Spring 2015

DOD HPC INSIGHTS

A publication of the Department of Defense High Performance Computing Modernization Program

Advanced Wear Modeling

for Improved Aircraft Tire Life

HPCMP CREATE™

Growing Impact of HPC on Acquisition Engineering

HPC Help Desk

Navy Cray Phi Usage Guide

Contents

First word

Dr. David Horner 1

HPC @ Work

Advanced Wear Modeling for Improved Aircraft Tire Life 2

New to the HPCMP 8

The HPC Systems 10

HPCMP CREATE™: The Growing Impact of HPC on Acquisition Engineering 12

ARL DSRC Supports Network Integration and Evaluation Exercise 15

ARCADE: Advanced Research, Collaboration, and Application Development 16

The HPC User

HPC Portal Update 21

Navy Cray Phi Usage Guide 22

QA in the Era of "Always On" Systems 24

HPC Help Desk 24

At the Centers

MHPCC Management Team 25

From the Director's Desk: ARL DSRC 26

From the Director's Desk: Navy DSRC 27

AFRL DSRC Innovation Laboratory: Incubating Discovery 28

Cover Photo:
Photos Courtesy of:
Defense Media Activity

First word

Dr. David Horner
HPCMP Director

It is my distinct honor to have been chosen as the Director of the High Performance Computing Modernization Program (DOD HPCMP). Let me take this opportunity to express my profound gratitude to my predecessor, John West, for the superlative work he did for the program and the nation, and to Dr. Roy Campbell for his outstanding contributions as Acting Director and Deputy Director.

My background in high performance computing was forged from my experience as a user of HPCMP resources, while performing research for the US Army Engineer Research and Development Center (ERDC), so a user focus is what I bring to the director's desk. The loss of our ability to communicate deeply with, and among, HPCMP users through our annual User Group Conference has eroded the once strong bond between the user base and the service delivery base. I intend to restore this bond through regional meetings; look for more on this topic in coming months.

I also strongly hold that in order for the program to continue its arc of achievement, we must clearly and forcefully articulate our impact to the warfighter. We must evaluate new opportunities to support high-impact service and agency requirements. We must capture and advertise the impact of our many successes across the entire program. We must develop a portfolio approach to balance our program investments among impact, innovation, transition, and current and future computing environments.

The program is strong. Our financial footing is solid. We have new, world-class supercomputers coming online this summer. Our future is bright. I look forward to meeting and working with each and every one of you. Thank you for the opportunity to serve.

DOD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM

Advanced Wear Modeling for Improved Aircraft Tire Life

Andrew Zakrajsek, Jonathan Childress, Dr. Michael Bohun
United States Air Force, 96TG/OL-AC Landing Gear Test Facility, Wright-Patterson Air Force Base

Dr. Samir Naboulsi
Contractor, High Performance Computing Modernization Program Air Force Research Laboratory DOD Supercomputing Resource Center (HPCMP, AFRL DSRC)

Photo Courtesy of: Defense Media Activity

Achieving a desired tire life for a new aircraft has always been difficult, especially for fighter aircraft. Initial aircraft tire life has been lower than expected for certain aircraft variants which results in unanticipated cost to the program. Tire life issues are a direct result of increasing performance requirements, varying weight requirements, and a lack of understanding regarding tire wear mechanics. To address this, the 96th Test Group, Aerospace Survivability and Safety Operating Location (96TG/OL-AC), in conjunction with the Air Force Research Laboratory (AFRL) DOD Supercomputing Resource Center is developing advanced computational modeling capabilities for tire wear design, testing and evaluation[1]. As observed in previous aircraft acquisition programs, significantly increasing tire life early in the acquisition cycle can benefit the logistical, environmental, and financial aspects of the program. Improving tire life for certain aircraft can save hundreds of millions of dollars over an aircraft's lifetime.

AFRL's tire wear testing effort focuses on predictive wear testing using the 96TG/OL-AC Landing Gear Test Facility's (LGTF) 168-inch internal drum dynamometer (168i), shown in (Figure 1). We focus on linking upper level system parameters (tire size, weight, aircraft specs, brake specs, etc.) to fielded tire life. To do this, the 168i's internal drum is fitted with ductile concrete tiles (Figure 2) cloned from actual Air Force runways to run mission representative test profiles. The 168i can test tire performance with simulated runway speeds up to 350 MPH, wheel loads up to 150Klbs, yaw angles of ±20 degrees, and camber angles of ±10 degrees. These physical testing capabilities have recently been complemented with the use of computational models to further study tire wear (Figure 3). The computational tools create a predictive tire qualification testing capability with the potential to simulate full aircraft missions.

*Photo Courtesy of:
Defense Media Activity*

[1] *"Improved Aircraft Tire Life through Laboratory Tire Wear Testing and Computational Modeling," Andrew J. Zakrajsek, Jonathan M. Childress, Michael H. Bohun, and Martin G. Vogel, Samir K. Naboulsi, Brett J. Fiepke, Ryan N. Vogel, Ryan M. Bena, Stephen J. Howell, and Clarence B. Alsobrook, paper number AIAA-2015-1125, SciTech 2015 Conference.*

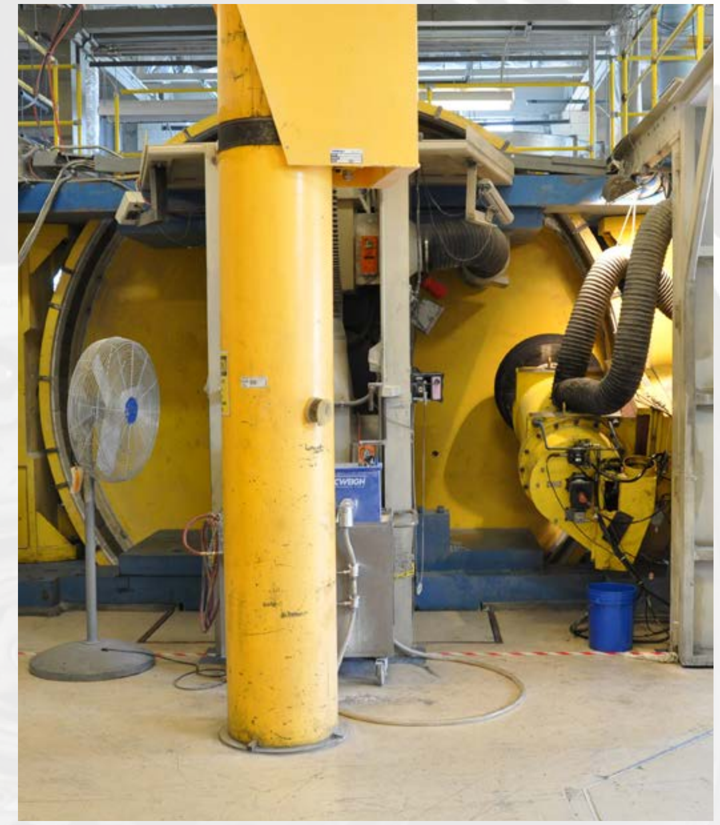

Figure 1

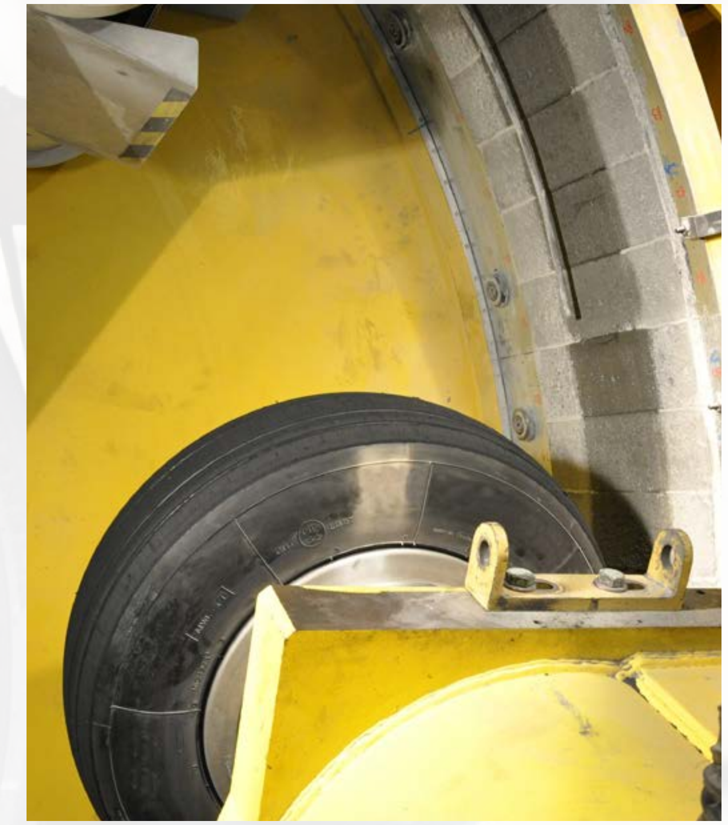

Figure 2

"FEM has become a powerful analytical tool for understanding and improving tire design and development."

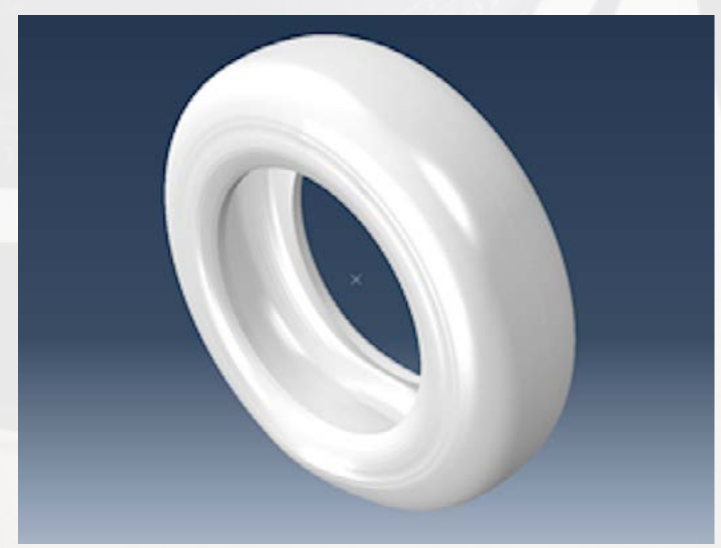

Figure 3

*Figure 1
USAF 168-inch Internal Drum Dynamometer with a tire instrumented for comparative wear testing.*

*Figure 2
Close up of tire and ductile concrete tiles on the Internal Drum Dynamometer.*

*Figure 3
Computational simulation of a deformed tire.*

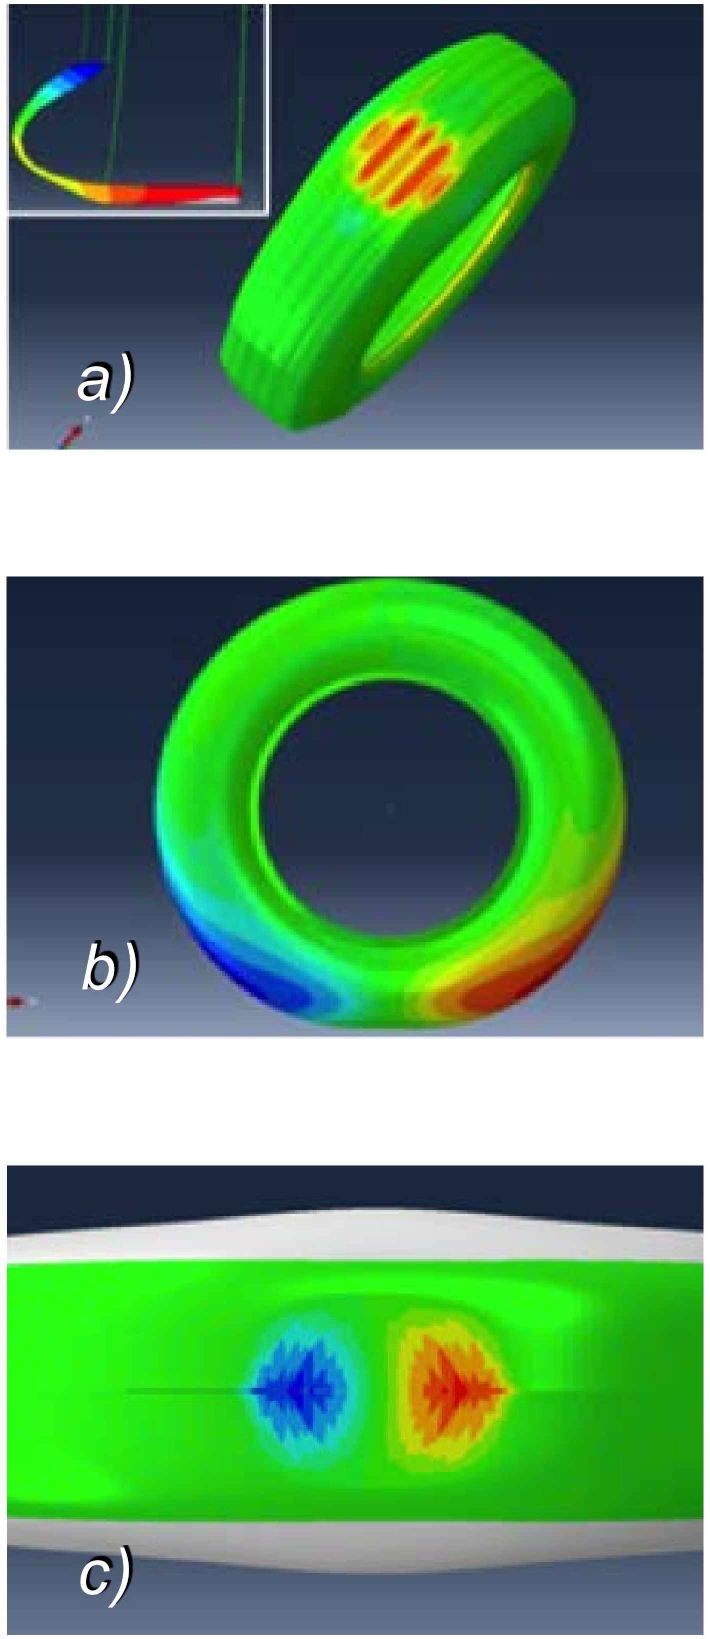

Advanced Wear Modeling with Finite Element Modeling and Missionized Testing

Finite Element Modeling (FEM) has become a powerful tool for understanding and improving tire design. With supercomputing resources, high-fidelity FEMs can efficiently simulate tire wear under any number of circumstances. We generated a preliminary generic aircraft tire model using Abaqus software as a proof of concept. Model inputs include numerous parameters affecting tire wear; such as runway surface conditions, aircraft operating parameters, and tire design parameters and simulates. This diverse set of input parameters help simulate real world tire wear accurately.

We started with a three-dimensional tire footprint which captured the static deformed shape of the inflated aircraft tire when aircraft weight was applied as vertical dead load. This established the contact patch surface area under different loading conditions. We then created a steady-state rolling simulation to capture the tire's response at a given rolling speed. From this, we obtained a free rolling equilibrium solution for a rolling tire. Using the model, we then performed a steady state rolling analysis under a variety of conditions, such as slip angle/yaw, camber thrust, and braking torque. Finally, we developed a tread wear simulation to calculate wear rate and the overall wear amount experienced in aircraft tires during takeoff, taxiing, and landing maneuvers.

Using the steady state solution and assuming a simple wear model based on the tire contact patch and slip state, we computed the nodal ablation velocity, which determines when tire wear occurs. The nodal ablation was then applied away from the footprint region to the periphery of the tire. Figure 4 shows results from the FEM model. Extrapolating over a given duration at a given traveling speed provides an approximate consideration of the transient process of tire wear for this steady state case. The simulation takes into account both the inertia and viscoelasticity effects. The results show trends consistent with known tire performance characteristics. Experimental data can now be used to further verify the model and assist in missionized profile development and predictive wear testing on the 168i internal drum dynamometer.

Figure 4 Sample tire wear result from FEM model:
a) Contact shear stress along the rolling direction of a tire at the end of the dead-load step using a scaling factor of three for the deformed shape;
b) Displacement of the tire at the end of the dead-load step;
c) Pressure contour plot of the tire for straight case after wear.

Achievements and Summary

The joint efforts of the 96TG/OL-AC LGTF and HPCMP resources have led to a series of substantial achievements that are bringing the realization of improved aircraft tire life closer to reality. With the recent addition of a replicated concrete runway substrate on the 168i, we are now, for the first time, able to conduct tests on aircraft tires indicative of actual field wear in the laboratory. This type of testing, coupled with the use of computational modeling for predictive tire wear analysis allows the USAF to work towards the development and implementation of a standardized qualification test for new aircraft tire designs before they are fielded.

Predicting aircraft tire life is challenging, but obtainable with the continued development of the USAF 96TG/OL-AC LGTF 168i's missionized testing program, and further development of predictive tire wear analysis with FEM computational modeling. This provides USAF aircraft acquisition programs a suite of invaluable tools to more accurately estimate operational performance and cost of support for the war fighter.

Photo Courtesy of:
Defense Media Activity

NEW TO THE DOD HPCMP

High Performance Computing Portfolio

Five systems and over 21 petaFLOPS

A Cray XC40 and an SGI ICE X High Performance Computing (HPC) system have been installed at the Army Research Laboratory (ARL) and Army Engineer Research and Development Center (ERDC) DOD Supercomputing Resource Centers (DSRCs). Excalibur entered production in April 2015, and Topaz is expected to be available for allocated use in October 2015.

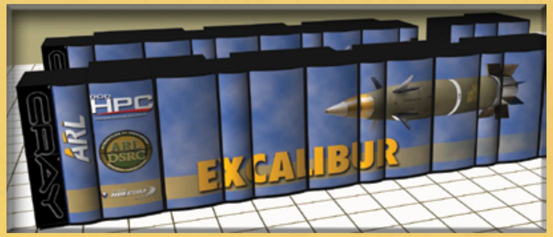

EXCALIBUR

(ARL DSRC)

CRAY XC40

SERVICE DATE 04/13/2015

Processor Cores	101,312
Memory (GB)	405,248
Disk (TB)	6,055
GPGPUs	32
teraFLOPS	3,774

Access to this resource (EXCALIBUR)... is having a significant impact in our fundamental research approach and is providing meaningful insights of fuel air mixture formation process in internal combustion engines"

Dr. Luis Bravo
Propulsion Division, Vehicle Technology Directorate, Army Research Laboratory

TOPAZ

(ERDC DSRC)

SGI ICE X

SERVICE DATE 07/01/2015

Processor Cores	125,440
Memory (GB)	443,584
Disk (TB)	17,496
GPGPUs	32
teraFLOPS	4,662

"With a theoretical peak performance of 4.6 petaFLOPS, Topaz will triple the ERDC DSRC's current compute capability. This will provide the DOD's RDT&E community with the resources needed to meet their ever-growing demand for HPC."

Robert M. Hunter,
Director, ERDC DSRC

Two Cray XC40s and an SGI ICE X have been installed at the Navy and Air Force Research Laboratory (AFRL) DOD Supercomputing Resource Centers. Conrad, Gordon, and Thunder are expected to be available for allocated use in October 2015.

CONRAD

(NAVY DSRC)

Named for Apollo 12 Astronaut Charles "Pete" Conrad.
The Cray XC40 system will contain 2.3 GHz Intel Xeon E5-2698v3 ("Haswell-EP") processors.

50,208 compute processors,
168 Intel Xeon Phi 5120D accelerators
205 terabytes of memory
3.04 petabytes of storage
and 2.0 petaFLOPS of peak computing capability

THUNDER

(AFRL DSRC)

The SGI ICE X system will consist of 125,888 compute cores, 356 NVIDIA GPGPUs, 356 Intel Xeon Phi coprocessors and 12 petabytes of usable storage.

Thunder will be collocated with the 8 M-Cell SGI ICE system, Spirit, deployed in 2012 and will feature 6 M-Cells — SGI's single largest M-Cell deployment. SGI M-Cell technology provides industry-leading power and cooling efficiency.

GORDON

(NAVY DSRC)

Named for Apollo 12 Astronaut Richard "Dick" Gordon.
The Cray XC40 system will contain 2.3 GHz Intel Xeon E5-2698v3 ("Haswell-EP") processors.

50,208 compute processors,
168 Intel Xeon Phi 5120D accelerators
205 terabytes of memory
3.04 petabytes of storage
and 2.0 petaFLOPS of peak computing capability

"This expansion in capability quintuples our pre-2014 capacity of 5.38 petaFLOPS to 26.4 petaFLOPS."

Christine Cuicchi,
HPCMP Associate Director for HPC Centers

"The power of Thunder will drive solutions to the most challenging problems facing our nation in today's volatile global environment."

Jeff Graham,
Director, AFRL DSRC

"Increased capabilities will allow our users to perform ground breaking science and realize previously impossible discoveries in DOD research."

Christine Cuicchi,
HPCMP Associate Director for HPC Centers

For more information on these systems, visit

centers.hpc.mil

DOD HPC MODERNIZATION PROGRAM SYSTEMS

The High Performance Computing Modernization Program (HPCMP) provides over 656,000 cores and 16 petaFLOPS to its users in the pursuit of improved scientific research, weapons design, force protection, and software development for the Department of Defense

Army Research Laboratory DSRC
Aberdeen Proving Ground, Maryland

PERSHING
IBM iDataPlex
Service Date: 01/07/2013

Processor Cores	20,160
Memory (GB)	45,696
Disk (TB)	2,370
teraFLOPS	419

EXCALIBUR
Cray XC40
Service Date 04/13/2015

Processor Cores	101,312
Memory (GB)	405,248
Disk (TB)	6,055
GPGPUs	32
teraFLOPS	3,774

Engineer Research and Development Center DSRC
Vicksburg, Mississippi

GARNET
Cray XE6
Service Date: 07/18/2013

Processor Cores	150,912
Memory (GB)	301,824
Disk (TB)	3,125
teraFLOPS	1,509

COPPER
Cray XE6m
Service Date: 11/19/2012

Processor Cores	14,720
Memory (GB)	29,440
Disk (TB)	443
teraFLOPS	138

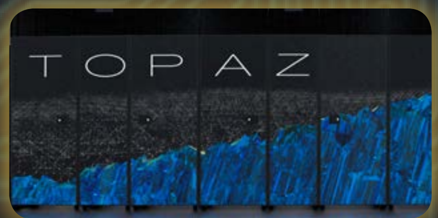

TOPAZ
SGI ICE X
Service Date: 07/01/2015

Processor Cores	125,440
Memory (GB)	443,584
Disk (TB)	17,496
GPGPUs	32
teraFLOPS	4,662

Maui High Performance Computing Center DSRC
Kihei, Maui, Hawaii

RIPTIDE
IBM iDataPlex
Service Date: 05/07/2013

Processor Cores	12,096
Memory (GB)	24,192
Disk (TB)	2,150
teraFLOPS	253

Air Force Research Laboratory DSRC
Wright-Patterson Air Force Base, Ohio

SPIRIT
SGI ICE X
Service Date: 03/25/2013

Processor Cores	73,440
Memory (GB)	146,880
Disk (TB)	2,458
teraFLOPS	1,528

PREDATOR
SGI UV
Service Date: 11/21/2013

Processor Cores	1,004
Memory (GB)	4
Disk (TB)	88
teraFLOPS	22

LIGHTNING
Cray XC30
Service Date: 09/01/2014

Processor Cores	57,000
Memory (GB)	152,704
Disk (TB)	3,520
GPGPUs	32
teraFLOPS	1,281

Navy DSRC
Stennis Space Center, Mississippi

SHEPARD
Cray XC30
Service Date 09/01/2014

Processor Cores	30,592
Memory (GB)	82,752
Disk (TB)	1,560
Coprocessors	124
GPGPUs	32
teraFLOPS	822

ARMSTRONG
Cray XC30
Service Date 09/01/2014

Processor Cores	30,144
Memory (GB)	81,728
Disk (TB)	1,560
Coprocessors	124
teraFLOPS	786

KILRAIN
IBM iDataPlex
Service Date: 01/07/2013

Processor Cores	19,812
Memory (GB)	40,832
Disk (TB)	24
teraFLOPS	407

HAISE
IBM iDataPlex
Service Date: 01/07/2013

Processor Cores	19,812
Memory (GB)	40,832
Disk (TB)	2,473
teraFLOPS	407

HPCMP CREATE™

The Growing Impact of HPC on Acquisition Engineering

Dr. Kevin Newmeyer
CREATE, Senior Operations Director, Phacil Inc.

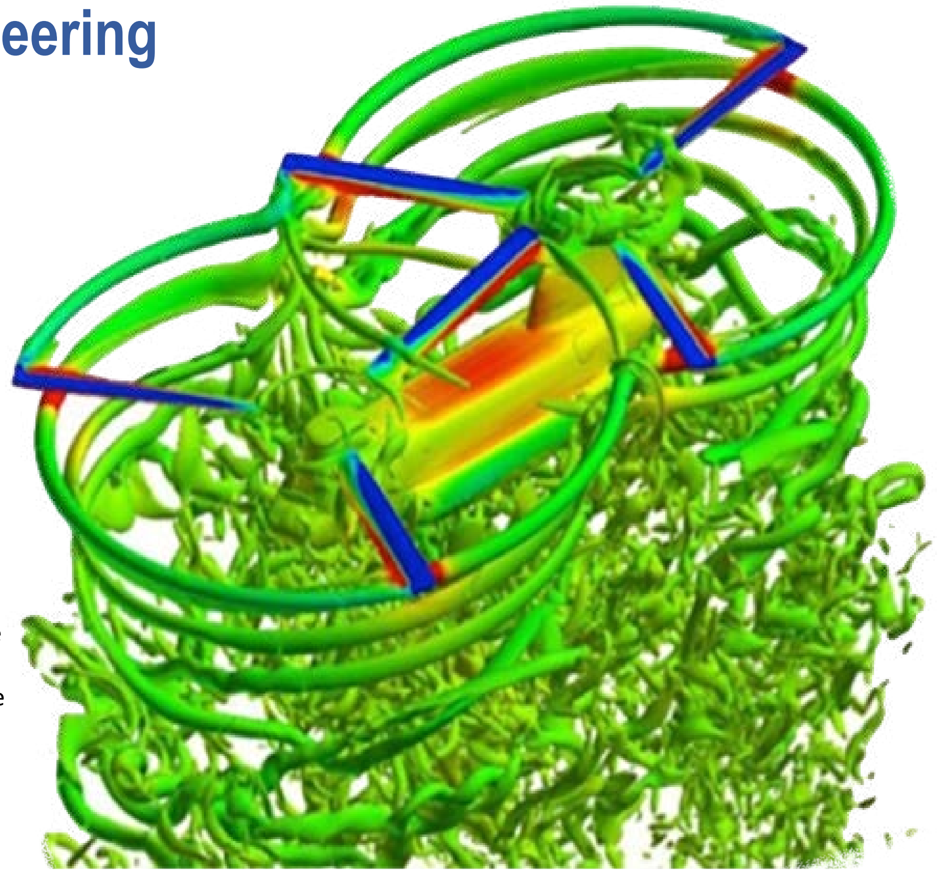

Figure 1:
CH-47 Rotor Vortices in Helios

Amid the rapid development of the High Performance Computing Modernization Program Computational Research Engineering Acquisition Tools and Environments (HPCMP CREATE™) projects, the defense acquisition community is taking greater advantage of the Department of Defense (DOD) High Performance Computing Modernization Program's (HPCMP) assets. While modeling and simulation in war-gaming has an extensive history in the defense arena, the multiphysics tools available to government engineers, via the CREATE program, are just beginning to have real impact on acquisition programs. High performance computing is currently able to positively influence the acquisition of major defense programs, provide government engineers the ability to evaluate competing proposals, and test virtual prototypes of weapons systems, while there is still sufficient time to substantively affect program decisions.

Traditional defense acquisition programs do not possess detailed performance analysis capabilities until relatively late in the DOD acquisition process; usually between milestones B and C. By this point, taxpayers have expended a significant amount of money, and severe performance problems may become apparent. Virtual prototyping using CREATE tools on HPCMP supercomputers allows for high-fidelity assessments of designs prior to building physical prototypes.

HPCMP CREATE™ History

After program approval in 2006, the CREATE teams began development of software to perform physics-based analysis of aircraft, ship, and electromagnetic sensor projects for engineers within the Department of Defense. The CREATE-Aviation, CREATE-Ships, and CREATE-RF projects are all supported by the meshing and geometry project, CREATE-MG. The aviation and ship projects consist of multiple tools for rapid-design space exploration, and detailed mutiphysics-based analysis of the integrated performance characteristics of virtual models.

CREATE Tools

CREATE-AV contains three software projects centered on improving the acquisition process for both fixed- and rotary-wing aircraft in support of Army, Air Force, and Navy aviation. The DaVinci program is a concept design tool that allows aircraft designers to evaluate several aspects of an aircraft design project. For example, designers can adjust the overall length of the fuselage, wing size and location, or desired thrust to optimize a project design. The Helios program for rotary-wing and the Kestrel program for fixed-wing vehicles allow advanced physics-based modeling of the airflow, airframe, propulsion integration, and control-surface effects. These tools allow for high-accuracy predictions of an aircraft throughout its flight envelope so that design problems can be identified and corrected before cutting metal. Data derived from the virtual prototype can also be used to streamline wind tunnel and operational test plans.

The CREATE-Ships tools provide similar capabilities for naval architects. The Rapid Ship Design Environment (RSDE) allows ship designers to explore, literally, thousands of configuration options for a ship in a short period of time.

This tool allows engineers to make informed decisions about tradeoffs between potential designs in terms of ship size, endurance, lift capacity, and speed, among other factors. A second tool in the Ships suite, Integrated Hydrodynamics Design Environment (IHDE) , allows for simplified hydro-dynamic analysis of ship designs. This information is crucial to assess ship performance in anticipated operating areas, as well as provide projections of fuel consumption integral to assessing ship endurance characteristics.

To support higher-fidelity analysis in ship design, CREATE-Ships developed two additional products, Navy Enhanced Sierra Mechanics (NESM) and the NavyFOAM program. Navy Enhanced Sierra Mechanics calculates the effects of external shocks on ships. The NavyFOAM product provides high-fidelity assessments of the hydrodynamics of ships and submarines.

NESM supports the acquisition process in several ways. For example, NESM supports the Live Fire Test and Evaluation (LFT&E) requirements for the new Ford Class nuclear-powered aircraft carriers. All ship classes are required to be tested for shock damage resistance. Historically, testing included only external explosive charges. After testing, engineers recorded the effects, and, if necessary, designed and implemented modifications for the ship. This process requires a ship's dedication to the lengthy test process, as well as dedication to retrofitting any shock resistance deficiencies into the ship's design. In addition, NESM allows for a computational assessment of shock on the ship and equipment while still in the design stage, when designers can incorporate fixes prior to construction.

Similarly, NavyFOAM provides high-fidelity, physics-based assessments of how a ship or submarine will perform hydrodynamically. This tool allows naval architects and engineers to assess how variations in hull forms or control surfaces impact ship performance across a large spectrum of operational conditions.

The CREATE-RF SENTRi project provides electromagnetic engineers and radar designers with a high-fidelity tool for assessing the expected performance of multiple radars and antenna systems mounted on aircraft, ships, and ground vehicles. This capability is of vital importance, as military platforms are modified during their service life. As anyone with a cell phone knows, electronics technology is constantly changing and growing in capability. For a military platform with a 20- or 30-year service life, changes in its communications and sensor gear is inevitable. SENTRi allows engineers to assess the potential performance and impact of an antenna design change on other systems. The latest versions of SENTRi are capable of assessing problems with more than 3 million unknowns. Likewise, the latest versions of SENTRi allow users to assess higher frequencies on larger platforms mathematically, using a supercomputer, that previously could only be assessed experimentally on specially instrumented platforms and ranges.

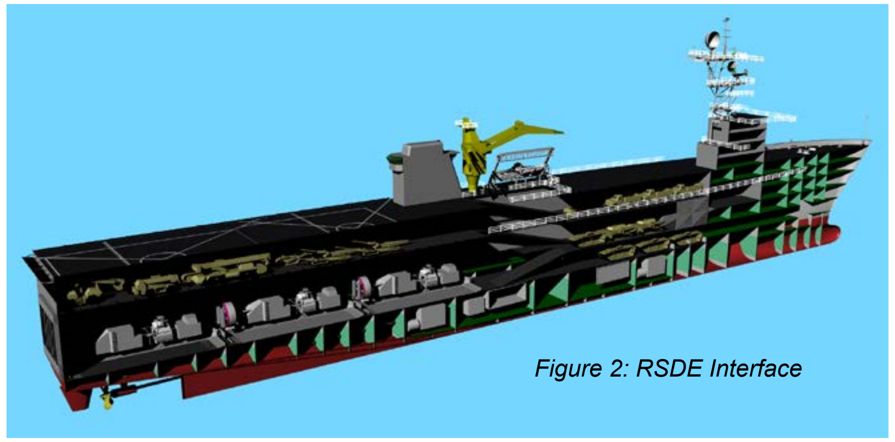

Figure 2: RSDE Interface

Figure 3: Live Shock Test

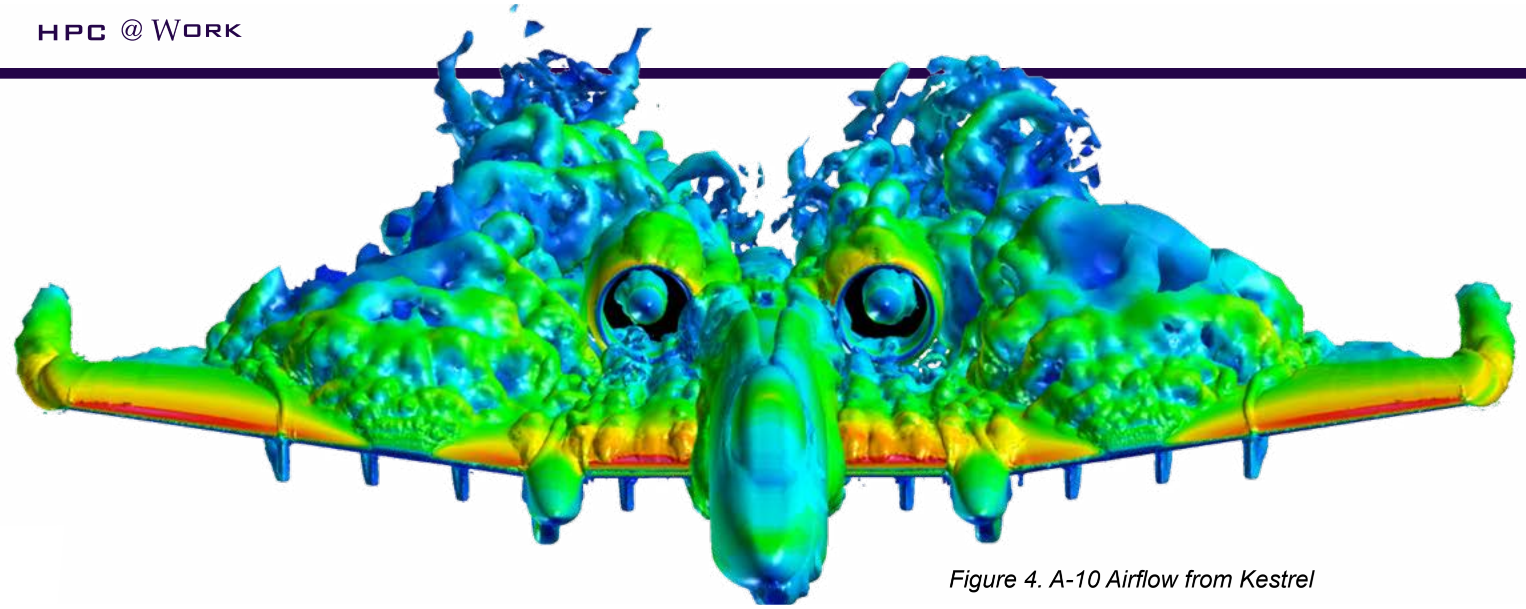

Figure 4. A-10 Airflow from Kestrel

Recent CREATE Acquisition Program Contributions

The U.S. Army uses CH-47 Chinook helicopters for battlefield mobility and resupply. The ability to lift more weight improves the combat and rescue capabilities of our forces, as well as their utility in humanitarian missions. Maintaining flight performance is a key element in the survivability of the aircraft in a hostile environment. To achieve additional lift capability, with no loss of flight performance, Boeing and the U.S. Army Aeroflightdynamics Directorate developed a new set of rotor blades for retrofit on the CH-47 helicopters.

Army aviation engineers and Boeing used CREATE-AV Helios software and HPCMP computers to confirm Boeing's predictions of improved rotor performance of the new blade design, and then for the first time, were able to validate computationally the complex aerodynamics and installed performance of the new rotors and the airframe. The Helios tool provided the unique capability of accurately assessing two sets of counter-rotating tip vortices and their multiple interactions with each other and the fuselage. This increased the engineers' confidence that the rotorblade retrofit would successfully meet its goal to provide 2,000 additional pounds of load capacity for the Chinook, without a significant negative impact on forward flight performance.

The CREATE-AV Kestrel software and HPC assets were used to assess performance impacts from the complex airflow interactions between engines and airframes for possible modifications to the A-10 Thunderbolt II "Warthog" close air support aircraft. Kestrel has also been used to assess the interactions between airframes and weapons to identify potential risks during the weapon launch sequence.

The A-10 close air support mission requires operations at high angles-of-attack and quick maneuvering, which increases the likelihood of engine air inlet distortion and potential engine stalls. Engineers used CREATE-AV Kestrel to model the potential changes to the wing and engine inlet to determine the effect on engine airflow throughout the plane's mission envelope. Kestrel enabled the A-10 Program Office to make informed decisions in a timely manner using virtual prototypes.

In a recent study, the Navy used RSDE to examine the possible design characteristics for new warships to provide physics-based answers to questions of "what can it do?" and "how much will it cost?"

The study examined a range of options for the next class of amphibious assault ships dubbed the LX(R). This type of ship is used to support the U.S. Marine Corps in its mission to conduct expeditionary operations from the sea. The ships typically carry large numbers of ground troops, helicopters, short-takeoff fixed-wing aircraft, and various types of landing craft. During previous studies of this type, the design team could only develop, analyze, and price-out five to twenty different ship concepts. With RSDE, the team was able to develop and evaluate more than 22,000 concept designs. This allowed the team to generate a spectrum of options on troop capacity, aviation capabilities, vehicle deck configurations, etc. The ship designers provided information to decision makers concerning the impact of ship cost and operational capabilities of different requirements for speed, endurance, and military capability of the concept vessels.

With these and many other efforts, the HPCMP CREATE™ program is positively impacting the acquisition process by providing timely scientific answers to important questions. The virtual prototyping concepts capable with the CREATE software demonstrate the utility of supercomputing to the acquisition community for more than war-gaming. Similar efforts in the private sector, with companies like Goodyear and Caterpillar, improved the reliability of products and significantly reduced time-to-market. The DOD is beginning to leverage the capabilities of HPCMP CREATE™ software to improve its acquisition of platforms.

Hercules System at ARL DSRC

Supports Army Network Integration and Evaluation (NIE) Exercise

Ken Renard, ARL DSRC

The Hercules system at ARL DSRC proved a vital resource in the most recent Army Network Integration and Evaluation (NIE) effort. The Army NIE is a semiannual exercise that focuses on the testing of complex and powerful Army tactical network systems. During these exercises, engineers use mobile communication platforms to collect significant amounts of valuable network performance data. Aggregation and analysis of the resulting large data sets yield valuable insight into the capabilities and effectiveness of the systems involved. As part of ARL's strategic partnerships with other Army tenants at Aberdeen Proving Ground, ARL and the Aberdeen Test Center (ATC) leveraged the Hercules system to address the NIE's requirements for computational processing, data storage, and high-speed network transfer for ATC's time-dependent, large-data analysis effort.

During an NIE exercise, data collections from radio systems can produce 1-2 terabytes of data in a 12-hour test period. ARL worked with ATC to gain Army and High Performance Computing Modernization Office (HPCMO) approval for a Dedicated Support Partition (DSP). Using Hercules and SDREN, the ARL-ATC teams successfully transferred and processed terabytes of data, originating at White Sands Missile Range, in a fraction of the time the task would have normally taken. Execution of the DSP began in October and continued through November. Using over 1.5 million Hercules hours, the NIE's data reduction effort was a success. Likewise, the requirement for a quick turn-around in network data reduction and analysis revealed a complementary relationship between HPC resources and a DSP. The codes used on Hercules were the result of a two-year joint effort between the organizations to rewrite the data reduction and network analysis software, while cultivating a successful HPC environment. The framework and algorithms developed during that period have been used during several other events. The ARL-ATC teams conducted significant updates between, and during events, in response to changing test designs and data formats. The HPC resources used during the 2.5-month period surrounding this year's official test event demonstrated the capabilities of the High Performance Computing Modernization Program (HPCMP) networking and computing infrastructures; resulting in an average turnaround time of less than 12 hours to process the data sets. The previous, non-HPC software was averaging a 2.5-day turnaround for much smaller, less complex data sets.

Advanced Research Collaboration and Application Development Environment

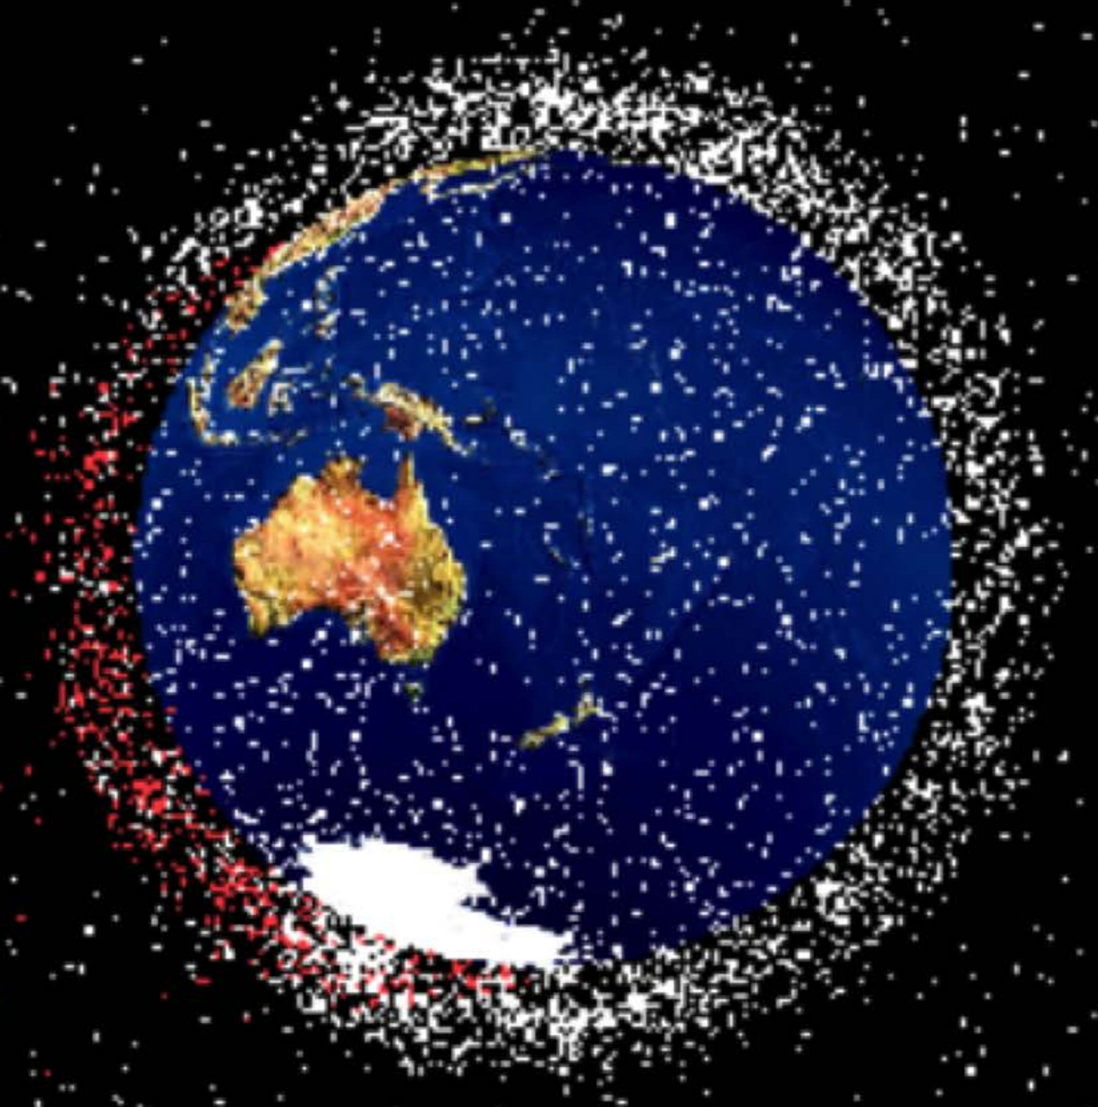

Chris Sabol, *U.S. Air Force Research Laboratory, RDSM, HSAI-SSA Director*
R. Bruce Duncan, *The Boeing Company*
Betty Duncan, *MHPCC DSRC*

Introduction

The Air Force Research Laboratory (AFRL), in partnership with the High Performance Computing (HPC) Portal project and the Air Force Space Command Space and Missiles Systems Center (SMC), is developing the Advanced Research, Collaboration, and Application Development Environment (ARCADE). ARCADE provides technology developers with all of the necessary components to take an idea and develop it into a service for transition into the Joint Space Operations Center (JSpOC) Mission System (JMS). This includes: software development tools, assisted application-to-service wrapping capability, access to relevant data sets, benchmarked system-of-systems testing with relevant scenarios, JMS documentation, test servers, User Defined Operational Picture (UDOP) development kit, and JMS Service-Oriented Architecture (SOA) implementation with access to key services. A high-level concept graphic of ARCADE is shown in Figure 1. The capabilities embodied in ARCADE provide a transformational approach to accelerate technology development and transition capabilities to the JSpOC. This article provides a brief description of ARCADE and status and achievements of its HPC-backed component, the ARCADE Portal.

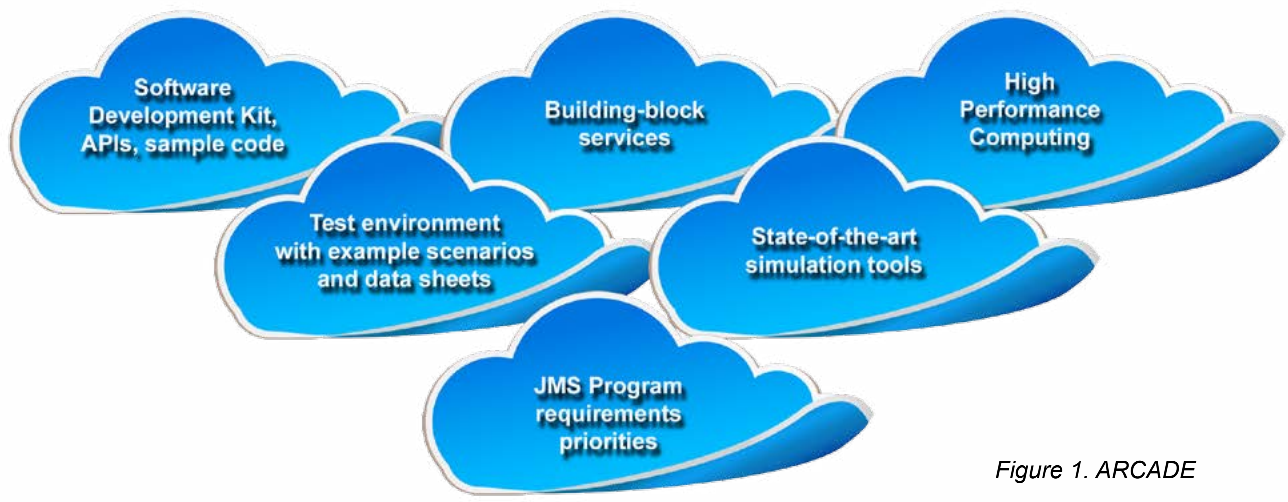

Figure 1. ARCADE

Background

The JMS acquisition program of record is being developed by SMC to deliver an integrated, net-centric Space Situational Awareness (SSA) and Command & Control (C2) capability to provide the following primary functions:

- Replace the legacy Space Defense Operations Center (SPADOC) and Astrodynamics Support Workstation (ASW)
- Rapidly detect, track, and characterize objects of interest: High-Accuracy Catalog
- Produce User-Defined Operational Picture (UDOP) and Space Order of Battle (SOB)
- Perform space threat analysis such as predict and report orbital conjunctions
- Conduct C2 of space forces in dynamic environment

Researchers are developing ARCADE to provide the JMS program, a DOD Acquisition Category I program, and its developers with a transformational application development capability to test and benchmark algorithms and services, accelerate technology development, and transition important functions to the JSpOC and other entities. Traditional acquisition strategies are cumbersome for JMS since SSA and C2 challenges are broad, and expertise extends throughout the country. Capability gaps are rapidly evolving, making agile acquisition a requirement; resources are extremely limited, making leveraging of activity in the Research, Development, Test and Evaluation (RDT&E) community a must. Given the broad spectrum of challenges, the JMS acquisition strategy should seek to remove bottlenecks to technology transition. ARCADE is being developed to mitigate these risks.

ARCADE provides a software application/capability development and evaluation environment effort to mature capabilities through various Technology Readiness Levels (TRLs) in support of the JMS program. Likewise, it is the initial step in the transition pathway for JMS developers. ARCADE is critical to rapid prototyping of future JMS capabilities. It will help support seamless transition in the FY16+ time frame, and provide data sources and services in a JMS relevant environment for all developers to collaborate and build upon. In addition, it will usher in a paradigm that could significantly save the government money to achieve JMS objectives.

ARCADE will overcome challenges to JMS, as no relevant environment exists to develop, mature, and evaluate new applications. It will provide the following key capabilities:

- Web access to software development tools and resources
- Assisted Web-service generation from executables
- Processes to evaluate and transition to JMS program
- End-to-end performance testing in user "sandboxes" with benchmarked scenarios
- Access to key JMS services for use in development and testing
- R&D instantiation of JMS SOA(ARCADE SOA)
- Documentation and help
- Relevant real and simulated data
- Streamlined Information Assurance process for hosting service in ARCADE SOA

ARCADE is comprised of two major components:

ARCADE Portal

AFRL Directed Energy Directorate (AFRL/RD) hosts a software development and test environment at the Maui High Performance Computing Center DOD Supercomputing Resource Center (MHPCC DSRC) using Riptide and Utility Server platforms, which is available through the HPC Portal. The ARCADE Portal provides rapid RDT&E capability that is primarily geared toward providing developers the tools they may need to develop their products, and the acquisition community the insight into identifying high-payoff performers; thereby dramatically reducing performance risk. The facility offers users access to the many common development tools offered by the HPCMP, and a Linux environment to develop, test, and run their code within. A workflow tool is included to allow users to conduct benchmarked system-of-systems testing with relevant scenarios for refinement and evaluation of services.

ARCADE SOA

AFRL Space Vehicles Directorate (AFRL/RV) hosts the ARCADE JMS Service-Oriented Architecture (SOA) at their Battlespace Evaluation and Assessment SSA Testbed (BEAST) facility at Kirtland Air Force Base, NM. The ARCADE SOA dramatically reduces integration risk by offering a true representation of the JMS SOA from both a software and hardware perspective. AFRL/RV provides the SOA on various enclaves by providing for rapid integration of these capabilities into the JMS SOA, and early exposure of R&D services to JSpOC analysts. Streamlined Information Assurance (IA) processes developed and approved for allowing service exposure on key networks, along with pertinent data sources are the key enablers to rapid fielding of new capabilities.

Service interfaces are designed to be consistent between the ARCADE Portal and SOA, to allow easy transition between the two. Combined, ARCADE provides all of the basic capabilities a developer needs to create, test, and prepare a service for transition to the JMS program.

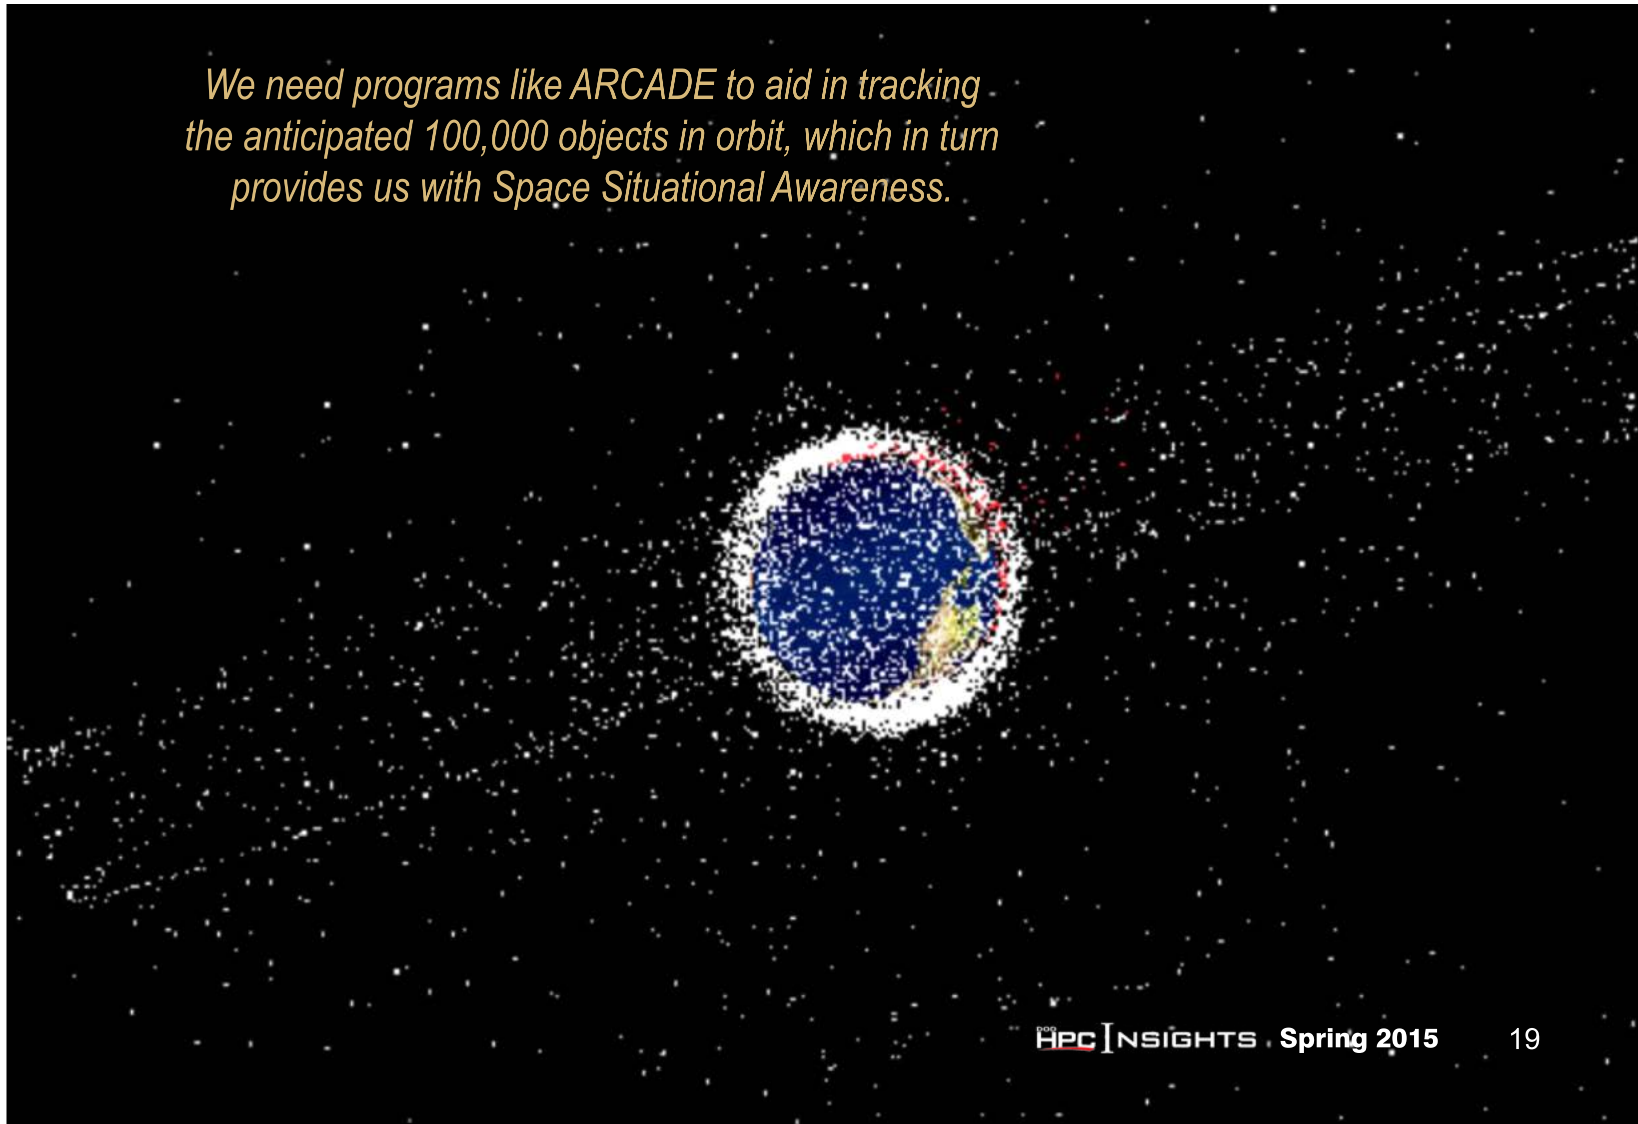

ARCADE Portal Status and Achievements

The ARCADE Portal has been supported through the HPC Software Applications Institute for SSA (HSAI-SSA) and the HPC Portal program, with most development occurring at the MHPCC. Several significant achievements and capabilities have been realized to date in the development of the ARCADE Portal environment. These include:

- Virtual Development Environment providing access to a suite of software development, test, and management tools through the browser
- Workflow integration control and parallel execution tools
- Uncorrelated track (UCT) resolution optical processing prototyping allowing for closed-loop catalog building exercises and demonstrations
- Correlation and orbit update astrodynamics code integration
- Assisted service wrapper generation
- Covariance-based Metric Tasker (CoMeT) capability
- Continued refinement of the design, development, and implementation of pregenerated benchmarked scenarios that preset workflow services with a portal graphical user interface
- Development of dynamic, event-/time-driven parallel execution workflow development
- Porting most of the functionality to the ARL Hercules system

Given these capabilities, a user can take an algorithm and:

- Create an executable through one of the virtual applications or Web shell
- Generate a JMS-compatible service through the assisted Web-service wrapping capability
- Conduct end-to-end, system-of-systems performance testing in a dynamic, closed loop scenario of relevance to the JMS program

The HPC Portal development team has designed, developed, and posted an informational public domain website at:

http://www.mhpcc.hpc.mil/portal/arcade/

This site provides more information on how to obtain an HPC Portal account and ARCADE Portal account. The informational website opening home page is shown in Figure 3 (right).

Summary and Acknowledgments

ARCADE:

- Is critical to rapid prototyping of future JMS capabilities
- Has a repository of documentation to help developers understand JMS needs, processes, and practices
- Houses development tools that assist scientists and engineers in creating applications and services for transition to JMS
- Provides data sources and services in a JMS-relevant environment for all developers to collaborate and build upon
- Contains workflow tools with scenarios and benchmarks to allow for performance evaluation of new approaches in a system-of-systems way
- Ushers in a paradigm that could significantly save the government money to achieve JMS objectives

Figure 3. ARCADE Home Web Page

Significant contributions to the ARCADE Portal have come from Jason Addison, Mark Elies, Brad Farnsworth, Tom Gemuend, Randy Goebbert, Jason Hirata, Jeff Houchard, Adam Mallo, Tar Nusso, Raphael Pascual, Glenda Ramos, Vicki Soo Hoo, Steve Tengolics, and Ron Viloria. Dr. Jeremy Murray-Krezan of AFRL/RV leads the ARCADE program, and Dr. Chris Atwood leads the HPC Portal project.

HPC Portal Update:

Towards Productive Supercomputing for Access-Constrained DOD Environments

***Randy Goebbert**, MHPCC DSRC,*
Pacific Defense Solutions, LLC,
a Subsidiary of Integrity Applications Inc.

The High Performance Computing (HPC) Portal development team, under the direction of the High Performance Computing Modernization Program (HPCMP), and in collaboration with HPCMP Centers, Security, and DOD Supercomputing Resource Centers (DSRCs), continues to expand HPC Portal DOD supercomputing access, while developing new capabilities and applications.

Key goals for the team are to provide transparent, secure, and user-friendly access to computational and storage resources for communities of practice that have not traditionally benefited from access to supercomputing. The HPC Portal opens opportunities for DOD agencies and users who, in the past, could not access the HPCMP DSRCs because of the requirement to install Kerberos authentication software to access the HPC resources. As part of this process, the team is interested in providing services that make it easy to develop new applications, or adapt existing applications for the HPC Portal framework. For the fiscal year 2014 HPCMP survey, the HPC Portal was ranked second in usage among HPC program services.

During 2014, the HPC Portal team helped expand access to supercomputing for the DOD science and engineering community by successfully launching the HPC Portal at the Navy, AFRL, and ERDC DSRCs. The team has also made fundamental improvements to the stability and performance of the core portal infrastructure. The HPC Portal is located at https://portal.hpc.mil. For convenience, all users with existing portal to Information Environment (pIE) accounts are automatically enabled for access to the HPC Portal without any additional requirements. At the HPC Portal site, users authenticate via the HPCMP OpenID capability using only a CAC or HPCMP YubiKey. The HPC Portal enables supercomputing and visualization access from a Web browser, and requires no user-installed software on the local workstation.

The HPC Portal helps bring the engineering workflow closer to increasingly large datasets, and establishes new ways for HPC to increase project cycle effectiveness for user communities. Currently, the Portal provides file and job management, access to the “command line,” access to the HPCMP CREATE™-AV Kestrel and Helios Web applications, “Distributed” MATLAB, and Utility Server-hosted COTS and GOTS (commercial and government off-the-shelf) applications (e.g., Abaqus, Ansys®, HPCMP CREATE™ Capstone, LS-DYNA, ParaView, Pointwise®), which are delivered as an HTML5 Virtual Network Computing (VNC)-based remote display.

The updated HPC Portal File Manager (Figure 1) now provides access to user files via the browser on the Utility Server, and DSRC supercomputer file systems. Users may upload, download, create, copy, edit, delete, modify permissions, and transfer files between HPC resources or the Utility Server at any given DSRC. This capability has greatly expanded during the last year in both functionality and performance.

The HPCMP office is actively developing and seeking new HPC Portal applications. Likewise, the HPC Portal team has also developed a downloadable virtual machine (VM)-based software development kit (SDK) that allows developers to quickly stand up new applications. Using the Software-as-a-Service (SaaS) paradigm, the workload is reduced since the applications need to be built, tested, tuned, and deployed only once to a DSRC, versus the inherent problems of supporting many disparate departmental clusters. Currently, the HPC Portal team is finalizing the “ParaViewWeb” application, a visualization tool that leverages HPC infrastructure to deliver interactive 3D Web-based visualizations. Read more about the development of the HPC Portal at http://www.mhpcc.hpc.mil/portal.

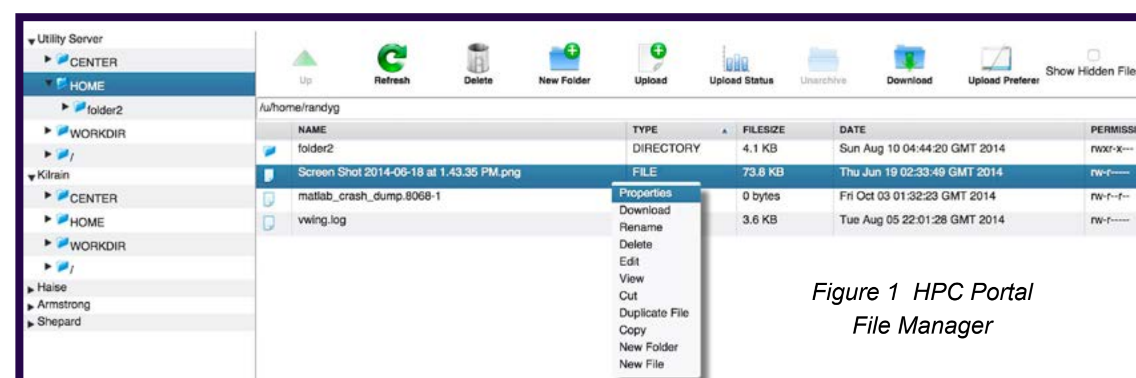

Figure 1 HPC Portal File Manager

Basic Usage Guide
for the
Intel Xeon Phi Nodes
On the Cray XC30 & XC40 Systems at the Navy DSRC

***Bryan Comstock**, Chief Technologist, Navy DSRC*
***Alan Minga**, Senior Presales Analyst, Cray Inc.*

The Cray XC30 systems, Armstrong and Shepard, delivered as part of an FY13 HPC system acquisition, as well as the Cray XC40 systems, Bean, Conrad, and Gordon, acquired as a result of the HPCMP's FY15 procurement activity, contain nodes with Intel Xeon Phi coprocessors. Armstrong and Shepard each contain 124 nodes with one Intel E5-2670v2 Xeon host processor, 32GB of memory, and one Intel 5120D Phi coprocessor. The 5120D is a 60-core model with 8GB of internal GDDR5 memory per coprocessor. Communication with the host E5-2670v2 processor is via the PCI Express (PCIe) interface. Cray uses the compact SXM form factor to fit each Phi on a blade of nodes, but the interface is still PCIe. The XC30 and XC40 systems support two Phi modes: Native mode, the mode that executes all code on the Phi; and Offload mode, where numerically intensive parts of the code run on the Phi, and the rest run on the host.

Notably, these systems currently do not support the Symmetric mode, where executables compiled specifically for the host and the Phi run on both at the same time. Each Phi core has two dual-issue pipelines for scalar instructions. Both scalar pipelines utilize 64-bit data paths with a one-clock latency. The scalar unit supports up to four hardware threads per core. The instruction decoder is fully pipelined, but designed as a two-cycle unit that allows a significant increase to the maximum core frequency. However, the core cannot issue instructions from the same context in two adjacent cycles. You must use at least two threads or cores to utilize all of the available cycles of the scalar unit. Each Phi coprocessor core also has a functional multithreaded vector unit that is 512 bits wide. There are thirty-two 512-bit vector units available per context.

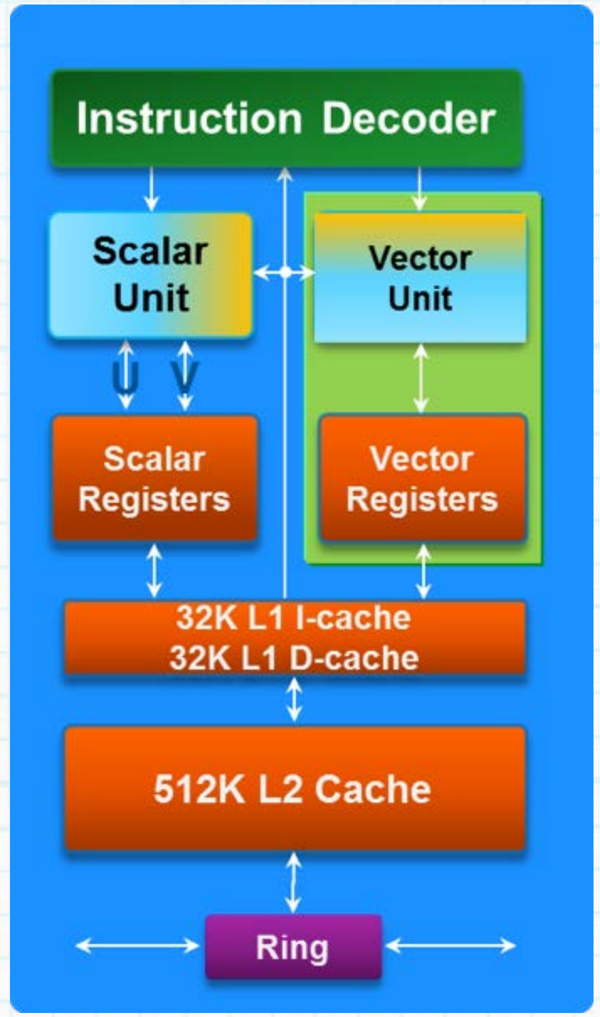

Performance Tips for Phi

1) **Always allow one core on the Phi to run the operating system. The Phi coprocessors on the XC30s and XC40s contain 60 cores each, therefore the user code should execute on 59 cores for the best performance.**
2) **Utilize four threads or cores when possible.**
3) **Use OpenMP environment variables, *KMP_PLACETHREADS* and *KMP_AFFINITY* to set specific thread affinity.**
4) **Use *aprun* option *–cc none*, to disable default *aprun* behavior and ensure proper thread affinity on the Phi.**
5) **In addition, Intel compiler version 15, available on all Cray systems, performs automatic data alignment for Phi, and will yield better performance on codes in which the user has not manually completed data alignment.**

Running in the Native Mode

Users can execute most codes on a Cray XC30 or XC40 Phi in native mode with ease, since the Phi uses an x86-64 instruction set. The code should use OpenMP, and scale to at least 240 OpenMP threads for the best performance.
To build a code for native mode usage, follow the steps below:

1) **Swap to the Intel Programming Environment**
   module swap PrgEnv-cray PrgEnv-intel
2) **Unload the ATP and libsci modules**
   module unload libsci atp
3) **Set the target processor to Phi**
4) **Add *–mmic* and *–openmp* to the compiler options**
5) **Add *–k* to your *aprun* command line to run code on the Phi portion of the node**

Example Cray XC30 or XC40 Phi Native Mode PBS Batch Script

```
#!/bin/sh
#PBS -q phi
#PBS -l select=1:ncpus=10
#PBS -A <PROJECT>
#PBS -L walltime=01:00:00
echo "Running the saxpy with aligned data."
echo "Running with 240 threads on 60 cores"
echo "Leaving uppermost core free for OS
threads"
export LD_LIBRARY_PATH=/tmp
export OMP_NUM_THREADS=236
export KMP_AFFINITY="granularity=fine,compact"
export KMP_PLACE_THREADS="59c,4t"
aprun -cc none -k ./fsaxpy_opt
```

Running in the Offload Mode

Phi's x86-64 instruction set makes it easy to run the Phi in native mode with almost any code. The code used should include OpenMP, and should scale to at least 240 OpenMP threads for the best performance.
To build a code for offload mode, follow these steps:

1) **Swap to the Intel Programming Environment**
   module swap PrgEnv-cray PrgEnv-intel
2) **Unload the ATP and libsci modules**
   module unload libsci atp
3) **Add *–openmp* to the compiler options**

Example Cray XC30 or XC40 Phi Offload Mode PBS Batch Script

```
#!/bin/sh
#PBS -q phi
#PBS -l select=1:ncpus=10
#PBS -A <PROJECT>
#PBS -L walltime=01:00:00
echo "Running the saxpy with aligned data."
echo "Running with 240 threads on 60 cores"
echo "Leaving uppermost core free for OS
threads"
export LD_LIBRARY_PATH=/tmp
export MIC_ENV_PREFIX=PHI
export PHI_OMP_NUM_THREADS=236
export PHI_KMP_AFFINITY="granularity=fine,compa
ct"
export PHI_KMP_PLACE_THREADS="59c,4t"
export OFFLOAD_REPORT=2
aprun -cc none ./fsaxpy_opt
```

Example SAXPY Code

```
! scale the calculation across threads requested
! need to set environment variables OMP_NUM_THREADS and KMP_AFFINITY

!next directive needed for offload mode only
!dir$ offload target (mic : 0)
!$OMP PARALLEL do PRIVATE(i,j,k,offset)
 do i=1, numthreads
    ! each thread will work its own array section
    ! calc offset into the right section
    offset = i*LOOP_COUNT

    ! loop many times to get lots of calculations
    do j=1, MAXFLOPS_ITERS
        ! scale 1st array and add in the 2nd array
        !$omp simd aligned( fa,fb,a:64)
        do k=1, LOOP_COUNT
            fa(k+offset) = a * fa(k+offset) + fb(k+offset)
        end do
        !$omp end simd
    end do
 end do
```

QA in the Era of “Always On” Systems

Dan Bigelow, *Quality Assurance Lead,*
MHPCC DSRC, University of Hawaii

Not long ago, “24 by 7” was a term used among service providers to describe their requirement to keep computer servers and networks available all of the time. In today’s fast moving and interconnected world, our utilities, phones, and network-based services are simply expected to ‘be there,’ anytime we need to access them – 24x7. This presents an interesting challenge for contemporary testing and quality assurance teams, who have typically been able to validate and verify system changes offline, before the software product or service was released for general access. Today, rather than being replaced, the systems discussed almost morph from version to version, showing little, if any, downtime.

Take the new HPC Portal installed at each DOD Supercomputing Resource Center (DSRC). The portals are each comprised of a general framework, where users can launch a variety of Web-based applications. Not only does the framework stand on its own as a customized application, but each of its sponsored applications (e.g., Kestrel, MATLAB, WebShell, FileManager, etc.) interact with that framework, as well as supporting their own software functions. Any changes and upgrades to a portal’s operating system and/or framework that might cascade into the various applications require that those applications be (re)tested and verified, to ensure that any unexpected side effects can be quickly identified and resolved.

While the timeless goals of quality assurance teams (availability, responsiveness, reliability, etc.) still remain, the traditional tools and models no longer apply. In a Web-based world, tracking user authentication, browser variations, network congestion, and security mechanisms all result in the need for an upgraded approach to quickly, efficiently, and deterministically validate these dynamic systems. The five deployed instances of the portal each support four or more applications and must be tested across eight browser families, which means the portal can no longer be efficiently tested using manual methods.

Fortunately, the tools and techniques used to identify system issues have matured right along with the systems they validate. The practical solution is twofold: extending classic unit-testing techniques to pretest small code segments ‘closer’ to the user interface, and the use of automated interface testing tools.

In order to keep up with these rapidly changing systems, we are relegated to using software to test our programs. In addition to using well-established unit testing systems like J-Unit, we now have automated testing toolkits such as the Selenium Framework (www.seleniumhq.org), and PhantomJS (www.phantomjs.org - a simulated Web browser). These toolkits allow us to prerecord user actions, convert those actions into test codes, and run those codes on servers that don’t need (or even have) monitors, keyboards, or mice. In this way, we can quickly and repetitively run complete sets of test simulations.

Because humans still have to record, refine, and maintain the many test sequences that ensure coverage of the key elements of our ever-changing framework and its applications, assuring the quality of the portal is not quite as simple as we would like it to be. However, it’s much more efficient, interesting, and productive to design and refine these automated testing systems, versus manually clicking through multiple variations of applications and browsers to determine if the framework still works properly.

CCAC Has Become HPC Help Desk

In order to more accurately describe the services provided, the Consolidated Customer Assistance Center (CCAC) has been renamed the “HPC Help Desk.”

This change became effective 18 May 2015. As part of this change, access to the HPC Help Desk ticket system has moved from https://help.ccac.hpc.mil to the new

https://helpdesk.hpc.mil

Accordingly, the email address to contact the HPC Help Desk has changed to help@helpdesk.hpc.mil. The old URL and email address will remain available with auto-forwarding and redirection for a period of time, but you are encouraged to update your bookmarks and contacts now.

The phone number, 1-877-222-2039, will remain the same.

AFRL Management Team

at Maui High Performance Computing Center DSRC

Betty Duncan, *MHPCC DSRC*


Ms. Laura Ulibarri

It is with great pride that the Maui High Performance Computing Center announces its new Acting AFRL Director, Ms. Laura Ulibarri. Ms. Ulibarri received bachelor and master of science degrees in physics from the University of New Mexico. Her most recent assignment was as the Deputy Chief of the Space Electro-optics Division, where she managed more than 400 personnel and executed a $116 million annual portfolio of Air Force, Department of Defense, and national research and development programs. Since joining government service, Laura has performed strategic planning for the AFRL Directed Energy Directorate’s Space Situational Awareness portfolio, and led the $450 million Air Force Maui Optical and Supercomputing (AMOS) site through a significant transformation, realigning the research and support operations into a cost-effective and highly-valued mission partner. Her extensive strategic planning, supervisory, and program management experience is a welcomed addition to the MHPCC team. She will work closely with HPCMP and AFRL management to redefine MHPCC’s role as an HPC Center. Ms. Ulibarri is excited about the opportunities presented by HPC to transform RDT&E for the DOD, and she is grateful to be a part of the team.


Dr. Chris Sabol

Dr. Chris Sabol is the Acting AFRL Technical Director for MHPCC. Dr. Sabol is a nationally recognized Space Situational Awareness (SSA) expert. Likewise, he is a trailblazer whose strategic vision enabled him to spearhead groundbreaking projects years before technical needs were formally recognized. These projects span from basic research to operational demonstrations. Dr. Sabol’s efforts have greatly enhanced the Air Force Research Laboratory’s ability to deliver breakthrough technologies. Dr. Sabol received a bachelor of science degree in aerospace engineering from the Georgia Institute of Technology, a master of science degree from the Massachusetts Institute of Technology, and a doctorate in aerospace engineering sciences from the University of Colorado. Dr. Sabol’s published works are international, and he has received many awards from the AFRL.


Major Micah O’Neal

Major Micah O’Neal is the AFRL MHPCC Deputy Director. Major O’Neal received a Bachelor of Arts degree from Saint Leo University in Florida, and a master of business administration degree from Strayer University in Washington, DC. With more than 20 years of service to the AFRL, Major O’Neal has earned many outstanding and distinguished honors and awards during his career.


Captain Issac Putnam

Captain Isaac Putnam is the AFRL MHPCC Operations Director. Captain Putnam is also the team leader for the Air Force telescopes operations at the Air Force Maui Optical & Supercomputing site. He earned a bachelor of science degree in electrical and computer engineering from the Virginia Military Institute and a master of science degree in electrical engineering (electro-optics) from the Air Force Institute of Technology at Wright-Patterson Air Force Base.

From the Director's Desk ARL DSRC

Dr. Raju Namburu, Director, ARL DSRC

Excalibur, the Cray XC40 procured during FY14 and installed at the ARL DSRC, arrived in mid-November, bringing enthusiasm and excitement with it. The Excalibur system will replace the IBM iDataPlex system, known as Pershing. It will be the first High Performance Computing Modernization Program (HPCMP) system that includes flash disk capability, and the unique feature of 120 terabytes of Solid-State Drive (SSD) storage, which possesses access speeds nearly 100 times faster than a standard hard drive. This increased speed can be significant in boosting the efficiency of programs that access large amounts of data. The DSRC and Cray teams worked together to bring the 101,184-core Cray XC40 system through acceptance and into production on April 13, 2015.

One of the key goals for Excalibur is to support the users across the entire spectrum of DOD communities, which include emerging research communities such as large-data analyses, cognitive sciences, and computational biology that are not typically users of more traditional High Performance Computing (HPC) systems.

Articles in this edition of HPC Insights feature a profile of the Excalibur system. Also featured is an update of our partnership with the Army Test and Evaluation community. This update includes our plans to address their computational and Input/Output (I/O) intensive requirements that include the demand for 'big data' analytics for the Army's Network Integration Evaluation (NIE) exercises. Also noted is our work with the Army NIE 15.1 exercise conducted this past quarter. Another article features news about the IBM iDataPlex system Hercules at ARL DSRC. The Aberdeen Test Center (ATC) leveraged Hercules extensively using a Dedicated Support Partition (DSP) to accelerate the processing and analyses of extremely large-field data sets.

As always, we are poised and committed to support the DOD user community by providing computational support to solve the research challenges most critical to our national defense.

From the Director's Desk Navy DSRC

David Cole, Director, Navy DSRC

In December 2014, I was truly fortunate to receive notification of my selection to succeed Tom Dunn as the Navy DSRC Director as he prepared to retire after 46 years of government service. During this period, he served in multiple key roles in support of the Department of Defense High Performance Computing Modernization Program (DOD HPCMP). Tom's first contribution to the HPCMP began with his leadership role in support of the Commander, Naval Meteorology and Oceanography Command (COMNAVEMETOCCOM); in an effort to develop and submit a proposal to establish a DOD Major Shared Resource Center (MSRC) at the Naval Oceanographic Office (NAVOCEANO). Due to his efforts, the NAVOCEANO supercomputing center was chartered as an MSRC in 1994. From May 1996 until July 1999, Tom served at the DOD High Performance Computing Modernization Program Office in Arlington, VA. Tom's tenure in the Arlington office encompassed several assignments including the DOD Supercomputing Resource Centers' Project Manager, Financial Manager, Technical and Deputy Director, and eventually Director of the HPCMP. Likewise, from October 2006 to December 2014, he served as the Director of the Navy DSRC.

One of Tom's most notable achievements as the DSRC Director was his implementation of an HPC architecture at the Navy DSRC. The new architecture addressed the needs of high-priority projects characterized by time-critical job activity occurring on a regular or recurring basis. Beginning with the DOD HPCMP combined FY11 and FY12 procurements, the Navy DSRC implemented a strategy to acquire two unclassified HPC systems of the same architecture for each HPC procurement cycle. This approach provides optimum availability by having one of the two systems available in the event of scheduled or unscheduled outages. Consistent with this approach, the Navy DSRC will field two large Cray XC40 HPC clusters into the unclassified environment during the summer of 2015. These two new systems will provide a significant increase in computational capability that will support time-critical HPC requirements.

Tom's numerous accomplishments and dedication to HPC significantly contributed to the high performance computing capabilities that are provided by the HPCMP today. In the time-honored Navy tradition, we wish him fair winds and following seas.

AFRL DSRC Innovation Laboratory: Incubating Discovery

Kevin Schoen, *AFRL DSRC*

In 2012, the Air Force Research Laboratory DOD Supercomputing Resource Center (AFRL DSRC) launched a 24-month project to evaluate the benefits of providing a High Performance Computing (HPC) Development Test Bed to our users. Our goals were to: (a) support innovative work in the HPC user community, (b) explore ways for nontraditional research groups to apply HPC to their workloads, and (c) encourage development of new capabilities in our High Performance Computing Modernization Program (HPCMP) systems. We named our Development Test Bed "Talon."

Talon sprang from a small, orphaned, 144-core cluster with 6 project nodes for dedicated applications and services, 12 conventional compute nodes, and 35 terabytes of shared disk with an Infiniband interconnect. The Talon project demonstrated the value of a Development Test Bed to HPCMP management. We are expanding this concept with Talon as part of a new Innovation Laboratory. The Innovation Lab helps transition outside DOD technology developments back to the DOD by adapting them to our environment. The Innovation Lab also supports DOD laboratory IT systems research and HPC technology research. It allows developments that would be disruptive or impossible on production HPCs. A project can evaluate, develop, and demonstrate new HPC capabilities. Users can experiment with tools for nontraditional HPC science domains like big data, data mining, human effectiveness, and bioinformatics.

The Innovation Lab using Talon is an innovation incubator. Innovation Lab projects receive system administration and application support, along with dedicated system resources. We support software like MySQL, Apache, PHP, Java, and Tomcat. We even offer security and engineering review and support for other applications if users need it. However, project owners must accept the risks of interruptions for configuration changes, crashes, debugging, and other similar processes.

AFRL DSRC is currently working on three key projects with hopes of developing new abilities for use in an HPCMP production environment. The first project involves evaluating models of the human mind for the U.S. Air Force 711 Human Performance Wing. This particular project, referred to as Mind Modeling, gives researchers a flexible framework to upload a model, and then to define various input parameters and simulation objectives using a simple Web interface. Next, this system dispatches millions of processing tasks to HPC nodes. As groups of tasks finish, it gathers the results and determines the next set of values to explore. This framework allows researchers to track execution metrics, gives them on-demand status reports, and supplies interactive tools for real-time data analysis. Likewise, researchers no longer have to resort to Linux command lines, nor do they need to copy and stage their data, create job scripts, or monitor job log files. This project is generalizing the framework so other users and projects can easily exploit massive parallelism.

A second project involves hosting the Microscope Image Analysis and Bioinformatics system for the 711 Human Performance Wing. This big data project processes millions of microscope images. HPC nodes calculate over 11,000 characteristics for each cell identified in each image. It uses HPC and machine learning to mine billions of data sets for desired results. The biologist or geneticist uses the project's Web interface to input experiment metadata, launch the HPC processing tasks, and generate analysis reports. This project is evolving to support collaboration with other government laboratories.

The third project involves an experiment to ensure the privacy of DNA sequencing on shared computers. The goal of this project is to launch and manage the HPC processing tasks, while ensuring privacy and integrity. This project also enables self-healing in areas where failed or suspect processing tasks can be restarted. Researchers working on this laboratory experiment intend to convey their research results to an accredited DOD system. Results of this research will stimulate discussion and allow members of the HPCMP information assurance staff to evaluate the capabilities of the project.

Anyone desiring an Innovation Lab Project may request one by following these easy steps: (1) submit a two-page project proposal with your system and software requirements; (2) describe your effort and identify impediments to your effort inherent in production HPCs; and (3) include your estimates of required system and labor resources, and identify the value or impact of these requirements to the DOD.

Talon brings a unique environment to our HPC users. It creates a place to experiment and explore, to push the boundaries of HPC, to try new things, to stretch the imagination, and perhaps to discover something valuable for the DOD, and ultimately for the men and women defending us.

For information about Innovation Laboratory projects:
email: sp-proposal@helpdesk.hpc.mil

For details on the Talon system, see the Talon User Guide at
http://www.afrl.hpc.mil/hardware

DOD HPC INSIGHTS

is a semiannual publication of the High Performance Computing Modernization Program.

Editorial Staff

Managing Editor
Rachelle Hintson

PETTT
George Stathopoulos
Lisa Diletti

DREN
Stephanie Parker

WD
Pam Gilliam

CREATE
Dr. Kevin Newmeyer

HPCMPO
Cathy McDonald

DAAC
Dr. Mike Stephens

HPCMPO Editor
Denise O'Donnell

UAG
Soumya Patnaik

Design / Layout
Ron Compton

Associate Director for HPC Centers
Christine E. Cuicchi

DSRCs

AFRL
Gregg Anderson

MHPCC
Betty Duncan

ARL
Brian Simmonds

Navy
Bryan Comstock

ERDC
Scotty Swillie
Nikkita Mays

Questions or Comments
feedback@centers.hpc.mil

POWERFUL RESOURCES

The High Performance Computing Modernization Program (HPCMP) provides the Department of Defense supercomputing capabilities, high-speed network communications, and computational science expertise that enable DOD scientists and engineers to conduct a wide range of focused research, development, and test activities. This partnership puts advanced technology in the hands of U.S. forces more quickly, less expensively, and with greater certainty of success.

Today, the HPCMP provides a complete advanced computing environment for the DOD that includes unique expertise in software development and system design, powerful high performance computing systems, and a premier wide-area research network. The HPCMP is managed on behalf of the Department of Defense by the U.S. Army Engineer Research and Development Center.

To learn more about the HPCMP, please visit our website at: hpc.mil.

To learn more about our High Performance Computing (HPC) centers, please visit: centers.hpc.mil.

HPCMP User Customer Service Center
1.877.222.2039
centers.hpc.mil

AFRL DSRC

US Air Force Research Laboratory
Wright-Patterson AFB, Ohio
http://www.afrl.hpc.mil

ARL DSRC

US Army Research Laboratory
Aberdeen Proving Ground, Maryland
http://www.arl.hpc.mil

ERDC DSRC

US Army Engineer Research & Development Center
Vicksburg, Mississippi
http://www.erdc.hpc.mil

NAVY DSRC

US Navy
DOD Supercomputing Resource Center
Stennis Space Center, Mississippi
http://www.navo.hpc.mil

MHPCC DSRC

Maui High Performance Computing Center
Kihei, Maui, Hawaii
http://www.mhpcc.hpc.mil