGA being irradiated with 105 MeV protons for a total dose of 17.0 krad [8] (170 Gy). No functional problems were observed.

In another experiment, two Kintex 7 devices were irradiated in an attempt at finding the device failure point [12]. The first broke down after receiving 340 krad (3400 Gy) and one other still functioned after receiving 446 krad (4460 Gy), after which the test was aborted. Both tests were performed using high energy (180 MeV) protons.

As shown in Section IV-A, the expected total mission dose is less than 10 krad (100 Gy). This is more than 30 times lower than the lowest observed failure point of a Kintex 7.

Considering these results, it is reasonable to assume TID effects in the FPGA can be ignored in the NCLE mission.

C. Configuration upsets

Single event upsets (that is: an incidental flip in the state of an element) can be especially troublesome in an FPGA, as upsets can affect the state of configuration bits. This means the logic as composed by the digital gates functionally changes. Depending on the bit that was struck, this can lead to faulty data output or bring the system into an erroneous state.

The main technique to combat configuration upsets is “scrubbing”. This means to constantly compare the active configuration bits with a protected (duplicated) reference and to reconfigure blocks where necessary, or by achieving similar functionality using error correcting code (ECC) bits.

The Kintex 7 used in the NCLE has an on-board configuration scrubber, which is able to correct single-bit errors in one word, and detect up to two-bit errors. Additionally, a more advanced single-error mitigation IP core is available, which is able to correct up to two-bit adjacent errors in one word, and detect any larger odd-count bit errors in one word, as well as some larger even-bit errors [13]. In further analysis, it is assumed this core is used.

Advanced two-bit adjacent error correction relies on storing additional Cyclic Redundancy Check (CRC) bits in the device BRAM. For designs that, unlike the NCLE DRS, use (almost) all BRAM tiles in the FPGA, this option is not available. In such cases, all multi-bit upsets should be considered uncorrectable.

It is also important to note that the scrubber does not fix upset bits immediately. The SEM core has a scrubbing latency of 12.9 ms for the specific FPGA in the NCLE [13], although this can be increased to save power. There is also a correction latency of 0.6 ms. This means that whenever a correctable upset occurs, the system is stuck in an imperfect state for up to approximately fifteen milliseconds.

D. Expected rate of configuration upsets

Providing a reliable estimate on the amount of expected errors is not entirely trivial. While there are many papers available in which relevant radiation test results are presented, the type and amount of radiation is different from the expected radiation at the L2 point, to a partially unknown degree.

Test results from literature have shown that for a Kintex 7 device in a 105 MeV proton beam, the effective configuration upset cross-section is $5.21 \cdot 10^{-15} \text{ cm}^2/\text{bit}$ [8]. This proton beam translates to a LET spectrum where approximately 40% of the events have a LET of at least $1 \text{ MeVcm}^2/\text{mg}$ and 10% has a LET of at least $8 \text{ MeVcm}^2/\text{mg}$ (similar to the one shown in [14, Fig. 2]). Comparing this to the expected mission spectrum shown in Fig. 3, it is clear that the proton tests are not entirely representative of the expected radiation environment. In the expected mission spectrum, only one in ten thousand events have a LET of at least $1 \text{ MeVcm}^2/\text{mg}$ and less than one in a million events have a LET of at least $8 \text{ MeVcm}^2/\text{mg}$.

TABLE II
EXPECTED CONFIGURATION UPSETS AND UNCORRECTABLE MBUS

Situation	Configuration upsets	Uncorrectable upsets
average	4.4 / week	0.11 / year
worst week	2.5 / hour	0.21 / week
worst day	11.5 / hour	0.14 / day
worst 5 minutes	3.6 / 5min	0.0018 / 5min

Literature has shown that for high LET ($5 - 20 \text{ MeVcm}^2/\text{mg}$), the configuration cross-section is reasonably linear with the event energy [4]. However, extrapolating these results to low energy events might not be valid; there is insufficient evidence supporting this relation behaves the same way at low LET.

Applying the 105 MeV proton cross-section as the cross-section average would mean ignoring the discrepancy between the beam spectrum and the expected mission spectrum. This is equivalent to pretending the system impact of particles hitting the device is larger on average than can reasonably be assumed. This will result in inflated error estimates; while unfortunate, it is better to stay on the side of “unrealistic worst-case” than to end up with an estimate that is significantly too low. As such, in this paper, the cross-sections measured using high-energy protons will be considered as valid.

In the concept design of the NCLE digital receiver system FPGA implementation, approximately 20Mbit of configuration SRAM is used. This translates to an effective configuration cross-section of $1.1 \cdot 10^{-7} \text{ cm}^2$. Combining this result with the flux from Section IV results in the average upset-rates found in Table II.

E. Multi-bit upsets and correctability

The earlier mentioned configuration scrubber in the Kintex 7 can correct up to two-bit errors, and detect all odd-number bit errors in a single word. While the bits of different words are physically interleaved to combat multi-bit upsets in a single word, these do happen occasionally. In literature, an upset in multiple bits across different words is sometimes called an MCU (multi-cell upset), whereas the term MBU (multi-bit upset) is reserved for upsets which flip multiple bits in a single word.

Fairly extensive testing has been done to characterize the multi-bit upset behavior in the Kintex 7 [15]. The tests with the lowest energy ions used nitrogen and oxygen ions with an energy of 200 MeV, impacting with an average LET of $1.16 \text{ MeVcm}^2/\text{mg}$ and $1.54 \text{ MeVcm}^2/\text{mg}$ respectively. Out of the presented tests, these should be most representative of the expected mission environment. All but one in ten thousand particle strikes during the mission is expected to have a LET of less than the average of the ion strikes in these tests, meaning the test results convey an absolute worst case.

In the nitrogen and oxygen ion tests, average incident rates of 0.4% and 0.05% were found for 2-bit adjacent MBUs and 3-bit MBUs respectively, as a fraction of all configuration upsets. The amount of ≥ 4 -bit MBUs and non-adjacent 2-bit

TABLE III
ESTIMATED RESOURCE UTILIZATION

Resource	Utilization (% of available)
CLB slices	11250 (44%)
BRAM tiles (36k)	210 (65%)
DSP blocks	400 (67%)

TABLE IV
CROSS-SECTIONS OF VARIOUS SLICES/BLOCKS

Resource	Cross-section per unit	In concept design
CLB slices	$4.22 \cdot 10^{-14} \text{ cm}^2/\text{slice}$	$0.47 \cdot 10^{-9} \text{ cm}^2$
BRAM tiles	$4.81 \cdot 10^{-11} \text{ cm}^2/\text{tile}$	$10.1 \cdot 10^{-9} \text{ cm}^2$
DSP blocks	$9.88 \cdot 10^{-13} \text{ cm}^2/\text{block}$	$0.40 \cdot 10^{-9} \text{ cm}^2$

MBUs was negligible. The average incident rate of MCUs was found to be 1.7%. These rates result in a total uncorrectable configuration cross-section of $5.5 \cdot 10^{-11} \text{ cm}^2$ and a negligible undetectable cross-section. The expected rate of uncorrectable upsets during the mission can be found in Table II.

F. Data upsets

Apart from configuration upsets, there are several other possible single event faults: errors in BRAM, propagated transients in multipliers and upsets in flip-flops. These errors are all related to user-data and it is not possible to detect them, unless logic is specially generated for that purpose.

As with the used configuration memory, the exact utilization of the BRAM-blocks, DSP-blocks and flip-flops in the final design is not yet fully known. Estimates were made from the concept design, which can be found in Table III. Non-listed sites (such as distributed RAM blocks and Muxes) have a sufficiently low expected usage that their cross-section contribution was deemed insignificant. These estimates are effectively the resource utilization of the concept design, rounded up slightly.

Cross-sections in a proton beam for BRAM upsets, DSP blocks and CLB slices (which contain LUTs and Flip Flops) are available from literature, obtained in similar conditions as the configuration memory cross-section results used in Section V-D.

The cross-sections for BRAM upsets, DSP blocks, and logic slices for a Kintex 7 7K325T were determined in literature as $2.17 \cdot 10^{-9} \text{ cm}^2/\text{device}$ (logic slices), $0.83 \cdot 10^{-9} \text{ cm}^2/\text{device}$ (DSP blocks) and $21.4 \cdot 10^{-9} \text{ cm}^2/\text{device}$ (BRAM) for full utilization of those respective parts [8]. The device under test in this paper was a larger FPGA than the 7K160T in the NCLE system, so some conversion is necessary.

Results from converting the cross-sections to per unit and full design cross-sections are shown in Table IV. Note that while it provides a convenient intermediate step in calculating the design cross-section from device cross-sections, converting the per device cross-sections to per logical slice cross-sections is somewhat fictitious, as many of these components are not truly separate blocks on the FPGA.

Adding all of these cross-sections together results in a total user data cross-section of $1.09 \cdot 10^{-8} \text{ cm}^2$, which happens to

TABLE V
EXPECTED RATE OF DATA UPSETS

Situation	Data upsets per day
average	$6.3 \cdot 10^{-2}$ ($\cong 1.9$ per month)
worst week	6.1
worst day	27.8 ($\cong 1.2$ per hour)
worst 5 minutes	103.6 ($\cong 4.3$ per hour)

be almost exactly 10% of the effective configuration cross-section. The resulting expected upset rate in the user data for various situations is given in Table V.

VI. FAULT IMPACT

An important note to keep in mind is that various upsets can have significantly different effects. For example, a configuration upset may add a data wire to an unconnected block, or a BRAM bit may upset right before data has been written to it. Neither upset has any effect on the system function. Such errors can be considered non-critical.

A. Critical cross-section

The effective cross-section for upsets which impact the system results in any way is considered as the critical cross-section. This is composed of the critical cross-sections for both configuration and data upsets.

The fraction of essential configuration bits (i.e., the bits which influence the system functionality when upset) as percentage of the total amount of configuration bits is called device vulnerability factor (DVF). This can be determined from a design using Xilinx tools. The DVF can be assumed to never exceed 10% [1]. As there is not yet a functional design available at the time of writing, the worst case is assumed.

This means that for every ten configuration bits flipped, only one will have any functional effect. As shown in Section V-E, the incident rate of multi-bit upsets is only 1.7%, meaning the percentage of upsets causing essential configuration bits to flip is approximately equal to the DVF. The resulting critical cross-section for configuration upsets is $1.1 \cdot 10^{-8} \text{cm}^2$.

The data upset cross-section is dominated by BRAM upsets, but the critical cross-section for these upsets could not be determined: the fraction of upsets that are critical heavily depends on the application running on the FPGA. No relevant testing has been performed to determine this fraction for the NCLE application, as it is not yet complete or functional at the time of writing. For this reason, all data upsets are assumed to be critical. The resulting critical cross-section for data upsets is $1.09 \cdot 10^{-8} \text{cm}^2$, identical to the total data upset cross-section given in Section V-F. This results in a total critical cross-section of $2.2 \cdot 10^{-8} \text{cm}^2$.

B. Semi-critical versus severely critical cross-section

While non-critical upsets are completely uniform, there is a scale of severity in critical upsets. Most critical configuration upsets that are corrected by the scrubber after several milliseconds might influence several data points, but these errors will likely disappear in averaging the large amount of data points.

However, uncorrectable configuration upsets and correctable upsets in certain parts of the control logic will cause the system to enter and stay in an incorrect state.

The semi-critical cross-section can be defined as the cross-section for upsets that result in a minor system error. What classifies as “minor” is unique for each system. For the NCLE DRS, any error which does not interrupt or impede the science modes and which does not significantly influence the gathered results can be considered a minor error. This translates to all correctable critical configuration upsets in most of the system, with the exception of small control parts.

Conversely, a “severe” error is any error which does significantly influence gathered science data. For uncorrectable upsets, this means faulty data points being accumulated until the device is reconfigured.

For critical upsets in control logic, it is likely that the science mode run produces some sort of invalid data. It is also possible the run finishes prematurely, never finishes at all, or even overwrites data gathered in previous science modes stored in flash because of a fault in an address calculation. While all of these errors are severe, the latter two are potentially catastrophic. It is unreasonable to assume these catastrophic errors form a significant portion of the severe errors, but the possibility of their occurrence should be taken into account.

Putting a number on the amount of configuration bits that would result in a severe error when upset is difficult, as it is extremely application specific. A reliable way would be extensive testing using fault injection, which is not possible without a semi-final design. For the NCLE DRS design, correctable critical upsets inside of the FFT and filter calculations would not result in a severe system fault, and hardware for these calculations make up the grand majority of the area on the FPGA. Because of this, the assumption is made that no more than 25% of the critical upsets result in a severe system error.

For data upsets, a similar problem exists: far from all critical data upsets result in a severe system error, but the fraction is hard to estimate without extensive testing. A small amount of data consists of matters like filter coefficients and loop control variables, which could result in severe propagated errors when upset.

Most of the BRAM tiles (which have the highest contribution to the data upset cross-section by far) are used for storing partially processed sample data. There is not enough insight about the final design available at the time of writing to give a precise estimate on what the severely critical fraction is, but preliminary investigations show that 10% should be a reasonable worst-case figure.

C. Calculation of semi- and severely critical cross-sections

The semi-critical cross-section, if defined as a superset of the severely critical cross-section, encompasses all critical configuration upsets ($1.1 \cdot 10^{-8} \text{cm}^2$) in addition to all data upsets (also $1.1 \cdot 10^{-8} \text{cm}^2$). The resulting approximate cross-section is $2.2 \cdot 10^{-8} \text{cm}^2$.

The severely critical cross-section encompasses all critical vital configuration upsets and critical uncorrectable upsets

TABLE VI
EXPECTED CRITICAL UPSETS (CONFIGURATION & DATA)

Situation	Critical upsets	Severely critical	Undetectable
average	3.8/month	8.1/year	1.9/month
worst week	6.0/day	2.1/day	3.0/day
worst day	2.3/hour	9.9/day	1.2/hour
worst 5 minutes	0.72/5 min	0.13/5 min	0.36/5 min

($2.8 \cdot 10^{-9} \text{cm}^2$), as well as all critical vital data upsets ($1.1 \cdot 10^{-9} \text{cm}^2$), resulting in a total cross-section of approximately $3.9 \cdot 10^{-9} \text{cm}^2$.

The resulting incident rates can be found in Table VI. While a significant part of these results is based on *very* rough estimates, they can reasonably be considered worst-case numbers.

D. Critical undetectable cross-section

The critical cross-section can be further split up into two parts: the detectable and the undetectable critical cross-section. This distinction is useful, as even severe functional problems can largely be mitigated as long as they are detectable, simply by marking the data produced in the rest of the science mode run as (possibly) invalid. However, for undetectable upsets, this is not an option, meaning they can spoil data without the system being able to mark spoiled data as such.

The undetectable configuration upsets are all MBUs that flip more an even number above three bits in one word. As mentioned in Section V-E, the cross-section for this event is negligible.

The data upsets are all assumed undetectable. This means the undetectable cross-section is the same as the critical data upsets for both the semi-critical undetectable cross-section ($1.1 \cdot 10^{-8} \text{cm}^2$) and the severely critical undetectable cross-section ($1.1 \cdot 10^{-9} \text{cm}^2$).

VII. NECESSITY OF ADDITIONAL FAULT TOLERANCE TECHNIQUES IN THE NCLE DRS

The presented expectations for upset rates and fault rates can be used to answer the question whether any additional fault tolerance techniques in the FPGA are strictly necessary for successful completion of the science objectives of the NCLE mission.

A. Practical problems and solutions

Considering Table VI, the answer to the above question is not a simple yes or no. While the data conveys a near-worst-case scenario, the estimated number of severe critical upsets is significant: approximately 40 over the total five-year mission duration. This is the amount of times an entire science mode run is expected to be corrupted or interrupted. As long as none of these events are catastrophic errors which lock up the DRS or destroy significant portions of the measurement data, this should not be enough to compromise any of the science goals.

While known to be small, the exact fraction of severe errors causing such catastrophic faults is unknown, so implementing

the system without any sort of measures to protect against these events would pose an unacceptable risk.

Preferably, measures should also be taken to mitigate the impact of semi-critical events. Upsets that flip the most significant bits of only few samples could have a significant impact on results of long accumulation. With in the worst case 230 of these events, this could influence the science data to an unacceptable degree.

B. Possibilities for semi-critical upset impact mitigation

While, as shown in Section V-E, virtually all configuration upsets are detectable and almost all are correctable, properly recovering from a configuration upset is not trivial. Any data that passed through a struck gate between discovery and the time of the last scrubbing pass can no longer be trusted.

The SEM core has a scrubbing latency of 12.9 ms for the FPGA in the NCLE [13]. To recover, the system would have to either be able to fully roll back to a state from 14ms earlier, or employ (triple) redundant calculations to fall back on in case of a fault. Considering the tight area and power budgets, neither of these is a viable option.

Most science modes accumulate large amounts of samples into single data points. These could store intermediate accumulations (over a time period that is significantly smaller than the expected upset rate) into separate memory locations, as opposed to keeping a single updating data block in flash. The storage budget is fairly large: storing intermediate accumulations for every separate minute would not be prohibitive. By generating metadata for each intermediate accumulation on whether upsets were detected during its collection, faulty intermediates can be filtered out.

It is possible to do this on a more local scale. Due to their accumulating nature, regularly dropping a few milliseconds of data is not a problem in most science modes. A possibility would be to accumulate results from several milliseconds in an intermediate accumulator, only adding them to the final result if no upsets are detected during that time. The disadvantage is that the data is actually thrown away, instead of marked as possibly faulty.

The SEM core provides functionality for classifying upsets by checking them to a list of essential bits stored in external storage. This requires the SEM core to interface with the flash memory, and implementing it is a significant amount of work. Due to the short lead-time of the project and low priority of this feature, it will possibly not be implemented. This means that when using the above approach, approximately ten times too many data windows would be dropped. As shown in Table II, non-critical configuration upsets have an expected occurrence rate below once a minute during the worst five minutes of the mission, meaning dropping several milliseconds of data ten times too often is not a large problem.

For science modes where a high temporal resolution (order of milliseconds) is necessary, the previous approach is not an option. These modes can not run for much longer than a few seconds, as the bandwidth to get that much data back to Earth is not available. A realistic alternative is to simply assume

the system does not experience any upsets in relevant bits during that time. Even during the worst possible hour of the mission, the average time to a critical upset is approximately 25 minutes; significantly longer than these measurements take.

As with storing intermediate accumulations, each run can be marked with metadata on whether upsets were detected during the collection of the data. Even data collected while the system was influenced by an upset is likely valuable to some extent (especially since most upsets are non-critical) and the decision to use it can be made on Earth.

C. Severely critical upset mitigation

Due to the extremely low incident rate, an acceptable response to severely critical upsets in the system would be to drop all tasks and request a power cycle from the CDHS. However, classifying the severity of detected upsets possibly poses a problem and not all severely critical upsets are detectable in the first place.

It is possible to decrease the incident rate of severe incidents by using traditional fault-tolerance techniques such as TMR. With the incident rates of critical faults, getting simultaneous faults in two voters can be considered statistically impossible, as long as they are not placed adjacently.

The small area and power budgets are less of a problem in this case, as the severely critical sections use only a small fraction of both of these budgets. Implementing TMR in the severely critical sections would eradicate the grand majority of the severely critical faults from the system, both those caused by configuration upsets as well as those caused by data upsets.

D. Catastrophic event mitigation

As catastrophic events are mostly comprised of a small subset of severely critical upsets, selectively applied TMR would drive down incident rate significantly. While the catastrophic cross-section should be close to negligible, some possibility for catastrophic errors still remains, mostly in the Kintex 7 system bits that cause full functional interruption when upset.

The cross-section for functional interruption events in the Kintex 7 is significantly smaller than the cross-section for latch-ups [7], which was considered negligible in section V-A. However, due to the catastrophic nature of these events, some sort of mitigation is in order. An option is establishing a “heartbeat” line to the CDHS, which can reconfigure and power cycle the FPGA, to automatically trigger a reset whenever this heartbeat line goes flat.

To offer further protection, this heartbeat could carry additional information about the DRS system state: if a science mode run takes longer than it should, CDHS could respond similar to the heartbeat going flat.

VIII. CONCLUSION

It has been shown reasonable to expect a significant amount of upsets in the NCLE DRS during the mission, but few are expected to be critical, and even fewer are expected to possibly compromise the mission goals. When using the built-in scrubber of the Kintex 7 as the only radiation error mitigation,

there would be a small but non-zero risk of catastrophic events in the FPGA endangering the acquisition of valid science data for the NCLE mission.

The most troubling risks were addressed with simple, realistic, low-overhead mitigation options. Assuming the proposed mitigation options (or similar ones) are implemented, radiation related errors in the DRS FPGA are extremely unlikely to compromise the science goals of the NCLE mission.

A simple approach has been presented to analyse the expected radiation-related system faults in Kintex 7 FPGA devices when used in space missions. Furthermore, it has been shown that such an analysis can prove useful in saving power on computationally expensive fault tolerance techniques in situations where a small power budget is available.

IX. ACKNOWLEDGEMENTS

This research was supported by the Netherlands Space Office (NSO). The authors would like to thank their collaborators in the Netherlands-China Low-frequency Explorer (NCLE) team, as well as the other members of the Radboud Radio Lab, for their helpful input at various stages of the project.

REFERENCES

- [1] Xilinx, “Device Reliability Report, First Half 2016 (UG116),” Dec. 2016.
- [2] ECSS, “Space product assurance: Techniques for radiation effects mitigation in ASICs and FPGAs handbook,” Sept. 2016.
- [3] ECSS, “Space environment: engineering standard,” Nov. 2008.
- [4] D. Lee, M. Wirthlin, G. Swift, and A. Le, “Single-Event Characterization of the 28 nm Xilinx Kintex-7 Field-Programmable Gate Array under Heavy Ion Irradiation,” July 2014.
- [5] D. Heynderickx, B. Quaghebeur, and H. D. R. Evans, “The ESA Space Environment Information System (SPENVIS),” *COSPAR Scientific Assembly*, vol. 34, Jan. 2002.
- [6] S. Seltzer, “Updated calculations for routine space-shielding radiation dose estimates: SHIELDOSE2,” *NSTIR*, Jan. 1994.
- [7] V.-M. Placintă, “Kintex-7 Irradiation, test bench and results,” Sept. 2016. TWEPP 2016.
- [8] D. Hiemstra and V. Kirischian, “Single Event Upset Characterization of the Kintex-7 Field Programmable Gate Array Using Proton Irradiation,” *Radiation Effects Data Workshop (REDW)*, June 2014.
- [9] J. M. Armani, J. L. Leray, and V. Iluta, “TID Response of Various Field Programmable Gate Arrays and Memory Devices,” *IEEE*, Oct. 2015.
- [10] Lattice, “LatticeECP3 Family Data Sheet,” Mar. 2010.
- [11] N. S. Saks, M. G. Ancona, and J. Modolo, “Generation of Interface States by Ionizing Radiation in very thin MOS Oxides,” *IEEE Transactions on Nuclear Science*, vol. NS-33, pp. 1185–1190, Dec. 1986.
- [12] H. Takai, “Soft Error Rate Estimations of the Kintex-7 FPGA within the ATLAS Liquid Argon (LAr) Calorimeter,” Sept. 2012. TWEPP 2013.
- [13] Xilinx, “Soft Error Mitigation Controller v4.1 LogiCORE IP Product Guide,” Apr. 2017.
- [14] P. O’Neil, G. D. Badhwar, and W. Culpepper, “Risk Assessment for Heavy Ions of Parts Tested with Protons,” *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, vol. 44, pp. 2311–2314, Dec. 1997.
- [15] M. Wirthlin, D. Lee, G. Swift, and H. Quinn, “A Method and Case Study on Identifying Physically Adjacent Multiple-Cell Upsets Using 28-nm, Interleaved and SECDED-Protected Arrays,” *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, vol. 61, pp. 3080–3087, Dec. 2014.