Rad Hard By Software for Space Multicore Processing

David Bueno, Dave Campagna, Dave Kessler, and Eric Grobelny
Honeywell Inc., Space Electronic Systems
{david.bueno, david.campagna, david.kessler, eric.grobelny}@honeywell.com

Motivation

As tomorrow's space missions continue to demand ever increasing processing requirements, current and future processing engines fabricated via Radiation-Hardened By Process (RHBP) techniques are struggling to keep up in terms of efficiency, performance, and cost. While today`s most powerful radiation-hardened Single Board Computer (SBC) is capable of about 100 MIPS/Watt of performance/power efficiency, many modern software-programmable COTS processors are now capable of upwards of 2 GFLOPS/Watt. The performance and efficiency disparity between COTS processors and RHBP processors is so large that multiple order-of-magnitude performance advantages can be gained by employing a Rad Hard By Software approach on these COTS processors, even if a percentage of throughput must be sacrificed for fault-tolerance.

Rad Hard By Software Techniques

Traditional approaches to using COTS processors in space have relied on fine-grained, instruction-level redundancy techniques such as triple modular redundancy (TMR). Such techniques can be highly effective for operating through worst-case environment scenarios, but those usually represent only a small fraction of the operating life of the system. During the time between processor faults, the 66% (in the case of TMR) redundant processing and the power it consumes is essentially wasted. The wasted power is particularly problematic given that existing commercial high-performance embedded processors typically consume between 10W and 25W each. Being forced to sacrifice an extra 20W to 50W (and dissipate the resulting heat) to ensure correct operation puts TMR solutions out of reach of many applications.

Honeywell's Rad Hard By Software solution brings "elasticity" to the processing capability by recovering from faults at the operation level instead of the instruction level. This approach is conceptually similar to techniques that have been employed in data communications for many years. In a system where data needs to traverse an unreliable link, the standard technique is to break the data stream into a sequence of discrete transactions. In each transaction, a portion of the data is sent and the recipient checks for errors. If an error is detected, the transaction is repeated until it eventually succeeds. Bit errors result in reduced bandwidth and increased latency, but overall data integrity is preserved.

With RHBS, one can think of sequences of instructions as the "data" in this concept and an "unreliable processor" as the "unreliable link." The concept, however, remains largely unchanged. An upset in the processor that is not immediately correctable causes the transaction to be retried until it succeeds. This action results in reduced instruction throughput and increased operation latency, but the overall integrity of the software operation is preserved. Both effects may seem unacceptable at first, but consider that the baseline throughput is at least an order of magnitude greater than what is possible on a rad-hard processor. With such capacity, even a 90% reduction in throughput due to upsets would not put performance below that of a traditional rad-hard solution, and the COTS processor would still provide very significant cost and power advantages.

In situations where upset rates are considerably less than worst case, RHBS provides the benefit of a usable surplus of processing throughput. In cases where processing needs are elastic, this surplus can be used to do more overall work. In cases where processing needs are static, this surplus can be used to "race to idle" in order to reduce overall power consumption.

Defense and science data processing applications with tolerance for elasticity are often an ideal fit for RHBS. Total dose hardness of many COTS processors can be expected to fall in the 100Krad range, which makes them widely applicable to a variety of missions once combined with Rad Hard By Software techniques designed to combat the single event upset (SEU) problem. However, RHBS is not intended to be employed in satellite control or bus control applications where a single manifested upset or missed deadline could result in catastrophic consequences. Applications demanding "strategic" hardness levels and a high level of prompt-dose hardness are also not necessarily a fit for RHBS.

Multicore Dependable Multiprocessing

Over the last 3 years, Honeywell has developed the Dependable Multiprocessor (DM) infrastructure in partnership with NASA/JPL through the Space Technology 8 (ST8) New Millennium Program (NMP) [1]. DM leverages a fault-tolerant software middleware layer in combination with COTS data processing nodes and an RHBP system controller to provide high levels of performance, availability, and reliability for NASA science applications. DM provides a suite of fault-tolerance options for RHBS that may be mixed-and-matched to suit the needs of a particular mission or application. Examples of these techniques include checkpointing/rollback, temporal redundancy, spatial redundancy, and algorithm-based fault-tolerance. In addition, DM provides a comprehensive suite of cluster management tools and a standard, fault-tolerant MPI-based software model familiar to application developers. The DM project is set to conclude in 2008 with a TRL6 validation using COTS PowerPC 7447a processing engines.

While the 7447a processing engines used for the DM TRL6 validation provide substantial performance and efficiency improvements over current RHBP processors, they do not

rival today's most powerful and efficient COTS processors. In this work, we extend the current DM infrastructure to a heterogeneous testbed of multicore processors in Honeywell's Research and Development in Advancement of New Technology (RaDiANT) Lab. Using this testbed, we explore benefits and tradeoffs of mapping DM to multicore processors. Among the architectures to be incorporated in the cluster are the Freescale 8641D and the low-power PA Semi PA6T-1682M processor. Honeywell is also exploring considerations for integrating the Cell Broadband Engine into future DM systems. Of currently available multicore processors, the PA Semi processor is particularly interesting due to its extremely low power (~12W @ 2 GHz), high performance, and availability in a conduction-cooled form factor similar to that of the current DM 7447a data processors.

To evaluate the multicore DM testbed, Honeywell will use a variety of applications across several domains. The National Optical Astronomy Observatory (NOAO) is providing two cosmic ray elimination applications that are expected to strenuously exercise the many fault tolerance modes of the ST8 DM system for its June TRL6 validation. We plan to compare the performance, efficiency, and fault tolerance of these applications on the multicore DM cluster with results obtained from ST8 TRL6 system. In addition, we plan to include parallel SAR and Hyperspectral Imaging applications in the comparison. To provide a reference against the current state-of-the-art in space computing, we will also evaluate the performance of the selected applications on platforms representative of current Rad Hard By Process technology.

References

[1] J. Samson, et al., "NMP ST8 Dependable Multiprocessor (DM)," *Proc. Eleventh Annual Workshop on High-Performance Embedded Computing (HPEC)*, Lexington, MA, September 18-20, 2007.