

Radiation Tolerant Intelligent Memory Stack (RTIMS)

Authors: Tak-kwong Ng, Jeffrey A. Herath
Affiliation: NASA Langley Research Center
Email: t.ng@nasa.gov, jeffrey.a.herath@nasa.gov

Abstract

The Radiation Tolerant Intelligent Memory Stack (RTIMS), suitable for both geostationary and low earth orbit missions, has been developed. The memory module is fully functional and undergoing environmental and radiation characterization. A self-contained "flight-like" module is expected to be completed in 2006.

RTIMS provides reconfigurable circuitry and 2 gigabits of error corrected or 1 gigabit of triple redundant digital memory in a small package. RTIMS utilizes circuit stacking of heterogeneous components and radiation shielding technologies. A reprogrammable field programmable gate array (FPGA), six synchronous dynamic random access memories, linear regulator, and the radiation mitigation circuitries are stacked into a module of 42.7mm x 42.7mm x 13.00mm.

Triple module redundancy, current limiting, configuration scrubbing, and single event function interrupt detection are employed to mitigate radiation effects. The mitigation techniques significantly simplify system design. RTIMS is well suited for deployment in real-time data processing, reconfigurable computing, and memory intensive applications.

1. Introduction

NASA has identified many systems that will require on-board satellite data processing for future missions. Continuing themes for addressing these requirements include the need for ever-increasing resolution, improved data quality, and additional capacity for raw and/or processed

data. The requirement to efficiently handle large data sets necessitates the use of larger on-board memories. This paper discusses the development of a radiation tolerant memory, suitable for both geo-stationary (GEO) and low earth orbit (LEO) missions.

RTIMS has been designed, built, tested and integrated onto a 3U Compact PCI printed circuit board (PCB). RTIMS capitalizes on previous technology investments at NASA Langley Research Center (LaRC), NASA Goddard Space Flight Center (GSFC), ASRC Aerospace, Irvine Sensors and 3D Plus USA. RTIMS incorporates the circuit stacking technology of 3D-Plus USA, as well as package-level radiation shielding and novel radiation mitigation techniques developed at LaRC. By using reprogrammable FPGA technology for the memory controller, RTIMS can also be a key element in adaptive/reconfigurable computing applications.

2. Need & Benefits

The need for compact, high performance, radiation tolerant memories has been identified as critical to enabling our nation's future space missions. Improved memory technologies offer reductions in size, cost, mass, power, risk and system complexity, and can be applied to advanced data processing, high bandwidth communication, and meeting the data volume needs of the next-generation high-resolution sensors.

Advances made by the RTIMS technology provide the following key benefits:

1. Significant reductions in the size and mass of mission memory arrays.

2. A radiation tolerant memory suitable for both GEO and LEO space missions by using new package level radiation shielding technology and Triple Modular Redundancy (TMR) FPGA techniques.

3. Simplified interface to a large SDRAM memory array with built-in logic for timing reads, writes and refresh cycles.

4. Novel self scrubbing and Single Event Functional Interrupt (SEFI) detection that allow a relatively "soft" FPGA to become radiation tolerant without external scrubbing and monitoring hardware.

5. Efficient design, test, and validation by incorporating the radiation mitigation technology at the component level rather than requiring the designer to implement a design unique solution at the system level.

6. Increased system reliability by distributing the radiation mitigation structure to each component instead of a single point failure at the system level.

7. Added mission flexibility by operating the memory array in a TMR architecture with 1Gb of storage or in an EDAC (Error Detection And Correction) mode where part of the memory is used to detect and correct errors with 2Gb of storage (corrects single bit errors and detects double bit errors). This allows RTIMS to be used effectively on many types of missions because it can be configured for the "harshness" of the expected environment.

8. Tightly integrated solution that includes a Field Programmable Gate Array (FPGA), six synchronous dynamic random access memories (SDRAMs), linear regulator, and the radiation mitigation circuitry into a single stack or component.

9. In-flight reconfigurability using Static Random Access Memory (SRAM) based FPGA technology allows the design to overcome both hardware and software errors that may be detected after launch during mission operations. This reduces overall

mission risk which is increasingly important as flight system development times and budgets decrease. It also allows RTIMS to adapt to changing mission conditions.

10. Additional logic and processing resources. By reprogramming the FPGA a designer can take advantage of the "extra" elements in the FPGA to implement other portions of their design.

These features provide considerable engineering savings as well as advanced functionality. RTIMS is enabling, because it allows standardized hardware to be used for completely different missions achieving lower costs since the same hardware design can be reused.

3. RTIMS Module Construction

The RTIMS module is a major step forward in the state of the art for space based memory arrays. A block diagram of RTIMS is shown in Figure 1. The new stacking technology used for RTIMS allows a heterogeneous stack of electronic parts to be built into a single component. The electronic parts in the stack can be bare electronic die, packaged parts (plastic or ceramic), or chip passives (capacitors and resistors). This new stacking technology typically provides an 80% reduction in required volume or footprint for a given application.

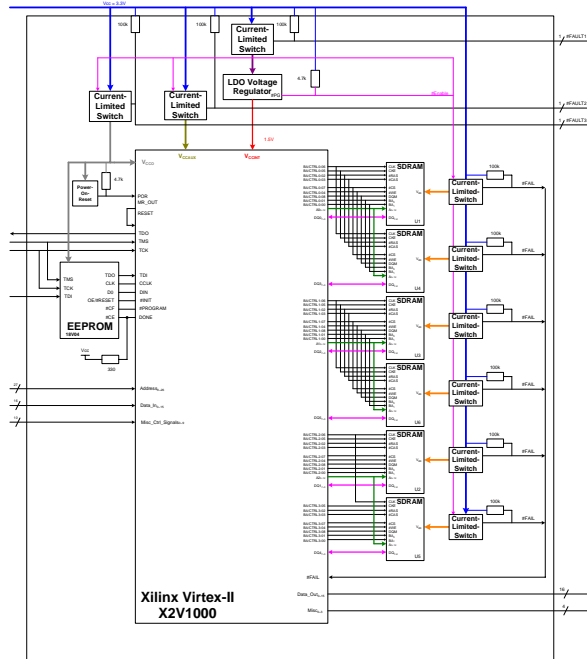


Figure 1. RTIMS Block Diagram

NASA GSFC, in conjunction with Centre National d'Études Spatiales (CNES) and the European Space Agency (ESA), completed a study on this stacking technology from 3D Plus. The study determined that the stacking technology is rugged enough and suitable for space applications [1].

The mechanical development of a 3D module is very similar to that used for a standard Multi Chip Module (MCM). The main difference from an MCM is the splitting of the design function onto several layers. These layers are essentially individual circuit boards allowing the stacking of electronic components placed on them. Once each layer (circuit board) of the module is designed, the layers can be fabricated, populated with components, tested, and burned in, all before final module assembly. This reduces the impacts of manufacturing errors and faulty components. The layers are then stacked together, aligned, and vertically spaced. A very pure epoxy resin from Dexter (HYSOL FP4450) is used to fill the spaces around the entire module and between each layer. After resin polymerization, the module is cut from the mold, exposing the external connections (flying-

leads) that will be used to vertically interconnect the layers. Nickel is then chemically and electro-chemically deposited onto the module. This effectively shorts all of the flying leads together. A laser is then used to create grooves that leave the appropriate vertical connections. Finally, a mission specific thickness of Tantalum is installed for radiation shielding.

The RTIMS module measures 42.7mm x 42.7mm x 13.0mm as shown in Figure 3.




Figure 2. RTIMS Cross Section




Figure 3. RTIMS Dimensions

The RTIMS module is designed with a 144 pin quad flat pack footprint and three active layers of circuitry. A cross section of the RTIMS module (without tantalum shield installed) is shown in Figure 2.

RTIMS also incorporates internal thermal drains and internal radiation shielding. A completed RTIMS module is shown in Figure 4.




Figure 4. RTIMS Module

4. Making the System Radiation Tolerant

4.1. Robustness vs. Capacity

A single event upset (SEU) may introduce one or more data bit errors in a SDRAM device. The RTIMS architecture provides the flexibility of organizing the six SDRAMs, each of 1 Gb capacity, into a 128 MWord (2 bytes per word) EDAC memory, or a 64 MWord TMR memory. This can be easily accomplished if all of the signal pins of the SDRAMs are connected to the FPGA. However, the Xilinx device selected for this project, XQ2V1000 in a BG575 package, does not have sufficient pins to support this solution.

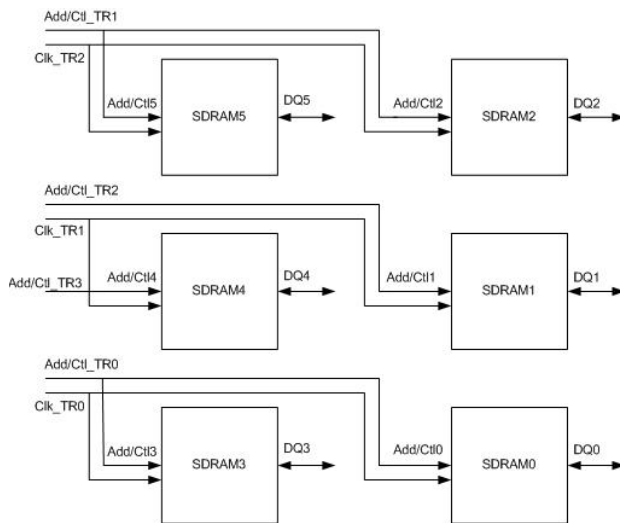


Figure 5. RTIMS FPGA SDRAMs Wiring

Figure 5 shows the RTIMS wiring scheme between the FPGA and the SDRAMs. The Clk_TR0, Clk_TR1, Clk_TR2 are copies of the SDRAM clock.

When operating in TMR memory mode, the control buses, Add/Ctl_TR3, Add/Ctl_TR2, Add/Ctl_TR1, and Add/Ctl_TR0, have the same signaling. During write operation, DQ5, DQ4, DQ3 are driven with the high order byte of the data word, and DQ2, DQ1, DQ0 are driven with the low order byte. During read, DQ5, DQ4, and DQ3 are voted to generate the high order byte of the data word, and DQ2, DQ1, and DQ0 are voted to generate the low order byte. There is no single point failure in this configuration.

When operating in EDAC memory mode, the SDRAMs are organized into two banks of 64 MWords each. Add/Ctl_TR3 and Add/Ctl_TR1 buses have the same signaling and control one bank of memory. Add/Ctl_TR2 and Add/Ctl_TR0 buses also have the same signaling and control the other bank. During write operations, either DQ5, DQ4, DQ2, or DQ3, DQ1, DQ0 are driven with the data word and the check bits. Similarly, during read operations, either DQ5, DQ4, DQ2, or DQ3, DQ1, DQ0 are selected for the data words and check bits. The control signals are single point failures. A SEU can induce multiple uncorrectable data bit errors.

4.2. Single Point Failure

The FPGA and the SDRAM are susceptible to SEU. The logic implemented in the FPGA utilizes the Xilinx TMR strategy, and the Xilinx TMRTOOL [2] is employed to triplicate the design. Due to the limited number of signal pins on the FPGA package, it is not possible to triplicate every i/o signal. The single pin i/o signals thus become single point failures.

The DataIn and DataOut buses of the module interface signals are not triplicated. If this is shown to be problematic, an EDAC scheme can be implemented to protect these buses. The twelve spare i/o on the module can be utilized for this purpose.

The interface signals between the FPGA and the SDRAMs are not triplicated. An upset to anyone of the control signals may result in multiple bit data errors. This may occur when RTIMS is operating in EDAC memory mode, possibly resulting in unrecoverable errors. For the TMR memory mode, this would only affect one of the three redundant words and would therefore be corrected.

An upset to one of the data signals may result in a single bit error. For the EDAC memory mode, this would result in a single bit error, but the data would be recovered (error corrected) via the EDAC circuit. For the TMR memory mode, it would only affect one bit of the three redundant words and would also be corrected.

4.3. SEU

The FPGA is susceptible to SEU. Xilinx XTMR scheme is employed to mitigate the problem. This scheme works based on the assumption that one SEU only induces one upset on a signal and its fan-outs. Furthermore, it is assumed that the upset is corrected before the next SEU. Current manufacturer test results indicate these assumptions are realistic.

4.4. Dynamic State Recovery

With redundancy, the circuit can function correctly with errors induced by one SEU. It is essential to correct the error states as quickly as possible, because the redundancy employed may not correct as errors accumulate. Xilinx XTMR scheme guarantees the correctness of the inputs of the flip-flop for a single error induced by a SEU.

An upset flip-flop can be corrected only when it is reloaded with the corrected value. The flip-flop primitive in the FPGA has a Clock Enable (CE) pin that controls the loading of the flip-flop. Hence an upset flip-flop is corrected only after the CE pin is made active. To mitigate this the usage of flip-flops with rarely active CE pins is eliminated in the RTIMS design.

4.5. Configuration Memory Refresh

The FPGA is programmed to implement an application by setting the states of the configuration memory cells. The configuration memory cells are sensitive to SEU. It is essential to correct a configuration memory upset as quickly as possible, because the logic redundancy employed may not correct as errors accumulate.

RTIMS utilizes a simple approach by refreshing the configuration memory cells at a regular interval that is selectable by the application. The default interval is 1 hour. The interval can be set to 10 minutes, 30 minutes, 1 hour, 4 hours, 24 hours, or off.

The configuration bit stream is stored in the non-volatile radiation tolerant EEPROM. Upon power-on the FPGA loads its configuration from the EEPROM. It is the design objective to store one configuration bit stream that would work for both initial configuration and periodic configuration refreshing.

Table 1 outlines the default configuration bitstream [3]. The default bitstream loads the entire FPGA with one command as indicated in step 2. This also overwrites the block ram. For configuration refreshing, the contents of the block ram are not overwritten, requiring instead a modified step 2. Steps 3 and 4 are also skipped during configuration refreshing.

Step	Function
1	Set up
2	Write Configuration
3	Initialize registers, activate all interconnects
4	Start startup sequence
5	Check CRC, Desynch to desynchronizes the configuration logic, 4 NOPS

Table 1. Default Bitstream

The modified configuration bitstream is shown in Table 2. During initial configuration loading, steps 2A-2C are functionally equivalent to step 2 in the default bitstream.

During configuration refresh, the content of the EEPROM is read sequentially, but steps 2C-5 are skipped by deasserting the CS signal on the FPGA. This effectively skips

the refreshing of the BRAM content, the initialization of registers, and the startup sequence. As steps are skipped, a revised CRC word is supplied. Step 5 is therefore also skipped. The revised CRC word is supplied in step 6. During initial configuration loading, though the CS signal is asserted, the Desynch command in step 5 signals the end of the configuration.

Step	Function
1	Set up
2A	Write Configuration Stream to GCLK, IOB1, IOI1, CLBs, IOI2, IOB2
2B	Write Configuration Stream to BRAM INT
2C	Write Configuration Stream to BRAM content AUTO CRC
3	Initialize registers, activate all interconnects
4	Start startup sequence
5	Check CRC, Desynch to desynchronizes the configuration logic, 4 NOPS
6	Check revised CRC, Desynch to desynchronizes the configuration logic, 4 NOPS

Table 2. Modified Bitstream

The configuration memory refresh controller can be implemented externally or internally. When it is implemented externally, it requires additional radiation tolerant devices. For this project, the configuration memory refresh controller is implemented in the FPGA that is internal to the RTIMS module. The internal controller does not increase the risk of failure as long as any SEU induced errors are corrected before the next SEU.

4.6. Detection of Refresh Failure

The configuration logic of the FPGA is hard-wired, and is not TMR protected. Though it has a small footprint, it has been shown to be susceptible to SEU. An upset here can only be corrected by a full initialization of the FPGA.

Previous methods suggest that when certain configuration control register reads fail, this implies that the configuration logic is not functioning properly. This method is indirect and inefficient, requiring triplicated configuration data bus pins on the FPGA (since the configuration refresh controller is

internal). An alternative efficient method for detecting configuration refresh failure is discussed next.

During configuration refresh, all distributed memory that is implemented with look-up tables within the Configurable Logic Blocks are initialized. This process can be used to detect the failure of the configuration refresh.

The configuration refresh control includes a 16x1 distributed memory. Just prior to configuration refresh, a known pattern is written into this memory. Configuration refresh, when successful, clears this memory to all zeros. Upon the completion of configuration refresh, this memory is read. Any non-zero content directly indicates the failure of configuration refresh.

4.7. DCM Failure Recovery

It has been shown that the functionality of the Digital Clock Manager (DCM) may not be recovered by configuration refresh alone. A reset to the failed DCM is required to recover its functionality. The approach to detecting and correcting this type of failure uses three counters. Each counter is clocked by one of the triplicate DCM clock outputs. The counters are started synchronously and, in normal operation, they will count in lock step. When one of the DCMs fail, the counter associated with this DCM will have a different count than the other two counters. Hence an error is detected and a reset to this DCM is generated.




Figure 6. DCM Checker

The DCM checker is shown in Figure 6. The checker itself is triplicated. A router macro is inserted between the DCM clock outputs and the DCM checker. A DCM reset signal is generated when either the Reset_DCM voter fails, or the DCM clock output fails.

4.8. SDRAM Scrubbing

The SDRAMs are also susceptible to SEUs. It is essential to prevent the accumulation of errors induced by multiple SEUs. Otherwise the error correction circuitry may not be able to recover the data.

RTIMS implements a scrubbing strategy where the contents of the SDRAMs are read sequentially at a fixed time interval. The interval is defaulted to one scrub every 30 us, on average. It takes approximately 32 minutes to refresh the entire TMR configured SDRAM array, and 64 minutes to refresh the entire EDAC configured SDRAM array.

The scrubbing process consists of reading the SDRAM. When an error is detected and the error is correctable, the SDRAM is updated with the corrected data.

Ideally the scrubbing is performed when the SDRAM is idling. Otherwise there will be contention between regular memory accesses and scrubbing memory accesses resulting in performance degradation.

This strategy is flexible and allows minimal contention by taking advantage of the fact that the time interval between scrubs can vary as long as the average interval is maintained. When normal accesses are sparse, the SDRAM scrubber waits until at least 8 scrubs are pending. It will then do scrubbing accesses until the scrubbing is ahead by 8. When normal accesses are dense, the SDRAM scrubber contends with normal memory access requests, resulting in a maximum 0.2% degradation in throughput.

5. Environmental Qualification

The RTIMS modules are currently installed on a 3U Compact PCI memory

board developed at LaRC. This board, in turn, will provide the structural support and electronics support during environmental testing for space qualification as dictated in MIL-STD-883 and MIL-STD-202. The RTIMS module was designed for a Total Ionizing Dose (TID) of 100Krad(Si) at 25°C, latch up immunity to at least 60MeV -cm²/mg at 25°C and an operating temperature of -40°C to +85°C. The types of environmental testing for RTIMS includes: thermal vacuum, vibration, accelerated life, and radiation.

The radiation testing of the RTIMS module investigates the performance of the module in a radiation environment. During initial testing we will take the finished modules to a proton test facility to better understand the proton Single Event Effects (SEE) sensitivity. Follow on testing will include performing a radiation transport analysis using a typical geosynchronous radiation environment to understand the modules' response to Total Ionizing Dose (TID).

5.1 Proton Testing

Piece-part heavy-ion SEE testing has been performed on all devices within the RTIMS module. This testing allowed for an analysis of destructive single event effects and individual component single event upsets and functional interrupts. But, system-level upset and functional interrupt events are possible that cannot be seen and the effectiveness of NASA LaRC's radiation mitigation technology can not be verified by this type of testing.

Ideally, a system is placed in a single-event-producing environment (similar to the mission) and monitored for these synergistic effects while studying mitigation effectiveness. The stacked nature of the module does not allow for ground-based heavy ion sources to penetrate through the entire module structure and impact every device as would higher energy space radiation.

In these cases, high energy proton sources are used which penetrate the entire module structure. This allows verification of the entire module's performance in a radiative environment. However, since proton-induced SEEs require a nuclear interaction, the event rate is significantly reduced, requiring higher fluences to see effects and larger module sample sizes due to high doses from these higher fluences.

5.2 Radiation Transport Analysis

Every component within the RTIMS module was tested with Cobalt-60 for TID. A couple of the components showed sensitivity to TID, so additional shielding was placed in the package. To better understand how the RTIMS module will perform in a relevant space radiation environment, a radiation transport analysis will be done with the radiation transport code NOVICE which calculates expected dosage for each component within the RTIMS module.

For this work, a GEO mission was selected. A year-long radiation environment will be calculated for two conditions - solar active and solar inactive (i.e., the radiation environment includes solar particle events for the solar active environment). These environments will be generated using the standard NASA GSFC space radiation environment models.

Next, the geometry of the RTIMS module will be translated from the schematics and layouts of the RTIMS module into geometry coding that the radiation transport code (NOVICE) can understand. If the calculated doses are less than the sensitivities for all of the components then the RTIMS module should survive greater than

100 Krads(Si) for a GEO mission (for solar active and inactive periods).

6. Summary

The objective of the RTIMS project is to develop and demonstrate an in-flight reconfigurable radiation tolerant stacked memory array based on state-of-the-art chip stacking, radiation shielding and radiation mitigation technologies. Upon completion of the environmental testing this objective is expected to be met.

RTIMS can be a complete computing module. Computing cores can be compiled into the on-board FPGA. The module can then support a distributed, reconfigurable computing architecture.

RTIMS is also suitable for high reliability computing at nuclear facilities. Automated or remote controlled nuclear waste handlers often have significant computing and/or data collection tasks. RTIMS' radiation tolerance makes it an excellent fit for these applications.

The radiation tolerant RTIMS are also suitable for data collection and/or computing applications on nuclear powered craft (aircraft carriers, submarines, and future nuclear powered spacecraft).

But, first and foremost, the RTIMS modules are designed for high performance space-based computing applications, including real-time data processing, reconfigurable computing, and memory intensive space-based systems. The RTIMS technology, which enables new measurements and information products, increases the accessibility and utility of data, and reduces the risk, cost, size, and development time of space-based systems.

[1] Plante Jeannette, *Evaluation of 3D Plus Test Structures*, Dynamic Range Corporation and NASA, pp 13, 12/12/2001.

[2] TMRTool User Guide, Xilinx, Inc., 5/30/2005

[3] Virtex-II Platform FPGA User Guide, Xilinx, Inc., pp 314-336, 3/2005