



Meeting the Demands of Robotic Space Applications with CompactPCI

The robotic tasks in manned and unmanned space applications need increasing sophistication, intelligence and autonomy, as well as radiation tolerance. CompactPCI-based hardware can provide superior performance under demanding conditions.

By Anthony Lai, Aitech Defense Systems

As space agencies around the world undertake more aggressive exploration missions, each space application requires increasing sophistication, intelligence and autonomy. Autonomous applications are performed mostly by sophisticated robotics through the use of precision control of mechanisms. High-bandwidth sensor and vision systems provide robots with intelligence, and data is processed with computation-intensive software algorithms. An obvious solution for meeting the software requirements of such tasks is the use of real-time, multi-tasking, embedded applications. But the complementary hardware used to fulfill those applications requires high levels of computing performance and the ability to operate in harsh orbital or terrestrial space environments.

Function	Requirement
Docking	Use of additional vision and ranging payloads to control the movement of one object toward another, demanding complex software and uninterruptible operation.
Autonomy	Non-parallel computing capabilities are used to handle multi-tasking software with artificial intelligence.
Precision Movement	Active control of mechanisms in real time requires high-data-rate sensor inputs.
Formation	Multiple instances of objects with reconfigurable software and hardware for executing a task.
Teaming	Requires processor module to be usable in multiple roles while providing capability to handle complex tasks based on situational awareness.
Servicing	Processor module must be able to survive and operate through severe launch and re-entry environments for multiple planets or multiple launch and re-entry phases to/from the same planet.

Table 1 The robotic tasks involved in space missions require different hardware capabilities, depending on the type of function needed.

Space Missions' Functional Requirements

Space exploration is typically a multi-step process, often involving initial remote-controlled investigation, subsequent human-controlled missions and additional servicing to sustain human missions. In each of these phases, robotics plays a key role in realizing a successful mission. The versatility of robotics is driven mainly by the addition of remote sensing inputs to control software executing on a high-end, space-ready processor. For example, a payload designed to provide range data requires heavy computation, processed in real-time, in order for the inputs to be useful for any robotic task. Typically, these tasks may include docking between two space-borne or terrestrial objects (such as a space vehicle and a transfer module, or a rover and a depot); autonomous maneuvering of spacecraft and rovers; precision pointing or tracking of communication systems; formation flying or teaming of multiple spacecraft, or driving of multiple rovers; and remote servicing of one spacecraft by another, or of one rover by another. Successful execution of these tasks depends upon unique software and hardware capabilities designed to accommodate the demands of

specific functions.

Hardware Feature	Capability
Microprocessor	Available PowerPC processors with inherent radiation-tolerant capabilities provide superior performance and low power consumption for space mission applications. ECC and parity-protected on-die L1 and L2 cache to provide mitigation to Single Event Effects.
Volatile Memory	SDRAM arranged in a triple-voting architecture (three bits per cell) ensures that the processor always receives the best available data and instructions.
Non-Volatile Boot Memory	Dual-redundant boot flash creates alternate paths to start a flight application in the case of corrupted firmware and/or application images.
Non-Volatile User Memory	ECC-corrected user flash memory provides real-time correction and detection to data corrupted by single event effects.
I/O Expansion	Conduction-cooled PCI mezzanine card (CPMC) slot delivers flexible and modular I/O expansion.
Bus Interface	Incorporating a standard bus interface ensures capabilities for additional redundancy and for integrating signals from outside the COTS enclosure.

Table 2 Different combinations of hardware features and capabilities are needed in a COTS 3U CompactPCI processor module used in a robotics space application.

Hardware Requirements

Satisfying the functional requirements of these various robotic tasks mandates specific hardware capabilities (Table 1). To achieve the cost, compatibility and lead-time advantages of a COTS approach, those capabilities must be accommodated within industry-standard form-factors. For example, in order to provide superior performance under demanding conditions, and survive the hostile space environment, a processor module designed in a conduction-cooled 3U CompactPCI form-factor might require different combinations of hardware features. These might include a microprocessor, volatile and non-volatile memory, provision for I/O expansion and certain bus interfaces (Table 2).

Radiation-Tolerant Processors for Space Modules

Uninterruptible operations are important in the computations performed for critical events in robotic tasks, such as docking, spacecraft maneuvering and pointing or scanning to maintain communication links. In order to ensure reliability in the space radiation environment, radiation-tolerant devices are critical for multiple system functions. These include the PowerPC system controller (with PCI bridge and memory controllers); user-programmable timers or counters; a safety watchdog management subsystem to an external radiation-hardened watchdog supervisor; reset mechanisms; a CompactPCI bridge; and all mitigation schemes.

A high-performance, low-power, silicon-on-insulator (SOI) PowerPC microprocessor has all the necessary attributes to offer unparalleled performance with a throughput of more than 1500 Dhrystone MIPs. In addition to the inherent radiation-hardness of the SOI process, the 750 PowerPC family of microprocessors built on this process has been tested to meet and exceed all levels of radiation hardness. Furthermore, its inherited latchup-immune feature provides mission-critical subsystems with uninterrupted operation throughout the mission and over an extended period of time.

But in order to maintain throughput exceeding 1500 Dhrystone MIPS, the processor's internal L1 and L2 cache memory must be enabled. Both L1 and L2 cache reside on the same SOI die as the processor, so they inherently have the same radiation characteristics as the processor does. The internal L2 cache has a capacity of 512 Kbytes and utilizes an 8-bit error check and correct (ECC) for every 64-bit word in memory to correct a majority of single-bit errors and to detect multiple-bit errors. The L2 cache tags also support parity and by-way locking. The L1 cache has 32 Kbytes of instruction cache and 32 Kbytes of data cache. Both types of cache are 8-way set associative, and the L1 cache tags also support parity.

Using PCI and CompactPCI bus interfaces in the module design makes available an industry standard I/O expansion slot via a PCI mezzanine card (PMC) site, and the access to a modular bus system via a cPCI backplane.

Designing Volatile RAM Memory Subsystems for Radiation Tolerance

Software computation for robotic applications usually involves processing large amounts of data. Therefore, a large amount of volatile memory is needed to enhance efficiency and give engineers more choices of the robotic control algorithms that can be used in various situations. Several techniques can be implemented in volatile memory to enhance the radiation hardness of a space-based processor board, and to ensure reliable mission operations with onboard flight software.

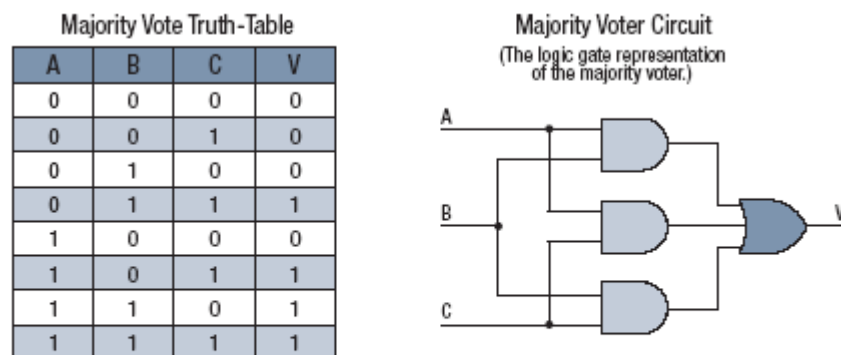


Figure 1 A majority-rule triple-voting mechanism can be used to implement a radiation-tolerant, 3U CompactPCI board's SDRAM volatile memory, thus enabling simultaneous writes to three memory banks in real time, and correcting data reads in real time.

All semiconductors are affected by radiation, including single event effects (SEE) of the processor and the side effects introduced by the instruction and data caches. Memory devices are particularly sensitive to space-based effects. For example, depending on the space environment, random access memories such as SDRAM can be particularly prone to flipping bits or erasure.

A number of strategies can be used to protect onboard volatile memory resources from radiation, as has been shown by proton and heavy-ion testing on various device types. For SDRAM applications, used often for providing instructions and data to the processor, one of the most effective techniques for maintaining data integrity takes advantage of triple redundancy with voting mechanism logic incorporated in a radiation-tolerant FPGA. For radiation hardness and reliability at the component level, the SDRAM controller can be implemented with a majority-rule, triple-voting mechanism in anti-fuse FPGAs, as opposed to SRAM-based FPGAs (Figure 1). Along with three physically separate banks of SDRAM, the volatile memory has been demonstrated to meet high-performance and radiation-tolerant operations. The SRAM-based FPGAs are used for the engineering development units for software development, and the anti-fuse FPGAs are used for the flight units for radiation hardness.

Protecting Non-Volatile Flash Memory

Multiple robotic tasks typically require different software modules and configurations. To maintain multiple boot images of applications and their configuration settings, a reliable, non-volatile user flash memory is needed to provide reconfigurability and adaptability. Protecting the integrity of the extensive firmware utilities stored in boot flash is essential to a successful mission operation.

As in the case of protecting non-volatile memory, boot flash capabilities can also be protected by dual redundancy—that is, two independent banks of boot flash—in combination with a watchdog mechanism. The watchdog mechanism comprises software-independent radiation-hardened circuitry with periodic service generated by the firmware or the flight software application. Providing redundancy in this way enhances the opportunity to boot successfully after initiated or environment-induced resets, such as a software reset, a power cycle reset or a single event functional interrupt.

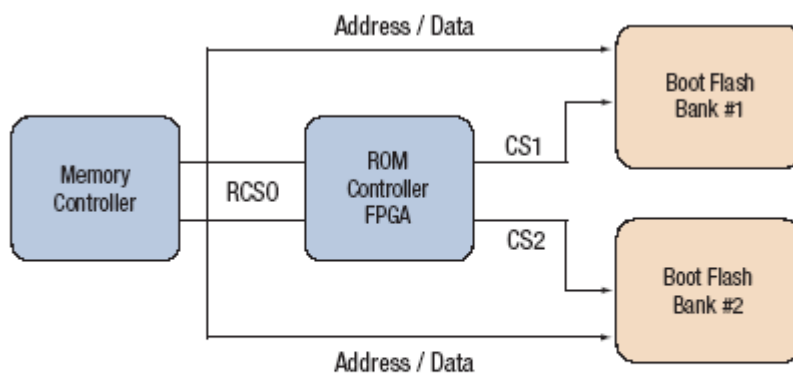



Figure 2 Dual-redundant boot flash memory can be controlled by a radiation-hardened watchdog supervisor function. The redundant flash is used to boot up if the watchdog expires without an initial service from the boot-up firmware and subsequent service by the flight application.

In this scheme, a defective boot flash can be overwritten by the contents of the intact boot flash, providing two identical copies for future startup operations. A record of the number of resets required to provide a successful reset and the bank of flash to boot up after the last reset can be maintained in separate hardware registers. For further assurance against soft errors, the control logic to implement the watchdog timer and the overwritten operations (Figure 2) are also implemented in an anti-fuse FPGA.

In addition to the boot flash, other non-volatile user memory is typically required to store the user's application along with data, such as static data tables or digital filter coefficients. Reliable NOR flash is often used to deliver the optimum performance for such random access scenarios. The NOR flash can be further enhanced with an ECC algorithm integrated with the flash memory controller as part of an anti-fuse FPGA. For example, a user's 32-bit data can be programmed into the two data devices, with the ECC syndrome calculated and stored in a third flash device (Figure 3). In this scenario, when the processor initiates a read request to the user flash, a CRC checksum is calculated and compared against the stored value. If the two checksums are different, a single-bit correction will be attempted or the flight software will be notified with a multi-bit error in the flash. Similar to the intent of dual redundant boot flash, the ECC-protected user flash is also designed to mitigate SEE.






Figure 3 User flash controller calculates and writes ECC syndromes to a third flash bank while data is written in 32-bit mode to the other two user flash banks. The ECC mechanism automatically detects multi-bit errors and corrects single-bit errors.

Accommodating Other System-Level Functions

As mentioned previously, radiation-tolerant systems include more than just the CPU and memory subsystems. Local expansion slots maintain the same system processor as that of a building block, while enabling additional mission- or application-specific I/O interfaces to be inserted into a single-slot solution. This can be an especially important feature when a robotic system possesses a high-speed I/O interface where high bandwidth traffic can be localized to a dedicated bus on the processor, instead of to an external bus system. A bus interface that allows multiple system processors to perform redundant or different tasks is also desirable for facilitating communication among other cards in an open architecture.

Furthermore, a bus system provides access for additional I/O functions that simply won't fit into a standard daughter card, such as a motor controller. With its high current and larger space-qualified components, it is typically too big to fit any industry standard I/O card form-factor.

A Rad-Tolerant COTS Board Solution

Several key hardware and system design elements are required to achieve a wide variety of robotic tasks for unmanned and manned space exploration. While many of the mitigation techniques discussed have been introduced in the past as individual products, they are now also available as integrated radiation-hardened solutions in a single board, such as the Aitech S950 (Figure 4).




Figure 4 The S950 radiation-tolerant board provides a COTS.

... provides a cost-
compatible solution via its
PowerPC processor and
standard 3U conduction-cooled
CompactPCI form-factor with a
conduction-cooled processor
mezzanine card expansion slot.

The incorporation of dual FPGA footprints for various components in the design produces an SBC with an engineering unit that is a form, fit and functional equivalent to the flight unit. The engineering design units allow for rapid prototype development of space missions, while maintaining software compatibility for the flight configuration. If the processor card and related I/O modules are offered in multiple radiation-tolerant levels for various flight configurations, this can accommodate engineers' environmental and operational requirements, as well as aggressive cost goals for space and terrestrial orbiting applications.

Aitech Defense Systems
Chatsworth, CA.
(888) 248-3248.
[www.rugged.com].

© 2005 RTC Group, Inc., 905 Calle Amanecer, Suite 250, San Clemente, CA 92673