SURVEY OF PROCESSORS FOR SPACE

Ran Ginosar
Ramon-Chips Ltd., 104 Galil Street, Nofit 36001, Israel
Email: ran@ramon-chips.com

1. INTRODUCTION

Satellites are expensive, computers are unreliable and if failing they will risk the mission. Past space computers were expensive and hard to get, hence the drive to make space computers based on COTS. This paper surveys various approaches to space computers and offers some recommendations.

2. REQUIREMENTS FOR USING PROCESSORS IN SPACE

Processors for space are required to be tolerant to the following radiation and environmental effects: TID (100-300 kRad for GEO and beyond, 10-50 kRad for some LEO missions), Latch-up, SEU, SET, SEFI, and temperature cycles. Another reliability issue is the complete and permanent failure of a processor or a critical sub-component. None of the methods surveyed here mitigate such failures, and the common mitigation method is based on deploying spare computers in combination with either a central reconfiguration circuit or a distributed recovery mechanism.

3. ARCHITECTURE OF PROCESSORS FOR SPACE

Processors for space applications are typically required to achieve the following targets: high performance, low cost, low power dissipation, and reliability. This is sometimes achieved by high integration (for high performance, low cost and low power dissipation) and high tolerance to radiation and environmental effects (for reliability). The problem is that most available processors and integrated systems-on-chip achieve only some of the targets and fail on others. This is indicated below when relative advantages and disadvantages are listed, and exemplified in later sections. The following sections introduce the different approaches to this tradeoff.

3.1 RH Processors

Certain processors are fabricated on dedicated RH processes. Advantages include: High tolerance to radiation effects, thanks to the RH process; In some cases, such processors achieve high performance. This can be especially true when using custom design methods similar to those employed for the design of COTS high-performance processors; Compatibility with similar COTS devices (example: PowerPC). This results in easy migration of codes and application to the space environment; This approach can offer high level of integration, including the inclusion of special I/O controllers dedicated to space applications.

The disadvantages of using RH processes include: High cost—RH processes have limited use and the high price of modern fab (in excess of one billion dollars) is amortized over a very small market; Not widely available—there are only about two RH fabs in the USA and no similar advanced processes elsewhere. Use of the RH processes in the USA is ITAR controlled and is not widely available to non-USA customers; Lags several generations behind COTS processors, in terms of performance and power—typical RH processes are based on 150nm CMOS, while high-end COTS processes belong to the 28nm generation, about six processing generations more advanced [31]; SEU rate getting worse—the RH process enables a fixed SEU per bit (about 10^{-11} errors/bit-day) but as the chips become more advanced and contain more memory and more flip-flops, the total SEU per chip is closer to 10^{-5} errors/bit-day.

RH processors include the following. Harris RH3000 32-bit Rad-Hard Controller was based on the MIPS 3000 architecture [19]. It was a rad-hard chip set (a CPU and a FPU), made in the 1990s in the USA. It achieved 20 MHz clock and 6 MFLOPS, and incorporated several peripherals including RH memory and some I/O controllers. TRW developed and Honeywell fabricated another rad-hard version of the MIPS 3000 architecture, named RH32. No data was available at the time of this report. The RAD6000, based on IBM System/6000 architecture, was fabricated in the late 1990s [20] by BAE Systems and was launched on the Mars Rover. It was subsequently replaced by the BAE Systems RAD750, a radiation-hardened processor based on the PowerPC 750. The RAD750 was released for purchase in 2001 [21][22]. The processor has 10.4 million transistors, is manufactured on 250 nm process and has a die area of 130 mm². It operates at 133–166 MHz, achieving 300 MIPS. The RAD750 carries a unit price tag in excess of US$200,000.

3.2 RHBD Processors

This family contains processors that are designed as ASIC and fabricated on commercial CMOS processes. Radiation hardness is achieved by design techniques in the layout, circuit, logic and architecture areas, hence the name Radiation Hardening by Design (RHBD). Advantages of this family include high tolerance to radiation effects (higher than RH processors), medium cost—more expensive than COTS processors, mostly due to low production quantities and high cost of qualification, but at the same time, they are less expensive than RH processors thanks to using a regular commercial fabrication process. Finally, RHBD processors can offer high integration (inclusion of I/O controllers dedicated to space applications) since they are designed as ASIC and since typically the CPU itself takes only a small portion of the silicon die. Disadvantages of RHBD processors—they are usually slower than COTS processors since they are designed as ASIC chips and not as custom processors.

Most RHBD processors are based on the successful European LEON architecture. The SPARC V7 ERC32 and TSC695FL are made by Atmel in France [23]. Originally a 3-chip set, it is now a single chip CPU. It has been used in a large number of satellites and systems in Europe and elsewhere. It achieves 12 MIPS / 6 MFLOPS at very low power (0.3W for the core excluding I/O). The OBC695A V7 SBC is offered by SSTL using Atmel's TSC695. Tiger V7 is another SBC offered by Saab Aerospace using Atmel's TSC695. Atmel LEON2 AT697 provides a major step forward from the TSC695FL SPARC V7 processor. It is a LEON2 SPARC V8 processor [26][27], based on IP core provided by Aeroflex Gaisler. The chip includes the processor, memory interface and a PCI interface and executes at 70-80 MHz. Astrium has developed a SoC derivative of LEON2 named MDPA (Multi-DSP/ µProcessor Architecture) featuring LEON2 at 70 MHz, I/O controllers for 1553, SpaceWire, and CAN [28][29]. The SoC contains a DSP module with modem, encoder and decoder (suitable for high speed TM/TC, 600 Kbps each direction). Aeroflex Gaisler has developed a SPARC V8 processor (LEON3) and SoC architecture and implemented it on ACTEL RTAX rad-hard FPGA [30]. This development revolutionized this field: With rad-hard FPGA, such processors are more widely available and can be customized to specific needs. The main limitation is that ACTEL RTAX FPGA parts are ITAR controlled. The second limitation is performance (clock rate below 25 MHz). Saab Aerospace has developed another SoC based on the LEON2FT code-named COLE [25], intended for their Panther SBC [24]. The SoC design demonstrates a high level of integration, incorporating multiple 1553, SpaceWire and CAN interfaces among others. Astrium has developed SCOC3, a SoC based on LEON3, incorporating a very large set of I/O cores [32] and fabricated on Atmel RH 0.18u process (ATC18RHA). Thales Alenia Space developed LEONDARE LEON3 SoC for space [33]. It is based on IMEC DARE RHBD library, which achieves density of 25Kgates/mm^2. The architecture combines LEON3 CPU with 3 SpW+RMAP and other cores, planned for 208-pin CQFP and fabricated on UMC commercial 0.18u process through Europractice shuttle service. Aeroflex has released UT699, another LEON3FT SoC, that combines PCI bus with several on-chip serial IO cores: 4 SpW, CAN and Ethernet. The ASIC is implemented on 0.25u CMOS, packaged in 352-pin CQFP, operates at 66MHz, and dissipates 5W. Aeroflex Gaisler and Ramon Chips developed GR712RC, a dual-core LEON3FT SoC fabricated on TowerJazz 0.18u CMOS. It contains multiple I/O cores and executes at clock rates higher than 100MHz, dissipating less than 2W. In order to support high reliability and high speed and avoid the risks associated with high pin count packages in space applications, a 240-pin CQFP package is used. The processor periphery architecture is based solely on serial I/O cores employing an I/O switch matrix to reduce the number of actually required I/O pins.

3.3 Single COTS Processor with no Fault Tolerance

A single COTS processor without any special provisions for fault tolerance may be suitable for space missions and tasks where continuous availability or high reliability are not required, or missions with benign radiation environment. COTS processors offer low cost and high performance, yet are susceptible to radiation and environmental effects. TI MSP430 is used on Cubesat nano-satellites [34][35]. Picosat, launched in 2001 by USAF Space Test Program (STP) to test a commercial satellite made by SSTL, contains a 80186 primary processor and a 80386 co-processor [11]. SSTL OBC386 is based on Intel 386 [12], while SSTL MPC8260 on-board recorder employs Freescale MPC8260 PowerPC [13]. SSTL DSP Module uses a single COTS TI TMS320C64xx [14], and AiTech S950 SBC uses PowerPC 750FX [15].

3.4 Single COTS Processor with Time Redundancy (SIFT)

In this approach, a single COTS processor is used together with Software Implemented Fault Tolerance (SIFT), which executes the entire software or certain software sections twice or more. There are two levels of granularity: *Instruction level redundancy*, where each instruction is executed twice and additional instructions compare the results, requiring compiler transformation of the code, and *procedure level redundancy*, where the programmer writes the code to invoke certain procedures twice, compare the results and use software for recovery in case of mismatch. The latter approach may also require some additional hardware to protect the critical data and the critical software. The main advantage of this approach is that it is relatively inexpensive. The principal disadvantage (in addition to the added complexity of programming) is the performance penalty: Executing some code twice, comparing the results and recovery in cases of mismatch typically incur substantial performance penalty, possibly nullifying the performance gained by using high speed COTS processors.

Three SIFT examples are known: CNES created DMT ('Duplex Multiplexed in Time'), time replication for COTS microprocessors [1]. Critical software sections are executed twice, compared and a recovery is initiated in case of fault. DMT uses SEE-protected storage to assure safe context. The replication overhead requires processor 4 times faster than the required processing capability. DMT is only a design idea, yet to be implemented in real hardware. Another example is a time replication technique based on an instruction level approach that has been developed by Politecnico di Torino (Polito) [6]. The third one is the ARGOS COTS vs. RH experiment: Stanford University researchers developed the EDDI technique (Error Detection by Duplicated Instructions) implemented in the USAF ARGOS satellite (launched in 1999), based on the IDT-3081 commercial processor (R3000 instruction set). Each instruction is executed twice. The COTS processor was compared with a rad-hard 10 MHz Harris RH3000. They execute the same instruction set, but the COTS processor is 20× faster. The RH3000 uses two CPUs, operating in parallel and cross-checking. The COTS board used no error detection or correction. Instead, the code is transformed to execute everything twice and check for matches. The code itself was checked with signature analysis using CFCSS (control flow checking by software signatures). In space, while the RH3000 exhibited several but rare crashes, the COTS systems required rebooting on average every 2-12 days [7].

3.5 Duplex COTS: DMR

This architecture employs two equal COTS processors (also known as dual modular redundancy, DMR), a matching hardware, and software for recovery from mismatches. There is no voting as there are only two copies of execution. On mismatch, computation is cancelled and repeated by software control. DMR offers high performance, thanks to using very fast COTS processors and thanks to the observation that SEU are rare. DMR is also relatively inexpensive: the COTS processors are inexpensive, and the principal cost item is the system board that needs to make two processors execute in lock-step, and contains the matching hardware. The disadvantages are that DMR requires special hardware and software for matching and recovery, and that modern COTS processors are sometimes unpredictable at the clock cycle level, due to methods of internal branch speculation and other algorithms that are designed to boost performance. Forcing two such processors to execute in lock-step every clock cycle may require significant slowdown of the processors.

We cite four examples: CNES DT2 (Dual Duplex Tolerant to Transients) targets 99% availability [1],[36], another paper design. The digital electrical flight control system of the AIRBUS A320/A330/A340 [2], built in the 1980s, consisted of eight computers in four redundant pairs, each pair constituting a COM-MON duplex system. The COM-MON pair managed flight commands, with the monitor computer (MON) monitoring the command computer. The platform computer of the BIRD micro-satellite developed at DLR (Germany, launched in 2001) is based on a PowerPC MPC623 micro-controller. It is implemented on a single board that is quadruplicated, with two of the four channels switched off and used as spare channels (in cold redundancy) in case of either a transient and not recoverable fault or a permanent failure on the two active channels operating in master-checker mode [5]. A Chinese paper [10] describes a duplex system with two ARM7TDMI-based ASICs, EDAC memory and running VxWorks.

3.6 Triple COTS: TMR at the system level

TMR (Triple Modular Redundancy) architectures combine three COTS processors and voting logic. The processors do not need to be stopped on SEU. TMR offers high performance (high end COTS such as the latest Pentium or PowerPC processors may be used) and high SEU tolerance –SEU errors, resulting in erroneous bits at the outputs of the processors, are fixed. The downside of TMR is high cost, requiring large area /

volume and power, as well as special hardware for voting and usually additional hardware and software for recovery from internal SEU errors (inside the processors) that cannot be fixed by voting and require scrubbing or reset.

One example of TMR is Maxwell SCS750 SBC, based on three IBM PowerPC750FX 1.6 GIPS (800 MHz) micro-synchronized microprocessors (i.e. working in lock-step) operating in TMR mode, with a centralized voter integrated in a rad-tolerant FPGA that is functionally immune to upsets. It may fly on the NPOESS constellation made of six satellites and dedicated to civil/military weather forecast (first launch delayed to 2013) [8]. The board is believed to cost in excess of $200,000 and to be ITAR controlled. Another example is EADS DMS-R, a computer board based on ERC32 (Atmel TSC695) processor, and a fault-tolerant TMR version using 3 boards for the ATV of the ISS (Autonomous Transport Vehicle of the International Space Station). It also added the old T405 Transputer processors for inter-computer links for voting. [no direct reference]. In Japan, the SERVIS-2 experimental large satellite, developed at USEF, Japan, includes the CRAFT TMR-like computer based on a COTS system on-chip processor [3]. The CPU is not disclosed. The computer contains 3 SoCs that operate in parallel. The purpose of SERVIS-2 is the verification of COTS parts and technologies in space environment. Finally, also in Japan, the INDEX micro-satellite developed at ISAS-JAXA (launched in 2005) is based on a TMR computer using Hitachi SH-3 commercial micro-controllers protected by a 'light' version of a triplex architecture (centralized voter integrated in a space-qualified FPGA) [4]. The computer is stopped for reconfiguration (i.e. for the faulty channel reinsertion phase) during about 2 seconds.

3.7 TTMR on COTS VLIW processors

COTS VLIW processors execute multiple instructions in parallel, and the parallel instruction streams are pre-programmed. SpaceMicro (USA) implemented this capability for “time TMR (TTMR),” where each instruction can be executed three times and the results can be compared and voted, all within the same VLIW processors. TTMR offers high performance (in fact, TTMR processors are the fastest available space processors today) and high SEU tolerance, thanks to embedded TMR mechanism, but it is expensive, is limited to VLIW processors, and is hard to generate code for. The only known examples are SpaceMicro Proton 100k and 200k [9]. The code executes two copies of an instruction, compares the result, on mis-match executes the same instruction the third time and compares for majority voting. Proton 100k SBC uses Equator BSP-16 [17] (up to 1200 MIPS) and Proton 200k SBC uses TI 320C64xx [18] (up to 4,000 MIPS fixed point or 900 MFLOPS). The result is MTBF of 30 years for unrecoverable error (10-4 errors/day). Proton 100k was launched aboard NASA TacSat-2 Micro Satellite in 2006, and is the instrument computer of the AFRL RoadRunner On-board Processing Experiment (ROPE), performing data management and processing of focal plane array data. A Proton 100k also serves as the core processor for medical equipment computers on the International Space Station (ISS).

4. SUMMARY OF ARCHITECTURE TYPES

Pignol suggested three levels of availability (as a result of SEE) [1] as a motivation to help select the appropriate cost—reliability trade-off: High-safety missions (e.g. manned missions, deep space probes) requiring very high fault detection / recovery performances whatever the architecture overhead may be (mass, power consumption, cost, etc.), missions that require medium level of fault tolerance (e.g. non-critical payloads) and may recover from SEE in a variety of ways and may sustain short periods of non-availability, and low-cost scientific missions not necessarily requiring high availability and therefore allowing the use of COTS components without fault-tolerance. However, given the increase in sensitivity of COTS components to radiation effects (due to smaller die geometries), it is possible that in the future some of these missions will also require SEE protections at the architectural level. In the following diagram, inspired by [1] and [9], the various architectures are arranged according to assumed level of SEE tolerance versus a hardware cost-function in terms of price and complexity.

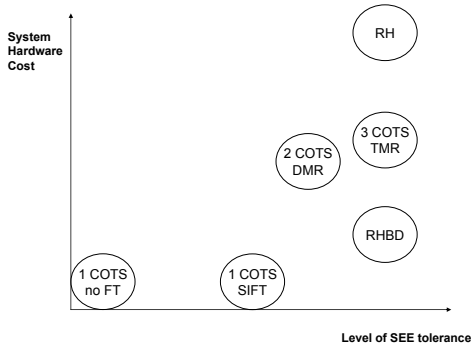

Figure 1: SEE tolerance versus cost of space processor architectures

Figure 2 compares the reported or estimated performance of the surveyed processors. As can be seen in this figure, TTMR and TMR provide the highest performance, followed by modern RH. The majority of the other processors are based on RHBD solution, which is more cost effective, as demonstrated in Figure 1.

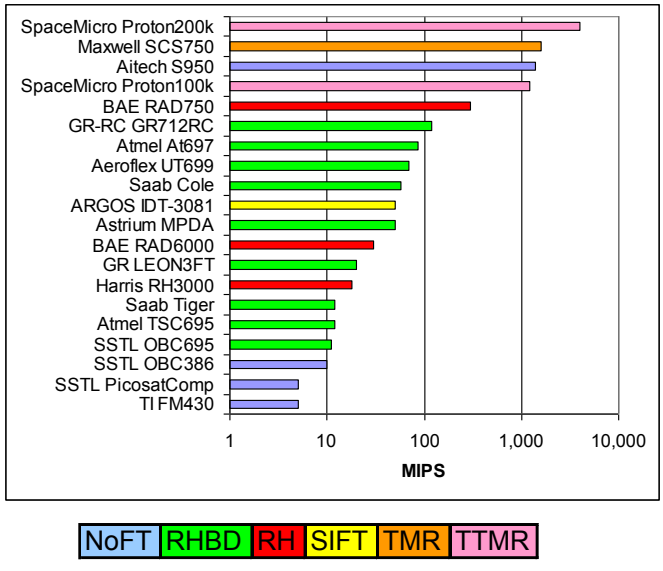

Figure 2: Performance of space processors

5. REFERENCES

[1] M. Pignol, "DMT and DT2: Two Fault-Tolerant Architectures developed by CNES for COTS-based Spacecraft Supercomputers," IOLTS, 2006.

[2] D. Briere, and P. Traverse, "Airbus A320/A330/A340 Electrical Flight Controls: A Family of Fault-Tolerant Systems," FTCS-23, 1993.

[3] H. Hihara, et al., "CRAFT: An Experimental Fault Tolerant Computer System for SERVIS-2 Satellite," ICSSC, 2003.

[4] H. Saito, M. Hirahara, S. Okano, et al., "INDEX: Piggy-Back Satellite for Aurora Observation and Technology Demonstration," IAF, 2000.

[5] P. Behr, W. Baerwald, K. Briess, and S. Montenegro, "Fault Tolerance and COTS: Next Generation of High Performance Satellite Computers," DASIA, 2003.

[6] M. Rebaudengo, M. Sonza Reorda, and M. Violante, "A new Software based Technique for Low Cost Fault-Tolerant Application," RAMS, 2003.

[7] M.N. Lovelette et al., "Strategies for Fault-Tolerant, Space-Based Computing: Lessons Learned from the ARGOS Testbed," IEEE Aerospace, 2002.

[8] R. Hillman, G. Swift, et al., "Space Processor Radiation Mitigation and Validation Techniques for an 1800 MIPS Processor Board," RADECS, 2003.

[9] D. Czajkowski, and M. McCartha, "Ultra Low-Power Space Computer Leveraging Embedded SEU Mitigation,". IEEE Aerospace, 2003.

[10] X. Wang and H. Sun, "The New Fault Tolerant Onboard Computer for Microsatellite Missions," J. China Universities of Posts and Telecom, 13(1):6-9, 2006.

[11] http://directory.eoportal.org/pres_PICOSat.html (2007).

[12] http://www.sstl.co.uk/documents/obc386_datasheet_2007.pdf.

[13] http://www.sstl.co.uk/documents/MPC8260 Solid State Data Recorder_1.pdf.

[14] http://www.sstl.co.uk/documents/DSP Real-Time Data Processing Engine.pdf.

[15] http://www.aitechlibrary.com/aitech/protected/s950.pdf, (2007).

[16] Saab Aerospace Panther SBC brochure.

[17] http://www.spacemicro.com/Products/Proton100k-Datasheet_v3.2-print.pdf.

[18] http://www.spacemicro.com/Products/Proton200k-DSP_Datasheet_v2.7.pdf.

[19] www.harris.com.

[20] http://www.baesystems.com/BAEProd/groups/public/documents/bae_publication/bae_pdf_eis_sfrwre.pdf.

[21] http://www.baesystems.com/BAEProd/groups/public/documents/bae_publication/bae_pdf_eis_rad750.pdf.

[22] http://www.aero.org/conferences/mrqw/2002-papers/A_Burcin.pdf.

[23] http://www.atmel.com/dyn/resources/prod_documents/doc4204.pdf.

[24] Saab Aerospace Panther processor board brochure, http://products.saabgroup.com/PDBWebNew/GetFile.aspx?PathType=ProductFiles&FileType=Files&Id=6216.

[25] http://microelectronics.esa.int/mpd2007/COLE-System-on-chip-MPD2007.pdf.

[26] http://microelectronics.esa.int/mpd2007/AT697F.pdf.

[27] http://www.atmel.com/dyn/resources/prod_documents/doc4226.pdf.

[28] http://microelectronics.esa.int/mpd2007/MDPA_Pres_March_2007.pdf.

[29] http://www.astrium.eads.net/equipment-product-catalogue/product-catalogue/products/communication-payload-control-and-processing-unit .

[30] http://www.gaisler.com/cms/index.php?option=com_content&task=view&id=196&Itemid=151.

[31] http://www.itrs.net/

[32] http://microelectronics.esa.int/mpd2007/SCOC3-MPD2007.pdf.

[33] http://microelectronics.esa.int/mpd2007/LEON3_DARE_ASIC_Rev1.pdf.

[34] http://focus.ti.com/lit/ml/slab034n/slab034n.pdf.

[35] http://www.cubesatkit.com/.

[36] M. Pigno *et al.,* "A Testbench for Validation of DST Fault-Tolerant Architectures on PowerPC G4 COTS microprocessors," DASIA 2011.