Lessons Learned Using COTS Electronics for the International Space Station Radiation Environment

John H. Blumer

Boeing Space & Communications Group
499 Boeing Boulevard
MC JW-63
P.O. Box 240002
Huntsville AL 35802
(256) 461-3640
John.H.Blumer@Boeing.com

Abstract. The mantra of *Faster, Better, Cheaper* has to a large degree been interpreted as using Commercial Off The Shelf (COTS) components and/or circuit boards. One of the first space applications to actually use COTS in space along with radiation performance requirements was the EXpedite the PRocessing of Experiments to Space Station (EXPRESS) Rack program, for the International Space Station (ISS). In order to meet the performance, cost and schedule targets, military grade Versa Module Eurocard (VME) was selected as the baseline design for the main computer, the Rack Interface Controller (RIC). VME was chosen as the computer backplane because of the large variety of military grade boards available, which were designed to meet the military environmental specifications (thermal, shock, vibration, etc.). These boards also have a paper pedigree in regards to components. Since these boards exceeded most ISS environmental requirements, it was reasoned using COTS mil-grade VME boards, as opposed to designing custom boards could save significant time and money. It was recognized up front the radiation environment of ISS, while benign compared to many space flight applications, would be the main challenge to using COTS. Thus in addition to selecting vendors on how well their boards met the usual performance and environmental specifications, the board's parts lists were reviewed on how well they would perform in the ISS radiation environment. However, issues with verifying that the available radiation test data was applicable to the actual part used, vendor part design changes and the fact most parts did not have valid test data soon complicated board and part selection in regards to radiation.

INTRODUCTION

The main purpose of the International Space Station (ISS) is to support payloads and testing in a near zero-g environment. Payloads are housed in racks, which in turn are mated to the inside of the ISS modules. Some larger payloads utilize an entire rack and design their own interface controller to interface with ISS power and data systems. NASA and Boeing research determined the majority of the payload users for ISS would fall into the category of only needing a fraction of the space provided by a rack. The research also indicated the smaller payload users did not want to go through the expense of designing a payload controller to interface to ISS. EXPRESS Rack was created to resolve the integration problem for the smaller payload users. It holds 10-15 payloads and subdivides the appropriate ISS resources to the individual payloads. The computer for EXPRESS Rack is the Rack Interface Controller (RIC). The RIC interfaces with ISS copper and optical communication buses and in turn translates ISS commands and protocol to more common and easier implemented interfaces. This allows the RIC to command and communicate with the payloads via common data links like 10BaseT Ethernet, EIA-422, and standard SMPTE-170M video (new specification equivalent to old RS-170A). Thus the payloads are isolated from the ISS interfaces. While the main command and control bus is the time proven MIL-STD-1553, some other ISS buses are fairly unique. The main payload data bus is the fiber optic High Rate Link (HRL) and the video bus is a Pulse Frequency Modulated fiber optic bus. A version of the 10BaseT Ethernet standard is also used. Figure 1 breaks out the I/O of the RIC.
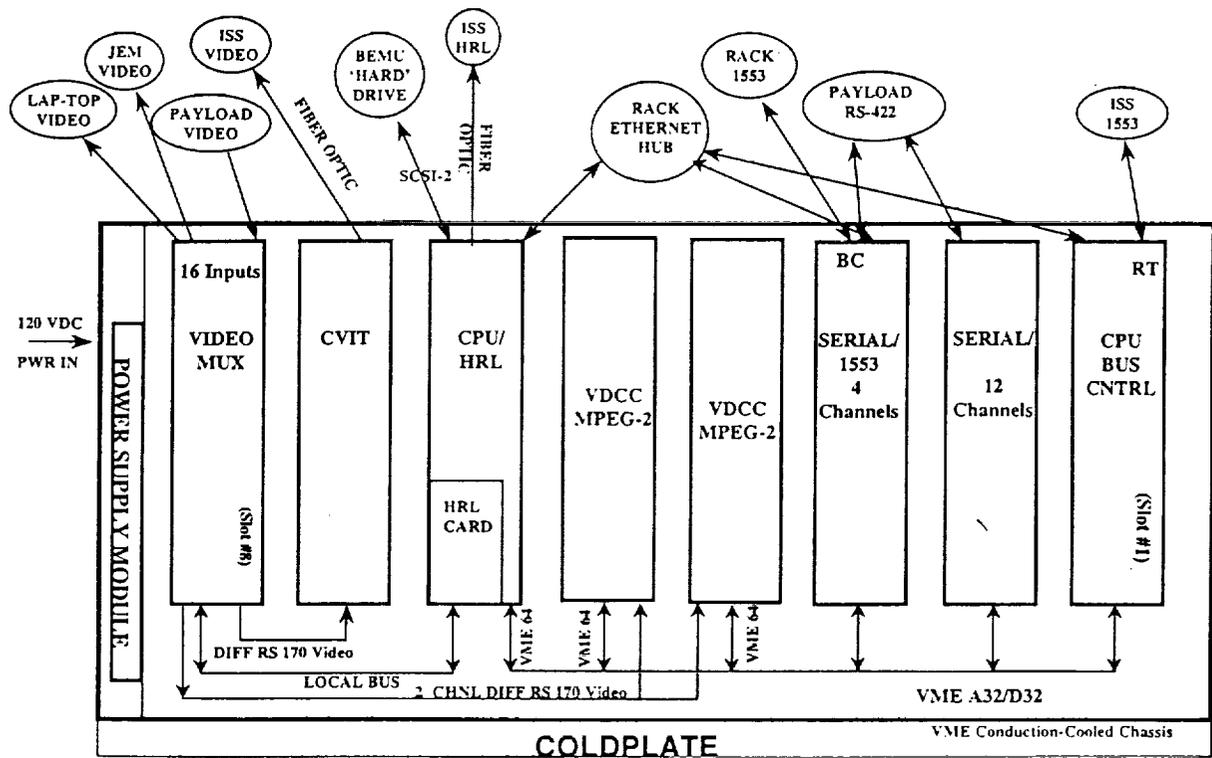
FIGURE 1. RIC VME Layout and I/O

DESIGN PHILOSPHY

As a result of being a new program in a new environment. numerous design requirements were not defined at the start of the program, with radiation requirements being one of the main ones. The philosophy used was to find the knee in the curve of cost vs. radiation tolerance. Rather than set a hard requirement, design goals were set and parts and equipment were evaluated against these goals.

RIC Design Philosophy and Requirements

Early in the program it was estimated that the RIC would have to borrow heavily from existing designs to meet the proposed schedule and funding level. Off the shelf space flight qualified computers were considered but traditionally they have long lead times and are expensive. Another even larger design obstacle to using traditional space flight controllers was most of the data buses, such as the multiple 10BaseT Ethernet ports, video, and fiber optic buses, etc. were not supported in traditional space flight controllers. Also, the ISS radiation environment is benign compared to traditional satellite environments, thus not requiring an expensive space grade computer system. With the geo-magnetic shielding provided by the low Earth orbit and shielding provided by the ISS module wall, surrounding rack equipment, and nominal chassis sidewalls the total dose environment was estimated to be 300 to 1000 rad/si for a 10-year mission, depending on component location. Thus, almost any component grade should meet the total dose requirements. However, upsets and latchups due primarily to the trapped proton belts surrounding the Earth and occasional heavy ion strikes were a significant probability and thus a design driver. The new, in early 1994, design philosophy of *Faster, Better, Cheaper* was dominant and COTS computers and boards were researched as way to meet the new philosophy. Military grade VME boards were researched from multiple vendors and several were found that could support the new Input/Output (I/O) requirements, while meeting or exceeding ISS thermal, shock, and vibration requirements. Note: for the purpose of this paper COTS is defined as a

catalog item or derivative of a catalog item, even if it meets military specifications. Also military grade VME components were readily available within the cost and schedule requirements. It was recognized that the unique ISS interfaces such as the fiber optic video bus, video switching requirements and HRL data bus were not available as COTS and would require custom designed boards.

Ionizing Radiation Impact for the Initial Test Flights

Parts lists were requested of applicable vendors so they could be evaluated in regards to radiation. Several vendors submitted parts lists for evaluation. The best cards, in regards to radiation, had applicable test data for 50-75% of their parts. Other cards had less than a quarter of their parts with applicable test data. The best candidates for the radiation environment were purchased for performance evaluation. After this round of elimination, a set of cards was selected for a test flight on the Space Shuttle. Three CPU (Central Processing Unit) boards were selected from two different vendors, plus a military grade VME computer chassis with power supply. One board functioned as the system controller with the other two boards handling payload and system I/O. Due to schedule and cost, only the items of greatest concern in regards to radiation were corrected. This consisted of simple vendor substitutions for identical type integrated circuits and replacement of the memory module in all three boards. The best radiation tolerant commercial SRAM (Static Random Access Memory), at this point in time, was a certain production run from Micron Semiconductor. The Micron SRAM was previously tested and had reasonable radiation tolerant numbers and thus was used in all three CPU card for the Space Shuttle flight.
The prototype RIC system was tested on two Space Shuttle flights (STS –83 and STS-94) in 1997 with no SELs (Single Event Latchup) observed. Note: an SEL is a condition that often causes additional current draw and usually needs to have power recycled to cure. Permanent damage often occurs in some parts if a SEL condition is not resolved in a timely manner. A system interruption did occur but it was not traced to a Signal Event Upset (SEU). Note: the most common SEU manifestation is a bit flip, where a component or memory location may change state from a "one" to a "zero" or vice versa.

Design Impacts Due to Loss of Mil-Spec Components

Although the Shuttle flight was a success it was recognized that a unit with better radiation numbers than the Space Shuttle unit was needed for the longer duration ISS mission. Due to operational usage, as well as ISS requirements, the ISS unit had more stringent requirements than the Shuttle unit. The 10 year mission life and higher reliability requirements drove the design to use more Mil-Spec components. Working against the higher reliability requirement was the higher probability of SEE (Single Event Effects – includes both SEU & SEL, as well as other effects) due to the higher 51.6-degree inclination orbit of the ISS. The higher inclination of ISS, as opposed to the standard Shuttle orbit inclination of 28.5 degrees, places ISS more in the tapped proton and electron belts or Van Allen Belts, thus increasing the probability of an SEE. Also at 51.6 degrees, ISS passes through the South Atlantic Anomaly (SAA) off the coast of Brazil in the Atlantic, which contains significant trapped protons and thus becomes a major contributor for increased SEE rates.

The initial parts evaluation for the cards used in the Space Shuttle unit was performed in 1994; thus the cards were designed prior to some I.C. (Integrated Circuit) vendors, like Motorola, pulling out of the military market. A second review of the parts for the proposed ISS VME cards approximately two years later showed significant component changes from the version used in the Space Shuttle test flight. With vendors pulling out of the Mil Spec market, the board vendors had made significant part and design changes in order to make use of the dwindling base of mil-spec components. Even parts that appeared to be the same had significant changes. One example is a Motorola 68302 serial controller, which was done originally on Motorola's military line using epitaxial wafers. A subsequent review after Motorola left the military parts business indicated the VME board vendor was using a Motorola 68302 die from Motorola's bulk CMOS (Complementary Metal Oxide Semiconductor) commercial line repackaged by a third party. The die was repackaged per military specifications using a ceramic package. Thus, while the part was still a ceramic military grade part (MIL-STD-883), the previous radiation analysis was invalid. The subsequent analysis on the boards indicated a considerably larger number of parts completely unknown in regards to radiation performance. An even larger concern were the parts which the original mil-spec vendor with known good radiation tolerance data

dropped out of the military parts business, with no equivalent grade drop-in substitute. One example of this is a FPGA (Field Programmable Gate Array) that is used in four places on one of the upgraded ISS boards. The previous Shuttle board used a radiation tolerant FPGA vendor that subsequently dropped out of the mil-spec business. The replacement FPGA, while initially unknown, was found later to have a significant destructive SEL risk. Thus, the design went from a known good FPGA in regards to radiation to a known bad one.

Component Upgrades & Approval

In tracking down all the unknown parts another change was noticed from the previous parts review, which was conducted almost two years earlier. A significantly higher percentage of the parts were commercial die from an I.C. manufacturer repackaged by a third party (E.g. 68302, SRAM, Flash memory, PowerPC, etc.). The practice of using a commercial die repackaged to mil-specification while a good idea in regards to most environmental issues further complicated the radiation analysis. Thus data had to be obtained from the vendor who packaged the die and then from the original die vendor to make an evaluation. Utilizing both local expertise and the radiation effects experts at BREL (Boeing Radiation Effects Laboratory); the vendor's part lists were re-evaluated. If known bad or suspect parts were identified, the first and most economical course of action was to find an exact replacement part with a known good radiation pedigree or applicable related test data. For several items, this was an acceptable solution (DRAM, FCT family drivers, etc.). However, for some parts like the DRAM, upscreening commercial grade components to the specified thermal and reliability requirements was utilized since no equivalent mil grade part was available. Also, many parts were approved by similarity, such as the FCT logic family parts. Test data was not available for the exact part needed, however parts from the same vendor and in the same FCT logic family were tested with satisfactory results. Thus, it was decided to approve the FCT parts by similarity. While this process does entail some risk, the risk was deemed small enough and was outweighed by the costs associated with redesign of the board to eliminate these parts or to test the parts in question. However, for several items, no test data could be found or test data found indicated poor radiation tolerance with no acceptable substitute part located. For some parts, it was determined the function was not needed or the function could be moved to another board location. (E.g. the RS-422 controller was depopulated from one board and the RS-423 channel was used instead).

Design impacts and associated NRE

In order to keep NRE (Non Recurring Engineering) costs low, radiation testing and board redesign were only used as a last resort. But for some parts, the function was crucial and no substitute could be found. One example that incurred NRE was to replace the function with programmable logic. UTMC's RadPal, a rad hard 22V10 PAL (Programmable Array Logic) was chosen for some functions. However, the two largest design impacts in terms of cost and schedule were the SEL issues related to the FPGA and to a lesser degree the 68302.

When the design for the serial board was selected, the military grade epitaxial version of the 68302 was available as the controller for the serial channel, which had adequate test data available. However, with the withdrawal of Motorola from the military business, the only substitute was the use of the repackaged version done by Thompson-CSF using a Motorola commercial bulk CMOS 68302. Test data obtained indicated the bulk CMOS version had a Single Event Latch-up (SEL) problem. Since the design impact was considerable to redesign the board to replace the part, it was decided to use a traditional workaround of adding circumvention circuitry to monitor for a latch-up condition. The circumvention circuit performed this function by monitoring current to the 68302 power pin. If the current exceeded a predefined threshold the circumvention circuit would, via a couple of transistors, open the power line to the 68302 power pin as well as pull the power pin to ground for a predefined time interval, which in principle halts the current flow and thus the SEL. It was decided to test and verify the circuit in conjunction with other components being tested by BREL (Boeing Radiation Effects Laboratory) with heavy ions at the Berkeley cyclotron. The circuit worked as designed during test. However, when the power was reapplied to the 68302 power pin, the current went back to SEL current levels. Subsequent research yielded the theory that even with the no power applied to the 68302 power pin, enough current was being sinked via the data and/or address lines to maintain a latchup condition. A second circuit redesign added high impedance tri-state drivers to the 68302 addresses and bus lines. In addition to switching a couple of transistors, the comparator circuit now also caused the tri-state drivers to

also switch to a high impedance state. A subsequent test at BREL using their californium (Cf-252) test chamber confirmed that the addition of the tri-state drivers corrected the problem.

Four of the FPGAs with a destructive SEL potential were used on one board. Also, the FPGAs had numerous power pins as well as data/address and I/O pins. When factored in with the board density, a circumvention circuit as used on the 68302 was deemed unworkable. The only solution, other than a large redesign effort, was to use another programmable device as a drop in replacement. Unfortunately, no drop in substitute FPGAs could be found that would match both the I/O pin out and meet the radiation requirements. The only solution that could be found was to use a radiation tolerant ASIC (Application Specific Integrated Circuit) that could replicate the same programmable logic as the FPGA. An ASIC was not initially the preferred solution due to higher NRE costs. An ASIC is a semi-custom design programmed at the vendor's factory vs. the end user programming an FPGA. A TEMIC radiation hard ASIC was selected as the FPGA replacement. TEMIC's Matra MHS division manufactured the ASIC. A process already in place by TEMIC allowed for the transfer of netlists from certain FPGAs to the radiation hard ASIC, with minimal NRE (as compared to a new development effort). Fortunately the FPGA in question was one that was supported by the TEMIC transfer process. While the process added cost and schedule impacts, it was significantly less expensive than a board redesign. After the ASIC completion, it was tested in the VME card and worked as a drop in replacement.

VDCC (Video Digitization and Compression Card)

Due to initial estimates and conservative ISS thermal data early in the program, the VME cards and power supply were specified to operate at +85C. This high operating range and the related reliability requirement more than any other requirement drove the use of military grade parts. For a second version of the RIC, a video compression requirement was added. The state of the art MPEG-2 compression algorithm selected and related components proved impossible to procure in military or even industrial grade components. Even upscreening the very high-density commercial packages available looked impossible. In fact one of the candidate MPEG-2 encoders, due to the density and clock rate, was only rated to +45C! Subsequent re-evaluation of the ISS thermal requirements led to reducing the upper end to +75C. Eventually a board design was selected that was based on a readily available commercial design using real COTS components (vs. mil grade components). Analysis and testing indicated components could be upscreened to meet the lowered thermal requirements. However, parts like the MPEG-2 encoder chip, which was rated up to 4 watts, required special care to ensure an adequate thermal path to the VME chassis sidewalls. In the end special thermal paths had to be used in addition to the normal thermal management layer.

The initial radiation analysis appeared to be even a larger driver than the thermal issue. The design used mostly state of the art components with most having no test history, neither direct or by similarity. Radiation testing seemed the only solution, but the normal test method of using a small vacuum chamber at a heavy ion test facility would limit testing up to a specific section of the board or a component at a time. Also due to penetration issues almost all parts have to be "delidded", thus removing the material over the die and exposing the die to the heavy ion beam. Due to these issues and the large number of parts needing to be tested, it was estimated the test would be long and expensive and greatly exceed the allocated budget. Testing at a proton facility was investigated as a possibility. Several facilities support testing with protons in an open facility that does not require delidding the components. However, because of the concern of a SEL induced by a heavy ion above the threshold of protons, the validity of testing only with protons was questioned. Fortunately, the Super Conducting Cyclotron at Michigan State University was opened up to non-academic testing at about the same time this problem was being evaluated. The Michigan State Facility facilitated testing the board with a very high-energy heavy ion beam in a large open air chamber. The facility also produced very high energy and thus highly penetrating heavy ions; thus delidding the parts prior to testing was not required. An X-Y positioning table was used during the test to position the part under test in the heavy ion beam in real time. Usage of the positioning table and the open air chamber allowed the use of a true COTS board as a test article, thus dramatically reducing test costs (no special test boards to produce). The open air chamber and positioning table also greatly reduced test setup and test time in the chamber as indicated in Figure 2.

Testing revealed several parts had radiation problems. For some parts, design workarounds or mitigation techniques were used. Traditional techniques like EDAC (Error Detection And Correction) was used for the commercial grade main memory, which suffered from a nominal SEU rate. Fortunately, EDAC was already supported in the

commercial grade memory controller. which passed radiation testing with minimal SEU concerns. The external L2 memory cache also had SEU concerns, but adding EDAC or other correction circuitry proved to be difficult without significantly slowing down the L2 cache and thus defeating the purpose of a high-speed L2 cache. Subsequent performance analysis indicated the performance requirement could be meet without the L2 cache, thus the L2 cache was eliminated from the design. The other section of memory, the FIFO (First In First Out) memory used for video buffering, also defied a logical mitigation technique, but the application was critical and could not be designed out. A part with an even greater SEE concern was the PCI controller, which the VME card used as a local bus controller. Test results indicated the PCI controller used had both SEU and SEL concerns. The PCI controller function could not be eliminated nor was a viable mitigation technique available. Since the card was already being redesigned to accommodate other design changes and thermal management, it was decided to use Actel's 54SX line of radiation tolerant FPGAs to replace the PCI controller and other logic devices used on the board. Actel also had available certified logic cores for their 54SX family. Fortunately, a PCI controller core was available from Actel, thus NRE to convert from the commercial PCI controller to the radiation tolerant FPGA was low. With the main SEL problems solved, the SEU concerns were addressed. Several parts, such as the FIFO memory, had significant SEU concerns with no apparent cost viable solution. The parts were evaluated in regards to their function and the effect a SEU would have not just on the particular part, but on the entire system. If the effect was a transient noise or disturbance in the video stream a large effort was not made to try and resolve the SEU. For example, an SEU in the FIFO between the video encoder and the host CPU (via local PCI bus) would cause a video artifact in the block of video pixels effected. However, the effect would normally only last for a few frames, often only occurring in a single frame. Thus SEUs in components like the FIFOs were deemed acceptable from a system viewpoint. The last remaining problem was the PowerPC 750 L1 cache. The L1 cache, which is located within the host CPU (PowerPC 750), is the source for approximately 99% of the CPU SEUs. Based on initial performance estimates, the L1 cache was disabled to meet the SEU goals. The predicted and measured loss in performance closely matched, with a measured decrease in CPU performance of almost 50%, due to loss of L1 cache.
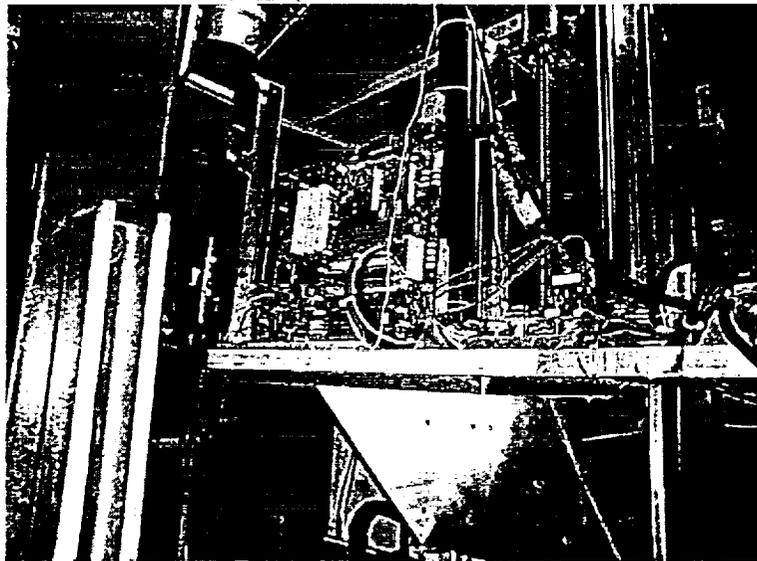

FIGURE 2 VDCC Radiation test Setup

AVDCC (Audio Video Digitization and Compression Card)

A third variation of the RIC added audio compression to the VDCC. The Analog Device's DSP (Digital Signal Processor) used for audio compression on the AVDCC was also tested at the Michigan State University Super Conducting Cyclotron and had significant SEL and SEU issues. Replacing the DSP was the initial design solution. However, an alternate DSP and software package could not be found that offered the same level of easy integration.

The final decision was to continue with the existing hardware and software design and try to resolve the radiation issues. For the SEL problem, a traditional circumvention circuit is planned. (Note: the updated AVDCC design is not yet in production). For the SEU effort, no traditional concept looked feasible. The DSP design used has two internal segments of memory and does not use external memory while operating. One segment is for the program or execution code and the other for data. Several design solutions were considered, but they all exceeded the cost and schedule targets. Once again, the design requirements were evaluated against actual performance. As in the video side, it was decided an occasional or random "pop" was acceptable. Thus, no effort was made to correct an SEU in the internal DSP SRAM (Static Random Access Memory) that contained and processed the audio data. In the program memory SRAM segment, it was decided an SEU could cause a longer lasting interruption or corrupted data stream and thus was a greater concern. However, the upset rate is low enough to be considered acceptable once the tolerance of the human ear is considered.

CONCLUSION

Several items worked against using even military grade COTS. First, program timing was bad with its start about the same time as several vendors withdrawing from the military grade component business. This accelerated the normal turn over in parts used. However, the problem is even worse when using pure commercial parts that have a production lifetime of only a few years at best. Thus, the design initially evaluated can have significant parts substitutions by the time the order is placed with a vendor. While most vendors control parts changes for their boards, they do sometimes make what they consider a transparent change and usually it is for most parameters. But it is seldom transparent for radiation effects. Components with identical electrical performance may not have identical radiation performance. Also, with the rise of third party vendors who repackage die to meet a higher thermal environment or military specifications the issue is even more clouded. In one case a third party vendor repackaged a SCSI interface chip. When the original die was no longer available, another die was used. While performance remained the same, the radiation tolerance could have changed considerably as the result of the die change. Thus, looking at the component on a board the component appeared to be identical to an earlier version with a totally different die. While most of the vendors were very helpful, they normally do not deal with radiation issues and thus initially part issues went unnoticed until subsequent part lists were issued. One lesson learned is that very close coordination with the vendor on components and parts used must be constantly maintained to avoid surprises. This is especially true if the vendor does not have design experience for radiation environments.
In several cases due to cost and schedule reasons (the reasons COTS was selected in the first place) traditional practices, like part replacement, could not be performed. More inventive solutions had to be utilized. In some cases the solution was simply looking at how the SEU manifested itself and determining if an SEU was a momentary transient in which the human eye or ear could easily tolerate.

What was an allowable SEE was also modified as the program progressed. Momentary anomalies in the audio and video streams were determined to be acceptable as the human ear or eye can accommodate such an interruption with no real loss of data. In several cases, such as the video buffer FIFO, the only alternative was an unacceptable redesign, which would have resulted in significant cost increase and a loss of performance.

The boards used were based on a COTS design, thus significant cost savings was realized in the form of using commercial software. However, the cost saving in terms of hardware was less than expected. While the recurring cost for the boards came close to original estimates, the non-recurring costs associated with the radiation enhancements exceeded original estimates. The lessons learned should help minimize such impacts in the future. However, the inherent risks of qualifying COTS, especially new state of the art components, for space use is largely educated judgement until the board/components in question are understood and tested.

Acknowledgments

I also would like to acknowledge the three military grade VME vendors Radstone (VME chassis), DY-4 (main I/O card in system), and Aitech (three serial I/O cards & VDCC/AVDCC). Even though the amount of design changes and paperwork required of the ISS program exceeded what all three were used to for a military program, they remained cooperative and provided effective economical suggestions on how to solve various problems.

REFERENCES