

A Radiation-Tolerant Low-Power Transceiver Design for Reconfigurable Applications

D. Weigand, M. Harlacher
ITT Industries Advanced Engineering & Sciences Division
1761 Business Center Drive
Reston, VA 20190-5337

Abstract—This paper describes the ongoing development of a radiation-tolerant version of ITT Industries’ Low-Power Transceiver (LPT). The LPT is a compact, flexible device that can be configured to perform custom communications and navigation functions in terrestrial, airborne, and space applications. The radiation-tolerant version of the LPT will maintain most of the LPT’s existing functionality and reprogrammability. This paper emphasizes the development of the radiation-tolerant digital modules, which contain reprogrammable FPGA and DSP devices, and their application to future space applications.

I. BACKGROUND

The vision of NASA’s Earth Science Enterprise involves the interconnection of future space instruments into a vast network, with each space instrument working individually yet collaboratively to achieve new levels of performance. This future network, which is referred to as the Sensorweb, allows multiple vantage points (e.g., LEO, MEO, and GEO) to provide data diversity. Within the Sensorweb, formations of micro- and nano-satellites perform autonomously while achieving reliability through redundancy. In addition, reconfigurable payloads reduce risk by allowing for context switching and instrument/algorithm upgrades and adaptations after deployment [1]. A key technology component that will enable these types of designs is a small, highly integrated, reprogrammable, multi-purpose communications and navigation payload that can withstand the radiation environments encountered over a variety of orbits.

II. THE LOW POWER TRANSCEIVER (LPT)

Over the last four years, ITT Industries and the National Aeronautics and Space Administration (NASA) have been developing the Low Power Transceiver (LPT). The LPT (Fig. 1) is a compact, flexible device that can be configured to perform custom communications and navigation functions in terrestrial, airborne, and space applications. Composed of multiple PC/104 modules, the LPT assembly is modular in nature and therefore suitable for implementing a wide variety of integrated functions (e.g., numerous simultaneous software receiver and transmitter channels over multiple frequency bands). The LPT performs signal-processing functions with reprogrammable FPGA and DSP devices. Additionally, the industry standard modules used in the LPT allow it to host application-specific and COTS modules that contain processors and interfaces.



Fig. 1. The Low-Power Transceiver (LPT)

Due to its flexible, programmable nature, the LPT is suitable to host virtually any form of modulation or demodulation. In existing implementations, the LPT firmware is designed to accommodate up to 16 receive RF bands, and to process up to 32 independent data channels from any one or all of the RF bands. In a typical configuration, twenty-four of these data channels are dedicated to processing GPS signals (e.g., two 12-channel GPS receivers), leaving eight channels for data communications. Each of these channels is physically identical, operates independently, and may be fed from any RF input band.

The demodulator typically implemented in the LPT is capable of processing both phase shift keyed (PSK) and more general linear phase modulated (PM) signals, either spread or non-spread, as required by NASA’s TDRSS and STDN communications systems and by GPS. Coherent tracking of symbols, carrier and, when necessary, PN code is accomplished for data rates from 50 bps to 1 Mbps. The most recent generation of the LPT includes a “*Dial-a-Channel*” feature that dynamically optimizes the LPT’s signal processing capability by allowing the user to increase the

nominal maximum data rate when the total number of receive channels is reduced. When this feature is combined with further exploitation of the FPGA's signal processing power, ITT expects that the LPT will be capable of receiving data rates in the range of 10s of Mbps to 100s of Mbps. The demodulator also includes Viterbi decoders and differential decoders for its communications channels. Channel metrics generated by the demodulator include frequency offset, signal level, Eb/No, lock detectors, and acquisition time.

The modulator is a dual-channel structure capable of modulating two independent data streams at data rates up to 5 Mbps each using BPSK, QPSK, OQPSK, and linear PM. A development path exists to increase the modulator's maximum data rates by further exploiting the programmable resources and clocking speeds of the FPGA. Like the demodulator, the modulated signal may either be spread or non-spread. In addition to the standard modulation types listed, a form of baseband filtered OQPSK is implemented in the modulator. This modulation affords a near-constant envelope waveform that may be configured to trade between channel bandwidth efficiency and RF power amplifier efficiency, yet is suitable for use with receivers that implement integrate-and-dump type matched filters. Further, the input data streams may independently be differentially encoded, convolutionally encoded, convolutionally interleaved and Manchester encoded, as desired.

In addition to the generic communications receiver, specific signal processing is incorporated in the LPT to process GPS signals using the civilian C/A codes. The existing LPT implementation performs both code and carrier phase measurements, and maintains an estimate of time. LPT also recovers and decodes each tracked spacecraft's ephemeris and almanac broadcasts. These phase measurements and broadcast messages may be combined with various forms of application-specific solution processors to form estimates for position, velocity and time as required by a particular mission. The most notable LPT GPS capability is its "time to first fix." On average, when provided with no a priori information regarding the GPS constellation relative to LPT position, and whether on the Earth's surface or in low Earth orbit (LEO), the LPT is able to search for all spacecraft, over the entire Doppler uncertainty region, and over all PN code offsets, in order to produce an estimate of position in approximately three minutes. This average time accounts for the randomness in acquiring the first four signals, for obtaining frame synchronization and for receiving the full ephemeris download from each of those four spacecraft.

As a highly integrated, multi-purpose communications and navigation payload, the LPT achieves a level of cost-reduction, miniaturization, reprogrammability, and performance that is attractive to space missions. A current LPT configuration integrates the functions of

communications, using both NASA's Space (TDRSS) and Ground Networks (STDN) for TT&C and science data relay, and navigation, using GPS. This specific capability will be demonstrated in orbit on an upcoming Shuttle flight in July 2002 (STS-107). The experiment will demonstrate simultaneous communications and autonomous navigation capabilities on orbit—critical requirements for both Space-Based Range Safety and Formation Flying applications. Additionally, the LPT will provide communication and navigation functions for the AFRL's one-year XSS-11 mission in 2004 and communication, navigation, and crosslink functions for the AFRL's one-year TechSat 21 mission in 2006.

III. THE RADIATION-TOLERANT LPT (rLPT)

Even though the first two generations of the LPT were focused more on proving functionality and performance and less on addressing environmental aspects, the intended operational environment of the LPT has always included the space environment, and the issues of high reliability and radiation tolerance are being addressed in the third generation of the LPT. ITT, NASA, and AFRL are in the process of working towards developing a radiation-tolerant LPT (rLPT) that has most of the functionality and reprogrammability of the current LPT with the ability to operate reliably in the space radiation environment.

The approach that is taken for developing a high-reliability, radiation-tolerant version of the LPT is to address the reliability and radiation tolerance of each LPT module independently. The first goal is to maintain the LPT's functional capabilities, performance, and reprogrammability in a radiation-tolerant design, and the second goal is to develop a common architecture that can be tailored for use in all applications—terrestrial, airborne, space—by populating the unit with components that meet the mission's reliability and radiation-hardness requirements. Three primary technology areas present themselves in the LPT with respect to addressing high reliability and radiation tolerance—power supplies, RF hardware, and digital hardware.

First, the power supply was most easily addressed by incorporating existing high reliability and space qualified components.

Second, because the materials commonly used in RF hardware are inherently immune to the effects of radiation, the components used on the RF boards can typically be qualified for space without significant risk. Hybrid modules are being developed to address the high-reliability aspects of these components. Each component's die is being selected for both its reliability performance and radiation tolerance.

Third, of greatest susceptibility are the digital components, specifically the A/D and D/A converters, Digital Signal

Processor (DSP), and FPGAs used in the LPT. Many A/D and D/A converters are susceptible to single event latch-up (SEL) and single event upset (SEU), so care must be taken in selecting and testing these components prior to use. In the LPT, many of these converters are being selected as a part of the hybrid designs, so they are being screened appropriately. The vulnerability of the DSP is addressed by using an existing, radiation-tolerant DSP chip.

The single largest “enabler” technology for the LPT is its heavy reliance on FPGAs. Fortunately, alternatives now exist for utilizing very high-density FPGA components in radiation environments. The third generation LPT uses a combination of FPGAs from both Actel and Xilinx. Actel FPGAs, which use anti-fuse technology, are available in a SEL immune version with TMRed flip-flops to ensure SEU hardness. The one-time-programmable Actel FPGAs will be used for all functions that must be available upon application of power to LPT and functions that are required to support the higher density Xilinx FPGAs. The third generation LPT uses Virtex-II FPGAs with “equivalent” ASIC gate densities as high as six million gates. The Virtex-II family of FPGAs has not been fully characterized in a radiation environment at the time of publication of this paper, but it is expected that the Virtex-II will behave similarly to the Virtex family, which has been characterized for use in a radiation environment.

IV. THE rLPT DIGITAL MODULE DESIGN APPROACH

NASA’s Earth Science Technology Office (ESTO) is funding research through the Advanced Information Systems Technology (AIST) program to design and develop the radiation-tolerant third generation LPT Digital module. This effort poses several design challenges, because it requires hardening the module’s FPGA and DSP devices, which perform the LPT’s advanced signal processing functions. Developing a fault-tolerant architecture for the SRAM-based reprogrammable FPGAs to mitigate SEUs and replacing the current DSP with a radiation-tolerant alternative are key elements of this effort. The resulting Digital module will be stacked with other radiation-tolerant modules to form the rLPT.

A. Fault-Tolerant, Reconfigurable FPGA Design

The second generation LPT Digital module (Fig. 2) uses two Xilinx Virtex-E FPGAs to provide the bulk of the LPT’s signal processing capability. In the third generation LPT, these FPGAs are replaced with radiation-tolerant Xilinx Virtex-II FPGAs. Xilinx uses an epitaxial silicon layer in its CMOS FPGAs to create FPGAs that are immune to SEL. These radiation-tolerant FPGAs preserve the reprogrammability and gate density that are necessary to host the types of reconfigurable advanced communications and navigation applications that are implemented in the LPT.




Fig. 2. The Second Generation LPT Digital module

1) *Radiation Effects on CMOS SRAM-Based FPGAs:* The two main categories of radiation effects are total ionizing dose (TID) and single event effects (SEEs). Testing has shown that the radiation-tolerant Xilinx Virtex-E FPGAs are immune to a total ionizing dose (TID) to at least 300 krads(Si), which exceeds the total accumulated radiation experienced on most long-duration satellite missions. It is expected that the Virtex-II will also exhibit such TID immunity.

There are two types of SEEs that affect CMOS SRAM-based FPGAs that have already been referenced in this paper—SEL and SEU. SEL is a potentially destructive effect, but testing has shown that the radiation-tolerant Xilinx Virtex-E FPGAs are immune to a linear energy transfer (LET) of up to 125 MeV-cm²/mg, which exceeds the expected LET for the space radiation environment. It is expected that the epitaxial layer that has been successfully used by Xilinx to create SEL immune FPGAs will be effectively used in the Virtex-II FPGAs.

However, these radiation-tolerant FPGAs are susceptible to SEUs, and there are three main categories in which SEUs are manifested—configuration upsets, user logic upsets, and architectural upsets, which are also known as single event functional interrupts (SEFIs) [2]. Configuration upsets alter the programmed function and internal signal routing of the FPGA, which may cause functional errors. User logic upsets affect bits that are dynamic during the FPGA’s functional operation (e.g., signal propagation, RAM contents). SEFIs are device unique upsets to the FPGA control elements. For example, the Virtex FPGA is susceptible to an upset to the power-on reset (POR) and the JTAG tap controller. Each of these SEFIs will result in a complete interrupt of device functionality and require external logic to recover operation.

The impact of the various manifestations of SEUs described above is dependent upon both the space radiation environment and the SEU mitigation techniques that are used,

because ultimately, the SEUs' impact on performance is quantified by the FPGA's functional upset rate with mitigation.

2) *Space Radiation Environment Analysis*: To quantify the effect of space radiation on the Xilinx FPGAs, ITT conducted an analysis of the radiation environment for a wide variety of orbits. Depending upon the orbit parameters (e.g., mission timeframe relative to sun's solar cycle, spacecraft trajectory, mission duration, and spacecraft shielding), the expected TID and SEU rate for the FPGAs vary by orders of magnitude. Space Radiation software was used to quantify the effects of radiation. The TID calculation is relatively straightforward. However, the SEU rate calculation and interpretation is more complex.

The amount of radiation encountered by a spacecraft is constantly changing as a function of the sun's solar activity and the spacecraft's position. To fully characterize the expected SEU performance, the average and worst-case radiation levels were calculated. The design accounts for these radiation levels depending upon the mission-criticality of the FPGA functions. For example, a mission critical function must have full fault-tolerance in the worst-case radiation environment, but a non-mission critical function may be built with fault-tolerance for a lesser radiation level. An availability analysis helps to determine the requirements for design robustness. Since the LPT is intended to provide communication and navigation functions throughout a spacecraft's mission, the design is targeted to mitigate SEUs that occur at the worst-case rate.

3) *Mitigation of SEUs in LPT FPGAs*: To develop a SEU-mitigating LPT architecture, ITT conducted a failure modes and effects analysis (FMEA) for the LPT design and developed tools to mitigate the effects of the failure modes. Table I shows the FMEA results for the LPT FPGAs.

Configuration bit upsets can be corrected using partial reconfiguration (PRC) of the FPGA's configuration bitstream, respectively, without having to pause the FPGA's operation, since the Xilinx Virtex FPGAs can be partially reconfigured without any interruption to the device operation [2]. PRC scrubbing continuously refreshes the FPGA's configuration bitstream, which corrects any bit errors that occur within the configuration bitstream. However, to mitigate any functional or signal routing errors that occur during the interval between a configuration bit upset occurrence and a configuration bit upset correction, triple modular redundancy (TMR) must be used throughout the design, as recommended by Xilinx [3].

The only way to mitigate user logic upsets is to use redundancy in the design. Although there are efficient error control techniques that can be used to mitigate these types of errors, the TMR that is already necessary to mitigate the functional and signal routing errors caused by configuration upsets will be used to mitigate the user logic upsets [3].

Architectural upsets are detected by looking for a unique fault signature. Recovery from a POR upset requires a full device reconfiguration, but fortunately, it has a very small cross-section ($<10^{-5}$ cm²), so it is expected to have a very low probability of occurrence on orbit—on the order of one POR every 170 years in a geostationary orbit. The JTAG tap controller upset has never been observed in testing, so it has an even smaller cross-section than the POR, and it can be detected and corrected within five clock cycles. The only way to completely eliminate the effect of SEFIs is to use hardware redundancy (i.e., multiple FPGAs). Fortunately, the cross-section of circuitry susceptible to SEFIs is so small that the probability of seeing either of them is exceedingly small (though it is quantifiable) [2].

TABLE I
SEU FAILURE MODES AND EFFECTS ANALYSIS

Type of Upset	Potential Effect of Upset	Mitigation of Upset
Configuration Upset	<ul style="list-style-type: none"> - No disruption (many SEUs have no effect on design function, because many designs do not use all configuration bits) - Corruption of routing and interconnects - Corruption of LUT values or logic 	<ul style="list-style-type: none"> - Detectable and correctible in configuration bitstream - Until upset is corrected, redundancy may be necessary to mitigate effect of upset
User Logic Upset	<ul style="list-style-type: none"> - Bit error - Flip-flop state error - RAM value (Block RAM or LUT-RAM) 	<ul style="list-style-type: none"> - Undetectable in bitstream readback - Some are self-mitigating or transient (e.g., upset value in shift register) - For persistent upsets (e.g., state machine), redundancy is necessary to correct upset
Architectural Upsets (Virtex FPGA)	<ul style="list-style-type: none"> - Power On Reset (POR) upset - JTAG tap controller upset 	<ul style="list-style-type: none"> - POR: Very low probability of occurrence—reconfigure device when upset signature is detected. If immunity to POR upset is desired, design must use multiple FPGAs - JTAG tap controller: Simple hardware configuration will mitigate this upset

Cross-section is the device's SEE sensitivity to ionizing radiation. It is the number of errors per ion fluence versus LET and it is expressed in units of cm^2 per device or per bit. Cross-section data for the radiation-tolerant FPGAs can be combined with the energy spectra for various radiation sources to estimate upset rates.

Based upon this analysis, ITT concluded that the LPT design requires a means for correcting static errors in the configuration bitstream as well as means for correcting dynamic errors in both the configuration bitstream and user logic. The LPT design does not intend to add hardware redundancy to mitigate the effects of the SEFIs, since the probability is acceptably small for most potential missions. Fig. 3 illustrates the judicious combination of TMR and PRC scrubbing that will be used to mitigate the effects of SEUs in the LPT.

In order to combat SEUs in the Xilinx parts, a number of techniques are available and have been thoroughly characterized by Xilinx in radiation test facilities. In essence, by following the manufacturers recommendations, including PRC and TMR, it is possible to eliminate the effects of configuration upsets and user logic upsets on device functionality.

Xilinx recommends certain techniques for using TMR throughout the VHDL design [3], and ITT plans to follow these generalized guidelines. These guidelines account for the use of general logic structures as well as the use of the Virtex FPGAs' special architectural features (e.g., BlockRAM, DLLs, and Arithmetics). However, these generalized guidelines recommend not using the FPGA's look-up tables (LUTs) for random access memory (RAM) functions, because PRC scrubbing in the Virtex FPGA will overwrite the dynamic bit values stored in the LUTs. Since the LPT design uses the FPGA's LUT blocks for RAM functions in several instances to improve design efficiency, it is necessary to develop a mitigation technique implementation that allows ITT the continued use of the LUTs as RAM.

The initial values of the LUTs are stored in the FPGA's configuration bitstream, so periodic PRC scrubbing of the FPGA will periodically overwrite the dynamic bit values stored in a LUT with their initial values. This is why Xilinx discourages the use of LUTs for dynamic RAM (LUT-RAM) and shift registers (LUT-SRL16) in a fault-tolerant FPGA design.

However, after a thorough analysis, ITT determined that LUT-RAM (and LUT-SRL16) can be used in a fault-tolerant design if the three redundant LUT-RAM modules are placed in independent columns of the configuration bitstream, as shown in Fig. 4. By spacing the redundant LUT-RAM modules throughout the configuration bitstream, the bit errors

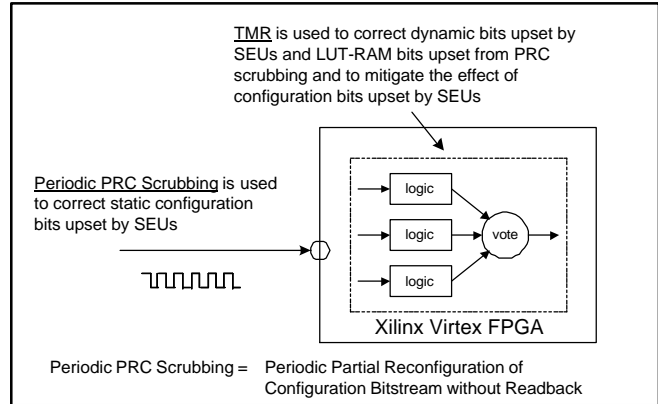


Fig. 3. High-Level Approach to SEU Mitigation

caused by partial reconfiguration in one of the LUT-RAM modules can be corrected with the other two redundant modules as long as they vote before the second module is corrupted from the PRC scrub.

Therefore, if the LUT-RAM bits are voted and refreshed at an interval that is greater than the shortest time between partial reconfigurations of any two redundant LUT-RAMs, the three redundant LUT-RAMs should be read, voted, and rewritten with the voted values before a second LUT-RAM is corrupted with PRC.

Using the approaches described above, the bit errors introduced through partial reconfiguration scrubbing will be corrected. However, the use of LUT-RAM opens the door to the chance that a SEU will not be corrected with TMR. Such an event would occur if a SEU occurs in a LUT-RAM at the same time that partial reconfiguration scrubbing corrupts the bits of one of the redundant LUT-RAMs. In this case, TMR voting will not correct both the SEU corrupted bit and the PRC corrupted bit. Fortunately, this is a very unlikely event. An analysis of the LPT design showed that the probability of such an event is very small—about one such error every 1473 years in a geostationary orbit.

Using this result and other lessons learned, ITT Industries has developed a set of guidelines that include general and LPT-specific design approaches for adding redundancy to the VHDL code and for implementing the PRC scrubbing. The design guidelines will continue to evolve as they are applied to the LPT architecture.

B. Radiation-Tolerant DSP

The third generation LPT Digital module uses a 400 MIPS Texas Instruments TMS320C5510 DSP. Unfortunately, there is not a radiation-tolerant version of this COTS DSP currently available, so an appropriate replacement processor must be used for applications that require a higher level of radiation hardness. In an effort to reduce the required DSP processing

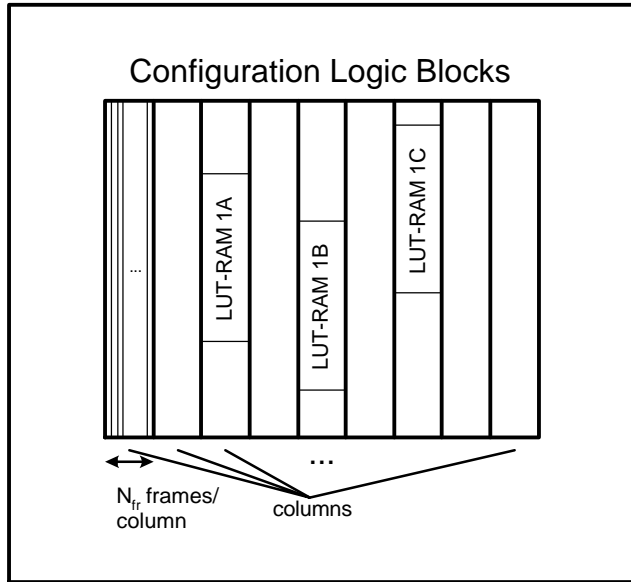


Fig. 4. Placement of LUT-RAM within FPGA's Configuration Bitstream

power for the radiation-tolerant DSP solution, many of the DSP's functions were moved into the FPGA.

The ideal radiation-tolerant processor is one that meets the LPT processing requirements and is low power, readily available, well supported with code development tools, and sized so that it can fit on the third generation Digital module. Both single-chip and soft-core solutions were considered for this application. Atmel's single-chip TSC21020F DSP is a functionally equivalent processor to the Analog Devices 21020 DSP. This processor provides 32-bit floating point processing capability and meets almost all of the requirements stated above, except that the DSP and its supporting memory will not fit on the third generation Digital module. This solution requires the development a new LPT module, named the "Rad-Hard DSP module." This new module, which will support the TSC21020F and all of its required program and data memory, will provide a plug-in solution for applications that require the radiation-tolerant DSP.

In the future, it is desired for the DSP to be embedded into the FPGA, but at this time, the FPGA technology does not contain sufficient resources to allow this approach within the LPT. The use of soft processor cores is attractive, since it moves the design one step closer to being a system-on-a-chip, and prepares the LPT for its future evolution.

V. THE FUTURE OF LPT

The generational evolution of LPT is far from over. ITT has already described the vision for a so-called fourth generation LPT, which will be termed the Miniature Multifunctional Integrated Terminal (M^2 inT). This new

generation will evolve the LPT packaging system into a miniature form factor while preserving or improving upon the core LPT capabilities. In its own right, the third generation LPT platform is a compact device measuring approximately 100 cubic inches in volume. However, the M^2 inT will strive to be only 8 cubic inches in volume—a cube measuring approximately two inches on a side. This will be accomplished by following the same formula used in the earlier generations of LPT—exploit latest-generation, commercially available technology while being mindful of the environmental requirements of space applications. This new, light-weight form factor will help revolutionize spacecraft design by allowing the transceiver to be placed near or inside antenna structures, virtually eliminating cable losses that plague existing spacecraft and limit the bandwidth available for science. Additionally, it will act as an enabler for a new generation of nano-satellites whose entire mass is less than a conventional transceiver/transponder.

New packaging schemes will be developed that expand upon the modular foundation of the first three generations of LPT, improving bottlenecks and trouble spots inherent to the existing design. Printed circuit boards will most likely be assembled using chip-on-board techniques in order to eliminate the size restriction imposed by chip packages. Likewise, innovative and ultra-high-density signal routing will be utilized in order to provide a dramatically improved bandwidth potential, allowing the new form factor to take advantage of future advancements in signal processing speed and density.

VI. SUMMARY

The successful conversion of the LPT to a functionally equivalent rLPT is critical for many space applications. Although this task is challenging, the recent development of radiation-tolerant COTS devices will facilitate this effort. ITT Industries plans to apply these new devices and technologies to the development of the rLPT. Furthermore, the M^2 inT promises to be a breakthrough technology for the emerging fleet of nano-satellites currently being planned.

ACKNOWLEDGMENTS

The authors would like to acknowledge Adesh Singhal, Bob Bauer, Steve Smith, and Mike Pasciuto from NASA's Earth Science Technology Office, Dave Israel and Roger Flaherty from NASA's Goddard Space Flight Center, and John Rush from NASA Headquarters for their support of the ongoing LPT development effort.

REFERENCES

- [1] "Exploring Our Home Planet, Earth Science Enterprise, Strategic Plan" NASA Earth Science Enterprise, Code Y, November 2000.

- [2] E. Fuller, M. Caffrey, A. Salazar, C. Carmichael, J. Fabula, "Radiation Testing Update, SEU Mitigation, and Availability Analysis of the Virtex FPGA for Space Reconfigurable Computing," MAPLD International Conference 2000, September 2000.
- [3] Carmichael, Carl, "Triple Module Redundancy Design Techniques for Virtex FPGAs", Xilinx Application Note XAPP197 (v1.0), November 2001.