Dependable Multiprocessor (DM) Architecture for Space Applications

Dr. John R. Samson, Jr.

Honeywell Aerospace, Defense & Space, Clearwater, FL

john.r.samson@honeywell.com

2nd Annual Fault Tolerant Space Computing Workshop

Sandia National Laboratory, Albuquerque, NM

27 May 2009

Honeywell

Outline

Honeywell

- **Introduction**
  - Dependable Multiprocessor* technology
    - Overview
      - -- Current technological shortcomings
      - -- COTS in space
      - -- Goals and Objectives
    - Hardware architecture
    - Software architecture
- **Current Status**
- **TRL6 Technology Validation**
- **Summary & Conclusion**

*** formerly known as the Environmentally-Adaptive Fault-Tolerant Computer (EAFTC); The Dependable Multiprocessor effort is funded under NASA NMP ST8 contract NMO-710209.**

Technological Shortcomings

- In terms of performance, current radiation-hardened technologies are 2-3 generations behind state-of-the-art devices designed and developed for terrestrial applications
- Long-held NASA and DoD desire to take advantage of Commercial-Off-The-Shelf (COTS) technology to increase science, surveillance, and autonomy capability in space
  - use of COTS parts is desirable due to high performance to cost ratio
  - BUT ...., COTS parts generally are designed for performance, not for power efficiency, fault tolerance, nor with (space) thermal issues and radiation effects susceptibility/vulnerability in mind
  - development and migration of applications from the laboratory to space is slow and costly
    - high non-recurring cost; incompatibility with standard cluster processing application software and parallel processing libraries
  - most COTS solutions are fixed, inflexible, and not power efficient, e.g., hard-wired self-checking or TMR (Triple Modular Redundancy)
- Need a technology and platform-independent solution that can incorporate techniques/technologies which allow us to overcome performance gaps with regards to throughput, power, mass, and cost

  e.g., ABFT (Algorithm-Based Fault Tolerance), FPGA, Rad Hard By Design (RHBD), Rad Tolerant By Software (RTBS)

Dependable Multiprocessor Technology

Honeywell

- Desire - -> ‘Fly high performance COTS multiprocessors in space’
  - To satisfy the long-held desire to put the power of today’s PCs and supercomputers in space, three key challenges, SEUs, cooling, & power efficiency, needed to be overcome
    - Single Event Upset (SEU): Radiation induces transient faults in COTS hardware causing erratic performance and confusing COTS software

      DM Solution { - robust control of cluster; - enhanced, SW-based, SEU-tolerance

    - Cooling: Air flow is generally used to cool high performance COTS multiprocessors, but there is no air in space

      DM Solution { - tapped the airborne-conductively-cooled market

    - Power Efficiency: COTS only employs power efficiency for compact mobile computing, not for scalable multiprocessing

      DM Solution { - tapped the high performance density mobile market

DM has addressed and solved all three issues

UNIVERSITY OF FLORIDA

NOAO

NMP Breakthrough Technologies

DM Goals and Objectives

- **Develop high-performance, COTS-based, fault tolerant cluster onboard processing system that can operate in a natural space radiation environment**
- **Demonstrate high throughput, low power, scalable, & fully programmable processing (≥300 MOPS/watt) (10x – 100x what is flying today)**
- **Demonstrate high system availability (≥ 0.995)**
- **Demonstrate high probability of timely and correct delivery of data (≥ 0.995)**
- **Demonstrate technology-independent system software that manages a cluster of high performance COTS processing elements**
- **Demonstrate technology-independent system software that enhances radiation SEE (Single Event Effect) tolerance**
- **Demonstrate ease of porting of applications from the laboratory to space (support MPI-based cluster processing)**
- **Develop and validate models that can predict system performance for a variety of applications in a range of radiation environments**

DM Hardware Architecture

Honeywell

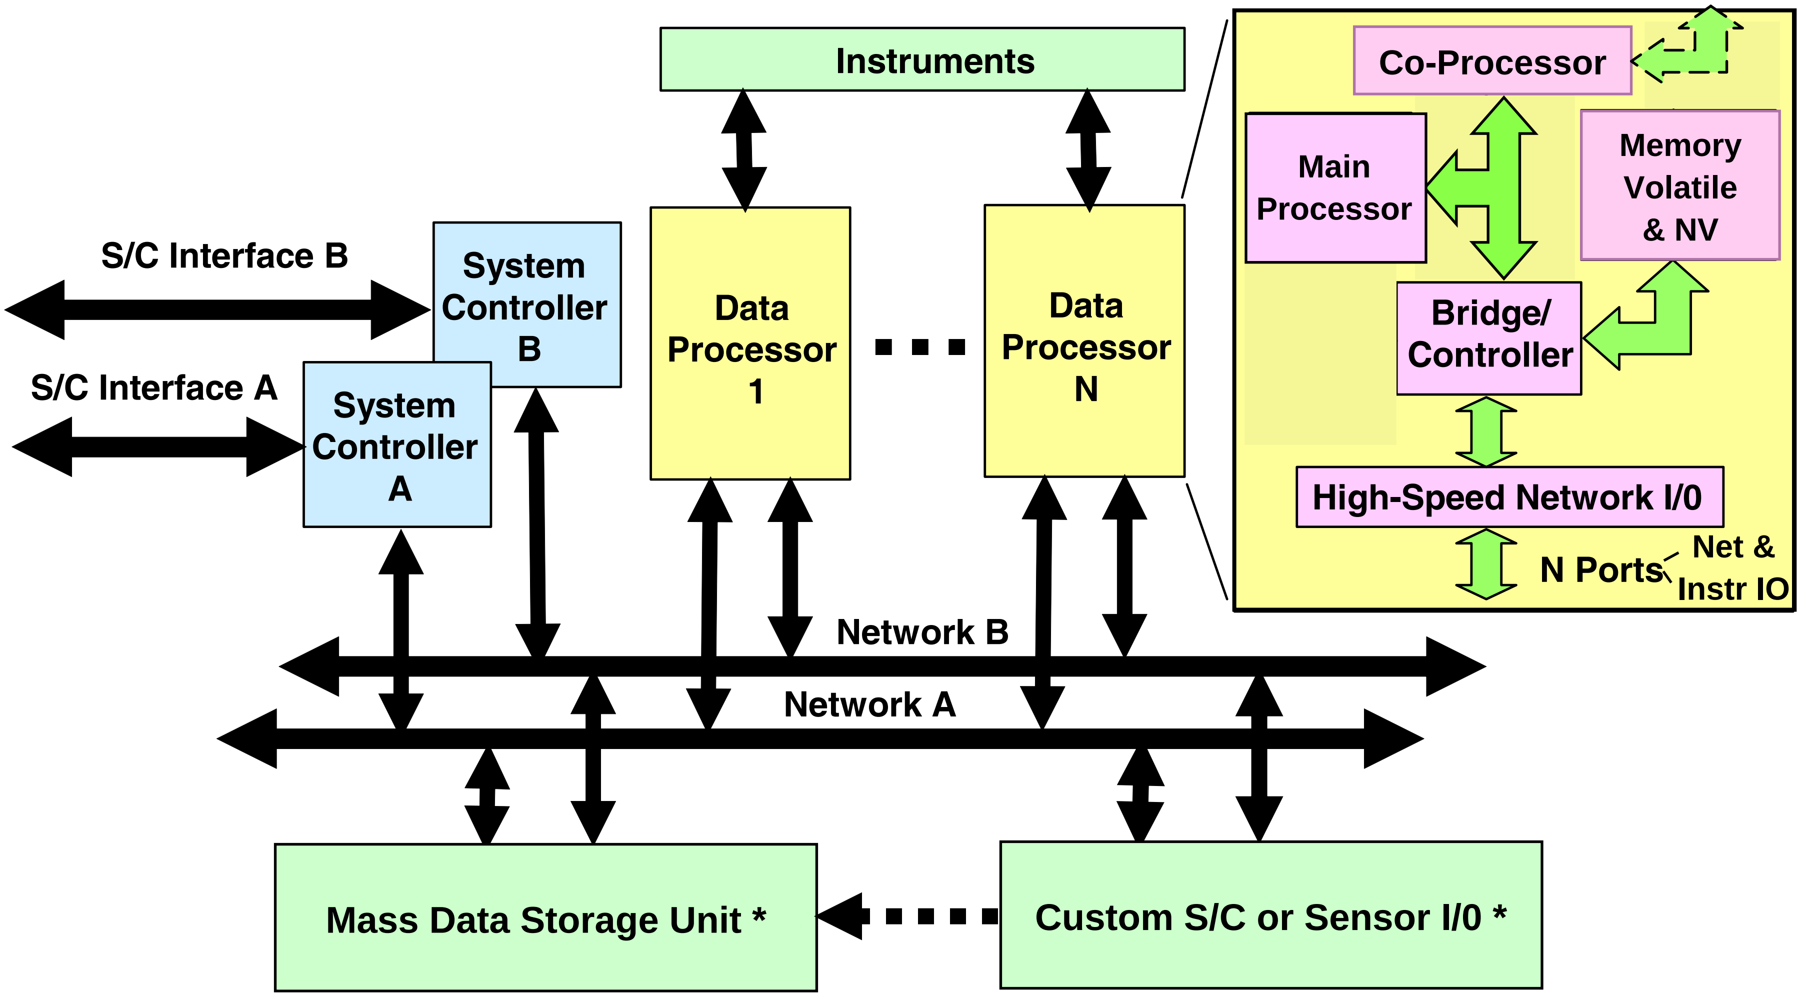

* Examples: Other mission-specific functions

DMM Top-Level Software Layers

Honeywell

DMM – Dependable Multiprocessor Middleware

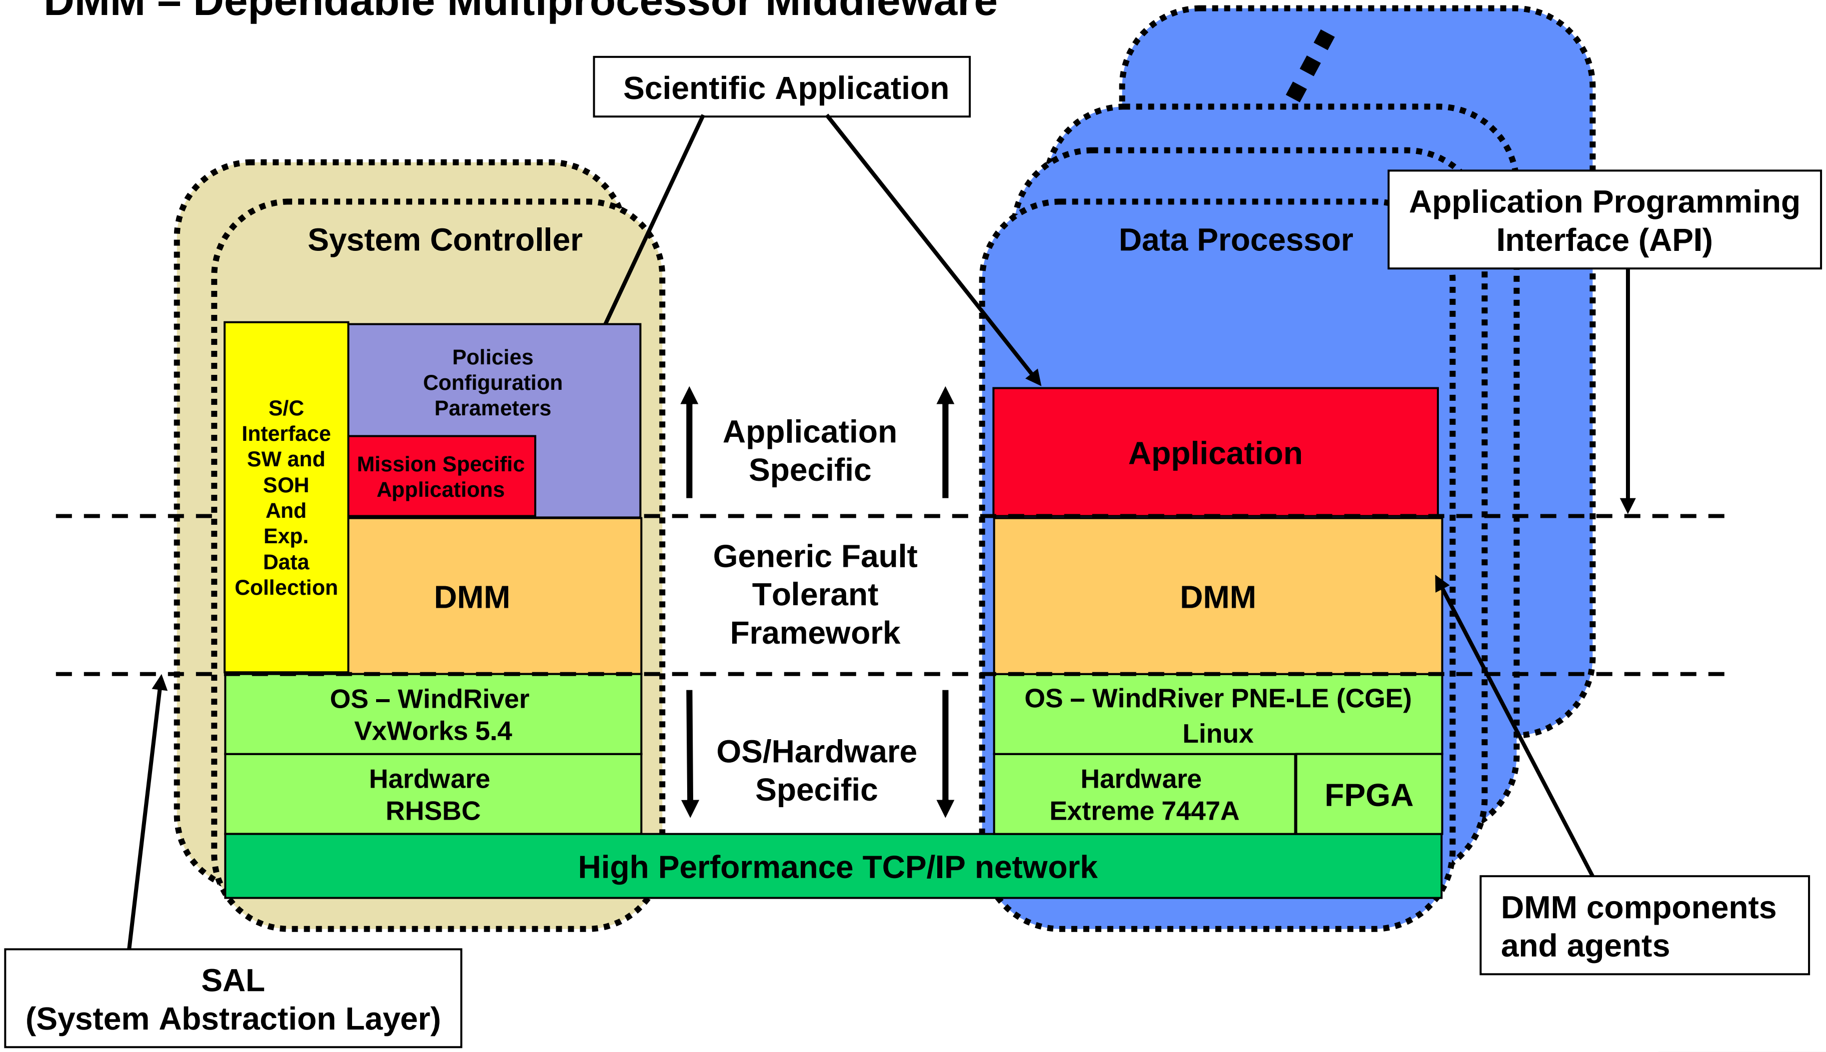

DM Top-Level Software Architecture

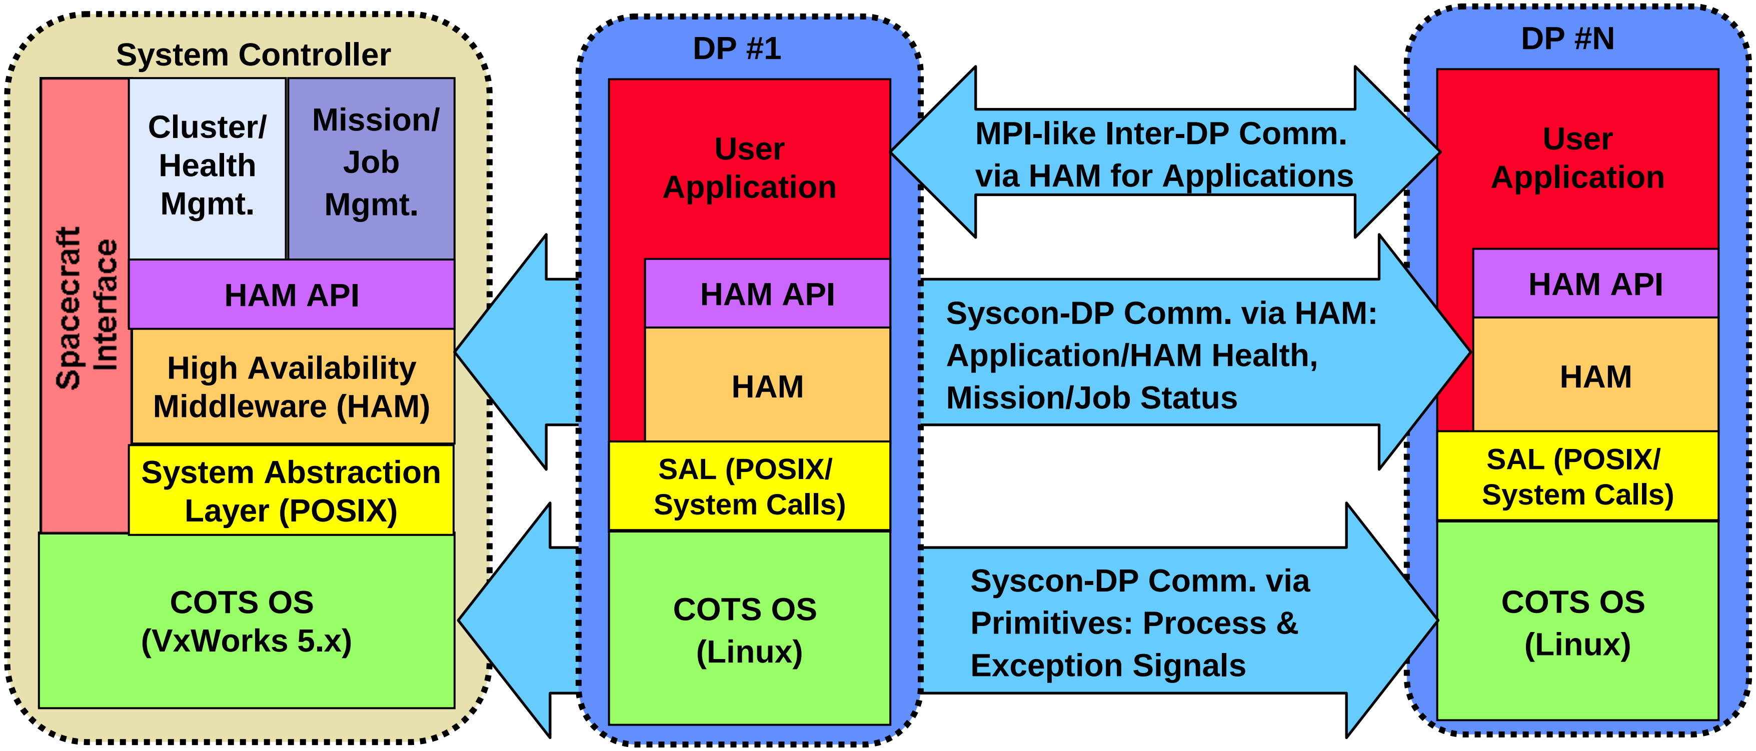

Dependable Multiprocessing Middleware

- **High Availability Middleware (HAM)**
  - Manages resources and application states
  - Provides cluster management including node discovery and network redundancy
  - Provides messaging infrastructure
- **System Services**
  - System level control of jobs and task management and failure detection and recovery schemes based on policy
- **Message Passing Interface (MPI) Services**
  - Provides interface for applications to transfer data using message passing protocol
  - Provides development environment for MPI-based applications
  - Provides job failover/restart/abort capabilities
- **Application Services (HAM APIs)**
  - APIs to communicate with DMM for application heartbeating
  - APIs communicate with DMM for recovery policy with user-defined fault detection
  - Mass data storage interaction for application data and check pointing

DMM System Services

Honeywell

- **Fault Tolerance Manager (FTM)**
  - **Maintains status table of system components and agents**
  - **Provides status and synchronization information to Job Manager**
  - **Detects faulty objects (applications, agents and nodes) through HAM**
    - Transition and client list handler in FTM monitors object state
    - Object start/stop handler allows FTM to manages DMM components execution state
- **Job Manager (JM)**
  - **Schedules, recovers and deploys (single task per processor) jobs**
  - **Cleans up check points and tasks on failure**
  - **Job Manager Agent (JMA)**
    - Forks tasks and relays status to FTM and JM
    - Detects application hangs and crashes
    - Receives fault detection from user-defined detection techniques (Algorithm Based Fault Tolerance (ABFT), replication and OS exception caused by application capture)
  - **Object start/stop handler allows JM and JMA to manage execution state of applications**
- **Mission Manager (MM)**
  - **Manages mission-level tasks and policies such as replication (spatial, temporal, simplex and parallel), periodicity, job scheduling and time outs**
  - **Cleans up data from replicas**

DMM Software Architecture Stack

Honeywell

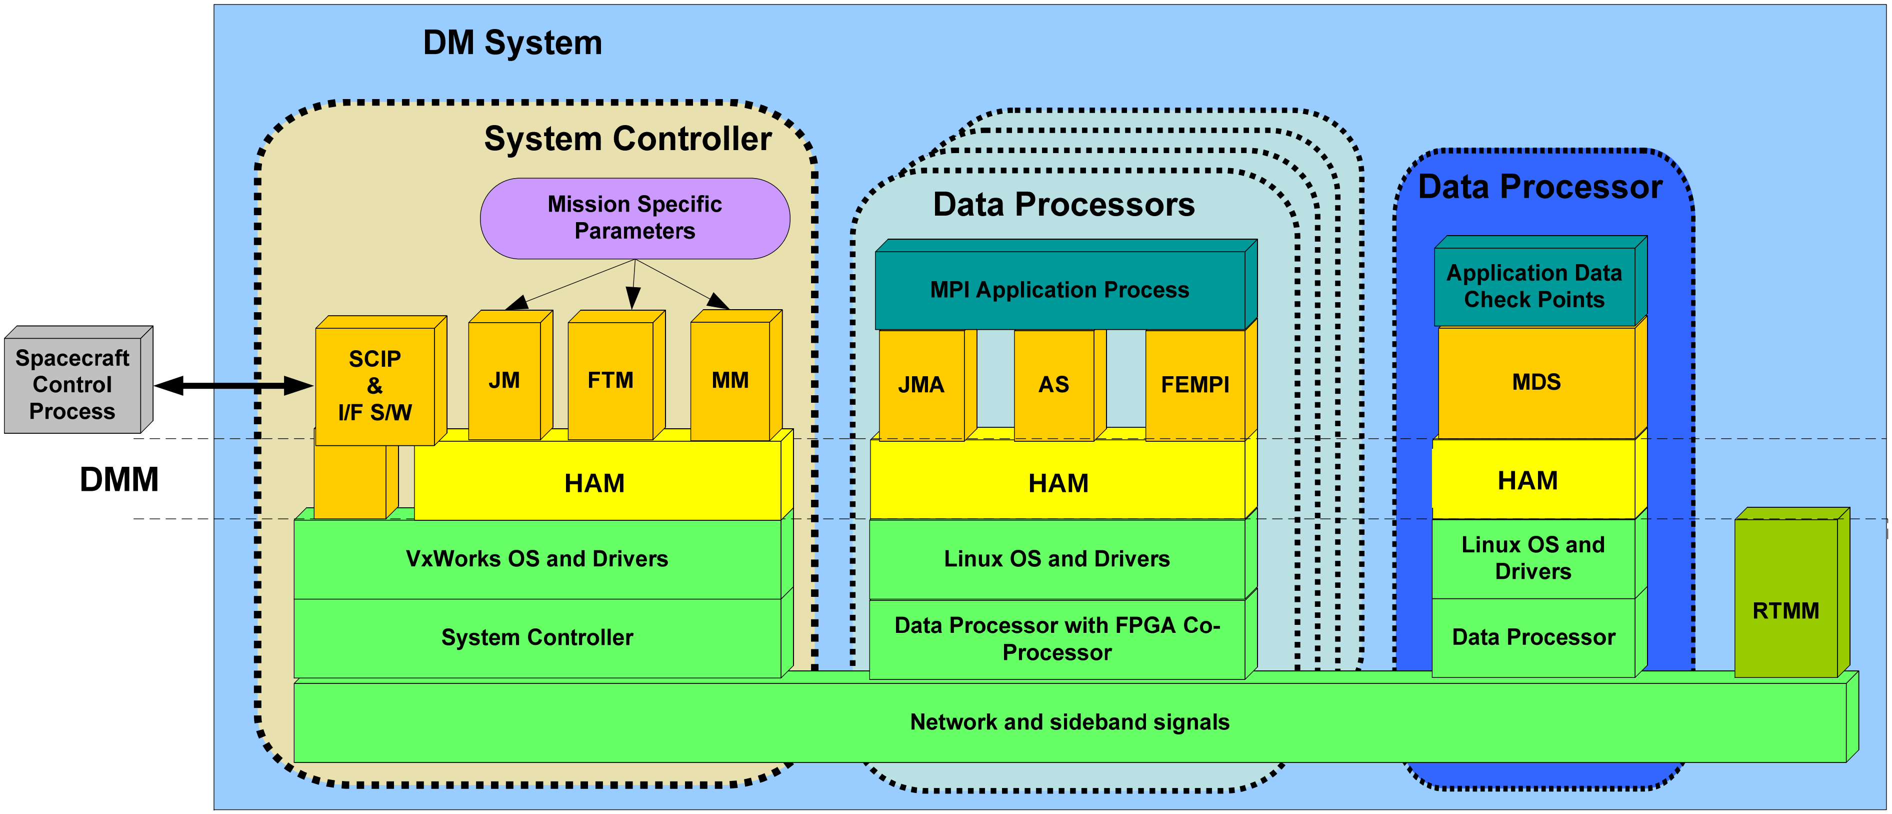

- HA Middleware
- Platform Components
- Application Components
- Mission Specific Components
- Dependable Multiprocessor MW Specific Components

JM – Job Manager
JMA – Job Manager Agent
MM - Mission Manager
FTM- Fault Tolerance Manager
FEMPI – Fault Tolerant Embedded Message Passing Interface
SCIP - Space Craft Interface Message Processor

AS – Application Services
MDS – Mass Data Storage
CMS – Cluster Management Services
AMS – Availability Management Services
DMS – Distributed Messaging Services
RTMM – Radiation Tolerant Mass Memory

Layers of Fault Tolerance

Honeywell

- **System-level fault tolerance**
  - **Heartbeat mechanism (between control and DP nodes)**
    - Detects node failures → allows FTM to perform hard node reboot
  - **Thread watch monitors (HAM implemented)**
    - Detects thread failures
- **Mission-level fault tolerance (defined in mission file)**
  - **Real-time deadline = job completes within allotted deadline, otherwise trash it and move on to next job**
    - Last resort mechanism to detect and recover from hung jobs
  - **Job selection**
    - Many jobs have multiple versions with varying degrees of fault tolerance
    - Performance vs. fault tolerance trade-off based on environment
  - **Priority and preemption**
- **Job-based fault tolerance (defined in job DAG file)**
  - **Temporal and spatial NMR**
    - Voting conducted by MDS unit
  - **Heartbeat mechanism (between app and JMA)**
    - Detects application hang or crash
- **Application-based fault tolerance**
  - **Algorithm-based Fault Tolerance (ABFT) = detection and (potential) recovery mechanism for data corruption**
  - **Inline-replication = temporal replication of code segments/functions with voting internal to application**
  - **Checkpointing**

UNIVERSITY OF FLORIDA NOAO

Examples: User-Selectable Fault Tolerance Modes

Honeywell

Fault Tolerance Option	Comments
NMR Spatial Replication Services	Multi-node HW SCP and Multi-node HW TMR
NMR Temporal Replication Services	Multiple execution SW SCP and Multiple Execution SW TMR in same node with protected voting
ABFT	Existing or user-defined algorithm; can either detector detect or detect and correct data errors with less overhead than NMR solution
ABFT with partial Replication Services	Optimal mix of ABFT to handle data errors and Replication Services for critical control flow functions
Check-pointing Roll Back	User can specify one or more check-points within the application, including the ability to roll all the way back to the original
Roll forward	As defined by user
Soft Node Reset	DM system supports soft node reset
Hard Node Reset	DM system supports hard node reset
Fast kernel OS reload	Future DM system will support faster OS re-load for faster recovery
Partial re-load of System Controller/Bridge Chip configuration and control registers	Faster recovery that complete re-load of all registers in the device
Complete System re-boot	System can be designed with defined interaction with the S/C; TBD missing heartbeats will cause the S/C to cycle power

DM Technology Advances to TRL7 Flight Experiment

Honeywell

Increasing system fidelity & integration

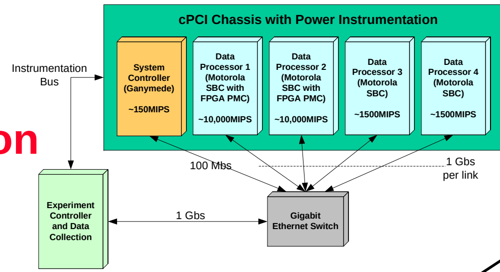

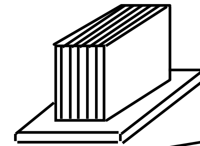

TRL7 Technology Validation

TRL6 Technology Validation

X*

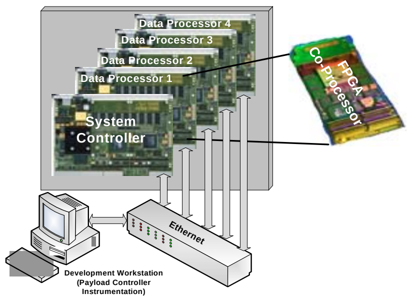

TRL5 Technology Validation

TRL7 Validation

- Demonstrate DM technologies in a real space environment
- Validate predictive models and predictive model parameters with experiment data
- TRL7 experiments will be identical to those performed and rung out during TRL6 demonstration and validation

TRL6 Validation

- Demonstrate enhanced DM technologies in a relevant environment on prototype flight hardware including exposure to radiation beam
- Validate and refine predictive models and predictive model parameters with experiment data
- complete set of canonical fault injection experiments

TRL4 Validation

- Demonstrated basic DM technologies in a laboratory environment on COTS hardware testbed including radiation source and sensor
  - Environment Sensor
  - Alert Generator
  - High Availability Middleware
  - Replication Services

NASA adds requirement for fault-tolerant cluster and fault-tolerant MPI capability

TRL4 Technology Validation

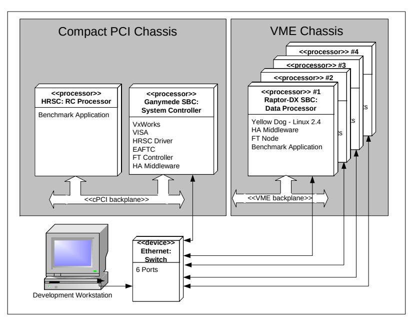

TRL5 Validation

- Demonstrate basic DM technologies in a laboratory environment on testbed hardware with partially integrated Fault Tolerance Services
- Develop predictive models
- Validate and refine predictive models and predictive model parameters with experiment data
- partial set of canonical fault injection experiments

*** In 8/07, NASA eliminated the ST8 flight experiment due to NASA budget issues**

UNIVERSITY OF FLORIDA

NOAO

DM Technology Readiness & Experiment Development Status and Future Plans

Honeywell

5/30/05
TRL4 Technology Validation
Technology Concept Demonstration

5/17/06
TRL5 Technology Validation
Technology in Laboratory Environment

10/27/06
NASA ST8 Project Confirmation Review

9/08 & 3/09 *
TRL6 Technology Validation
Technology Demonstration in a Relevant Environment *

X*

~~Launch 11/09 *~~
~~Mission 1/10 - 6/10 *~~

5/31/06
Preliminary Design Review
Preliminary Experiment HW & SW Design & Analysis

6/27/07
Critical Design Review
Final Experiment HW & SW Design & Analysis

X*

TRL7 Technology Validation
Flight Experiment

5/06, 4/07, & 5/07
Preliminary Radiation Testing
Critical Component Survivability & Preliminary Rates

5/08, 7/08, 8/08, 10/08, & 1/09
Final Radiation Testing
Complete Component & System-Level Beam Tests

Key: [ ] - Complete

- Per direction from NASA Headquarters 8/3/07, the ST8 project ends with TRL6 Validation;

  Preliminary TRL6 demonstration 9/15/08;

  Final TRL6 report 6/09

UNIVERSITY OF FLORIDA NOAO

DM TRL6 Testbed System

Honeywell

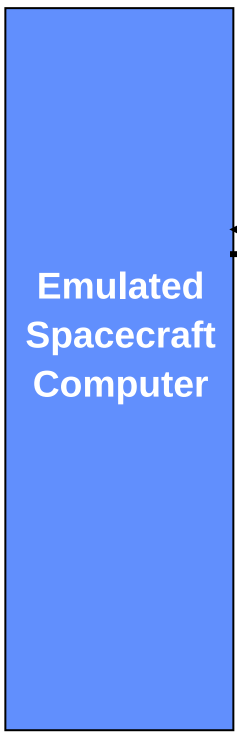

System Controller:
Wind River OS
- VxWorks 5.4
Honeywell Ganymede
SBC (PPC 603e)

RS422

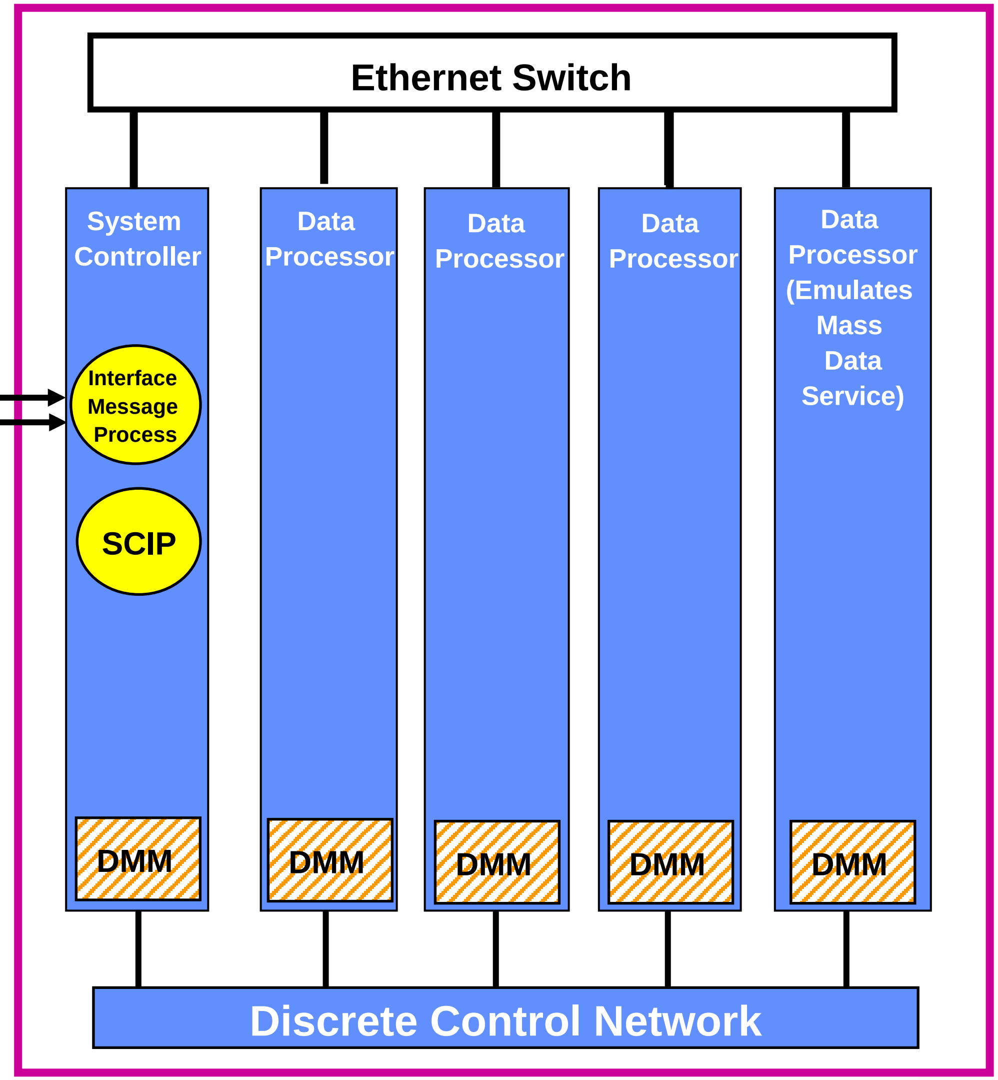

Data Processor:
Wind River OS
- PNE-LE 4.0 (CGE) Linux
Extreme 6031
PPC 7447a with AltiVec co-processor

Networks:
cPCI
Ethernet: 100Mb/s

SCIP – Space Craft Interface Process

UNIVERSITY OF FLORIDA NOAO

NMP Breakthrough Technologies

DM TRL6 (Phase C/D) Flight Testbed

Honeywell

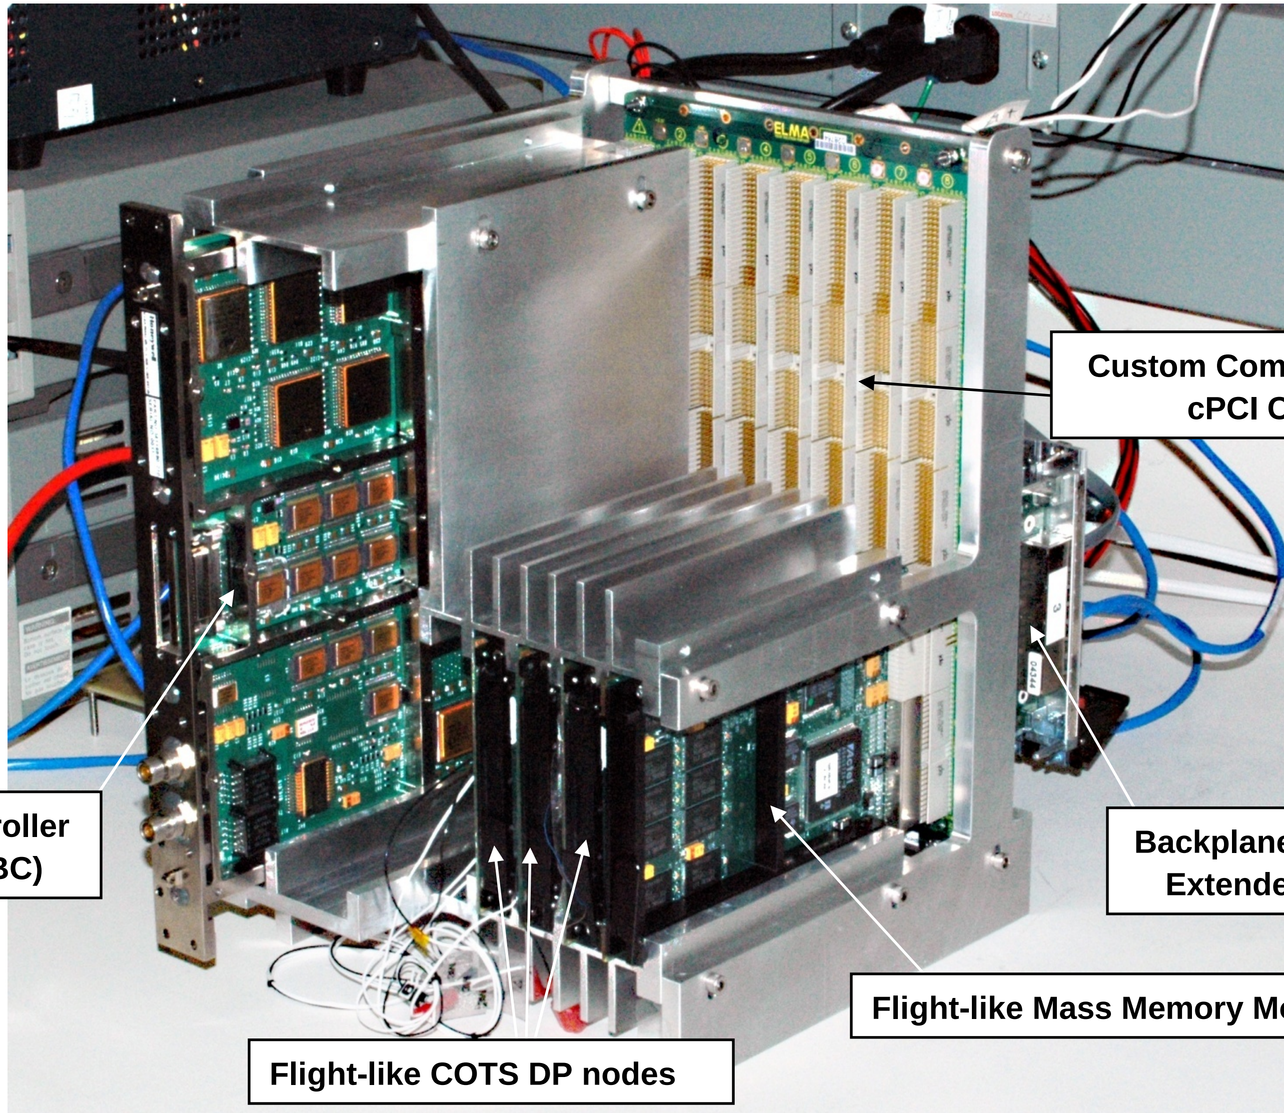

Custom Commercial Open cPCI Chassis

System Controller (flight RHSBC)

Backplane Ethernet Extender Cards

Flight-like Mass Memory Module

Flight-like COTS DP nodes

UNIVERSITY OF FLORIDA

NOAO

NMP

DM TRL6 Status – Key Elements

Honeywell

- **Radiation Testing**
  - Honeywell and JPL proton and heavy ion testing established SEE rates for all components on COTS DP boards
  - System-level testing performed with one COTS DP board exposed to proton beam while running the flight experiment application software suite
    - OS, HAM, DMM, application, instrumentation
  - DM flight experiment instrumentation including emulated ground station operated successfully
  - Post-experiment data analysis demonstrated
  - DM middleware performed as designed
  - DM system successfully recovered from all radiation-induced faults
- **DM Models (Markov and Discrete Event Simulator)**
  - Demonstrated DM predictive Availability, “Computational Consistency,” and Performance models
  - Models based on component-level radiation test results and SWIFI (Software-Implemented Fault Injection) campaigns
  - Extrapolated performance to various radiation environments, i.e., orbits, and other applications

UNIVERSITY OF FLORIDA NOAO

DM Models

Honeywell

- Developed Four Predictive Models
  - Hardware SEU Susceptibility Model
    - Maps radiation environment data to expected component SEU rates
    - Source data is radiation beam test data
  - Availability Model
    - Maps hardware SEU rates to system-level error rates (SWIFI)
    - System-level error rates + error detection & recovery times ➔ Availability
    - Source data is radiation beam test data and measured testbed detection & recovery statistics
  - Computational Consistency Model
    - Models number of erroneous datasets and late deliveries to the end user
    - Source data is radiation beam test data and the measured error detection & recovery coverage from testbed experiments
  - Performance Model
    - Based on computational operations, arithmetic precision, measured execution time, measured power, measured OS and DM SW overhead, frame-based duty cycle, algorithm/architecture coupling efficiency, network- level parallelization efficiency, and system availability
    - Source data is radiation beam test data and measured testbed performance and output of the availability model predictions

DM TRL6 Status – Key Elements

Honeywell

- Demonstrated ability to meet NASA level 1 requirements/goals
  - > 0.995 Availability (LEO environment)
  - > 0.995 "Computational Consistency," the probability of timely and correct delivery of data (LEO environment)
  - > 300 MOPS per watt
    - 308 MOPS/watt HSI application on 7447a processor with AltiVec - measured
    - 276 MOPS/watt HSI application on 7447a processor with AltiVec (including System Controller power) - measured
    - > 332MOPS/watt HSI application on new, industrial temperature range, low power 7448 processor with AltiVec (including System Controller power) – analytical
      - -- 7448 includes EDAC on L2 cache; drop-in replacement for the 7447a (helps DM Availability; 7447a only has parity on L2 cache)
    - 1077 MOPS/watt HSI application on PA Semconductor dual core processor with AltiVec - measured
- Demonstrated ease of use & low overhead
  - Independent 3rd party with minimal knowledge of fault tolerance ported two (2) diverse applications to DM testbed in less than three (3) days
  - Applications included scalable parallelization, hybrid ABFT/in-line replication, 2D convolution and median filter ABFT library functions, FEMPI, and check-pointing
  - Porting experience to DM << 1 hour/10 executable SLOC (TRL6 requirement)
  - Low DMM overhead (~6%) (same platform, same application with & without DMM)

UNIVERSITY OF FLORIDA NOAO NMP

DM Technology - Ease of Use

Honeywell

- Successfully ported four (4) real applications to DM testbeds
  - HSI (Hyper-Spectral Imaging) *
    - scalable MPI application
    - ~ 14 hours to port application to DM system with DMM, hybrid ABFT, and in-line replication
    - ~ 4 hours to implement auto-correlation function in FPGA
  - SAR (Synthetic Aperture Radar) *
    - scalable MPI application
    - ~ 15 hours to port application to DM system with DMM, hybrid ABFT, in-line replication, check-pointing
  - CRBLASTER (cosmic ray elimination application) **
    - scalable MPI application
    - ~ 11 hours to port application to DM system with DMM, hybrid ABFT, and in-line replication
    - scalability demonstrated ~ 1 minute per configuration
  - QLWFP2C (cosmic ray elimination application) **
    - scalable, fully-distributed MPI application
    - ~ 4 hours port application to DM system with DMM
    - scalability demonstrated ~ 1 minute per configuration
  - NASA GSFC Synthetic Neural System (SNS) application for autonomous docking *
    - ~ 51 hours to port application to DM system with DMM (includes time required to find a FORTRAN compiler to work with DM)

* Port performed by Adam Jacobs & Greg Cieslewski, doctoral students at the University of Florida and members of ST8 DM team

** Port performed by Dr. Ken Mighell, NOAO, Kitt Peak Observatory, independent 3rd party user/application developer with minimal knowledge of fault tolerance techniques, per TRL6 requirement

Summary & Conclusion (1 of 3)

Honeywell

- **DM TRL6 Technology Validation Demonstration**
  - System-level radiation tests validated DM operation in a radiation environment
  - Demonstrated high performance, high availability, high probability of timely and correct delivery of data, predictive models, and ease of use
  - SWIFI testing is continuing

- **Flying high performance COTS in space is a long-held desire/goal**
  - Space Touchstone - (DARPA/NRL)
  - Remote Exploration and Experimentation (REE) - (NASA/JPL)
  - Improved Space Architecture Concept (ISAC) - (USAF)

- **The problems and pitfalls of flying COTS in space are understood**
  - Prado, Ed, J. R. Samson, Jr., and D. Spina, "The COTS Conundrum," *Proceedings of the 2000 IEEE Aerospace Conference*, Big Sky, MT, March 9-15, 2003
  - Samson, Jr., John R., "SEUs from a System Perspective," Single Event Upsets in Future Computing Systems Workshop, Pasadena, CA, May 20, 2003

- **NMP ST8 DM project has brought this desire/goal closer to reality**

Summary & Conclusion (2 of 3)

Honeywell

- DM technology is applicable to wide range of missions
  - Science and autonomy missions/Landers/Rovers
  - MKV
  - UAVs (Unattended Airborne Vehicles)
  - UUVs (Unattended or Un-tethered Undersea Vehicles)
  - ORS (Operationally Responsive Space)
  - Stratalites
  - Ground-based systems & rad hard space applications

- Multiple applications have been successfully ported to and demonstrated on DM testbeds
  - SAR, HSI, NBF-SNS, CRBLASTER, QLWFP2C, Matrix Multiply, 2DFFT, LUD

- DM technology independence has been demonstrated on wide variety of platforms
  - x86, PPC clusters
  - PA Semi dual core processor, 8641D dual core processor
  - FPGAs
  - heterogeneous systems
    - -- HW
    - -- SW
  - VxWorks, Linux OS

UNIVERSITY OF FLORIDA

NOAO

Summary & Conclusion (3 of 3)

- To date, NASA has invested $12.2M in the development and demonstration of DM technology
- The DM project has further developed, refined, and demonstrated the process for migrating high performance COTS computing to space
- Validating DM technology in space is still needed
  - to establish that all-important space pedigree
  - to validate the process the process of migrating high performance COTS computing to space
  - to validate the predictive models in a real-space environment
- Since NASA eliminated the flight experiments from the ST8 project, DM has been looking for an alternative ride to space
  - currently looking for an advocate for a SERB (Science Experiment Review Board) flight experiment
  - with the exception of the sponsor page, SERB paperwork is filled out
- DM technology has potential applicability to common space architecture
  - two-chart summary of DM technology applicability to common space architecture will be presented at Architecture Working Group panel session on 5/29

Back-Up Material

Dependable Multiprocessor Experiment Payload on the ST8 "NMP Carrier" Spacecraft

Honeywell

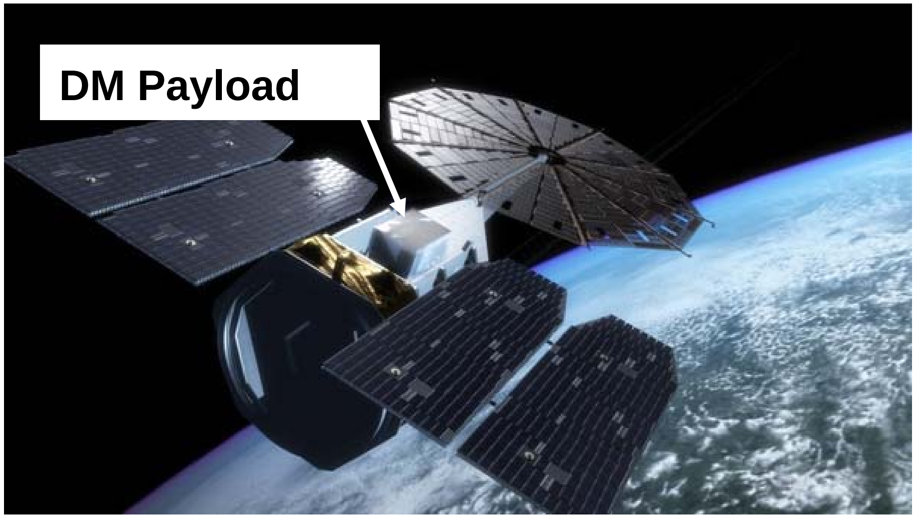

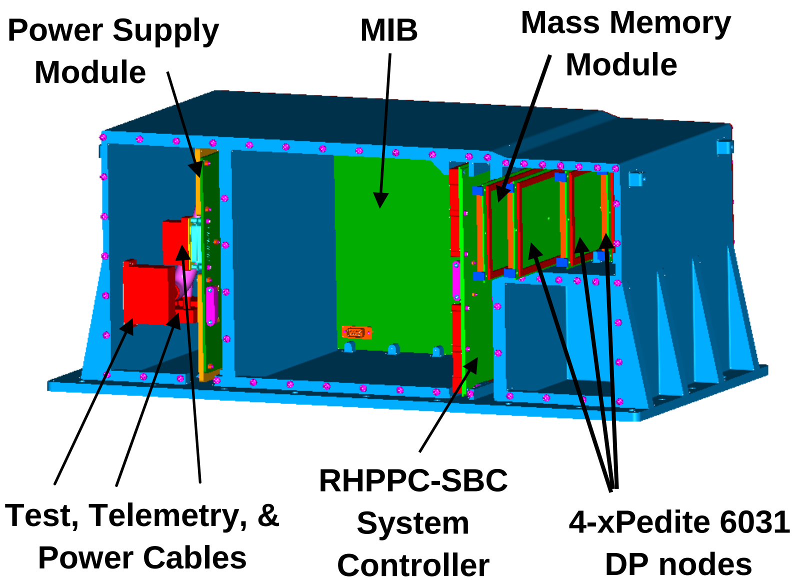

ST8 Orbit: - sun-synchronous
- 955 km x 460km @ 98.2^{o} inclination

Software
- Multi-layered System SW
  - OS, DMM, APIs, FT algorithms
- SEU-Tolerance
  - detection
  - autonomous, transparent recovery
- Applications
  - 2DFFT, LUD, Matrix Multiply, FFTW SAR, HSI
- Multi-processing
  - parallelism, redundancy
  - combinable FT modes

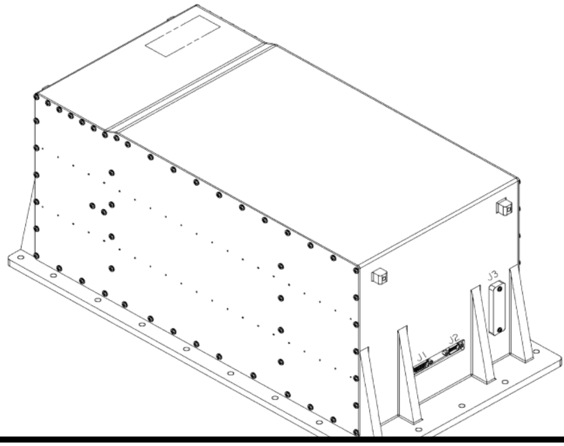

Flight Hardware
- Dimensions
  10.6 x 12.2 x 24.0 in.
  (26.9 x 30.9 x 45.7 cm)
- Weight (Mass)
  ~ 61.05 lbs
  (27.8 kg)
- Power
  ~ 121 W (max)

The ST8 DM Experiment Payload is a stand-alone, self-contained, bolt-on system.

UNIVERSITY OF FLORIDA NOAO