Achieving Multipurpose Space Imaging with the ARTEMIS Reconfigurable Payload Processor

Ian A. Troxel, Matthew Fehringer, Michael T. Chenoweth
SEAKR Engineering, Inc.
6221 S. Racine Circle
Centennial, CO 80111
303.790.8499
{ian.troxel, matt.fehringer, mike.chenoweth}@seakr.com

Abstract—The components that comprise the Advanced Responsive Tactically Effective Military Imaging Spectrometer (ARTEMIS) payload processor, to be deployed on the AFRL TacSat-3 satellite, provide a flexible and high-performance platform upon which space imaging applications can be deployed. The payload system consists of an FPGA camera interface and processing board as well as a G4-based single board computer among other components. The combination of FPGA and general-purpose processor provides resources attuned to strike a balance between compute-intensive sensor data extraction processing and administrative tasks such as health monitoring and information downlink. This paper provides a description of the system design and highlights the system's performance and flexibility. [1] [2]

TABLE OF CONTENTS

1. INTRODUCTION.....1
2. ARTEMIS PROCESSOR ARCHITECTURE.....2
3. ARTEMIS PROCESSOR FLEXIBILITY.....4
4. TACSAT-3 MISSION OVERVIEW.....5
5. CONCLUSIONS.....7
REFERENCES.....7
BIOGRAPHIES.....7

1. INTRODUCTION

NASA and other space agencies have identified the need to improve the processing capabilities onboard remote space platforms to support the complexity of future space exploration missions [1]. Increasing amounts of data collected coupled with limited downlink speeds makes the need to process data in-situ critical. The US Air Force also recognizes the coming "data crunch" and sees increasing spacecraft processing performance as vital to the success of future missions [2]. In addition to the need to improve raw computational performance, both military and civilian space agencies have determined spacecraft flexibility is key to meeting divergent mission objectives with minimal cost and risk. Multiuse payloads can support the needs of multiple users while reducing non-recurring engineering and increasing mission survivability.

To meet flexibility and resilience objectives, some researchers have chosen to develop adaptable software that executes on fixed hardware resources. For example, power-aware computing is a straightforward software management approach whereby the prioritized processing requirements of applications are matched against component power usage and reserves to optimize performance/Watt [3]. This approach provides a simple means for multiple applications to share hardware resources but lacks flexibility and fault tolerance.

Environmentally Adaptive Fault Tolerant Computing [4] is another software-based approach recently developed that focuses on mission fault tolerance along with resource management. The middleware deployed on these systems measures and predicts the radiation environment in which the spacecraft is traveling, prioritizes mission requirements based on this data, and then deploys applications in a manner that optimizes fault tolerance versus performance [5]. While this approach provides an improvement in system flexibility, board- and device-level hardware redundancy was necessary to achieve an acceptable level of fault tolerance in the deployed system [6]. A mixture of software mechanisms and hardware redundancy is likely to be the best overall solution for most programs but limiting the scope to which hardware is replicated is critical to avoiding unnecessary increases in weight, power, bulk, etc.

Several related projects have proposed new techniques that allow spacecraft to adapt to emergent and changing mission requirements by providing hardware flexibility and fault tolerance with limited redundancy. The so called Flexible Mission Spacecraft (FMS) concept has been proposed as a means to reduce design re-engineering by promoting system flexibility with minimal redundancy [7]. Some proposed FMS solutions focus on building plug-and-play subsystems [8] while others focus within the processing subsystem. The Adaptive Avionics Experiment (AAE) project, led by AFRL, focuses on the responsive space mission by proposing a design that includes a collection of signal processors whose input can be dynamically changed to receive data from one of several sensor suites [9]. Generic sensor interfaces provide a means to easily tailor the base architecture to a specific mission's sensors. Though the system did not have a successful flight, the AAE design served as a good first step toward an FMS in support of responsive space missions.

1

[1] 1-4244-1488-1/08/$25.00 ©2008 IEEE.
[2] IEEEAC paper #1202, Version 1, Updated October 19, 2007

Building on the AAE design, the Advanced Responsive Tactically Effective Military Imaging Spectrometer (ARTEMIS) is another payload processing system that is designed for the FSM mission. ARTEMIS processor supports application independent processing by coupling compute-intensive Field-Programmable Gate Arrays (FPGAs) with a powerful PowerPC-based general-purpose processor. ARTEMIS' adaptable architecture provides flexibility when tailoring the system design to meet the needs of various missions and allows for in-situ reconfigurability to vastly improve system fault tolerance. The ARTEMIS payload is scheduled to launch on the AFRL TacSat-3 mission in 2008.

This paper outlines the ARTEMIS payload processor and highlights its performance and flexibility. The organization of the remainder of the paper is as follows. Section 2 describes the ARTEMIS architecture and Section 3 highlights the flexibility of the ARTEMIS processor. A discussion of plans to deploy ARTEMIS on the upcoming TacSat-3 mission is given in Section 4 and Section 5 concludes the paper.

2. ARTEMIS Processor Architecture

The ARTEMIS processor is an AFRL funded project that seeks to develop a general-purpose signal processing platform that pushes the boundaries of payload performance, adaptability and survivability.

The key system objectives proposed by the program include:

- Support scalable processing from 9 to 400 GFLOPS
- On-orbit system reconfigurability
- Incorporation of open standards
- SEE and TID tolerant
- Flexible I/O architecture to ease mission tailoring
- User interface supporting rapid development

The main components of the ARTEMIS processor system architecture are shown in Figure 1. Interfaces between components within the ARTEMIS processor include power plane, Compact Peripheral Component Interconnect (cPCI) and Spacewire for command and data handling instructions (C&DH), and high speed serial to support fast data processing. The ARTEMIS processor interacts with other spacecraft components via several external interfaces including SV 28V power, Spacewire and RS422 for spacecraft C&DH, Gigabit Ethernet (GigE) for uplink commands and downlink data, and sensor connections that can interface via a wide array of standards using adaptable mezzanine connectors. Several custom interfaces have been developed and interfaced through the adaptable sensor connector including LVDS camera links for a Hyperspectral Imager (HSI) and a High Resolution Imager (HRI), other digital interfaces for a mass data storage board, and a focus mechanism, as well as an analog input for a position sensor. The ability to support such a wide array of devices is testament to the versatility of the adaptable sensor connector.

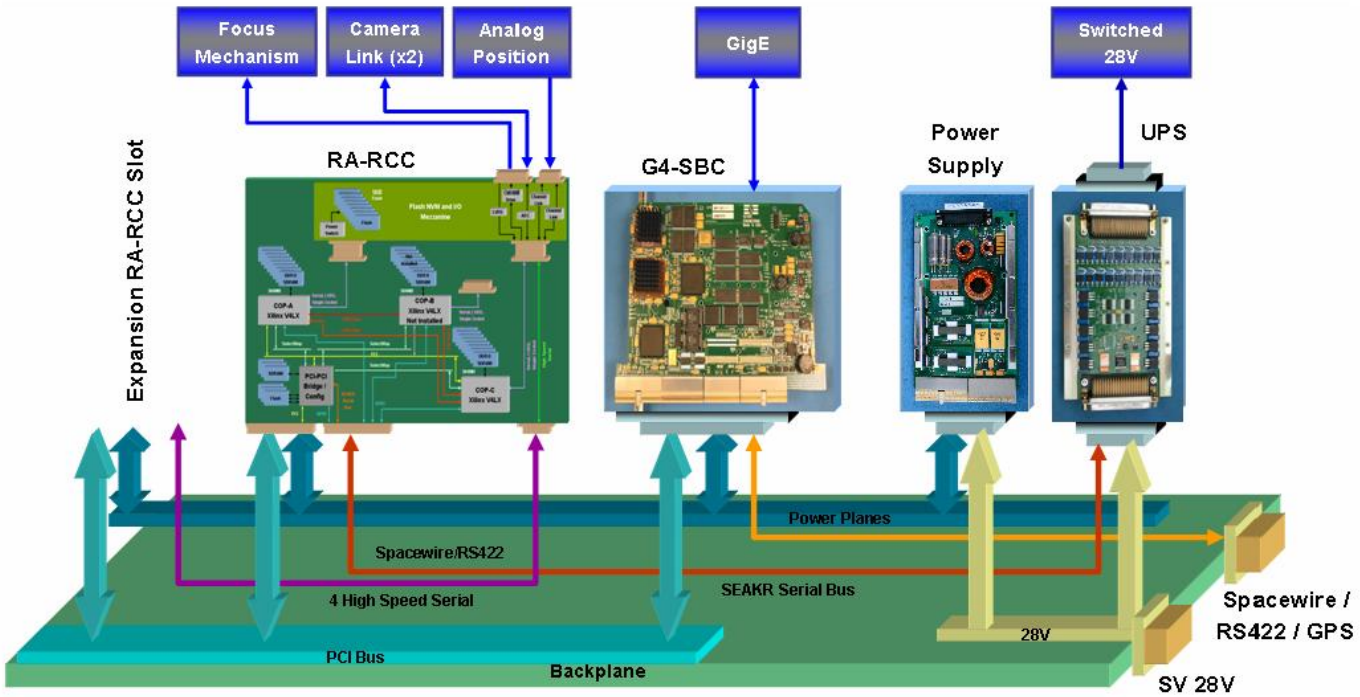

Figure 1 – ARTEMIS Processor System Architecture

Four types of boards compose the ATERMIS processor system including a power supply, a Universal Power Switch (UPS), a G4-based single-board computer (G4-SBC), and a Responsive Avionics Reconfigurable Computer (RA-RCC) board. The power supply receives the space vehicle 28V power and outputs regulated +3.3V, +5V, +15V, and -15V as required by the other boards in the system. The UPS takes in the space vehicle 28V power and receives commands via the SEAKR Serial Bus to switch the power to redundant payload power planes and individual payload sensor components such as cameras. Power switching commands are passed from the G4-SBC to the RA-RCC via the cPCI bus. The RA-RCC then passes the commands to the UPS via the SEAKR Serial Bus. Redundant power supplies and UPSs can be included if required by a mission's fault tolerance requirements. The total power budget for the ARTEMIS processor and sensors is 47W.

The G4-SBC manages external interfaces to the spacecraft and up/down links, controls system configuration and orchestrates data processing. A block diagram of the G4-SBC board architecture is shown in Figure 2. The board consists of the Freescale MPC7457 processor with 32KB of L1 cache, 256 KB of L2 cache and an internal Alta-Vec coprocessor; a PC107A memory controller and PCI bridge; a support FPGA that provides Error Detection and Correction (EDAC) protected interfaces to three types of memory and an RS422 LVDS interface; cPCI and SpaceWire interfaces; and a GigE PMC interface. The G4-SBC provides 600 DMIP sustained performance, 256 MB of SDRAM, 1GB of Flash memory, greater than 30Krad radiation tolerance and consumes a maximum of 15W.

The G4-SBC receives commands from the SpaceWire or RS422 interface and controls the rest of the system.

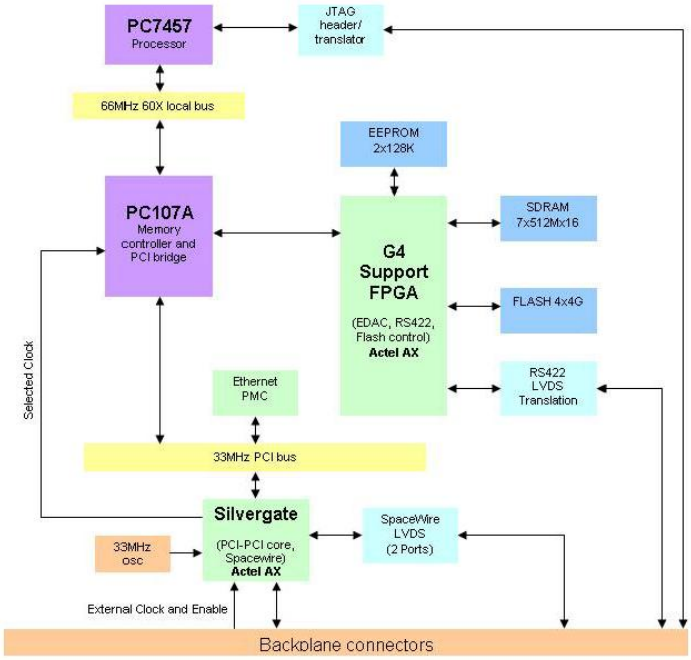

Figure 2 – G4-SBC Board Architecture

The primary functions of the RA-RCC are to control the payload sensor functionality, perform on-board processing of the sensor data, and to control power switching of the sensors and nonvolatile mass data storage. The RA-RCC board architecture is shown in Figure 3.

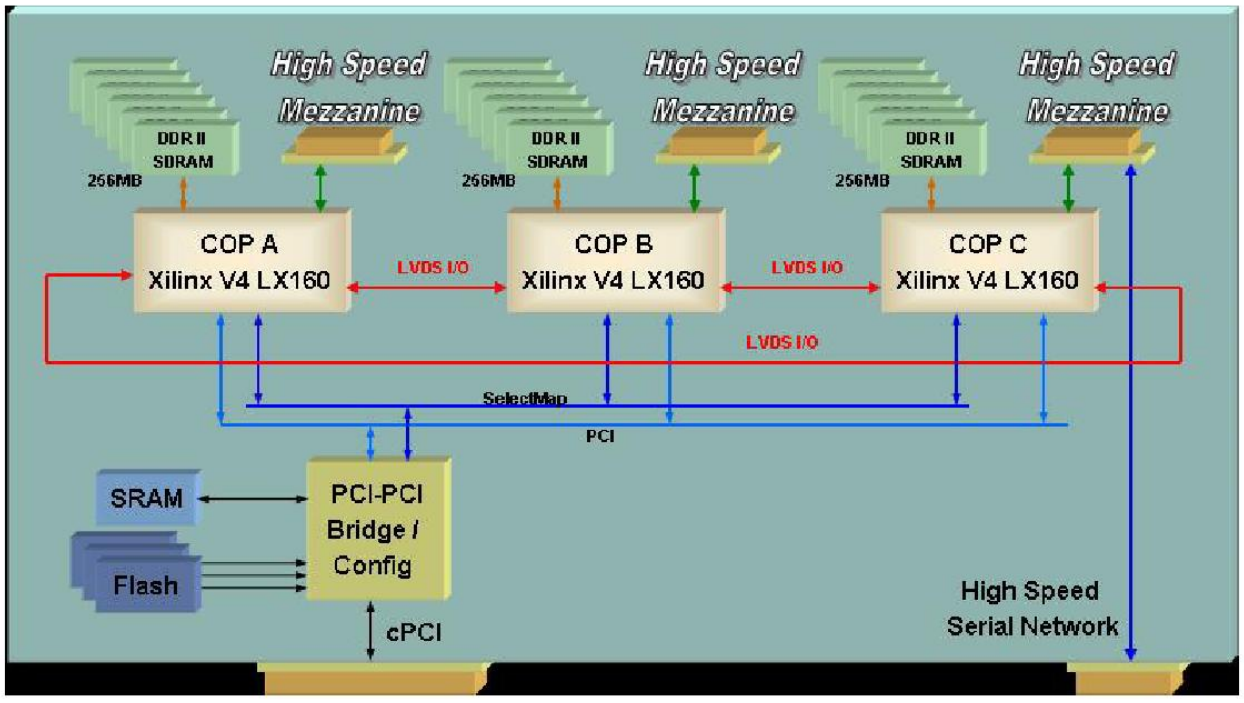

Figure 3 – Responsive Avionics Reconfigurable Computer (RA_RCC) Board Architecture

The board contains four FPGAs including one Actel RTAX2000 and three Xilinx V4 LX160 coprocessors (COPs). The radiation tolerant Actel provides a PCI-to-PCI bridge between the back plane and the local PCI bus interconnecting three COP FPGAs. The Actel also controls configuration management and scrubbing of the COPs via the SelectMap interface as well as the 4Gb bank of triple-redundant Flash memory used for configuration program storage and a 512Mb bank of triple redundant SDRAM used for system configuration management. The Actel also controls the serial bus interface for the UPS card which controls the power state of the payload sensors.

Each COP controls and has access to a 256MB bank of Reed-Solomon EDAC protected DDR2-SDRAM available for processing applications. The COPs also interface to adaptable high-speed mezzanine interfaces that can connect them to sensors, or additional memory or system interconnections. At present, the ARTEMIS processor connects to a 16GB bank of Reed-Solomon EDAC protected NAND Flash, two camera links for camera control and receipt of the high rate image data, current drive for opto-isolator control lines for the focus mechanism stepper motor controller, and analog position telemetry from the focus mechanism via mezzanine cards as previously described. Also, custom LVDS I/O interfaces are implemented between the COPs to improve inter-processor communication. Due to the flexibility of FPGAs, all of these interfaces may be enabled or disabled as required by the mission without changing board hardware layouts. These adaptable interfaces provide a cost-effective way to customize the ARTEMIS core components for a wide range of missions.

The COP FPGAs also provide front-end signal processing and/or payload processing in conjunction with the G4-SBC. FPGAs provide a fabric upon which a variety of high-performance signal processing algorithms can be deployed. Libraries of predefined functions are available or the developer can build custom functionality as required.

The flexible nature of these devices allow the system to adapt to late developing requirements or changing on-orbit conditions such as algorithm updates and degraded mode requirements due to catastrophic hardware failure.

3. ARTEMIS Processor Flexibility

Figure 4 shows several of the many I/O options the RA-RCC flexible mezzanine interconnect provides. Any one of the COPs can be connected to one or more mission-specific devices via tailored mezzanine cards and interface logic developed for the COP. The setup shown in Figure 4A shows one setup option whereby each COP is independently interfaced to one or more dedicated resources. This setup provides a simple design but does not guard against mechanical failures on the input channels unless some external redundancy mechanism is employed. A second option (Figure 4B) provides maximum redundancy and flexibility with all sensor inputs available to be switched to any of the COPs. This setup provides the ability to fully bypass a hardware failure with a COP or mezzanine card interface. Due to the reconfigurable nature of FPGAs, logic resources need not be wasted within the COPs to support each potential interface at all times. In-situ reconfiguration can occur on an as needed basis to support degraded modes. This option also facilitates fully replicated processing of input data.

Another option (Figure 4C) provides a hybrid between the two previous approaches supporting a mixture of redundancy as mission requirements dictate. Figure 4D shows another option where physical separation between redundant interfaces can be achieved with the use of an additional mezzanine card (or other means of interconnection). In this configuration the two COPs on the ends are physically isolated while the center COP can support redundancy or additional processing in either I/O domain.

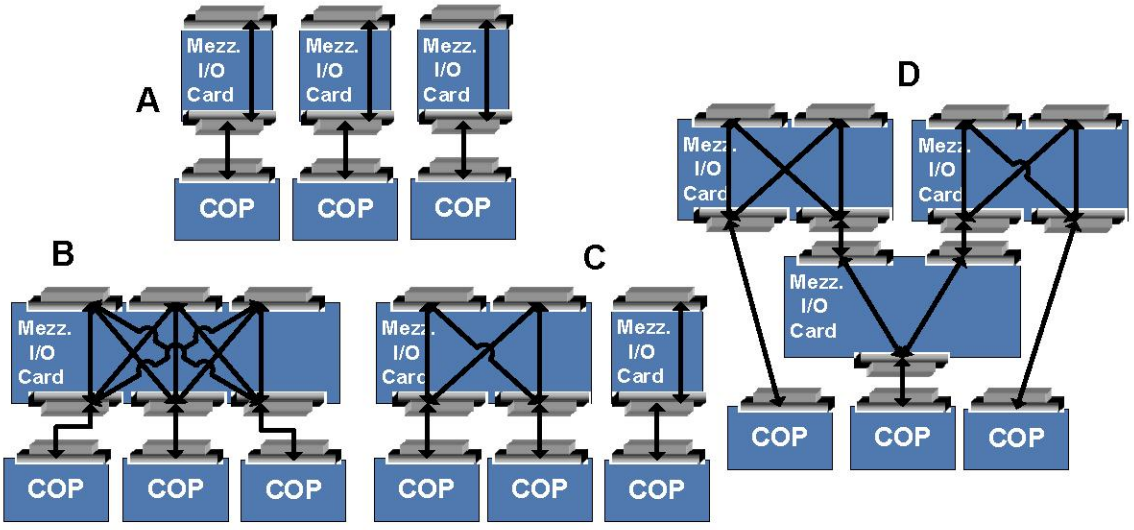

Figure 4 – Select COP I/O Fault Tolerance Options

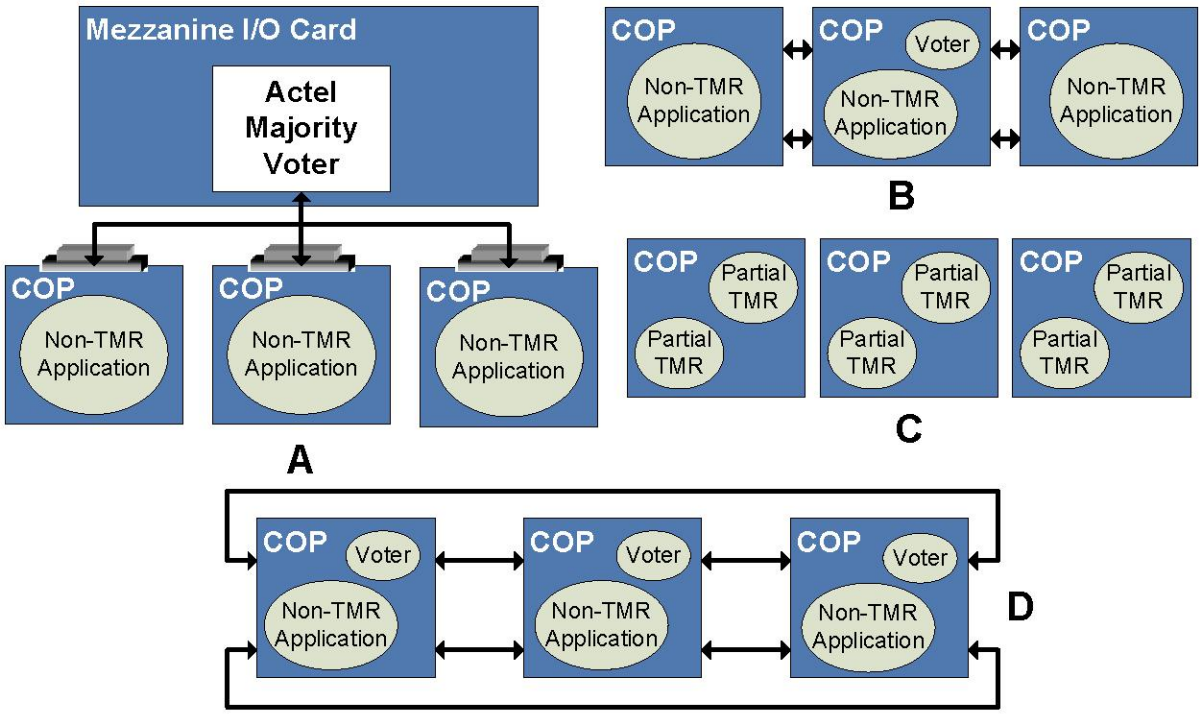

Figure 5 – Select COP Application Fault Tolerance Options

The use of high performance yet radiation susceptible Xilinx FPGAs as processing elements on the RA-RCC requires the use of period configuration scrubbing and other Single-Event Effects (SEE) mitigation techniques. The use of Triple Module Redundancy (TMR) with a radiation hardened voter is a common method deployed in aerospace systems but the flexible design of the RA-RCC provides the designer with some interesting options to consider (Figure 5). Using the flexible mezzanine I/O connectors, a fault tolerant voter could be deployed as shown in Figure 5A. For this option it is assumed that the voter would be integrated into one or more mezzanine cards such as those shown in Figure 4B-4D in a manner in which it doesn't interfere with the I/O interfaces to payload sensors. This mezzanine approach allows the designer to choose an option that is most appropriate for their algorithm rather than requiring the voter to be designed on the RA-RCC.

Other mitigation options exist that do not require the use of a mezzanine connector. Figure 5B shows an option whereby a radiation immune (i.e. triplicated) voter is placed on one of the COPs. The output data from each of the three COPs would be passed to the COP with the voter for verification. An alternative approach to the centralized voter is to have voters distributed in each of the devices (Figure 5D). This approach provides the benefits of having improved fault tolerance due to redundant voters which also provides the option to save COP resources by not requiring each individual voter to be radiation immune. However, this approach requires more data to be shared between the devices because each voter needs a copy of data from all other devices.

Figure 5C provides another TMR strategy known as selective TMR in which critical parts of an application are triplicated within the FPGA such that an external voter is not required [10]. Unlike the other options presented in Figure 5, all three COPs are not required to perform redundant computation and can work on separate applications or separate stages within the same application. The key to this strategy is to identify the critical control paths, data and other algorithm features to which TMR coverage can be applied such that the algorithm's fault tolerance is at an acceptable level. It should be noted that the flexibility of the RA-RCC allows applications to be dynamically switched between options if desired due to changes in the radiation environment or mission objectives or due to system failures. Such flexibility greatly improves system fault tolerance and reduces overall project risk.

4. TacSat-3 Mission Overview

The Tactical Satellite (TacSat) program is a joint AFRL and NRL demonstration program whose main goal is to develop the capability to field inexpensive space systems in time of crisis to augment and reconstitute existing capabilities or perform entirely new tactical theater support missions [11]. Some of the key criteria that will define success for the TacSat program in meeting the needs of responsive space missions are to deploy low cost ($20 million or less) mission-specific spacecraft rapidly (i.e. activated on orbit within six days of call up) and provide between six to twelve months of reliable mission operations [12].

The TacSat-1 mission took the first step toward improving responsive space capability and also expanded SIPRNET to space and the TacSat-2 mission provided a space-based common data link for tactical communications [13]. TacSat-3 will expand the program's capabilities by deploying a standard component interconnect bus, demonstrating initial component plug and play ability and including a land-focused Hyperspectral imaging (HSI) payload [14]. The ARTEMIS payload will provide the HSI sensor and application processing capability for the TacSat-3 mission scheduled to launch in 2008. Future TacSat missions will focus on improving payload capabilities, extending dwelling time in LEO and link to other satellite programs such as T-SAT.

The HSI payload onboard TacSat-3 consists of a telescope, an imaging spectrometer, and a high resolution imager that combine to produce raw high resolution images. This payload unit is captured by the *ATEMIS Sensor* box in Figure 6. The HSI data collected from the sensor is processed by the *Sensor Processor* shown in Figure 6 which consists of one ARTEMIS payload processor as shown in Figure 1. HSI allows for spectral match indication and identification which provides anomaly detection within a given scene that allows the warfighter to distinguish man-made materials from natural materials among other benefits. A fundamental capability of the ARTEMIS payload is to autonomously process HSI data and produce tactically relevant data for dissemination directly to the warfighter in the form of text or imagery as the situation requires.

The goals of achieving responsive space capabilities motivated all aspects of the ARTEMIS sensor and processor payload development. Design trades were carefully evaluated at each step with cost and schedule impacts of foremost consideration. The resulting sensor maintains technical performance while containing costs even with a rapid development schedule of twelve months [14].

ARTEMIS provides a foundation for high-performance and flexible payload system design that meets the TacSat development program's responsive space objectives. An innovative aspect used to control cost and schedule is to decouple payload processing capabilities from the spacecraft sensors. As previously described, the ARTEMIS processor employs mezzanine connectors with flexible hardware interfacing to connect sensors to processing resources. Also, by decoupling the sensor from the processor, the ARTEMIS payload can be deployed on divergent spacecraft architectures without redesigning the sensors – only a re-spin of the processor board to accommodate any additional interfaces (e.g. VME backplane) is required. Also, this decoupling allows for a smooth transition when future versions of the processing or sensor boards are developed. Such flexibility will allow future missions to quickly augment sensors and processing capabilities around the core payload architecture deployed on TacSat-3. The ARTEMIS payload processor is a key technology to enabling the responsive space mission and is actively being investigated for additional mission.

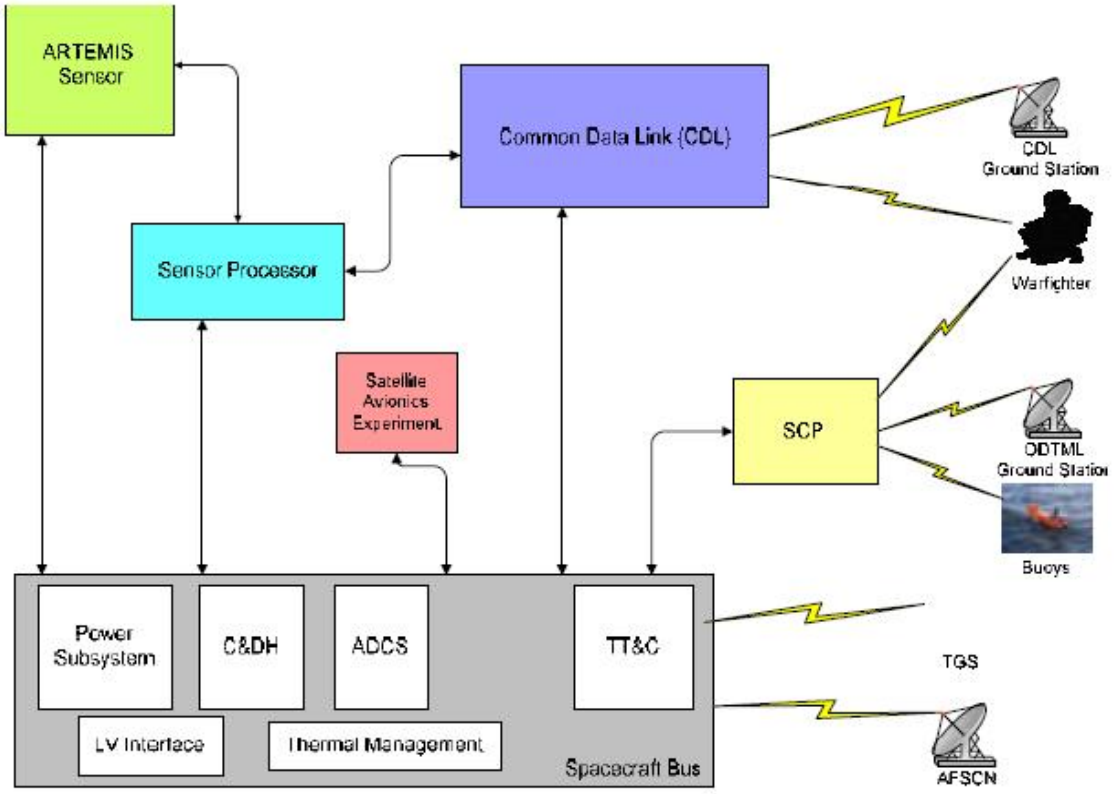

Figure 6 – TacSat-3 Architecture c/o AFRL [14]

5. CONCLUSIONS

The Advanced Responsive Tactically Effective Military Imaging Spectrometer (ARTEMIS) payload processor was motivated and described. Various means by which the ARTEMIS processor design improves system fault tolerance and flexibility and the system's ability to meet the challenges of the responsive space mission were demonstrated. Key features described include the flexible mezzanine interface that decouples the payload sensors and processors. This flexibility enables smooth transitions between successive generations of sensor and processor components. The mezzanine connectors also provide for a variety of fault tolerant I/O interconnect and radiation mitigation strategies as outlined in Section 3. The means by which ARTEMIS autonomously provides Hyperspectral Imaging intelligence to the warfighter from the TacSat-3 spacecraft and the mission's importance were briefly outlined in Section 4. Upgrades and additional deployment of the ARTEMIS payload processor are being actively pursued.

REFERENCES

[1] NASA, "2006 NASA Strategic Plan," NP-2006-02-423-HQ, NASA Headquarters, Washington D.C., February 2006.

[2] US Air Force, "Strategic Master Plan FY06 and Beyond," US Air Force Space Command, Peterson AFB, CO, October 2003.

[3] P. Shriver, S. Briles, J. Harikumar and M. Gokhale, "A Power-Aware Approach to Processing Payload Design," *Proc. Government Microcircuit Applications and Critical Technology Conference*, Tampa, FL, March 31-April 3, 2003.

[4] D. Brenner, J. Ramos, G. Galica and C. Walter, "Environmentally Adaptive Fault Tolerant Computing," *IEEE Aerospace Conference*, Big Sky, Montana, March 5-12, 2005.

[5] I. Troxel and A. George, "Adaptable and Autonomic Management System for Dependable Aerospace Computing," *Journal of Autonomic and Trusted Computing (JoATC) Special Issue on Autonomic and Trusted Computing Systems and Applications* (accepted and in press).

[6] J. Samson, J. Ramos, I. Troxel, R. Subramaniyan, A. Jacobs, J. Greco, G. Cieslewski, J. Curreri, M. Fischer, E. Grobelny, A. George, V. Aggarwal, M. Patel and R. Some, "High-Performance, Dependable Multiprocessor," *Proc. IEEE/AIAA Aerospace Conference*, Big Sky, MT, March 4-11, 2006.

[7] J. Bystroff, "System Architecting Challenges of Changing Missions for a Flexible Mission Spacecraft," *Proc. AIAA Responsive Space Conference*, Los Angeles, CA, April, 24-27, 2006.

[8] B. Jackson and K. Epstein, "A Reconfigurable Multifunctional Architecture Approach for Next-Generation Nanosatellite Design," *Proc. IEEE Aerospace Conference*, Big Sky, MT, March 18-25, 2000.

[9] D. Lanza, J. Lyke, P. Zetocha, D. Fronterhouse and D. Melanson, "Responsive Space through Adaptive Avionics," *Proc. AIAA Responsive Space Conference*, Los Angeles, CA, April 19-22, 2004.

[10] P. Samudrala, J. Ramos and S. Katkoori, "Selective Triple Modular Redundancy for SEU Mitigation in FPGAs," *Proc. International Conference on Military Application of Programmable Logic Devices (MAPLD)*, Washington, D.C., September 9-11, 2003.

[11] J. Raymond, P. Stadter, C. Reed, E. Finnegan, M. Hurley, C. Merk, T. Kawecki and C. Garner, "A TacSat Update and the ORS/JWS Standard Bus," *Proc. AIAA Responsive Space Conference*, Los Angeles, CA, April 25-28, 2006.

[12] S. Straight and T. Davis, "Tactical Satellite 3: Requirements Development for Responsive Space Missions," *Proc. Space Systems Engineering Conference (SSEC)*, Atlanta, GA, November 8-11, 2005.

[13] T. Duffey and M. Hurley, "Operationally Responsive Tactical Microsatellites," *Journal of Naval Research*, 2005.

[14] T. Davis and S. Straight, "Development of the Tactical Satellite 3 for Responsive Space Missions," *Proc. AIAA Responsive Space Conference*, Los Angeles, CA, April 24-27, 2006.

BIOGRAPHIES

Ian A. Troxel is the Future Systems Architect at SEAKR Engineering, Inc. Ian received his Ph.D. in Electrical and Computer Engineering from the University of Florida in 2006, where his graduate research focused on developing fault tolerance and application management

services for embedded space systems. Ian's research interests include fault tolerance strategies for next-generation processor systems, developing characterization methods for radiation acceptance testing of advanced processors and the design of optical interconnects for embedded space platforms.

Matthew J. Fehringer is a Program Manager at SEAKR Engineering, Inc. Matt received his Bachelor's Degree in Electrical Engineering from the South Dakota School of Mines in 1990. While with SEAKR, he has managed the development and production of solid state data recorders, non-volatile memory components and systems, mission data formatting products, and the ARTEMIS Sensor Processor.

Micheal T. Chenoweth is a hardware design engineer at SEAKR Engineering, Inc. Mike received his Masters in Electrical Engineering from the Colorado School of Mines in 1998, where his graduate research focused on developing an electromagnetic acoustic transducer and wavelet analysis system for the non-destructive diagnosis of power substation ground grids. Mike is the technical lead for the ARTEMIS processor.