UCRL-JC-107525
PREPRINT

Operation of Commercial R3000 Processors in the Low Earth Orbit (LEO) Space Environment

J. L. Kaschmitter
D. L. Shaeffer
N. J. Colella
C. L. McKnett*
P. G. Coakley*

This paper was presented at the
IEEE Nuclear and Space Radiation Effects Conference
San Diego, CA, July 15-19, 1991
and submitted for publication in *IEEE Transactions*, December 1991

August 9, 1990

Lawrence
Livermore
National
Laboratory

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

* JAYCOR Inc.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Operation of Commercial R3000 Processors in the Low Earth Orbit (LEO) Space Environment*

J. L. Kaschmitter, D. L. Shaeffer, and N. J. Colella
Lawrence Livermore National Laboratory
Livermore, CA 94550
and
C. L. McKnett and P. G. Coakley
JAYCOR Inc.
Santa Monica, CA

UCRL-JC--107525

DE92 006894

Abstract

Spacecraft processors must operate with minimal degradation of performance in the Low Earth Orbit (LEO) radiation environment, which includes the effects of total accumulated ionizing dose and Single Event Phenomena (SEP) caused by protons and cosmic rays. Commercially available microprocessors can offer a number of advantages relative to radiation-hardened devices, including lower cost, reduced development and procurement time, extensive software support, higher density and performance. However, commercially available systems are not normally designed to tolerate effects induced by the LEO environment.

Lawrence Livermore National Laboratory (LLNL) and others have extensively tested the MIPS R3000 Reduced Instruction Set Computer (RISC) microprocessor family for operation in LEO environments. We have characterized total dose and SEP effects for altitudes and inclinations of interest to systems operating in LEO, and we postulate techniques for detection and alleviation of SEP effects based on experimental results.

Introduction

A recent trend for satellite systems operating in the LEO environment is to increase the number of platforms in a constellation while reducing the cost and complexity of each platform. This has the effect of making the constellation more reliable by reducing its vulnerability to the failure of a single platform.

The proliferation of systems in LEO mandates cost reduction of individual platforms where possible. The market for processors radiation-hardened to meet requirements for dedicated operation in the LEO natural space radiation environment is miniscule compared to the total commercial market. This—and the time and effort required to produce specialized radiation-hardened processors—cause these parts to be very expensive, and they frequently lag the state-of-the-art in capability. Restricting costs and maintaining technological performance edges compel satellite system engineers to accept the use of commercially available parts without modification.

Meeting mission requirements with low cost satellites typically mandates elimination of redundancy for key subsystems and acceptance of reduced system reliability. Since these satellites may be less tolerant of single-point failures and typically rely on the processor for instantaneous attitude and navigation, they must be robust against the introduction of erroneous data or radiation-induced upsets to their operation.

Designers using commercial parts in LEO applications must exercise caution in qualifying parts to ensure that integrated circuit (IC) fabrication process changes have not degraded the effective radiation hardness of the parts. This problem can be obviated by performing lot-sample qualification of parts and inventorying sufficient parts to cover anticipated usage. This is feasible because commercial parts are much cheaper to buy, particularly in quantity, and satellite program parts volume will not require captive production runs.

Design Philosophy

Two effects dominate consideration in applying microprocessors in the low earth space environment: SEP and total ionizing radiation dose over the life of the mission. SEP includes upset due to protons and ions, and latchup, burnout, and gate rupture due to dielectric breakdown in CMOS junctions by ionizing interactions with charged particles.

The impact of these effects on the design process are summarized in Fig. 1. Alleviation of potential catastrophic failure due to high total dose mandates initial specification of shielding mass based on LEO orbit and mission life. Catastrophic failure due to total ionizing dose and single-event burnout and gate rupture must be prevented. Burnout and gate rupture can be prevented only by component selection. Total ionizing dose can be handled by parts selection, parts derating, and shielding. Once the maximum shielding sectional density has been fixed, further mitigation of latchup and single event upset (SEU) rates can be achieved by design techniques. This approach addresses a wide range of severity in environmental effects. Variances of one or two orders of magnitude in environmental effects will not cause a system failure—but at worst—will cause limited performance degradation.

Recent process developments in integrated circuit manufacturing continue to improve the prospects for use of commercial parts in the LEO environment. Fabrication of circuits on epitaxial layers has enhanced total dose tolerance substantially, and the increasing use of insulating substrates (SIMOX, SOS, or SOI) for very small feature sizes (<0.3 μ) will significantly improve upset performance.

* This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract no. W-7405-Eng-48.

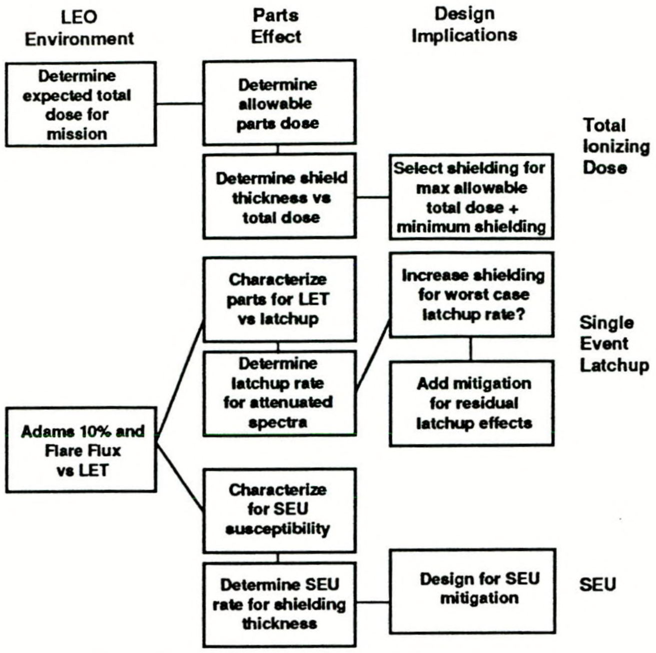
LEO Environment
Parts Effect
Design Implications
Determine expected total dose for mission
Determine allowable parts dose
Total Ionizing Dose
Determine shield thickness vs total dose
Select shielding for max allowable total dose + minimum shielding
Characterize parts for LET vs latchup
Increase shielding for worst case latchup rate?
Single Event Latchup
Determine latchup rate for attenuated spectra
Add mitigation for residual latchup effects
Adams 10% and Flare Flux vs LET
Characterize for SEU susceptibility
Determine SEU rate for shielding thickness
Design for SEU mitigation
SEU

Figure 1. Design process for use of commercial parts in LEO.

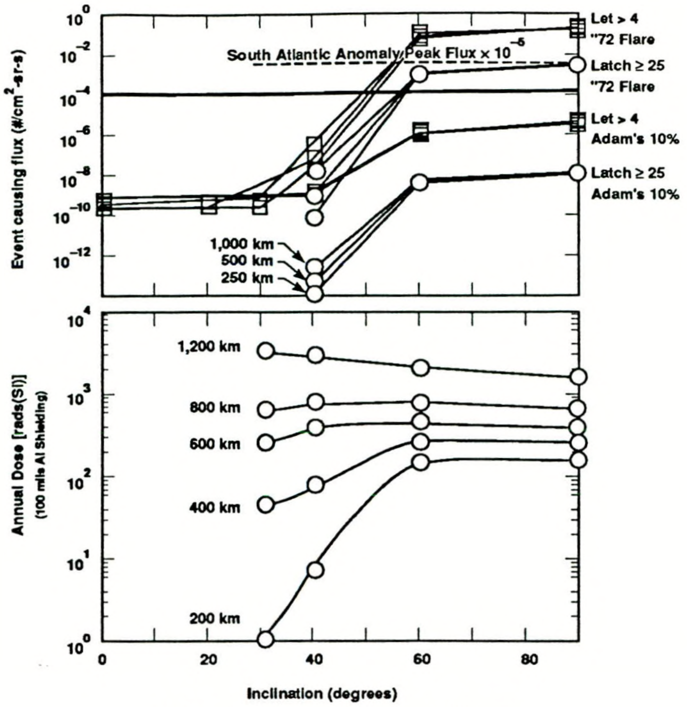

Figure 2 . LEO environments [1,2,5] (100 mils Al, 2π semi-infinite shielding).

The LEO Environment

Because of the influence of the earth's natural magnetic shielding, the severity of the LEO environment can vary by orders of magnitude, depending upon orbital inclination and altitude. Substantial differences also occur during periods of solar flare activity. Effects on electronics are often discussed for two levels which describe flare activity—flare maximum and Adams 10% worst-case. Normal, non-flare environments will be sized by the proton flux of the South Atlantic Anomaly. Figure 2 [1,2,5] summarizes SEP effects and total ionizing dose for a range of LEO orbits. LET values greater than 25 are used in a later section to predict the frequency of latchup.

Flare maximum estimates frequently reference energy levels and distributions for the 1972 King flare, which has been extensively analyzed and is frequently used as a relative worst-case environment. Preliminary reports indicate that the large 1989 flare was roughly comparable and that our conclusions would be valid for this environment as well.

Choice of an appropriate environment for performing system design trades is problematic in that the flare environment contains radiation effects many orders of magnitude more severe than the normal environment, but these effects occur only a very small percentage of the time. Flare activity tracks the 11-year solar cycle, during which actual flares may occur for periods ranging over several days.

The system designer must determine a tolerable rate for Single Event Phenomena by considering system performance requirements during flares and the mission lifetimes, and the risk to the mission if an SEP goes undetected or unmitigated. Factors which may influence this decision include the ability to restart under ground control, the acceptability of degraded operation during flares, and possible catastrophic consequences of upset—such as unintended attitude adjustments, motor firings, destruction of sensitive instruments by improper pointing (e.g., at the sun). Our analysis assumes that the system will operate through a major flare event, albeit with a minimal performance impact (≤10% of the nominal processing speed).

Radiation Effects on Commercial R3000/R3010 Microprocessors

We have designed and built an operational high-performance processor based on the MIPS R3000 RISC architecture. Reference [3] provides details of this design, which includes 32 K words of instruction and data cache memory, 512 K words of 80 ns main memory, 2 CMOS gate arrays for control and I/O, and a small amount of miscellaneous logic. The cache memory uses 25 ns 16 K × 4 CMOS high-speed SRAM; the main memory is based on 128 K × 8 CMOS SRAM. The computer is designed for multiprocessor operation and can be implemented with space-qualified high-density packaging [4].

We have tested both component parts as well as the completed, operational processor under various radiation effects. We propose a design methodology addressing those effects using a combination of shielding and mitigation techniques following the approach outlined in Fig. 1.

Total Dose

We have performed total dose testing (unpublished work) of key components of the R3000-based processor, including the Hitachi 1 Mbit SRAM and R3000/R3010 parts from various manufacturers. Results were obtained for Cobalt 60 (SRAM only) and for high-energy protons [7,8]. Absolute worst-case commercial part tolerance to total dose is greater than 1000 rads, so we have sized for this value (although some devices demonstrated hardness in excess of 100,000 rads).

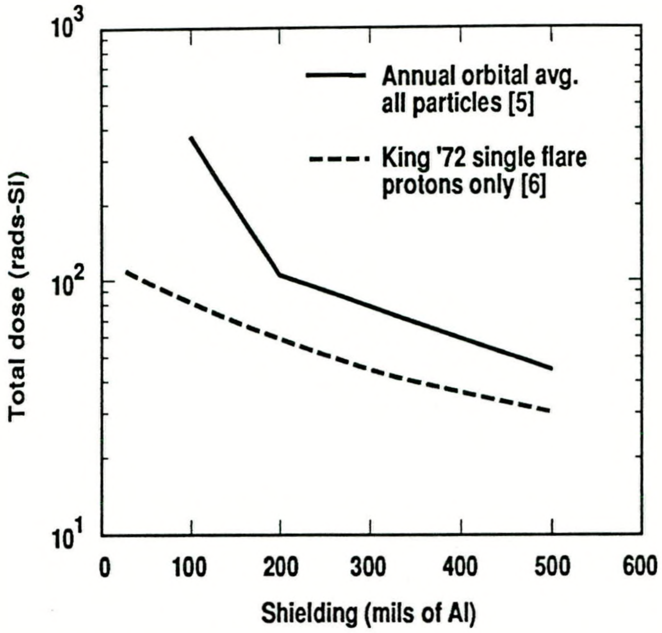

Figure 3. Total dose (rad-Si) vs shielding, 500 km 60° orbit.

Figure 3 details the total ionizing dose received for different aluminum shielding thicknesses at a nominal 500-km, 60° orbit and for both the average annual orbital natural environment and a single King 1972 event. For a 5-year mission at high inclinations, a total shielding of 200 mils of aluminum will limit the proton total dose to less than 1000 rads and prevent failure of any component within the processor. Assuming the processor is packaged in a 5-cm × 5-cm × 0.7-cm high-density multichip module (MCM) and that there are 100 mils of shielding inherent in the spacecraft structure and skin, the processor module weight would be less than 50 grams.

Single Event Upset

We and others [7–10] have performed multiple tests for single-event upsets on this processor and the R3000 family of parts for various proton energies and ion LETs. Figure 4 shows the calculated upset rate for the processor as a function of orbital inclination (at 500 km) and shielding thickness of aluminum. Results shown here are based on cross sections data measured for 256 MeV incident protons, [7] and heavy ion data [10], assuming single-bit error-correction in main memory. Errors in the cache memory are assumed not to be an issue because the cache memory incorporates parity error detection and will detect SEU errors and rewrite the cache block from main memory. At 60° inclination and the nominal 200 mils of shielding selected to reduce total dose, we will experience approximately 1 upset/min during a flare equivalent to Adam's worst case composition of the 1972 King flare [1]. SEU recovery (see software section below) will be less than 100 msec, so this error rate is not a problem. However, the performance degradation is substantial, so we elect to increase shielding to 300 mils, or 1 upset every few minutes. This will increase our processor MCM weight to less than 100 grams.

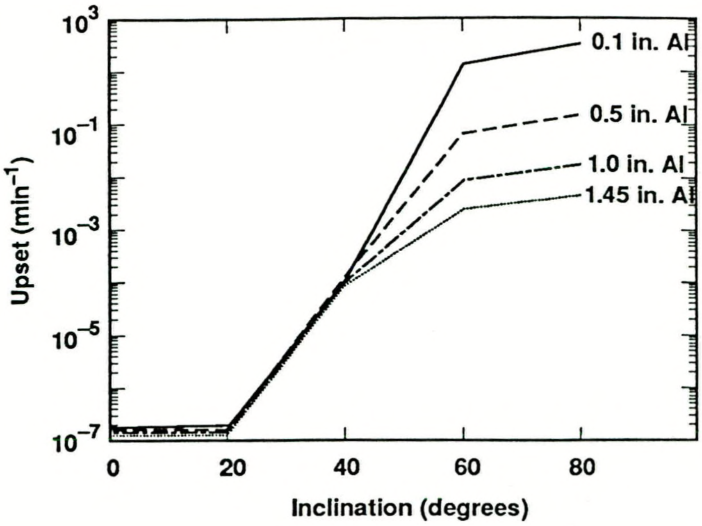

Figure 4. R3000-based processor upset rate vs inclination. King 1972 flare with trapped protons (500 km).

Single-Event Latchup

Single-Event Latchup thresholds were also measured for the R3000/R3010. The effective latchup rate can be roughly determined by examining the flux of particles with LET > 25. This can be determined from Fig. 1 because large CMOS parts—such as the R3000—are roughly 1 cm^2. Assuming that the R3000 bits have a sensitive volume $2 \times 2 \times 2$ μm^3 and an upper bound of 100,000 sensitive bits (or a total sensitive cross section of 4×10^{-3} cm^2 in agreement with reference 10), this would mean that the R3000 part would see approximately one latchup every few days during a peak event such as the 1972 flare. If accumulated over all parts in the processor this is a high enough rate to require mitigation.

Single-Event Mitigation Techniques

SEU can be detected and corrected in the main memory by the use of a single-bit correction, double-bit detection error detection and correction (EDAC) code on the the main memory, and memory scrubbing at a sufficiently high rate to prevent occurrence of double-bit errors. Multiple copies of critical recovery information must also be kept in the event an SEU occurs while writing memory. Errors in the cache are detected by parity, and cache is marked invalid and reloaded from main memory when this happens.

Figure 5 shows memory scrub times (access every memory location and correct any single-bit errors) required for operation during a peak 1972 flare for the Hitachi 1 Mbit SRAMs, for a 500-km, 60° orbit. Mean time to failure is the time after which a correctable single bit-error will become an uncorrectable double-bit error (assuming a nominal amount of shielding). Acceptable error rates are determined by mission duration and risk tolerance. Multiyear missions can require failure rates <1/10^8 s, which can mandate scrub times of every 0.1 s and can have an impact on system performance if not implemented transparently in hardware. Alternatives are to put critical orbital maintenance code/data in an SEU-hardened, lower-density SRAM.

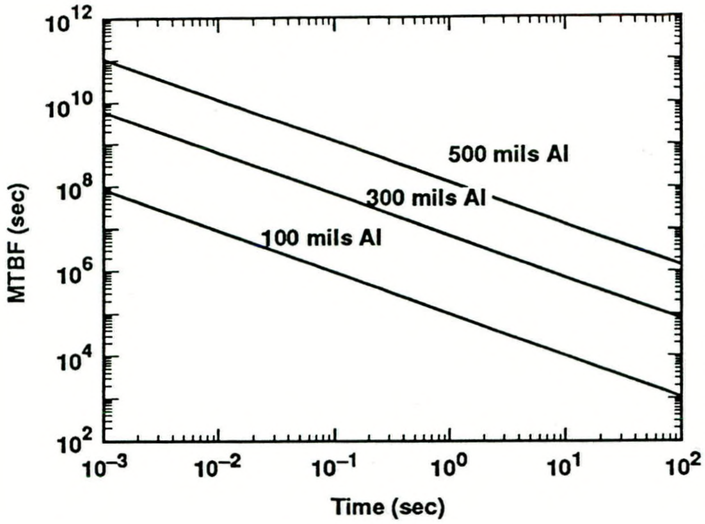

Figure 5. Adam's worst-case '72 flare; 500 km, 60" orbit. Memory scrub time.

This may be a reasonable alternative if such memory is required to deal with restart after a latchup.

Table 1 contrasts different system approaches to SEU-handling in the processor itself; these are based on impacts on system software, size, power, weight, performance, and fault coverage. In state-of-the-art systems, software development often dominates program effort, risk, and complexity. Reducing complexity in the software can have a major beneficial impact on program success. All approaches—unless reliant on a zero or very low upset rate realizable by rad-hard devices or by multiple, redundant CPUs (which necessitate voting and control handling)—require software checkpointing so that, in the event of an SEU, recovery can be performed quickly and without significant loss of state.

Increases in size and weight caused by the use of redundant processor chips can be ameliorated by employing state-of-the-art high-density packaging. In the case of multiple CPUs or computers, increases in power can be offset by powering down hardware during noncritical calculations.

High fault coverage, ease of programming, minimum cost, and hardware impact are best achieved by operating a pair of processors simultaneously, executing identical code, and comparing the outputs ("lockstepping") to detect upsets. Correction can be handled without major complications to the task of writing application software by saving and restoring state in a "warm" start, as described below.

We have examined techniques for lockstepping the R3000. Although the R3000 is amenable to lockstepping, it is not the ideal candidate. Other architectures designed for fault-tolerant transaction processing and with better power/performance ratios may be better suited; these are the subject of on-going investigations.

Other areas of the R3000 which require careful consideration are the translation lookaside buffer (tlb) and configuration registers. Both of these structures are relatively static. Consequently, errors in these structures may not be detected quickly enough by some of the techniques listed in Table 1. This is another strong motivation for using a lockstepping approach: tlb and configuration errors will show up as soon as any operation using these structures attempts execution. Operations which rely on self-checking may proceed through several correct mathematical calculations before encountering damaged data structures.

Processor Operation and Design

Hardware

Lockstepped operation of the R3000 requires that two additional functions be designed into the computer: (1) clock synchronization and (2) output compare, write inhibit, and restart control. These functions are best accomplished by implementing logic external to the R3000/R3010 chips due to the complexity and cost of modifying the basic chip designs.

A proposed clock synchronization circuit design is shown in Fig. 6. The SYSOUT clocks are first tested at power up to determine if they are in or out of phase with each other. If they are out of phase, the phases are aligned by halting the 2X clocks to the slave CPU for one 2X clock cycle. After basic phase alignment, the clocks are kept in phase by using a phase comparator to adjust the input phase of the 2X clock to the slave CPU.

We have identified a scheme for lockstepping the R3000 which is relatively easy to implement and minimizes design effort, schedule time, and cost. We propose to incorporate the compare logic with the R3020 write buffer function [11]. Address and data busses are routed through the write buffer, so

Table 1. Estimated impacts of SEU mitigation techniques.

SEU mitigation	Software impact	Performance degradation	Physical/cost	Fault coverage
Lockstep CPUs (2)	Checkpointing required	None or small [a]	Two CPU chip sets Moderate cost	Very high
Redundant CPUs (>2)	Voting and control handling	None or small	More H/W and cost control H/W	Very high
Cross checking	Frequent S/W cross checks	Slight	Two independent computers Moderate cost	Uncertain
Rad-Hard process	Possibly none or slight	Small	More chips. Extreme costs. Custom design/ process	High
Software checking	Extreme, bug prone	Moderate	None. Long S/W dev.	Uncertain (poor)
Hardware checking	Checkpointing only	Small to moderate	More gates/chips. High costs. Custom design	Fair to good

[a] Little or no impact if implemented in CPU design, on chip.

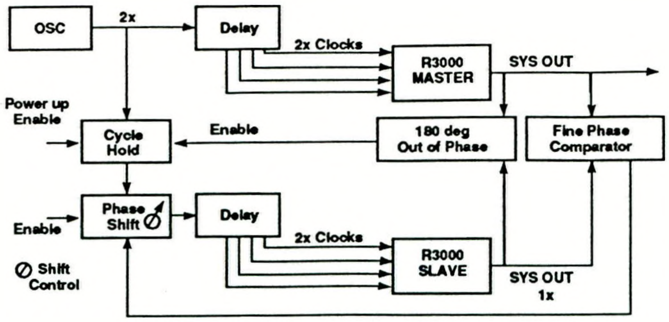

Figure 6. R3000 lockstep clock synchronization.

these lines already exist there. In addition, the R3020 is a relatively simple part (<3000 gates) which already exists as a gate array macro; thus, it can be easily incorporated with the compare function. Also, writes to main memory can be easily inhibited in this part, on chip, without off-chip timing penalty.

The output compare and restart logic would include the functions associated with the R3020 write buffer since all accesses to I/O and memory are accomplished through the buffers. Versions of the R3020 have been implemented in gate array logic, and one approach is to add features for SEU detection to this chip. Upon detection of a miscompare, the logic will generate a high priority restart to the processor, causing a "warm" boot, as described below. Figure 7 is a block diagram of our approach.

A small possibility exists for writing bad data to memory through an SEU occurring on the output of the write path after the checking circuitry. If the compare is continued throughout the write operation, this bad write condition can be detected. This issue is addressed by maintaining multiple copies of critical data.

Increased power consumption incurred by adding the slave CPU can be reduced by powering down the slave during periods of reduced processing activity or in more benign radiation environments (neither flare nor South Atlantic Anomaly). This can be accomplished by holding the slave CPU in a reset state and enabling the master CPU to continue operation.

Software

The SEU mitigation scheme described in the preceding section detects only SEUs. Handling of SEUs in hardware would require three or more processors ("triple modular redundancy"), with voting between processors determining the correct outcome of an instruction. This approach can potentially minimize software overhead, but it exacerbates the power problem.

If two processors are used, handling the detected occurrence of an event must be done in software routines. We have simulated this approach and determined that performance overhead and software-implementation impact are acceptable. The programmer will have to generate structured code which allows for state backup between routines, or within routines if the code affects inter-routine states. This is compatible with our existing coding approach; the complex operating system interactions can be written once and made transparent to the user.

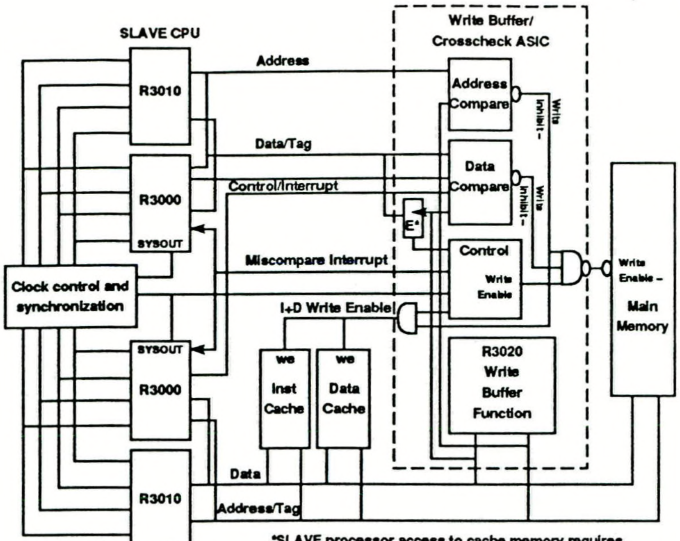

*SLAVE processor access to cache memory requires 18 nsec SRAMS for 33MHz performance.

Figure 7. Block diagram of lockstepped R3000/R3010.

This approach may be characterized as a "backward error recovery with recovery point" technique [12].

To minimize the time required to recover from an event and retain control of the spacecraft, we have designed an approach based on doing a "warm" start instead of a complete "cold" reboot. The advantage is that the system can restore its state and continue operation without having to completely reload memory and reconstruct its status. This implies that the critical memory sections will not be power cycled. Latchup immune memory must be used for the warm start to prevent failure due to SEL. We have identified high density, latchup-immune, commercial memories. Estimates are that a warm start can be easily completed in 10 to 100 ms.

In order to do a warm start after an SEU, the software must back up system state at periodic checkpoints. Types of data which would be logged might include ephemeris, attitude and sun vector, command status, mission data, critical housekeeping—such as battery state and communications status, and state information needed for system restoration. In the event of an SEU, all data within the CPU or FPU would be assumed suspect and would be reconstructed by restoring state and reverting program execution to the most recent checkpoint.

These requirements have implications for system software design and operating system implementation. In particular, the system must save state information in specially implemented structures that are not disturbed by a system restart. Two copies of critical structures will be kept, and state variables within the structure will be defined to indicate structure validity. Preliminary estimates indicate the performance overhead to be less than 10%.

MALLOC (memory allocation) types of operations must also be modified to handle upsets. If memory is left in an indeterminate state following an upset, recovery would not be possible. Memory allocation operations must also update state structures that will allow reconstruction of the state of memory and level of depth of calls.

Writing software using state machines requires careful front-end design to structure the system. This structuring has a beneficial side-effect in providing a well-organized and consequently easily-supported set of code. Once this framework is established and programmers become familiar with it, writing code with checkpointing places minimal burden on the application programmer.

Latchup Circumvention

Quick and reliable detection of a latchup condition is required for circumvention to be effective and to prevent burnout and gate rupture. This can be problematic in the event of partial latchup because current consumption in CMOS logic varies widely with circuit usage and may be difficult to differentiate from conditions of heavy processing. Discrimination can be accomplished by placing the processor in a known processing state and sampling the power-supply current monitor. An algorithm determines the average current and compares the sample to the average. If the sample exceeds the average by a predetermined threshold multiplier, latchup circumvention is initiated. The sample is compared to a long-term average to prevent increases in supply current, caused by total dose effects, from triggering the latchup circumvention logic.

In the event of latchup detection and subsequent power circumvention, the processor will power cycle. A circuit for providing a rapid power down and recovery to latchup-susceptible logic is shown in Fig. 8. This circuit has been built and demonstrated to perform circumvention within 0.4 µs.

Conclusions

Use of commercial parts in processors operating in a Low Earth Orbit environment requires careful consideration of total dose effects and special techniques to deal with Single Event Upset and Latchup. Total dose can be diminished by shielding; the resultant mass penalty is reduced by relying on newer electronic packaging technologies (i.e., multichip modules) to minimize the volume that needs to be shielded. SEU is handled in higher inclination orbits by cross checking between synchronously operating processor chips, with miscompares generating an exception to the software and requiring a reboot and reload of all data in the processor chip(s). Cache memory is protected by parity checking: parity errors force a reload from main memory. Main memory is protected by error checking and correction and regular "scrubbing" to prevent single-bit errors from becoming double-bit errors. Latchup and potential device damage are dealt with by sensitive overcurrent detection and immediate circumvention. All of these techniques can be implemented with commercially available parts, and they provide a high certainty of reliable operation in the LEO environment.

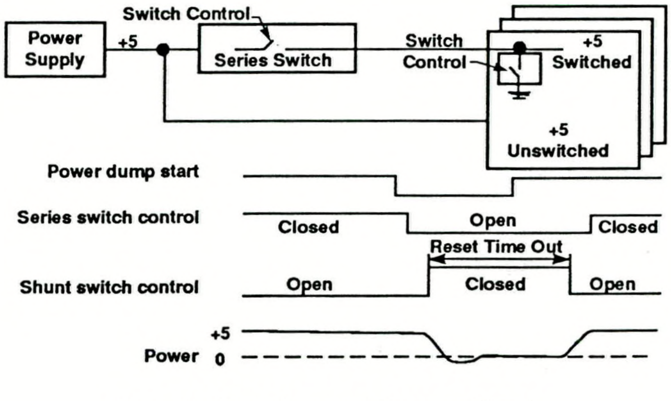

Figure 8. Power Down and Recovery Circuit.

Acknowledgments

The authors thank Dave Butner of LLNL for help in software design for SEU recovery and Joe Kimbrough and other members of the Spacecraft Survivability Group at LLNL for obtaining radiation effects data on commercial R3000 parts.

References

[1] J. H. Adams, Jr., "Cosmic Ray Effects on Microelectronics, Part IV," NRL Memorandum Report, 5901 (1986).

[2] E. Petersen, "Single Event Upsets in Space," Tutorial Notes, IEEE Short Course (17 July, 1983).

[3] J. L. Kaschmitter, F. E. Rubarth, L. H. Capots, and W. S. Scott, "Wafer Scale RISC Processor," Government Microcircuit Applications Conference (GOMAC), Vol. XV, Session 2.6, page 27 (1989).

[4] A. T. Barfknecht, D. B. Tuckerman, J. L. Kaschmitter, and B. M. McWilliams, "Multichip Packaging Technology with Laser-Patterned Interconnects" *Multichip Modules, System Advantages, Major Constructions and Materials Technologies*, pg. 162, IEEE Press (1991).

[5] E. G. Stassinopoulos and J. M. Barth, "Transport and Shielding Analysis of the Non-Equitorial Terrestrial Low Altitude Charged Particle Radiation Environment, Volume 1: Solar Minimum," NASA X-601-84-6, NASA Goddard Space Flight Center (January 1984).

[6] J. R. Letaw, author: *SPACE RADIATION CODE*, Severn Communications Corp., Millersville, MD (1990).

[7] D. L. Shaeffer, N. J. Colella, R. W. Davis, S. M. Denton, J. L. Kaschmitter, J. R. Kimbrough, J. W. Wilburn, and D. Holtkamp, "High Energy Proton SEU Test Results for the Commercially Available R3000 Microprocessor and R3010 Floating Point Unit," submitted for publication in IEEE Transactions on Nuclear Science, Vol 2 (December 1991).

[8] J. S. Browning, "Results of the Proton Beam Tests of the DMS Board Components at LAMPF WNR," Sandia Internal Report (July 9, 1990).

[9] J. H. Adams, Jr., "The Variability of Single Event Upset Rates in the Natural Environment," IEEE Transactions on Nuclear Science, Vol NS-30, No. 6 (December, 1983).

[10] D. Vail, "Estimating the On-Orbit Single Event Upset Behavior of a MIPS R3000 Microprocessor," Harris Corporation Internal Paper (February, 1991).

[11] L. Vainsecher, and M. Gavrielov, LSI Logic Inc., personal communication.

[12] B. Randell, P. Lee, and P. Treleaven "Reliability Issues in Computing System Design," *Comput. Surv.* 10, 2, 123–164 (June 1978).