Advanced Hybrid On-Board Science Data Processor - SpaceCube 2.0

TOM FLATLEY
HEAD – SCIENCE DATA PROCESSING BRANCH
NASA GODDARD SPACE FLIGHT CENTER
GREENBELT, MD 20771 USA

Abstract - Many of the missions proposed in the Earth Science Decadal Survey (ESDS) will require 'next generation' on-board processing capabilities to meet their specified mission goals. Advanced laser altimeter, radar, lidar and hyper-spectral instruments are proposed for at least ten of the Decadal Survey missions, and all of these instrument systems will require advanced on-board processing capabilities to facilitate the timely conversion of Earth Science data into Earth Science information. Both an "order of magnitude" increase in processing power and the ability to "reconfigure on the fly" are required to implement algorithms that detect and react to events, to produce data products on-board for applications such as direct downlink, quick look, and "first responder" real-time awareness, to enable "sensor web" multi-platform collaboration, and to perform on-board "lossless" data reduction by migrating typical ground-based processing functions on-board, thus reducing on-board storage and downlink requirements. The convergence of technology collaborations between the NASA Goddard Space Flight Center (GSFC), the Air Force Research Lab (AFRL) and the Naval Research Lab (NRL) is being leveraged to produce a SpaceCube 2.0 on-board data processing system that will meet the on-board processing needs of the ESDS missions.

I. INTRODUCTION

SpaceCube is a Field Programmable Gate Array (FPGA) based on-board science data processing system developed at the NASA Goddard Space Flight Center (GSFC). The goal of the SpaceCube project is to provide 10x to 100x improvements in on-board computing power while lowering relative power consumption and cost. The SpaceCube design strategy incorporates commercial rad-tolerant[1] FPGA technology and couples it with an upset mitigation software architecture to provide "order of magnitude" improvements in computing power over traditional rad-hard flight systems (see TABLE 1).

SpaceCube is based on the Xilinx Virtex family of FPGA processors, which include CPU, FPGA and DSP resources. These processing elements are leveraged to produce a hybrid science data processing platform that accelerates the execution of science data processing algorithms by distributing computational functions among the elements, allowing each type of processor to do "what it's good at". This approach enables the implementation of complex on-board functions that were previously limited to ground based systems, such as on-board product generation, data

TABLE 1
PROCESSOR COMPARISON

Notes: 1 – typical, 35 MIPS at 15 watts 2 – typical, 300 MIPS at 15 watts 3 – processor device only, total board power TBD 4 – 3000 MIPS at 7.5 watts (measured) 5 – 5000 MIPS at 10 watts (calculated)	MIPS	Power	MIPS/W
MIL-STD-1750A	3	15W	0.2
RAD6000	35	10-20W	2.33[1]
RAD750	300	10-20W	20[2]
SPARC V8	86	1W[3]	86[3]
LEON 3FT	60	3-5W[3]	15[3]
SpaceCube 1.0	3000	5-15W	400[4]
SpaceCube 2.0	5000	10-20W	500[5]

reduction, calibration, classification, event/feature detection and real-time autonomous operations. The system also is fully reconfigurable in flight, through either ground commanding or autonomously in response to detected events/features in the instrument data stream.

II. SPACECUBE EVOLUTION

The SpaceCube 1.0 system is based on the Xilinx Virtex 4 (FX60) FPGA technology and first flew as part of the Relative Navigation Sensors (RNS) experiment during Hubble Servicing Mission 4 in May 2009. The Hubble unit (SpaceCube 1.0a) interfaced with three cameras and ran autonomous rendezvous and docking algorithms to simulate a robotic servicing capture and release. The system operated on-orbit for approximately 60 hours, and met all RNS experiment goals. The second SpaceCube experiment (SpaceCube 1.0b) flew as part of the Naval Research Lab (NRL) Materials on the International Space Station Experiment 7 (MISSE7) payload and was installed on the Space Station in November 2009. The purpose of the MISSE7 SpaceCube experiment is to test and validate our "radiation hardened by software" (RHBS) mitigation techniques. The RHBS strategy involves a combination of traditional "memory scrubber" functions and a "program execution error detection and correction" software architecture that runs in the background to find and fix single and multi-bit radiation induced upsets. The goal of RHBS is to enable the reliable use of commercial rad-tolerant processing elements in space and to "asymptotically

[1] susceptible to radiation induced upsets, but not to destructive failures

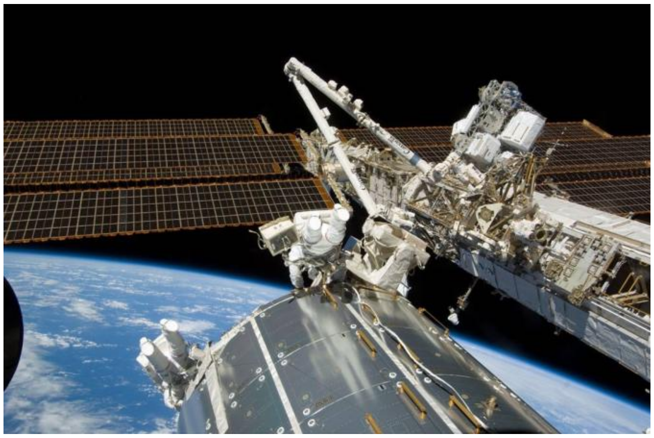

FIG 1
MISSE7 BEING INSTALLED ON STATION

approach" the reliability of true rad-hard systems while providing more computing power for less cost. The MISSE7 SpaceCube experiment will remain on the Station for the next 2-5 years and will serve as an on-orbit testbed for current and future RHBS technology development efforts. Early upset detection/correction results from the MISSE7 SpaceCube experiment are shown in FIG 2. To date, the MISSE7 SpaceCube has operated continuously on orbit for over 160 days and has experienced 60+ upsets (total upsets in four FPGAs), but all of the upsets have been successfully mitigated and no functional errors have occurred.

SpaceCube 1.5 is being designed to support a quick-turnaround sounding rocket flight and serves as a stepping stone to the next generation SpaceCube 2.0 system currently being developed under funding from the Earth Science Technology Office (ESTO). SpaceCube 1.5 moves from the Virtex 4 to the Virtex 5 family of FPGAs and adds several industry standard plug-and-play interfaces such as Gigabit Ethernet and SATA-II. The SpaceCube 1.5 system

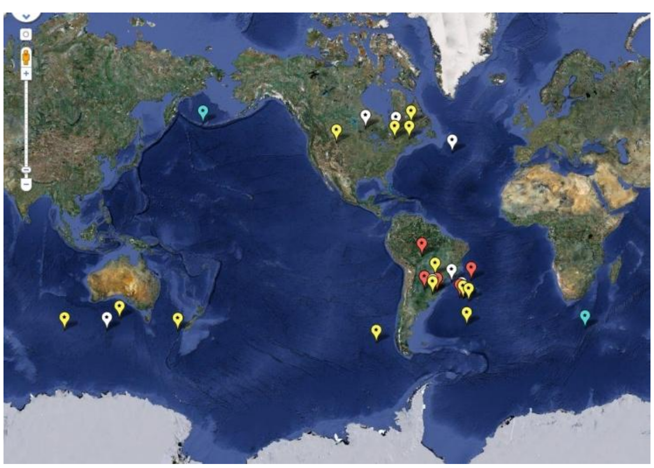

FIG 2
SPACECUBE UPSET DETECTION/CORRECTION

is scheduled to launch in the October-December 2010 timeframe and will likely lead to a series of sounding rocket technology demonstration flights in collaboration with the Department of Defense (DoD) Operationally Responsive Space (ORS) program.

ESTO is currently funding the development of the SpaceCube 2.0 system under their Advanced Information Systems Technology (AIST) program, to provide advanced on-board science data processing capabilities to the upcoming series of Earth Science Decadal Survey (ESDS) missions. The SpaceCube 2.0 system is based on Xilinx Virtex 5 technology and adds Spacewire and cPCI interfaces, in addition to the Gigabit Ethernet and SATA-II interfaces developed for the SpaceCube 1.5 system. SpaceCube 2.0 also moves from custom packaging that was used on SpaceCube 1.0 and 1.5 to an industry standard 3U form factor, and is fully scalable to support a wide variety of mission requirements (See FIG 3).

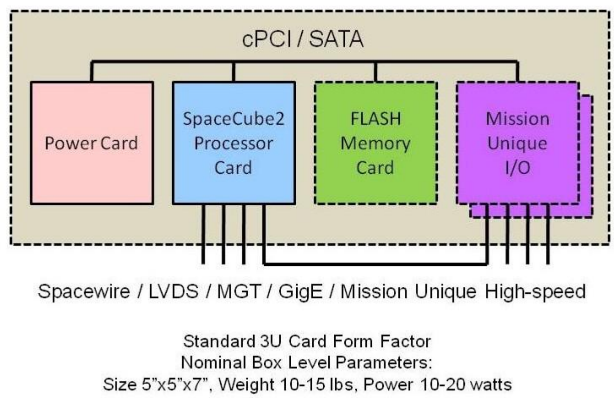

FIG 3
SPACECUBE 2.0 BLOCK DIAGRAM

A SpaceCube 2M (SpaceCube Mini) system is also being developed to support Unmanned Aerial Vehicle (UAV), in-situ/embedded sensor and robotic/rover applications, and will incorporate a significant subset of the full SpaceCube 2.0 features and functions in a small (3"x3"x3") package. The SpaceCube 2.0 and SpaceCube 2M processor cards are being designed to accommodate either the commercial Virtex 5 FX130T part (with embedded PowerPCs) for "high-performance" applications or the radiation hardened Virtex-5QV FX130 SIRF part (developed by the Air Force Research Lab, implemented without embedded PowerPCs) for "high-reliability" applications.

III. SPACECUBE APPLICATIONS

A number of SpaceCube on-board science data processing applications have been developed and demonstrated under NASA/GSFC Internal Research and Development (IRAD) funding during the past five years. This includes applications such as Synthetic Aperture Radar (SAR) processing and hyper-spectral data product generation. The general approach for developing SpaceCube applications involves obtaining algorithms (usually in Matlab or IDL) and sample data sets from science team collaborators, converting the algorithms to 'C', testing and validating results on the Virtex CPUs (PowerPCs), benchmarking the compute intensive functions, and then implementing a variety of Virtual Hardware Description Language (VHDL) and other library functions to accelerate the processing. Typical "hybrid processing" results have yielded a 25x to 50x speed-up in compute intensive functions vs. CPU-only performance.

One of the key mission enabling capabilities provided by the SpaceCube is the computing power to perform on-orbit data reduction, which allows science instruments to collect significantly more data without increasing on-board storage or downlink bandwidth requirements. Dramatic data volume reductions can be achieved by moving typical ground-based processing functions on-board, without the use of compression or other "lossy" techniques. Our SAR "Nadir Altimetry" application achieved a 6:1 data volume reduction by simply moving the first stage ground processing functions on-board (see FIG 4), and our SAR "Delay-Doppler Mapping" application achieved a 165:1 data volume reduction by producing the image products directly on-board (see FIG 5). These results can enable an instrument to run for hours rather than minutes or continuously (respectively), and provide data products to the users in real-time.

Similar results have been achieved processing hyper-spectral data. Fire detection/characterization, cryospheric classification and ocean color (chlorophyll) applications have demonstrated on-board product generation capabilities that can be used by Earth Science data customers such as emergency first responders and researchers in the field who would benefit from real-time direct broadcast data products (see FIG 6).

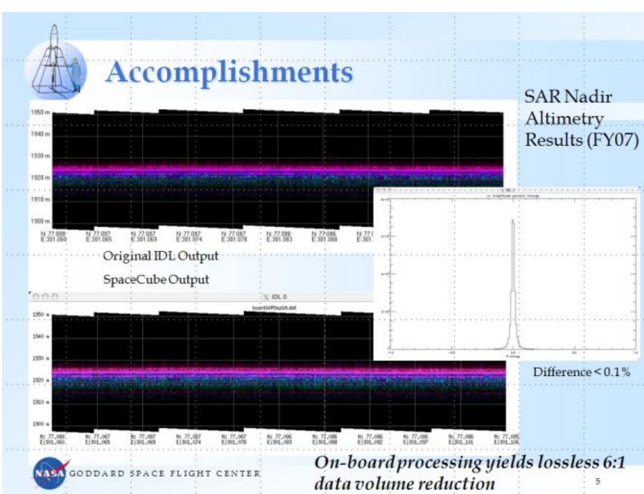

FIG 4
SAR NADIR ALTIMETRY RESULTS

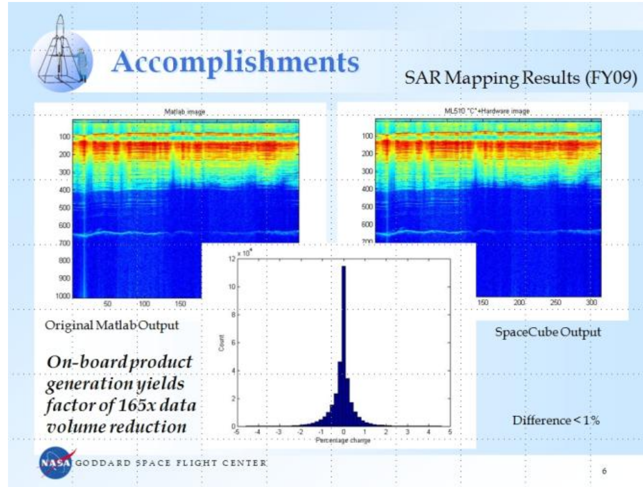

FIG 5
SAR DELAY-DOPPLER MAPPING RESULTS

On-board product generation can also be leveraged to autonomously detect and react to events, such as changing data collection modes when storms/hazards are detected, or changing observation targets when planned targets are obscured by cloud cover. Algorithms to process a variety of events/targets (wildfire, volcano, hurricane, flood, algal bloom, oil spill, etc) can be stored in SpaceCube FLASH

FIG 6
HYPER-SPECTRAL PRODUCT GENERATION

memory and the "active" algorithms can be reconfigured in real-time to process the current event.

IV. CONCLUSION

Many of the missions proposed in the Earth Science Decadal Survey will require "next generation" on-board processing capabilities to meet their specified mission goals. Advanced laser altimeter, radar, lidar and hyper-spectral instruments are proposed for at least ten of the Decadal Survey missions, and all of these instrument systems will require advanced on-board processing capabilities to facilitate the timely conversion of Earth Science data into Earth Science information. Both an "order of magnitude" increase in processing power and the ability to "reconfigure on the fly" are required to implement algorithms that detect and react to events, to produce data products on-board for applications such as direct downlink, quick look, and "first responder" real-time awareness, to enable "sensor web" multi-platform collaboration, and to perform on-board "lossless" data reduction by migrating typical ground-based processing functions on-board, reducing on-board storage and downlink requirements.

The goal of the SpaceCube program is to deliver a true flight-worthy science data processing system that approaches the on-orbit reliability of full radiation hardened systems, with 10x to 100x the processing capability, at a fraction of the cost. Initial results from the SpaceCube experiments conducted during 2009 have been extremely promising, and by coupling commercial processing elements with RHBS technology we hope to soon offer the science community the choice between "perfect data using a RAD750" or "100x more data plus next generation capabilities, with (maybe) an occasional bad pixel, using SpaceCube 2.0".

V. References

http://en.wikipedia.org/wiki/SpaceCube