

• Technology Validation: NMP ST8 Dependable Multiprocessor Project ^{1,2}

John R. Samson, Jr., Jeremy Ramos
Honeywell Aerospace, Defense and Space
13350 U.S. Highway 19 North
Clearwater, Florida 33764-7290
john.r.samson@honeywell.com
jeremy.amos@honeywell.com

Minesh Patel
Tandel Systems, LLC
12401 62nd Street North
Clearwater, Florida 33773
mpatel@tandelsys.com

Alan D. George
Department of Electrical and Computer Engineering
University of Florida
Gainesville, FL 32611-6200
george@hcs.ufl.edu

Raphael Some
Jet Propulsion Lab, California Institute of Technology
4800 Oak Grove Drive
Pasadena CA 91109
rsome@jpl.nasa.gov

Abstract—Current and future space-based processing applications are requiring, and will require, increasing amounts of onboard processing capability. One way to achieve a high level of processing capability is through the use of COTS (Commercial-Off-The-Shelf) high-performance processors. While state-of-the-art COTS processors are exhibiting adequate Total Integrated Dose (TID) performance to meet the requirements of the natural space radiation environment, Single Event Upsets (SEUs) caused by Galactic Cosmic Rays and Solar Protons will remain a problem. Traditional approaches to mitigate the SEU problem involve fixed redundancy schemes such as Self Checking Pairs (SCP) or Triple Modular Redundancy (TMR). While effective in mitigating the effects of SEUs, use of these techniques comes at a high price, 100% overhead for SCP, and 200% overhead for TMR, particularly when such a level of protection is not needed. In such cases, it would be beneficial to be able to convert that overhead into useful mission processing capability or power reduction. The idea behind Dependable Multiprocessor is to sense the environment and configure the processing system appropriately to maximize the computational performance density, i.e., the computational performance to power ratio, available to the mission.

The goal of the Dependable Multiprocessor project is to provide spacecraft/payload processing capability 10x – 100x what is available today, enabling heretofore unrealizable levels of science and autonomy. To date, Dependable

Multiprocessor (DM) technology has been developed as part of NASA's New Millennium Program (NMP) ST8 (Space Technology 8) project. The objective of this NMP ST8 effort is to combine high-performance, fault tolerant, COTS-based cluster processing and fault tolerant middleware in an architecture and software framework capable of supporting a wide variety of mission applications. Dependable Multiprocessor development is continuing as one of the four selected ST8 flight experiments. The focus of the current phase of the Dependable Multiprocessor project is two-fold: 1) to meet the TRL5 Technology Maturity Assessment technology validation requirements, and 2) to complete the plans for the TRL7 flight experiment and demonstration. This paper describes the validation experiments, demonstrations, and performance achieved to date, and the plans for Dependable Multiprocessor flight validation.

TABLE OF CONTENTS

- 1.0 INTRODUCTION**
- 2.0 TECHNOLOGY OVERVIEW**
- 3.0 TECHNOLOGY VALIDATION PLAN**
- 4.0 TRL4 VALIDATION**
- 5.0 TRL5 VALIDATION**
- 6.0 TRL6 VALIDATION**
- 7.0 TRL 7 FLIGHT VALIDATION EXPERIMENT**
- 8.0 STATUS**
- 9.0 SUMMARY & CONCLUSION**

¹ The project formerly was known as the Environmentally-Adaptive Fault-Tolerant Computing (EAFTC) project.

² Paper # 1510 Copyright 0-7803-9546-8/06/\$20.00©2006 IEEE; This paper has not been published elsewhere and is offered for exclusive publication except that Honeywell reserves the right to reproduce the material in whole or in part for its own use and, where Honeywell is obligated by contract.

1. INTRODUCTION

Many next-generation space missions will require onboard high-performance processing for science payload, as well as autonomous data analysis. Current space-qualified computing systems, built around radiation-hardened processors, cannot provide sufficient performance, i.e., throughput, or performance-density, e.g., throughput per watt, to meet these requirements. In terrestrial laboratories, science data processing is performed on parallel processing cluster computers. Similarly, the complex models envisioned for future highly autonomous robotic systems also need high-performance, parallel or supercomputer architectures to meet near real-time requirements. A cluster computer comprises a set of single board computers, interconnected by a high speed switched network, running a file-oriented multi-threading operating system and a “middleware” which controls and coordinates parallel processing applications. A typical system might consist of 10 to 20 Motorola G4 based single board computers, interconnected via a gigabit Ethernet, running the LINUX operating system and an MPI middleware. The parallel processing applications are typically written in a version of FORTRAN, C or C++ and are supported by parallel math libraries such as ScaLAPACK or PLAPACK. In the most advanced architectures, Field Programmable Gate Arrays (FPGAs) are used to implement the algorithms directly in hardware. FPGAs allow configuring of hardware “on the fly”, and provide the most power and time efficient implementations of mathematical routines.

Over the past few generations, Commercial-Off-The-Shelf (COTS) computer components have become extremely resistant to the debilitating effects of radiation. Many commercial parts can withstand many 10s of kilorads of Total Ionizing Dose (TID) and are immune to catastrophic Single Event Latchup (SEL). The primary issue preventing the deployment of a COTS-based spaceborne cluster computer is their continued susceptibility to Single Event Upsets or SEUs, (a.k.a. soft errors). SEUs however, unlike TID and SEL, entailing only a bit flip from 1 to 0 or 0 to 1, do not cause permanent damage. Further, in the latest generation of computer electronics, SOI CMOS (Silicon on Insulator Complementary Metal Oxide Semiconductor) has proven to be approximately an order of magnitude less susceptible to SEU than previous bulk CMOS. If we can withstand a few errors per day per processor, without unduly impacting system dependability, it would be possible to fly essentially commercial cluster computers. Not only would this provide mission enabling performance and performance-density levels, but it would significantly lower the cost of development as standard laboratory science codes could be easily ported to these

systems without the expensive and error prone process normally associated with moving complex codes from the lab to a new platform.

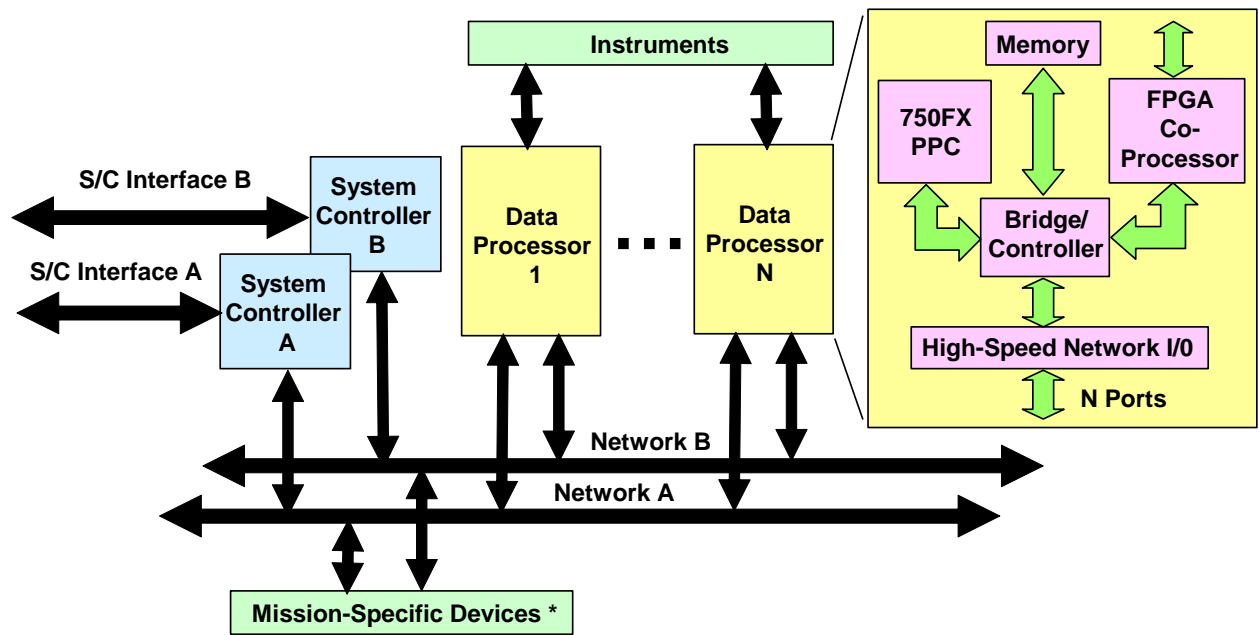
The Honeywell Dependable Multiprocessor experiment will validate the technological concept, the architecture, the fault tolerance techniques and the associated performance, reliability, and availability models behind this technology. Supplementing ground-based testing, the in-space validation will test those aspects of the technology which cannot be effectively exercised on the ground. This includes the ability to withstand concurrent omni-directional, multi-species, multi-energy, and extremely high energy radiation while meeting required reliability and availability levels. The experiment will also provide the data required to calibrate the associated models and to allow scaling of the models to radiation and computing environments well beyond the ST8 LEO/MEO environments.

2. TECHNOLOGY OVERVIEW

Dependable Multiprocessor technology comprises four key elements:

- An architecture and methodology which enables the use of COTS-based, high-performance, scalable, multi-computer systems in a space environment, incorporating reconfigurable co-processors, and supporting parallel/distributed processing for science codes, and accommodating future COTS parts/standards through upgrades.
- An application software development and runtime environment that is familiar to science application developers, and facilitates porting of applications from the laboratory to the spacecraft payload data processor.
- An autonomous and adaptive controller for fault tolerance configuration, responsive to environment, application criticality, and system mode, that maintains required dependability and availability while optimizing resource utilization and system efficiency
- A methodology and tools which allow the prediction of the system’s behavior in the space environment, including: predictions of availability, dependability, fault rates/types, and system level performance

Figure 1 depicts the Dependable Multiprocessor hardware architecture. The basic architecture consists of a redundant radiation-hardened system controller which acts as the controller for a parallel processing cluster of COTS-based, high-performance, data processing nodes, and a redundant



* Mass Data Storage Unit, Custom Spacecraft I/O, etc.

Figure 1 – Dependable Multiprocessor Hardware Architecture

network interconnect. In the current implementation, the radiation-hardened system controller is a Honeywell 603e PPC Rad Hard SBC, and the high-performance data processing nodes are PPC750FX compute nodes with FPGA co-processor accelerators. The interconnection of the parallel processing cluster is via Gigabit Ethernet. The system can be augmented with mission-specific elements, including mass storage, custom interfaces, and radiation sensors, as required.

A top-level overview of the Dependable Multiprocessor software architecture is illustrated in Figure 2. A key feature of this software architecture is the incorporation of a set of generic fault tolerant middleware techniques implemented in a software framework that is independent of and transparent to the specific-mission application, and independent of and transparent to the underlying platform (HW and Operating System). This independence and transparency is achieved through well-defined, high-level, application interfaces, an API (Application Programming Interface) to support mission-specific application needs, and an SAL (System Abstraction Layer) which isolates the remainder of the software system from the underlying platform, simplifying the porting of this software system to other platforms and allowing the generic fault tolerance middleware services to be available to future mission applications on future onboard processing platforms.

More information on the Dependable Multiprocessor and related technologies can be found in references [1] – [14].

TECHNOLOGY BENEFITS

The goal of the Dependable Multiprocessor project is to provide spacecraft/payload processing capability 10x – 100x what is available today. Figure 3 depicts the potential benefits of Dependable Multiprocessor Technology applied to the IOMI (Indian Ocean Meteorological Instrument) project which was performed in conjunction with the NMP EO3 GIFTS (Geosynchronous Imaging Fourier Transform Spectrometer) effort. A comparison of performance and performance density for a 1K complex FFT benchmark is provided for today’s technology shown above the dotted line and Dependable Multiprocessor technology shown below the dotted line. The FFT example was chosen because it is a familiar benchmark and is a function found in many science applications. One of the key elements of the Dependable Multiprocessor implementations is the high Reliability and high Availability provided by the Dependable Multiprocessor Fault Tolerant Middleware and supporting fault tolerance techniques such as Replicated Services and Algorithm-Based Fault Tolerance (ABFT).

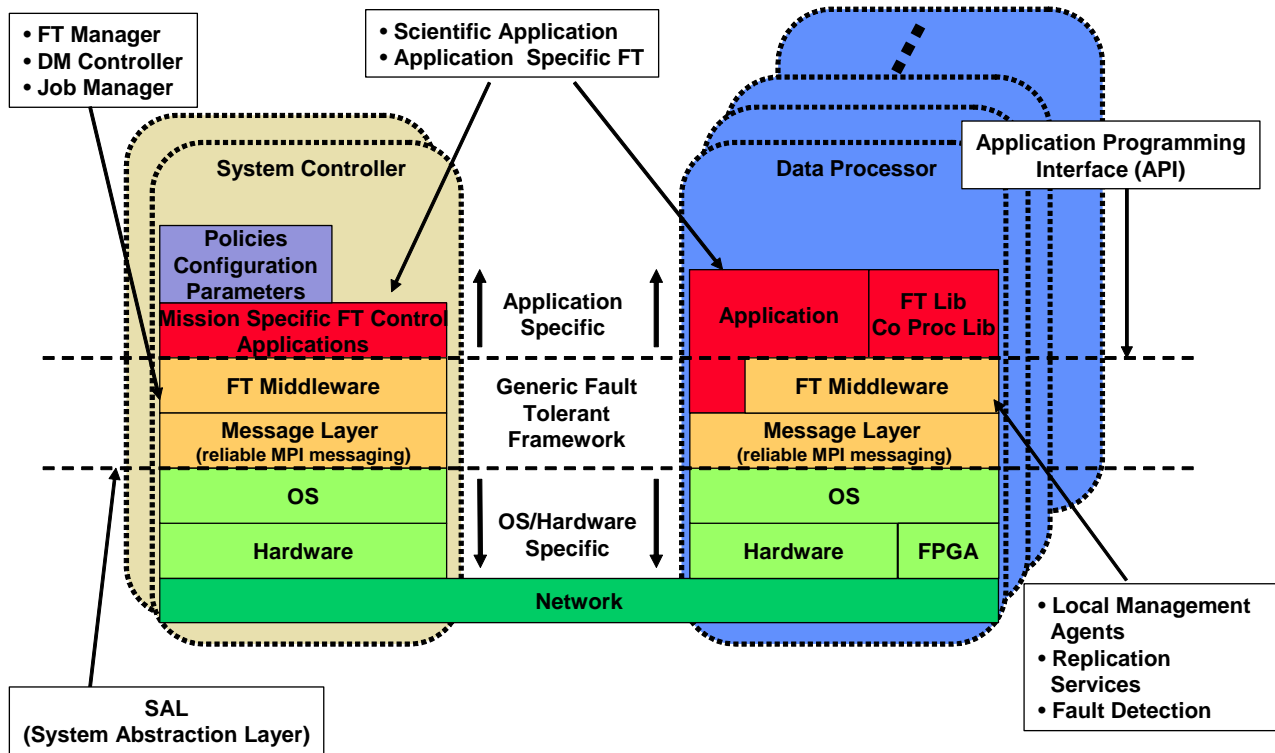




Figure 2 - Dependable Multiprocessor Software Architecture

**NMP EO3 Geosynchronous Imaging Fourier Transform Spectrometer Technology
Indian Ocean Meteorological Instrument (IOMI) - NRL**

Radiation Tolerant 750 PPC SBC		133 MHz ~ 266 MOPS ~ 1.2 kg	1K Complex FFT in ~ 448 μ sec ~ 13 MOPS/watt
Radiation Hardened Vector Processor		DSP24 @ 50 MHz ~ 1000 MOPS ~ 1.0 kg	1K Complex FFT in ~ 52 μ sec ~ 45 MOPS/watt

NMP ST8 Dependable Multiprocessor Technology


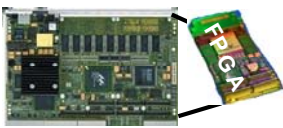
Dependable Multiprocessor 750FX PPC SBC only		1 GHz ~ 1500 MOPS ~ 1.2 kg	1K Complex FFT in ~ 45 μ sec ~ 75 MOPS/watt
Dependable Multiprocessor 750FX PPC SBC with FPGA Accelerator		1 GHz ~ 10000 MOPS ~ 1.4 kg	1K Complex FFT in ~ 6 μ sec ~ 400 MOPS/watt

Figure 3 – Dependable Multiprocessor Technology Benefit Example: Comparison of NMP ST8 Dependable Multiprocessing Technology and Technology That Would Be Flying Today on NMP EO3

3. TECHNOLOGY VALIDATION PLAN

The overall Dependable Multiprocessor technology validation plan is depicted in Figure 4, starting with the TRL4 validation at the end of Phase A, the Concept Formulation Phase. This is followed by the TRL5 validation at the end of Phase B, the Formulation Refinement Phase, the TRL6 validation at the end of the Phase C, the Design and Build Phase, and culminates with the TRL7 validation in the Flight Experiment Operations Phase. Each new validation level is characterized by increasing system fidelity and integration. One of the key elements of TRL5 is the development and validation of models which can be used to predict the performance, reliability, and availability of the Dependable Multiprocessor flight experiment and future NASA missions. The TRL6 and TRL7 experiments will refine and validate the models and the parameters used in the models. After a successful TRL4 demonstration at the end of Phase A, which proved the underlying environmental monitoring and reconfiguration capabilities of the system, NASA requested that the TRL5 effort focus more on high-performance, fault-tolerant cluster processing with fault-tolerant MPI (Message Passing Interface) capability to

provide a software development environment that is familiar to NASA science application developers.

Models

One of the key Dependable Multiprocessor project deliverables is the set of models which can be used to predict Dependable Multiprocessor performance in future NASA missions in different radiation environments, in different orbits, and with technology upgrades, and descriptions of how to use them. The objective of the TRL5, TRL6, and TRL7 technology validation experiments is to validate the models and the parameters used in the models. There are five (5) basic models: the Canonical Fault Model, the Radiation Effects and Hardware SEU Susceptibility Model, the Availability Model, the Reliability Model, and the Performance Model. Figure 5 depicts the Dependable Multiprocessor modeling flow and the inputs and outputs of each one. The Canonical Fault Model identifies the faults which are used as the basis for the Hardware SEU Susceptibility Model and against which the fault tolerance performance of the Dependable Multiprocessor will be evaluated. The Radiation Effects

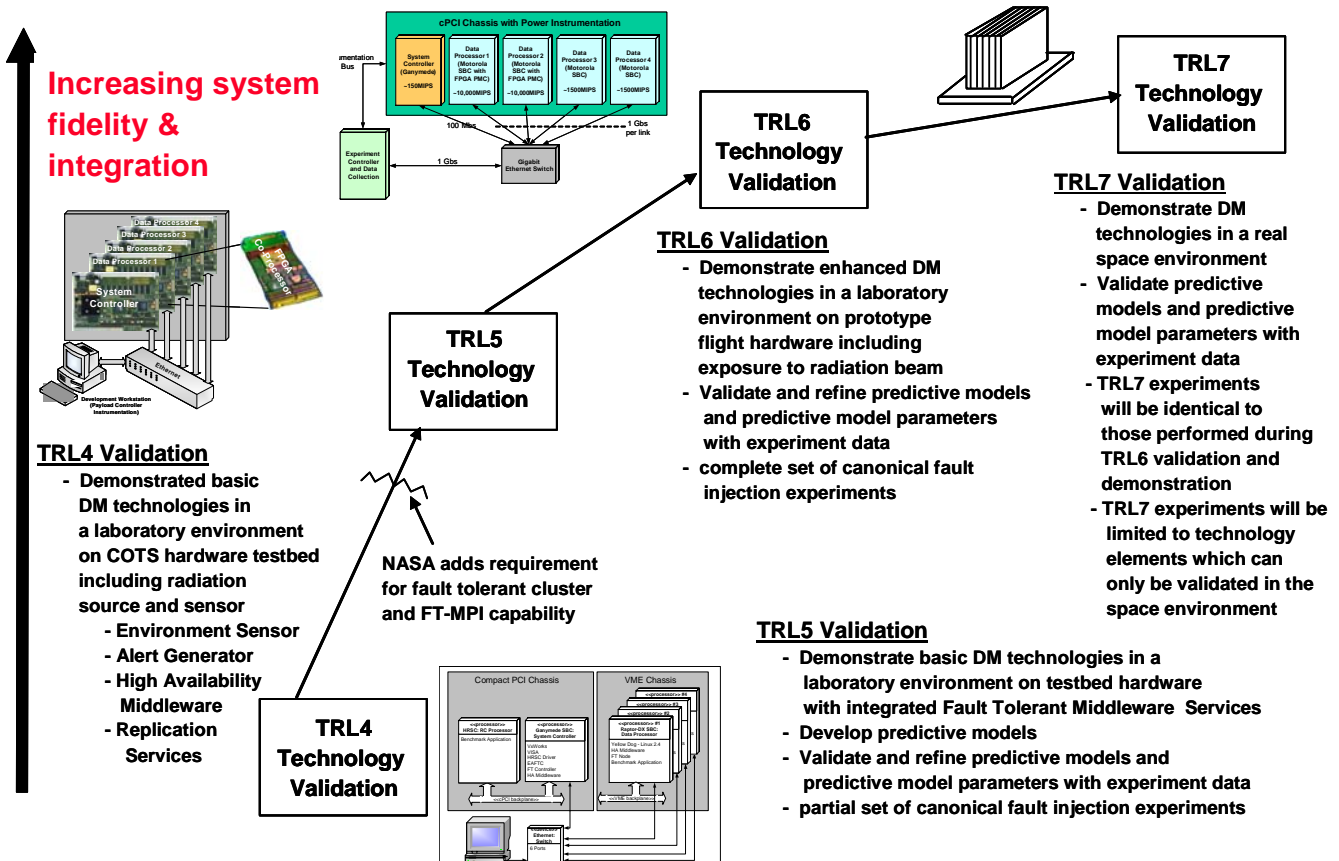


Figure 4 – Dependable Multiprocessor Technology Validation Plan

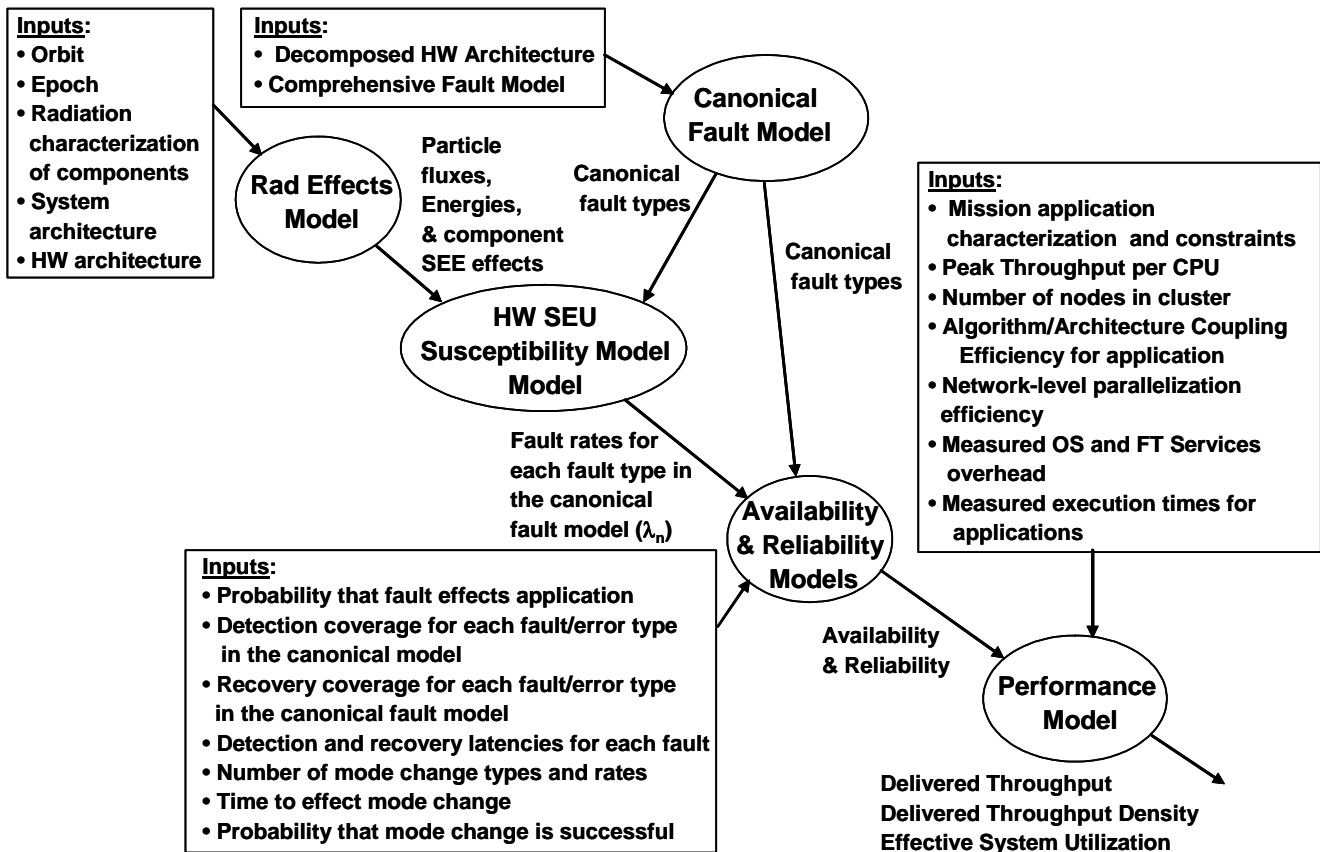


Figure 5 – Dependable Multiprocessing Model Flow

Model takes into account the Dependable Multiprocessor system architecture, the Dependable Multiprocessor hardware architecture, the mission orbit, the mission epoch or time frame, and the radiation characterization of the components. The Radiation Effects Model outputs the expected particle fluxes, energies, and component SEEs (Single Event Effects) for the given orbit. The Hardware SEU Susceptibility Model outputs the fault rates for each fault type in the Canonical Fault Model. The outputs of the Hardware SEU Susceptibility Model are combined with the detection coverage for each fault/error type in the Canonical Fault Model, the recovery coverage for each fault/error type in the Canonical Fault Model, the detection and recovery latencies for each fault/error type in the Canonical Fault Model, the probability that a particular fault affects the application, the number of expected mode changes for the mission, and the time to effect the mode change to predict the Availability and Reliability of the Dependable Multiprocessor in the particular mission application. The outputs of the Reliability and Availability Models are fed into the Performance Model which takes into account the mission application, the peak throughput of the CPUs in the high-performance data processing nodes, the algorithm/

architecture coupling efficiency for the application, the number of nodes in the cluster, the network-level parallelization efficiency, the measured OS and FT services overhead, and the measured execution times for the applications to determine the effective delivered throughput (MOPS), the effective delivered throughput density (MOPS/watt), and the effective system utilization for the mission.

4. TRL4 VALIDATION

At the TRL4 TMA (Technology Maturity Assessment), which was conducted at end of Phase A – the Concept Formulation Phase, the basic environmentally-adaptive technologies were demonstrate on COTS testbed hardware including a radiation source and sensor. This demonstration comprised the functionality of the environment sensor, the environment alert generator, high availability middleware, and high-level replication services, e.g., SCP (Self-Checking Pair) and TMR (Triple Modular Redundancy). The Dependable Multiprocessor system demonstrated the capability to switch from simplex operation to SCP and TMR operation and back as the radiation level was varied.

5. TRL5 VALIDATION

The Dependable Multiprocessor project is currently in the Phase B, Formulation Refinement Phase. The Refinement Phase will culminate with the TRL5 demonstration. The focus of the Formulation Phase is on high-performance, fault tolerant, cluster processing to meet the needs of future science missions. The COTS testbed hardware used in TRL5 is depicted in Figure 6.

The TRL5 system consists of four (4) high-performance COTS data processing nodes, two (2) COTS processors to emulate the redundant radiation hardened system controllers, and redundant Gigabit Ethernet switches. Two of the data processing nodes have FPGA coprocessor accelerators. One of the data processor nodes is used to emulate a payload mass data storage element. The clock rate of the controller nodes are reduced to match the performance of the radiation-hardened controllers in the flight system. LINUX OS is used on all nodes.

A top-level overview of Dependable Multiprocessor software and the partitioning and mapping onto the Dependable Multiprocessor hardware is depicted in Figure 7.

The Dependable Multiprocessor software architecture includes the middleware layers which provide fault tolerance for the cluster and a thin isolation layer which makes porting between platforms a minimal and straightforward process. Dependable Multiprocessor Fault Tolerant Middleware includes COTS High Availability Middleware (HAM), and the Dependable Multiprocessor Job Management Services (JMS), Fault Tolerance Management Services (FTMS), Fault-tolerant Embedded Message Passing Interface (FEMPI), Environment Sensor Manager (ESM), FPGA Co-Processor Services (FCPS), Replication Services (RS), and Checkpoint and Rollback (CR) functions. The Job Management Services function consists of the Job Manager (JM) which executes on the system controller node and the Job Manager Agents (JMAs) which execute on the high-performance data processing nodes. Correspondingly, the Fault Tolerance Management Services function consists of the Fault Tolerance Manager (FTM) which executes on the system controller and the Fault Tolerance Management Agents (FTMAs) which execute on the high-performance data processing nodes. The Environment Sensor Manager (ESM) combines spacecraft ephemeris, environment sensor measurements, detected error types and rates obtained from the Fault Tolerance Manager, and operational history of the spacecraft to generate environmental alerts. The environment alerts, combined with the established mission rules and policies, guide the Job Manager in configuring the

Dependable Multiprocessor for the given environment. In addition to adapting to the radiation environment, the Dependable Multiprocessor can also adapt to different mission operation modes and functional criticality. The COTS HAM functions include the basic cluster management services, availability management services, replicated data base services, and data messaging services including reliable communications. The Dependable Multiprocessor Fault Tolerant Middleware components execute on top of the LINUX OS.

In addition to validating the basic Dependable Multiprocessor functionality, the TRL5 experiments will include measurement of parameters needed to validate the performance, reliability, and availability models. These parameters include: the detection and recovery coverage for each fault/error type in the Canonical Fault Model, the detection and recovery latencies for each fault/error type in the Canonical Fault Model, the probability that a particular fault affects the application, the number of expected mode changes for the mission, the time to effect the mode changes, the peak throughput of the CPUs in the high-performance data processing nodes, the algorithm/architecture coupling efficiency for the application, the network-level parallelization efficiency, the OS and Fault-Tolerance services overhead, and the measured execution times for the applications, with and without FPGA Co-Processor acceleration, to determine the system reliability, the system availability, the effective delivered throughput (MOPS), the effective delivered throughput density (MOPS/watt), and the effective system utilization for the mission.

6. TRL6 VALIDATION

The TRL6 technology validation will be performed on increasingly higher fidelity hardware. For the TRL6 validation, the COTS-emulated System Controller will be replaced by a prototype Rad Hard System Controller. In addition, the TRL6 system will be taken to a radiation beam test facility where one of the high-performance data processing nodes will be exposed to a particle beam as shown in Figure 8 to further validate the technology and the flight experiment design.

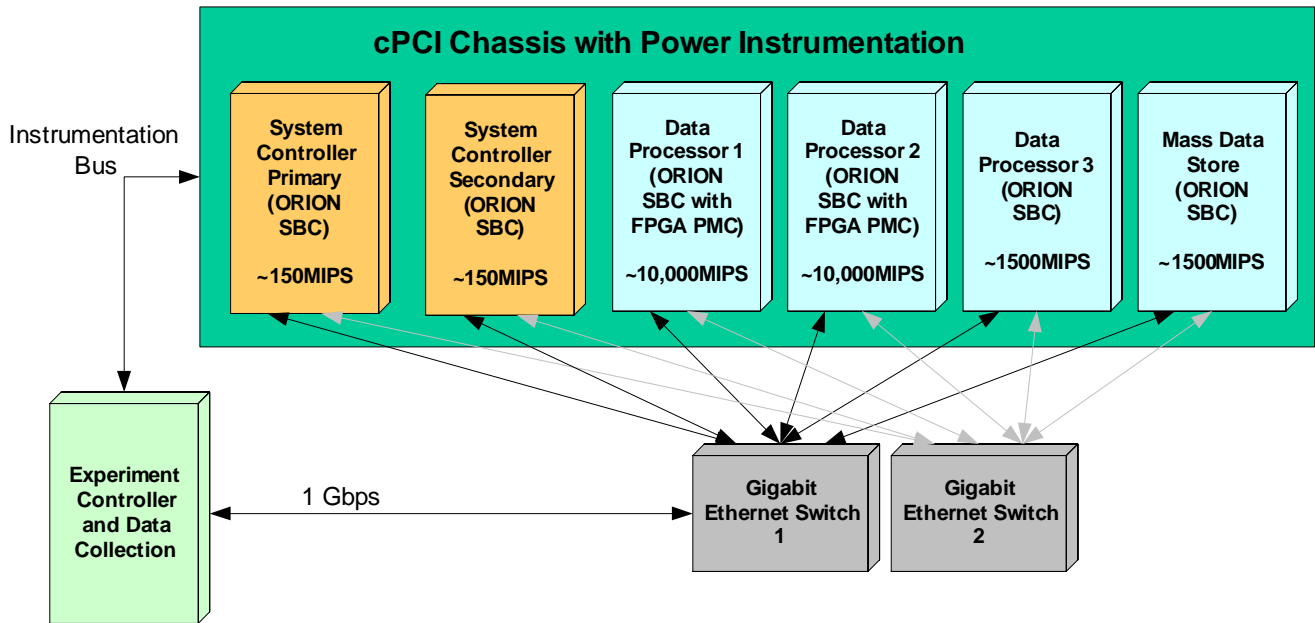
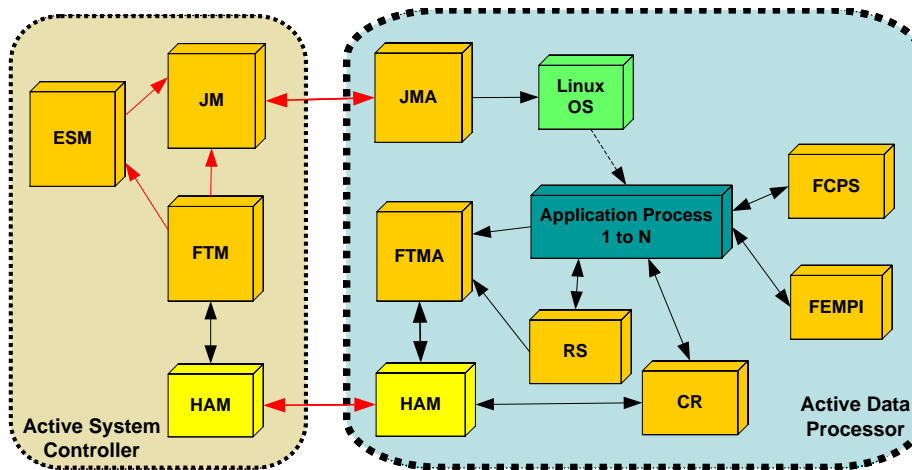


Figure 6 - TRL5 Hardware Testbed System

- **Dependable Multiprocessor Middleware Components**

- Environmental Sensor Manager (ESM)
- Job Management Services (JMS)
 - Job Manager (JM) + Job Management Agent (JMA)
- Fault Tolerance Management Services (FTMS)
 - Fault Tolerance Manager (FTM) + Fault Tolerance Management Agent (FTMA)
- High Availability Middleware Services (HAM)
- Fault-tolerant Embedded Message Passing Interface (FEMPI)
- FPGA Co-Processor Services (FCPS)



Links in Red are HAM DMS based communication links.

Figure 7 – Dependable Multiprocessor Middleware Components, Partitioning, and Mapping

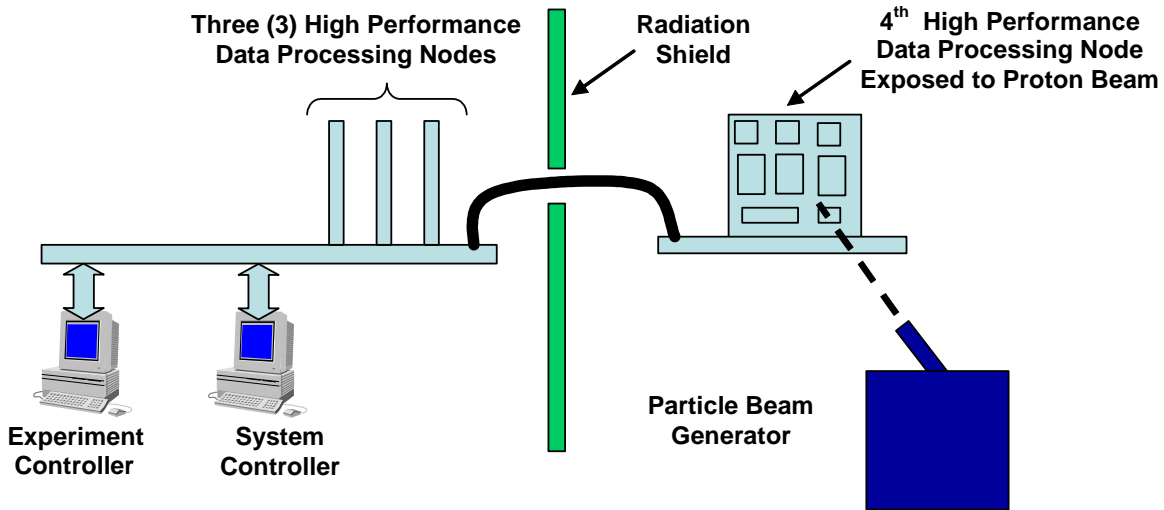


Figure 8 – TRL6 Particle Beam Testing

7. TRL 7 FLIGHT VALIDATION EXPERIMENT

The proposed ST8 spacecraft is depicted in Figure 9. The Dependable Multiprocessor (DM) experiment will share the spacecraft bus with three (3) other ST8 experiments: 1) the NGU (Next Generation Ultra-flex) deployable solar array experiment, 2) the MLHP (Multiple Loop Heat Pipe) experiment, and 2) the deployable SAILMAST experiment. Orbital Sciences Corporation is under contract to provide the ST8 spacecraft bus.

The Dependable Multiprocessor experiment configuration is depicted in Figure 10. The Space Segment comprises the S/C bus and the Dependable Multiprocessor experiment payload. The Ground Segment comprises the NASA ST8 mission ground facility, which consists of two elements, the USN (Universal Space Network) which will provide the communication link between the ground and the S/C, and the MOC (Mission Operations Center), which also will be provided by OSC, and the Experiment Control facility at Honeywell. Experiment command requests will be forwarded to the spacecraft through the Mission Operations Center. Experiment telemetry and data received from the spacecraft will be transmitted over an Internet link to Honeywell where data reduction and analysis will be performed.

The objectives of the Dependable Multiprocessor flight experiment are four-fold:

- 1) to expose a COTS-based, high-performance processing cluster to the real space radiation environment,
- 2) to characterize the radiation environment,
- 3) to correlate the radiation performance of the COTS components with the environment, and
- 4) to assess the radiation performance of the COTS components and the Dependable Multiprocessor system response in order to validate the predictive Reliability, Availability, and Performance models for the Dependable Multiprocessor flight experiment and for future NASA missions.

The highly-inclined (98.5°), elliptical ST8 mission orbit with a planned apogee of at least 1300 km and a perigee of 300 km was selected to maximize the data collection capability for the Dependable Multiprocessor experiment.

Except for some power-up and initialization testing, whenever the Dependable Multiprocessor payload is powered up, Dependable Multiprocessor operation is expected to be a free running experiment, collecting radiation environment characterization and radiation event (SEU) data, correlating the environment and detected events with S/C ephemeris, and monitoring and reporting Dependable Multiprocessor response. The Dependable Multiprocessor experiment is expected to be run continuously for four of the six month ST-8 mission to maximize the amount of data collected.




Figure 9 – ST8 Spacecraft and Experiment

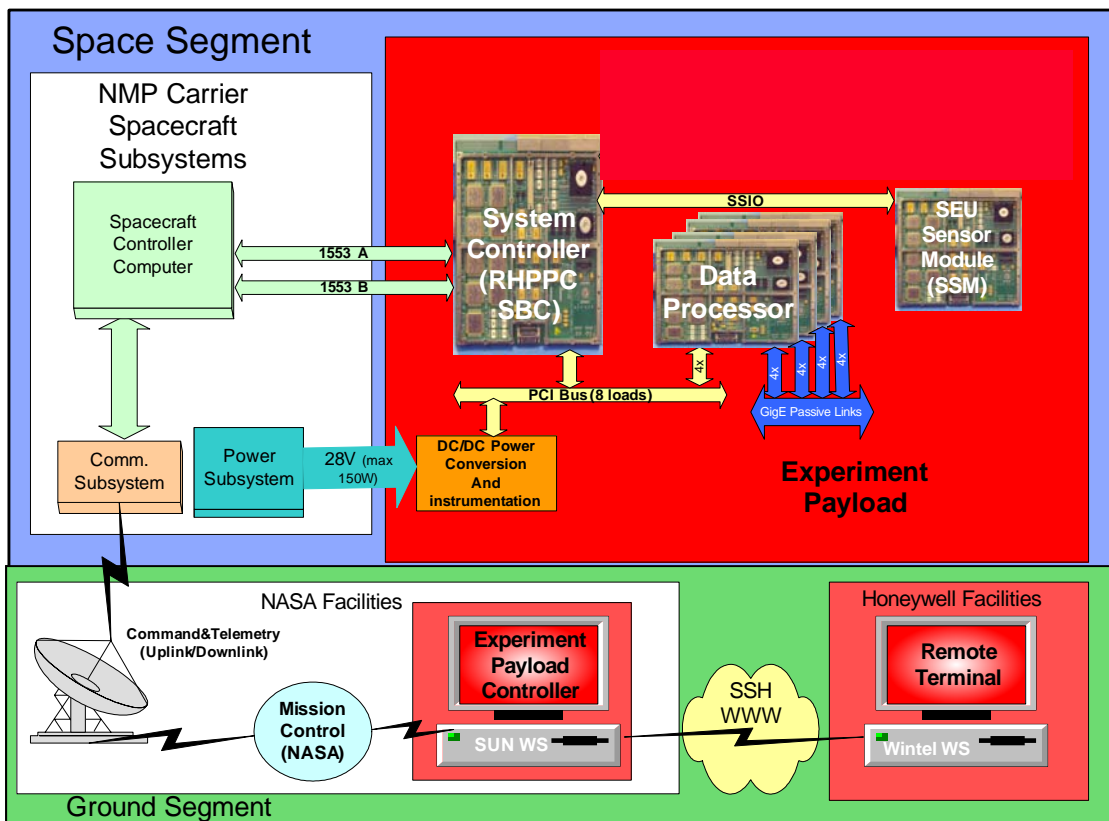


Figure 10 – Dependable Multiprocessor Experiment Configuration

The Dependable Multiprocessor flight experiment will encompass measurement of component and system parameters that can only be validated in a real space environment. Primarily, these are the component fault/error rates due to radiation, and the accuracy of the predictive fault/error model. Table 1 shows the primary data to be taken during the flight experiment. These data will support the experiment objectives identified in the previous paragraphs. The spacecraft ephemeris will be used to correlate the radiation performance of the COTS components with the orbit location. Other technology validation data, cluster performance, cluster performance density, error detection and recovery latencies, Operating System overhead, and Fault Tolerant Middleware overhead, will be collected in the TRL5 and TRL6 ground-based technology validation experiments. These parameters do not need to be re-validated in space because they are not expected to change from the values measured during the ground-based experiments.

Figure 11 is a high level depiction of the flight experiment mode and operational flow. The environmental sensor data will be sampled continuously throughout the course of the experiment. However, in order to minimize the data storage and downlink bandwidth requirements, the data will be stored in a circular buffer. The data will only be captured for downlinking upon detection of an event, detection of an alarm, a sensed change in the radiation environment, a specific ephemeris of interest, e.g., over the poles or in the vicinity of the South Atlantic Anomaly (SAA), or upon request from the ground. The circular buffer will contain the most recent frames of data to correlate the event with the radiation event that led up to the event.

There is no real science mission instrument in the Dependable Multiprocessor flight experiment. Synthetic

science application data will be processed continuously. The processed output will be compared with known correct output to determine if an error occurred which was not detected by the Dependable Multiprocessor system. The data collected will be stored and downlinked when the S/C is in view of one of the ground stations.

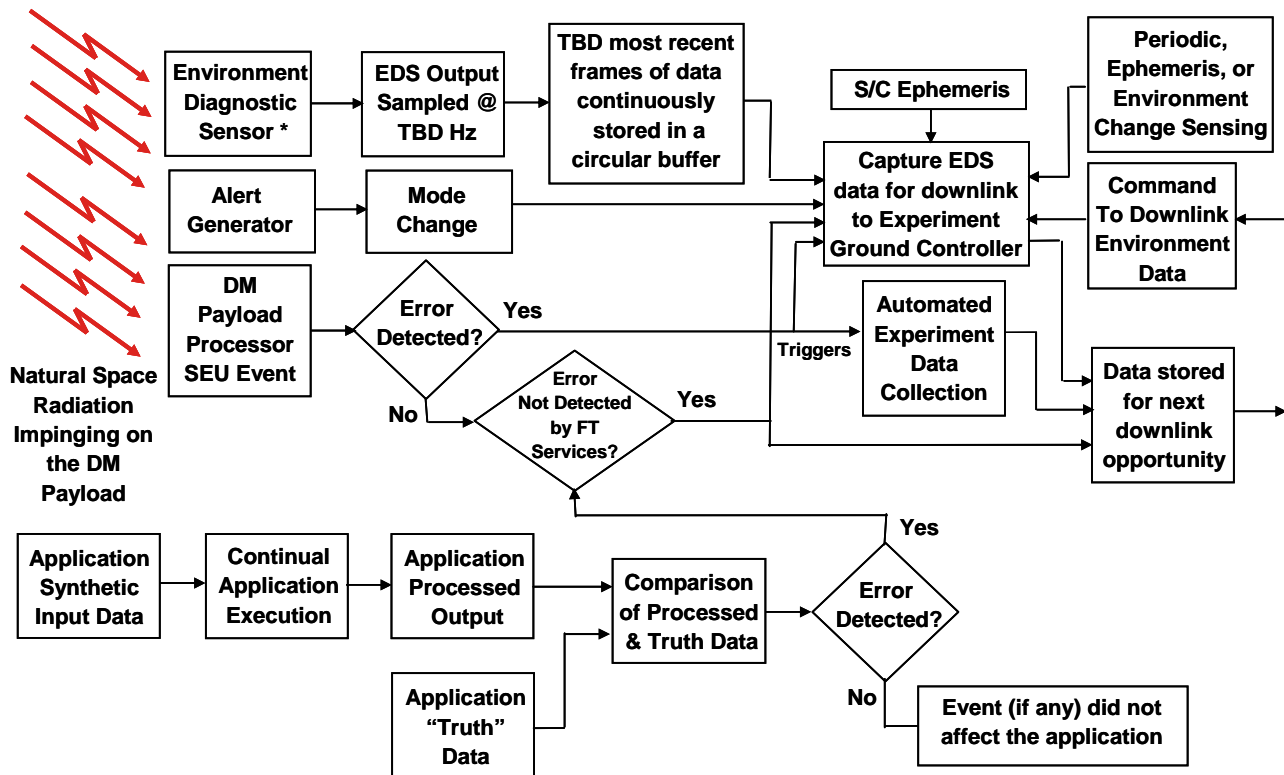
Although the Dependable Multiprocessor flight experiment is years away, planning for the flight experiment is underway to understand the data storage and downlink requirements. In order to validate the reliability and availability models, the basic data measurements that need to be collected during the Dependable Multiprocessor space experiment are identified in Table 2.

8. CURRENT STATUS

The Dependable Multiprocessor project is currently in the Phase B – Formulation Refinement Phase. The Phase B effort will culminate in the TRL5 technology validation demonstration in May of 2006. Currently at the mid-point of the effort, the JMS (Job Management Services), the FTMS (Fault Tolerance Management Services), the FEMPI (Fault tolerant Embedded Message Passing Interface), and the FPGA Co-Processor Services (FCPS) functions have been demonstrated in an integrated system environment. The remaining tasks to be done in Phase B include: the incorporation and demonstration of ABFT (Algorithm-Based Fault Tolerance) and Replication Services functions, and the execution of a comprehensive SWIFI (Software Implemented Fault Injection) campaign which will be used to analyze and evaluate system performance, and to validate the predictive reliability, availability, and performance models.

Table 1 - Experiment Data Categories

Experiment Data Categories
• Spacecraft Ephemeris
• Fault rates, types, and error detection and recovery coverages
• SEU Alarm Measurements
• Environment Diagnostic Sensor Measurements



* The Environment Diagnostic Sensor is not part of the Dependable Multiprocessor technology validation. It is needed for correlation of the occurrence of SEU events and the radiation environment, and for calibration of the Radiation Effects/HW SEU Susceptibility Models

Figure 11 - Experimental Mode/Operational Concept

Table 2 – Flight Experiment Data

Radiation Characterization Data	SEU Event Data
Time Stamp	Time Stamp
S/C Ephemeris	S/C Ephemeris
Heavy Ion Count per Energy Bin	Detected Memory Upset
Proton Count per Energy Bin	Detected Error Type/Location
	Recovery Time
	Undetected Application Error

9. SUMMARY AND CONCLUSION

The goal of the Dependable Multiprocessor project is to provide spacecraft/payload processing capability 10x – 100x what is available today, enabling heretofore unrealizable science and autonomy. Dependable Multiprocessor technology is a key enabler for future NASA science missions including increased autonomy for remote exploration, landing support, and lunar or Martian surface rovers. Over the past few generations, Commercial-Off-The-Shelf (COTS) computer components have become

more resistant to the debilitating effects of radiation. Many commercial parts can withstand many 10s of kilorads of Total Ionizing Dose (TID) and are immune to catastrophic Single Event Latchup (SEL). The primary issue preventing the deployment of a COTS-based space-borne cluster computer is their continued susceptibility to Single Event Upsets or SEUs, which cause only soft, transient errors, not permanent hardware failures. Further, the latest generation of computer electronics, SOI CMOS, has proven to be approximately an order of magnitude less susceptible to SEU than previous bulk CMOS. If Dependable

Multiprocessor technology allows a system to withstand a few errors per day per processor, without unduly impacting system dependability, it will be possible to fly, essentially commercial, cluster computers. Not only would this provide mission enabling performance and performance density levels, but it would significantly lower the cost of development, as standard laboratory science codes could be easily ported to these systems without the expensive and error prone process normally associated with moving complex codes from the lab to a new flight platform.

Migrating high-performance COTS processing to space is not a new idea. A key element of the Dependable Multiprocessor project, which distinguishes it from previous attempts to migrate COTS to space, is that NASA is providing the ride. NASA has already issued contracts for the spacecraft, the launch vehicle, and the ground facilities. The Dependable Multiprocessor experiment only needs to get through the remaining NASA and NMP gates to realize the goal of flying COTS high-performance computing in space.

REFERENCES

- [1] Ramos, Jeremy, J. Samson, M. Patel, A. George, and R. Some, "High Performance, Dependable Multiprocessor," *Proceedings of the 2006 IEEE Aerospace Conference*, Big Sky, MN, March 4-11, 2006.
- [2] Greco, James, G. Cieslewski, A. Jacobs, I. Troxel, C. Conger, J. Curreri, and A. George, "Hardware/Software Interface for High-performance Space Computing with FPGA Coprocessors," *Proceedings of the 2006 IEEE Aerospace Conference*, Big Sky, MN, March 4-11, 2006.
- [3] Samson, Jr. John R., J. Ramos, A. George, M. Patel, and R. Some, "Environmentally-Adaptive Fault Tolerant Computing (EAFTC)," 9th High Performance Embedded Computing Workshop, M.I.T. Lincoln Laboratory, September 22, 2005.
- [4] Ramos, Jeremy, and D. Brenner, "Environmentally-Adaptive Fault Tolerant Computing (EAFTC): An Enabling Technology for COTS based Space Computing," *Proceedings of the 2004 IEEE Aerospace Conference*, Big Sky, MN, March 8-15, 2004.
- [5] Samson, Jr. John R., "Migrating High Performance Computing to Space," 7th High Performance Embedded Computing Workshop, M.I.T. Lincoln Laboratory, September 22, 2003.
- [6] Samson, Jr., John R., "Space Touchstone Experimental Program (STEP) – Final Report 002AD," January 15, 1996.
- [7] Karapetian, Arbi, R. Some, and J. Behan, "Radiation Fault Modeling and Fault Rate Estimation for a COTS Based Space-borne Computer," *Proceedings of the 2002 IEEE Aerospace Conference*, Big Sky, MN, March 9-16, 2002.
- [8] Some, Raphael, W. Kim, G. Khanoyan, and L. Callum, "Fault Injection Experiment Results in Space Borne Parallel Application Programs," *Proceedings of the 2002 IEEE Aerospace Conference*, Big Sky, MN, March 9-16, 2002.
- [9] Some, Raphael, J. Behan, G. Khanoyan, L. Callum, and A. Agrawal, "Fault-Tolerant Systems Design Estimating Cache Contents and Usage," *Proceedings of the 2002 IEEE Aerospace Conference*, Big Sky, MN, March 9-16, 2002.
- [10] Lovellette, Michael, and K. Wood, "Strategies for Fault-Tolerant, Space-Based Computing: Lessons Learned for the ARGOS Testbed," *Proceedings of the 2002 Aerospace Conference*, Big Sky, MN, March 9-16, 2002.
- [11] Samson, John R., and C. Markiewicz, "Adaptive Resource Management (ARM) Middleware and System Architecture – the Path for Using COTS in Space," *Proceedings of the 2000 IEEE Aerospace Conference*, Big Sky, MN, March 8-15, 2000.
- [12] Samson, Jr., John R., L. Dela Torre, J. Ring, and T. Stottlar, "A Comparison of Algorithm-Based Fault Tolerance and Traditional Redundant Self-Checking for SEU Mitigation," *Proceedings of the 20th Digital Avionics Systems Conference*, Daytona Beach, Florida, 18 October 2001.
- [13] Samson, Jr., John R., "SEUs from a System Perspective," Single Event Upsets in Future Computing Systems Workshop, Pasadena, CA, May 20, 2003.
- [14] Prado, Ed, J. R. Samson, Jr., and D. Spina, "The COTS Conundrum," *Proceedings of the 2000 IEEE Aerospace Conference*, Big Sky, MN, March 9-15, 2003.

ACKNOWLEDGEMENTS

The authors would like to thank the following people and organizations for their contributions to the Dependable Multiprocessor effort: Sherry Akins, Dr. Mathew Clark, Lee Hoffmann, David Lupia, and Roger Sowada of Honeywell Aerospace, Defense & Space; Paul Davis and Vikas Aggarwal from Tandel Systems; a team of researchers in the High-performance Computing and Simulation (HCS) Research Laboratory at University of Florida led by Dr. Alan George. Members of the team at UF include: Ian Troxel, Raj Subrmanian, John Curreri, Mike Fisher, Grzegorz Cieslewski, Adam Jacobs, and James Greco; Brian Heigl, Paul Arons, Gavin Kavanaugh, and Mike Nitso, from GoAhead Software, Inc. Other members of the team are Dr. Ravishankar Iyer and Dr. Zbigniew Kalbarczyk from the University of Illinois and Armored Computing.

The Dependable Multiprocessor effort is funded under NASA NMP ST-8 contract NMO-710209.

BIOGRAPHIES



John R. Samson, Jr. is a Principal Engineering Fellow with Honeywell Aerospace in Clearwater, Florida. He received his Bachelor of Science Degree from the Illinois Institute of Technology, his Master of Science degree and the Degree of Electrical Engineer from the Massachusetts Institute of Technology, and his Ph.D. in Engineering Science with Specialization in Computer Science from the University of South Florida. During his 35-year career, John has worked at M.I.T. Lincoln Laboratory, Raytheon Company Equipment Division, and Honeywell Space Systems. His work has encompassed multiple facets of ground, airborne and space-based surveillance system applications. He has spent most of his career developing onboard processors, onboard processing architectures, and onboard processing systems for real-time and mission critical applications. He was Principal Investigator for a pioneering study investigating the feasibility of migrating high-performance COTS processing to space, a predecessor of the work described in this paper. Dr. Samson is the Principal Investigator for the Dependable Multiprocessor project. He is an Associate Fellow of the AIAA and a Senior Member of the IEEE.



Jeremy Ramos has been a Honeywell Aerospace employee since 1999. Mr. Ramos received a B.S. degree in Computer Science and Engineering from the University of South Florida in 1999. Prior to his engineering career Mr. Ramos served for 7+ years with the

United States Army as a Technician in the Army Ordnance Corps. Mr. Ramos is Technical Director of the Honeywell Dependable Multiprocessor project. Among his peers Mr. Ramos is considered as an expert in the areas of computer architecture, system simulation, and reconfigurable computing.



Minesh I. Patel is a systems and software architect and consultant with Tandel Systems in Clearwater, Florida. He received his BSEE and BSCpE in electrical and computer engineering and his MSCpE and Ph.D. in Computer Science and Engineering from the University of South Florida. His research and technical interests include software and system fault tolerance, artificial intelligence and machine learning, embedded and real-time systems and high-performance, parallel and distributed computing. Dr. Patel is Lead Software Architect for the Dependable Multiprocessor project.



Alan D. George is Professor of Electrical and Computer Engineering at the University of Florida, where he serves as Director and Founder of the High-performance Computing and Simulation (HCS) Research Laboratory. He received the B.S. degree in Computer Science and the M.S. in Electrical and Computer Engineering from the University of Central Florida, and the Ph.D. in Computer Science from the Florida State University. Dr. George's research interests focus on high-performance architectures, networks, services, and systems for parallel, reconfigurable, distributed, and fault-tolerant computing. He is a senior member of IEEE and SCS, and can be reached by e-mail at george@hcs.ufl.edu. Dr. George is the Principal Investigator for the Dependable Multiprocessor software research and development effort at the University of Florida.



Raphael Some is a program technologist at JPL for the New Millennium Program. He has served as Contract Technical Manager and Leader of the Technology Review Board for the ST8 Dependable Multiprocessor project. Prior to his involvement with the NMP ST8 project, Mr. Some was the Chief Engineer for the Remote Exploration and Experimentation Project at the Jet Propulsion Laboratory. Previously, at JPL, he formulated and managed the Smart Sensor project. His experience prior to JPL includes the development of fault tolerant space based supercomputers as well as a variety of avionics and signal processing systems for both commercial and military applications. He holds a BSEE from Rutgers University.