DoD HPC INSIGHTS

Spring 2013

SIKORSKY UH-60 BLACKHAWK
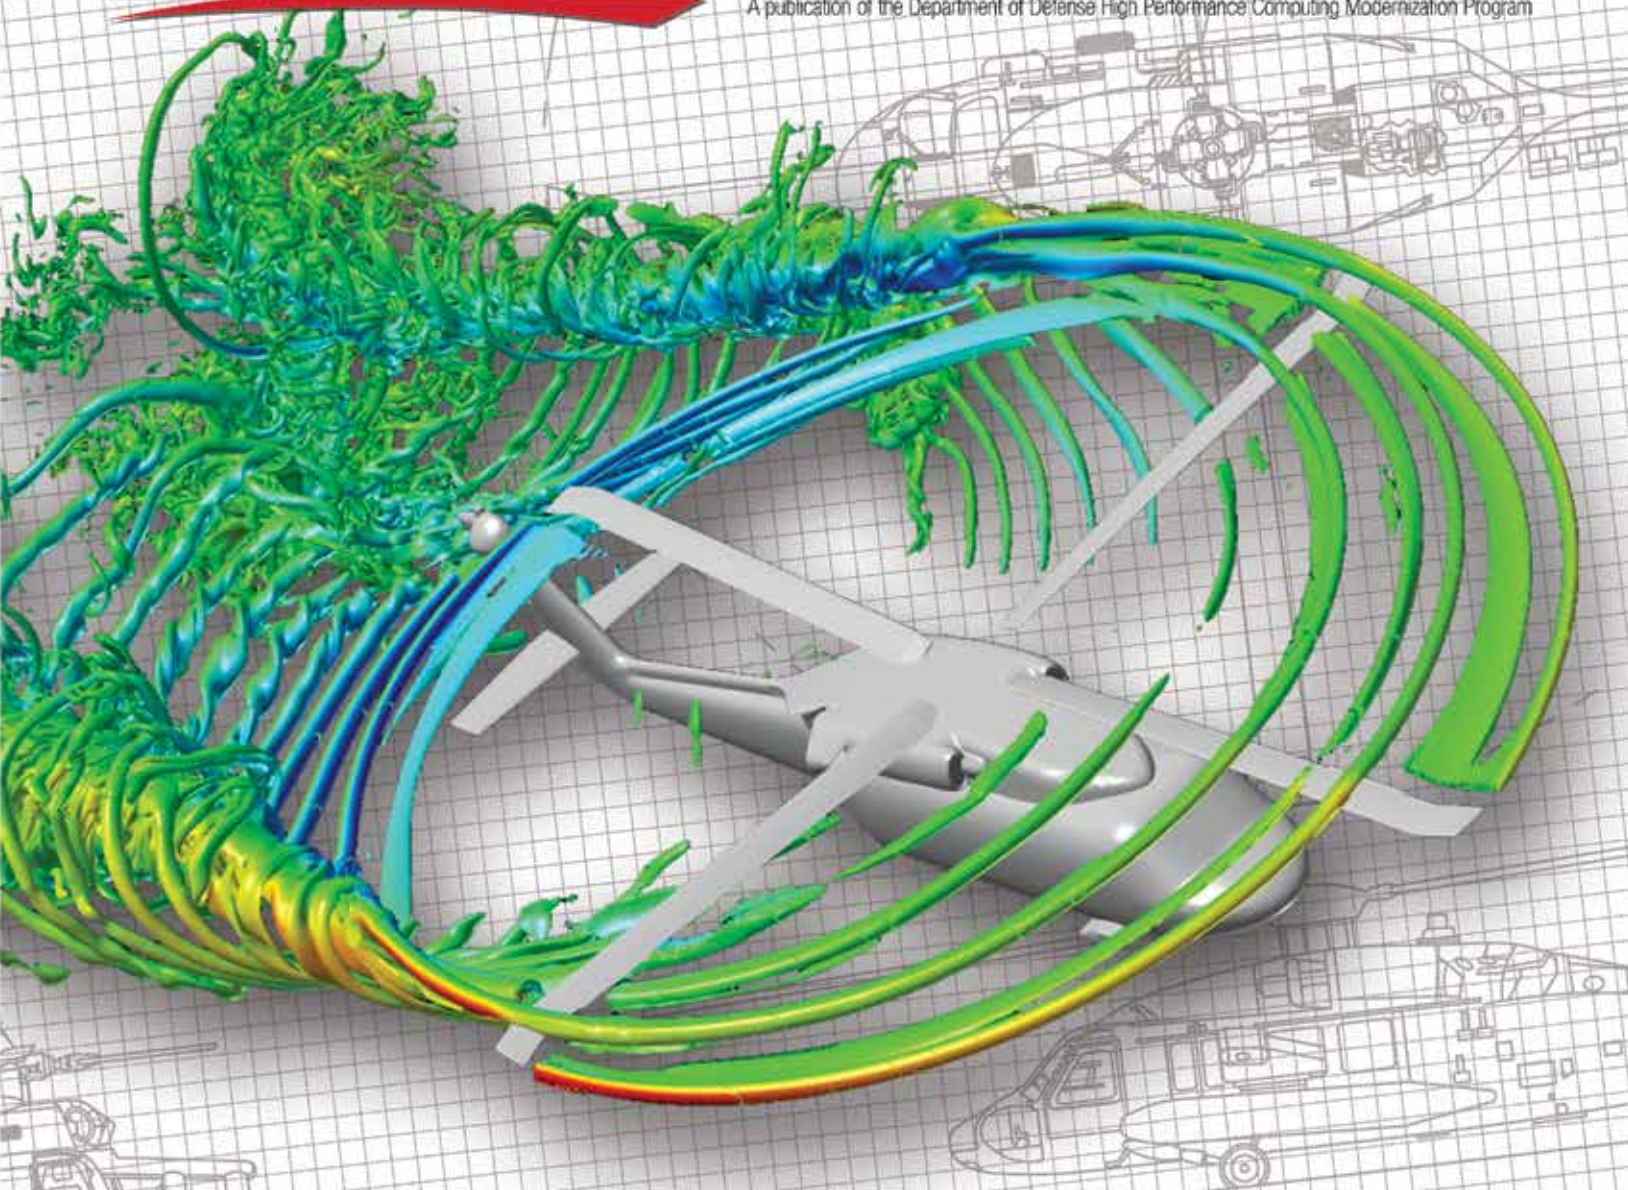
DoD HPCMP Computational Research and Engineering Acquisition Tools and Environments (CREATE) Program

HPC Tools Clear the Path for Unmanned Air Vehicles

CREATE Collaborative Development and Community Support

Contents

About the Cover: CREATE-AV Helios simulation of UH-60 in forward flight using adaptive wake vortex tracking. Simulation done by Mark Potsdam, AFDD. VIZ picture by Nathan Hariharan, CREATE-AV using FieldView. See story, page 7.

First Word

By John E. West, Director, Department of Defense High Performance Computing Modernization Program Office, Vicksburg, Mississippi

John E. West

In today's economic climate, I am acutely aware of the need to describe why HPC matters to our stakeholders—Congress, the Pentagon, and taxpayers. This is where your success stories play a crucial role. Distilling the power of HPC into something meaningful to American taxpayers and lawmakers can be difficult without being able to point to end products and services that demonstrate the value of HPC in their everyday lives.

Your successes send a powerful message to key decision makers: that HPC is, for many in the DoD, the tool of first resort for solving ever more complex technical problems. This message highlights the capabilities that HPC can bring to other areas of the Department, areas that constantly strive to streamline their scientific and technical discovery processes—namely the acquisition engineering and test and evaluation communities, where HPC is not yet as fully integrated as it is in the research and development community. Our consistent message is that HPCMP resources and expertise are enabling simulations of phenomena that are only possible on HPC machines, and we must continue to invest in a robust computational capability in order to maintain our Nation's technological advantage on the battlefield.

The DoD is in a period of substantial change. I am confident that many of these changes are needed to properly position the Department for the future; however, today we are being asked to make structural changes in the way we do business. Within the HPCMP provider community, we face significant challenges in planning and budgeting, travel, contracting, security, and in our interface to the IT enterprise. HPCMP users are facing similar changes in their home organizations.

Through it all, however, we must remember that the work we are doing together—and the work we make possible for HPCMP users—is part of a much bigger picture. Together we are improving the safety and effectiveness of our Nation's fighting forces. This is a job worth doing even when our circumstances make it more difficult than we would otherwise wish.

DoD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM

DoD HPCMP Computational Research and Engineering Acquisition Tools and Environments (CREATE) Program

By Dr. Douglass Post, Chief Scientist, DoD High Performance Computing Modernization Program

The DoD HPCMP Computational Research and Engineering Acquisition Tools and Environments (CREATE) Program is a long-term program launched in FY2008 by the Office of the Secretary of Defense. The objective is to develop and deploy multiphysics-based computational engineering software that, when used in conjunction with increasingly capable high performance computing systems, accurately predicts the performance of naval, air, and radio frequency antenna weapons systems. These tools are designed to enable trade space optimization of new and retrofit designs and to avoid costly (time and money) design flaws and rework. The ultimate goal of CREATE is to catalyze a revolution in weapons system design, and development methodology from reliance on building and testing physical prototypes to virtual prototype design and evaluation, followed by physical prototype validation for the DoD Government and industry research, engineering, and acquisition communities.

The CREATE Program is developing nine individual software applications. The CREATE Air Vehicles project consists of DaVinci—a conceptual design tool; and Kestrel and HELIOS (Helicopter Overset Simulations)—high-fidelity, full-vehicle, multiphysics analysis tools for fixed-wing and rotary-wing aircraft, respectively. The CREATE Ships project consists of the Rapid Ship Design Environment (RSDE)—or conceptual design of ships; NavyFOAM—or prediction of ship hydrodynamic performance; Navy Enhanced Sierra Mechanics (NESM)—for analysis of ship shock and shock damage; and the Integrated Hydro Design Environment (IHDE)—to facilitate use of naval design tools. The CREATE Radio Frequency Antenna project application, SENTRi, is a tool for analyzing and optimizing the performance of RF antennas integrated with DoD weapon platforms. A fourth project, Meshing and Geometry (MG), consists of Capstone, a set of tools for generating the geometries and meshes needed by the other CREATE tools. This collection of tools provides an integrated design and analysis capability that can be applied at all stages of the acquisition process, from requirements definition and conceptual design through sustainment and platform modernization. They are already being used to analyze existing DoD weapon systems and to design new systems. Led by the HPCMP,
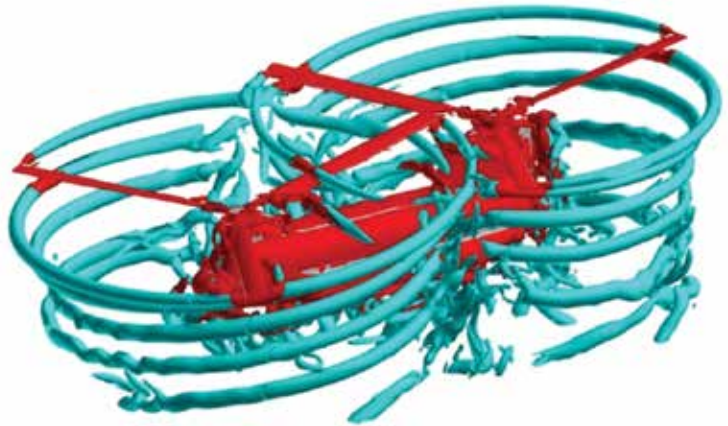
Helios calculation of interaction of vortices from rotor-tips on the CH-47.

this software is being developed by multi-institutional teams of Government and contractor personnel located at military laboratories and facilities with the relevant technical expertise and customer communities for each application area. For instance, NESM is being developed by a group at the Carderock Naval Surface Warfare Center that has the lead for shock vulnerability of naval vessels.

Each year, the CREATE teams release a new version of each software application with increased capability to meet technical and usability requirements of its customer base. For CREATE AV, DaVinci 1.0 provided an initial capability to construct a new airframe model from "parts" that can be combined into an integrated airframe model for analysis. This capability is being used to assess candidate designs for the next-generation Air Force cargo aircraft, among other applications. Kestrel v3.0 enables simulations with two or more bodies in relative motion with control surfaces, providing the ability to calculate store separation, together with the ability to achieve stated accuracy and parallel scaling goals. Kestrel is being used to assess flight worthiness and performance for a number of DoD existing and planned aircraft, including six UAV candidate systems. HELIOS v3.0 enables general multirotor and fuselage modeling for a complete rotorcraft and an automated adaptive mesh refinement scheme that enables accurate tracking of the vortices to study rotor and fuselage interactions. This capability is being used to assess the proposed CH-47 rotor retrofit that is designed to provide an additional 2000 pounds of lift for hover operation.

For CREATE Ships, RSDE is being used for a pilot

project for the DoD Engineered Resilient Systems Program to demonstrate the utility of design optimization of naval vessels. NESM v1.1 provides the capability to conduct a virtual "Full Ship Shock Trial" to analyze the vulnerability of equipment function in a naval vessel to underwater explosions. This capability is being validated for the Full Ship Shock Trial alternative for the CVN-78, the Navy's new class of aircraft carriers. NavyFOAM v3.0 provides the capability to assess ship hull resistance and maneuvering capability. An ongoing application is the design of the bow planes and maneuvering empennage for the Ohio Replacement Submarine, the Navy's new Ballistic Missile Launch Submarine. IHDE added an enhanced ability to analyze bare hull resistance, seaway loads, and seakeeping with a large number of simultaneous analyses. These capabilities are being used to analyze many different ships, including the DDG-51 Flight III bow bulb modification, the Medium Affordable Surface Combatant (MASC), and the Green Arctic Patrol Vessel (GAPV).

The latest version of SENTRi, v3.0, provides much faster algorithms for electromagnetic analysis that allow assessment of larger and more complex antennas and platforms and the ability to handle anisotropic and continuously varying materials. SENTRi can calculate the time-varying electric field patterns on full-sized aircraft such as the CV-22 Osprey Tilt-rotor and Apache helicopter, and it is being applied to other DoD weapon systems as well. Capstone v3.0 provides automated, near-body volume meshing with boundary layers, improved capability to generate unstructured surface meshes, and other improvements. It is being used to provide meshes for next-generation Air Force cargo plane designs generated with DaVinci and for other applications.

The CREATE customer and user communities continue to grow. Over 50 DoD programs are now using the CREATE tools to address acquisition program issues. Over 500 licenses for CREATE software tools have been issued (although not all of these are active). The growth of users has led to the need to support them. The HPCMP has provided funding to seed the establishment of a User Support Consortium (USC) by the DoD Aviation community. The Joint Aeronautical Commander's

Group (JACG), composed of the commanders of the AF Life Cycle Management Center, NAVAIR, and Army Aviation, formed a Tri-Service Integrated Project Team to oversee the USC and to develop a plan for the DoD Aviation community to fund the USC once the HPCMP seed money has been expended. The Air Force Arnold Engineering and Development Complex is hosting and setting up the USC. There is strong interest in the CREATE tools by the US Defense industry. Motivated by the promise of the CREATE software and other physics-based models, the NDIA has established a conference series on Physics-based Modeling for U.S. Defense that features the CREATE software. A number of major Defense industries (Boeing, Bell Helicopter, Raytheon…) are beginning to use the CREATE tools and are showing interest in adopting them into their product development workflow.

Increasingly restrictive computer security access requirements are reducing the ability of DoD scientists and engineers to access high performance computing resources (including the CREATE tools) from their workplaces. A likely endpoint of the present process is that all DoD computer users will have access to only Microsoft Office and a web browser, precluding use of scientific and engineering tools and limiting the use of the HPCMP computers. To allow engineers and scientists access to the HPCMP resources, the HPCMP and CREATE have begun developing a "portal" capability that provides secure access through a modern browser. The initial pilot project is being executed by the Maui DSRC and CREATE. It involves developing a portal that allows a remote user with only a browser to set up, run, and analyze Kestrel problems through a browser. Success with Kestrel is being extended to all the CREATE tools and will be available to facilitate the use of other applications on the HPCMP computers.
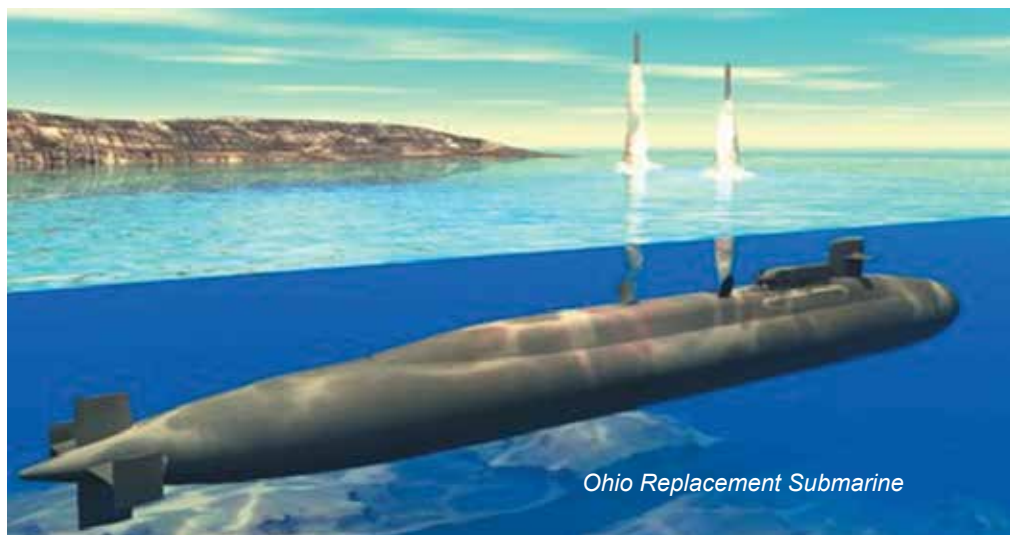

Ohio Replacement Submarine

Adoption of the use of CREATE software by the acquisition engineering community is essential for success. In addition to ensuring user support, the CREATE team is working with ERDC, the DoD, and the Defense industry to address intellectual property and technology transfer issues that need to be resolved to ensure adoption of the CREATE tools while protecting the sensitive information in the tools through export control mechanisms and other means. Among the options being explored are Cooperative Research and Development Agreements (CRADAs) with individual industries, Broad Area Announcements (BAAs), and other contract mechanisms. In addition, to facilitate adoption of the tools, CREATE is supporting use of nonsensitive versions of their CREATE tools by the U.S. Military Academies and by academic groups, such as naval architecture and aeronautical engineering departments of a number of major universities, as part of their course curricula.

The CREATE Program is approaching the midpoint of the planned development schedule. While the original CREATE software engineering strategies have been successful, the team has evaluated the effectiveness and is recasting them to improve the CREATE development process. Key process areas that have emerged as crucial are ensuring the sufficient focus on the development team, meeting customer needs, employing technically mature component technologies, identifying and employing successful agile development methodologies, and gathering requirements in an effective manner. While the features and technical capability of the software will continue to be improving, a major focus during the second half of the CREATE program will be improving usability, including scaling and improved user interfaces and user access, and increasing the focus on verification and validation together with the addition of methods for quantifying uncertainties.

CREATE staff have continued to receive national and international recognition and awards from professional societies and from the Services. These include two AIAA and an SNAME Fellow appointments, the National Academy of Sciences Gibbs Brothers Medal for Naval Architecture, the ASNE 2011 Gold Medal, the IDC HPC Innovation Excellence 2012 Award (given to the HPCMP for work done by a CREATE staff member), and The Technical Cooperation Program (TTCP) 2012 Scientific Achievement Award, as well as numerous internal DoD and Service awards.

The CREATE Program is successfully developing and deploying software with the new features needed by the DoD aircraft, Naval, and RF engineering communites. Customer growth is strong, both in terms of users and programs. CREATE software is already contributing to the analysis and design of important DoD systems, such as the CH-47 rotor-blade retrofit, the Ohio replacement submarine, the CVN-78 shock test, the NAVAIR UAV flight certification, the AF next-generation cargo plane, and many other systems. Progress is continuing in all areas, including user support and access, Intellectual Property and deployment issues, and Software Engineering.

Electric field on CV-22 calculated by SENTRi
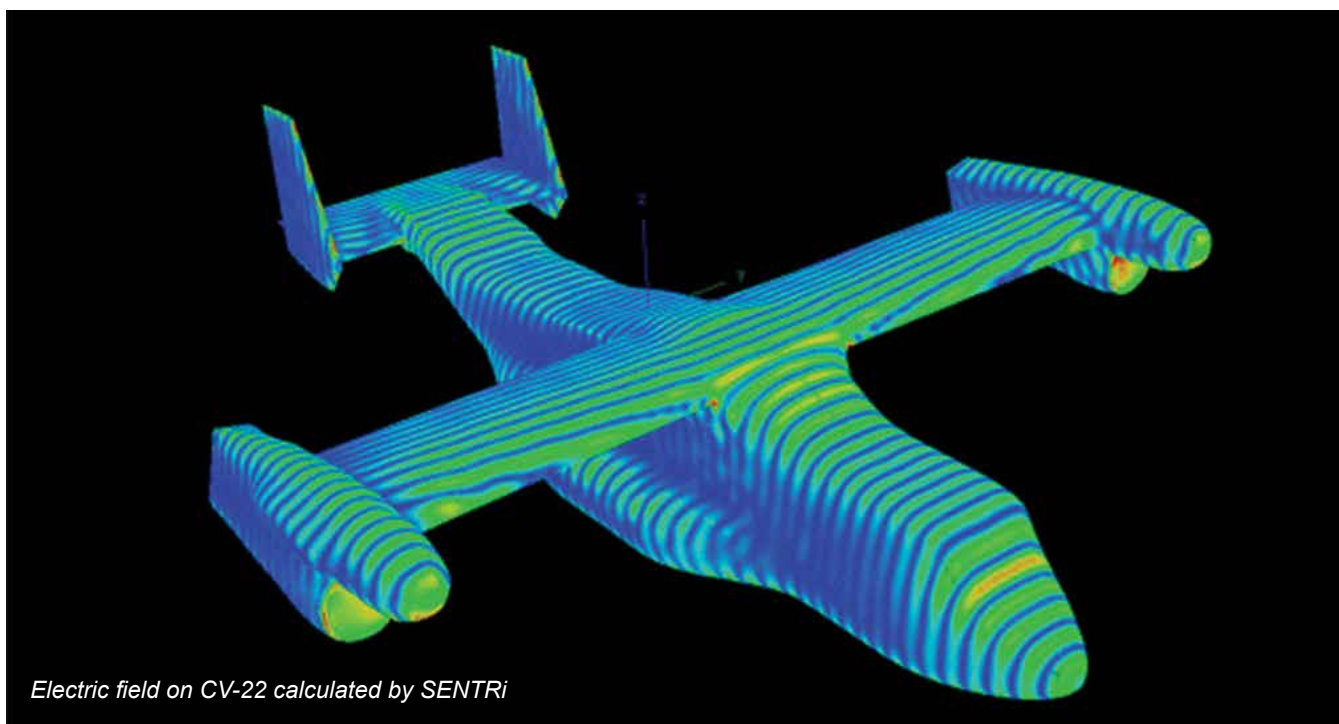
HPC Tools Clear the Path for Unmanned Air Vehicles

By Theresa Shafer, Naval Air Systems Command (NAVAIR 4.3.2.1), Applied Aerodynamic and Store Separation Branch

Problem

To respond to requests for proposal (RFPs) for unmanned aircraft vehicle (UAV) systems, companies need to demonstrate the performance of their proposed UAV systems. Also, Navy aircraft require air worthiness certification, which normally requires data from wind tunnel tests, flight tests, or empirical methods. Traditional aircraft design programs use computational fluid dynamics (CFD) analysis, wind tunnel testing, and flight testing to obtain flight clearances. However, small UAV development programs do not have the resources to obtain extensive databases needed to support full flight clearance decisions.

How

Shadow Ops used CFD to create a database that can be used in flight clearance decisions, which enabled small UAV programs to obtain flight clearance certifications. The CREATE-Air Vehicles (AV) created a standard, streamlined, validated process using enhanced methods to support the small UAV flight clearances using computational methods, which was applied to the small UAVs Exdrone and Aerostar. For the Exdrone UAV, the computationally-based engineering (CBE) model was created from 2D drawings, and six configurations/grids and 205 data points were generated and completed. In 2009, an original aerodynamic database for Aerostar was created. The engineers created the 3D CBE model by using a coordinate measuring machine and a flight vehicle. Nine configurations/grids were generated, and 176 data points were created.

Impact

The data generated through computational methods as part of the Aerostar UAV effort helped to mitigate risk in the flight certification process by providing engineers with sound data with which to make more informed air worthiness assessments. The small UAV program office (PMA-263) has since relied on this process to gain clearances for the Raven, Aerostar, and Exdrone vehicles. Databases in support of clearances for the SCANEAGLE and STUAS vehicles are currently in progress using the same process. Unbiased aerodynamic performance data were provided to the PMA-263 for aerodynamic assessment of vehicles enabling flight envelope definitions based on accurate, engineering data. Flight clearances have been defined based on the data created by the CREATE-AV program.

AeroVironment RQ-11 Raven is a small hand-launched remote-controlled unmanned aerial vehicle (or SUAV).

PMA-263 is NAVAIR's program office for Navy and Marine Corps Small Tactical Unmanned Aircraft System.

Using the established process of CFD analysis, differences in forces and moments on the Aerostar vehicle because of the presence of the antennas was calculated in a timely manner and communicated with the performance and flight dynamic subject matter experts. The data directly affected the flight clearance decision allowing these subject matter experts to make a decision on the flight clearance for Aerostar with these antennas. The same approach has been used two to three times a year for Aerostar with geometric modifications and was used on the Raven vehicle to provide decision makers with a flight dynamic and performance evaluation capability throughout its life cycle that was not possible prior to using these CFD tools.

Additional impacts include the small UAV flight clearance support. CFD was used to assess the aerodynamic effects of adding two minipods to Aerostar in two loading configurations. More than 90 runs were completed, and data and recommendations were supplied to the technical area experts in four days. CFD was also used to examine the aerodynamic effects of adding four antennas in various configurations. Sixty runs were completed; and data was supplied to technical area experts in less than 1 week, which was used to directly support the Greenbox flight clearance.

Metrics

There was an improvement in time and cost from the Exdrone to the Aerostar UAV:

• Exdrone UAV: 4 months to complete and cost about $80k, much less than the cost of a full wind tunnel test.

• Aerostar UAV: 6 weeks to complete and cost about $30k, reducing the time and cost because of the lessons learned from the Exdrone assessment.

Flight certifications for small UAVs have a tight turn-around time (2 to 6 weeks). For a single UAV, the ShadowOps team completed about 200 separate runs that use 128 to 512 processors each. In 2009, about 1 million computational hours were used for this project. In 2010, 1.8 million computational hours were used. For the Aerostar Greenbox clearance (adding antennas to Aerostar), 4 grids, 60 computational solutions, and postprocessed data were completed in 1 week. Completing this many runs in such a short time could not have been accomplished without HPCMP resources.

Who Did the Work

The program office (PMA-263) that supports the CREATE-AV ShadowOps has continued to fund the effort for small UAV clearance support since the Exdrone and Aerostar projects were completed. Dr. Theresa Shafer, a NAVAIR 4.3 aerospace engineer on loan, led the UAV effort.

What Tools Were Used

The Kestrel high-fidelity modeling tool for fixed-wing aircraft CFD, structural mechanics, propulsion, and control was used. The systems that were used to process the data were the following HPCMP resources: *Diamond*, *Harold*, *Mana*, *Hawk*, *Falcon*, *MJM*, *Raptor*, and *Garnet*.

What Is CREATE: A New Approach for Reducing Weapon System Acquisition Time and Cost

As a part of the HPCMP, the Computational Research and Engineering Acquisition Tools and Environments (CREATE) program is developing sophisticated multiphysics-based computational engineering design and analysis software tools. This brings the fruits of many years of research to bear on current and future engineering challenges. Used with the HPCMP high performance computers, these tools can identify and help eliminate design defects and integration problems much earlier in weapon systems design and test processes, before major schedule and budget commitments are made, resulting in reduced acquisition time and cost. CREATE addresses the dual need for (1) development and optimization of integrated conceptual

Other Impacted Systems

The methodology was applied to obtain flight certification data for the following small UAV systems being assessed by NAVAIR:

Aerostar

Exdrone

Raven

Scaneagle

STUAS
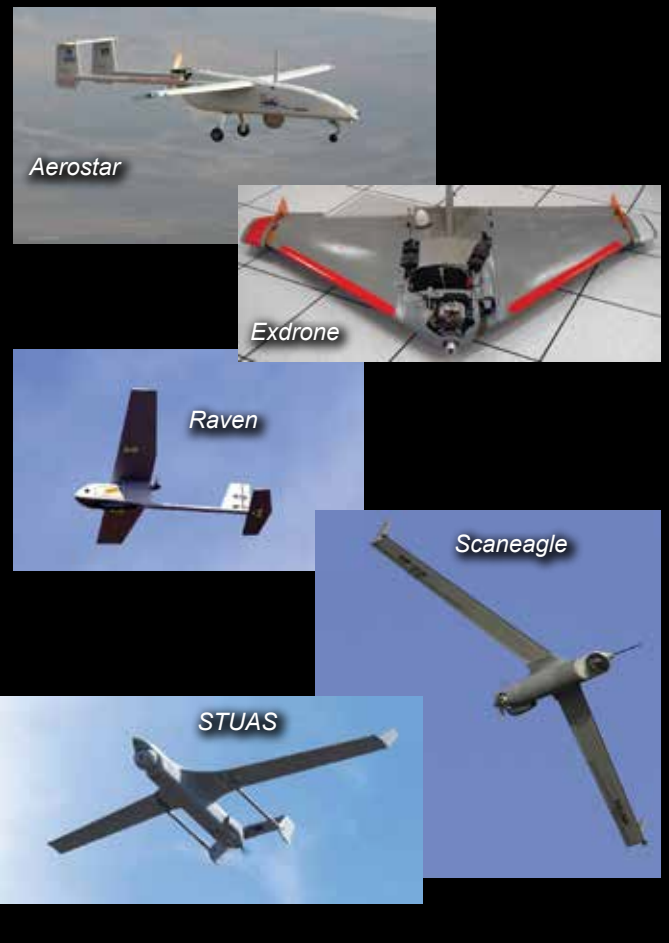
designs and (2) high-fidelity, detailed design analysis of DoD aircraft, ship, and radio-frequency antenna weapon systems. Engineers can use the tools from the initial stages of acquisition through sustainment. CREATE tools allow engineers to rely more heavily on computational prototypes that can be constructed less expensively and much earlier in the acquisition process than experimental prototypes, well before metal is cut. Testing can then focus on validating mature designs and on investigating the basic physical phenomena that determine system performance. The CREATE tools are being developed and deployed by teams hosted by Service research and development, engineering, and acquisition organizations (e.g., SEEK EAGLE Eglin Air Force Base and Naval Surface Warfare Center Carderock). While funded primarily by the HPCMP, the Services are investing time and energy introducing the new tools and processes into individual projects and planning for long-term support. The tools are already being used by over 50 DoD programs.

Building Large-Scale Software in the DoD

By Maria Murphy, MHPCC DSRC Sr. Applications Engineer, and Chris Atwood, CREATE, High Performance Computing Modernization Program Office

A few years ago, the CREATE development team was looking for the collaborative functions of GitHub, SourceForge, Forge.mil, or Google Code—but with secure authentication and access to thousands of cores. The extended team now includes several dozen developers and hundreds of users spanning five time zones. More formally, several studies note that while collaboration is not the sole independent variable determining the outcome of projects, it is crucial in addressing the key project factors: scoping mission, retaining stakeholder support, capturing requirements, scheduling, engaging clients, completing technical tasks, monitoring and controlling projects, troubleshooting and ultimately the client accepting the resulting software product (Atwood *et al.* 2010).

Software development teams have evolved to this highly-collaborative development model over the past few years. This evolution is exemplified by the growth in open-source efforts, which have resulted in a wide range of applications, including Linux, Apache, MySQL, and Android. These community efforts have driven the widespread adoption of tools and practices that enable collaboration with distributed members, each member potentially originating from differing organizations, geographies, or backgrounds. The distributed nature of these software product efforts is similar to the distributed Computational Research and Engineering Acquisition Tools and Environments (CREATE) software program (Post 2008). However, a significant complicating factor for DoD code arises from the security policies required for the development and deployment of International Traffic in Arms Regulations (ITAR) materials, which address issues of national interest. These necessary security elements include authentication, authorization, network access controls, server hosting constraints, and user client system privileges, each of which may vary by institution, management, and physical site.

This article summarizes the effort to support successful project practices while minimizing interruption to developers, which has resulted in the implementation of collaboration systems configured at http://create.hpc.mil to transparently facilitate, document, and institutionalize (1) communication among all the key roles in the project and (2) support and feature requests from the user community and from stakeholders from DoD Programs.

CREATE Program

The CREATE program is designed to improve the DoD acquisition process (AP) by developing and deploying three sets of advanced computational engineering design tools. Acquisition programs use these tools for the design of military air vehicles, ships, and RF-antennas. The military air vehicle (AV) design project will develop a design optimization tool that can simulate unsteady, separated flow, initially for individual aircraft components and ultimately for an entire aircraft. The military ship (SH) design project will develop and deploy accurate, physics-based models, for Navy vessels, that address three key capabilities: ship-shock response, hydrodynamics, and design space exploration for concept design. The RF-antenna design project will build an efficient electromagnetic design code that incorporates modern physics and computational algorithms for high performance computers. Since all three of these projects depend on discretization, a fourth project has been refactored, one focused on modeling, meshing and geometry (MG). This new generation of computational design tools will enable acquisition system engineers to rapidly produce optimized designs for complete systems and to make better design decisions than has been possible with prior design tools.

CREATE Community Services – Governance Model

In support of these software project efforts, a CREATE Collaboration Council (CCC) has been formed whose members include each of the CREATE Project Managers, members of the HPCMPO staff representing the DoD Supercomputing Resource Centers (DSRCs), and the Defense Research and Engineering Network (DREN). The CCC provides recommendations, with risks and mitigations, to the CREATE Program Office (PO) for Go/No-Go decisions. The CCC mission is to impact DoD Acquisition Programs by recommending collaboration tools and program processes for the benefit of CREATE software developers and of the AP engineering user community. Pursuant to the Program software engineering practices, the Council aims to provide easily-accessible, productive, and secure means of sharing practices across the software lifecycle; including code configuration, unit testing, continuous integration methodologies, techniques, Application Programming Interfaces (API) sharing, issues and defects tracking, software manuals,

video tutorials, training material to the AP workforce, and software support. From the recommendations of the CCC and with the approval of the PO, a collection of services has been hosted to support the application development lifecycle.
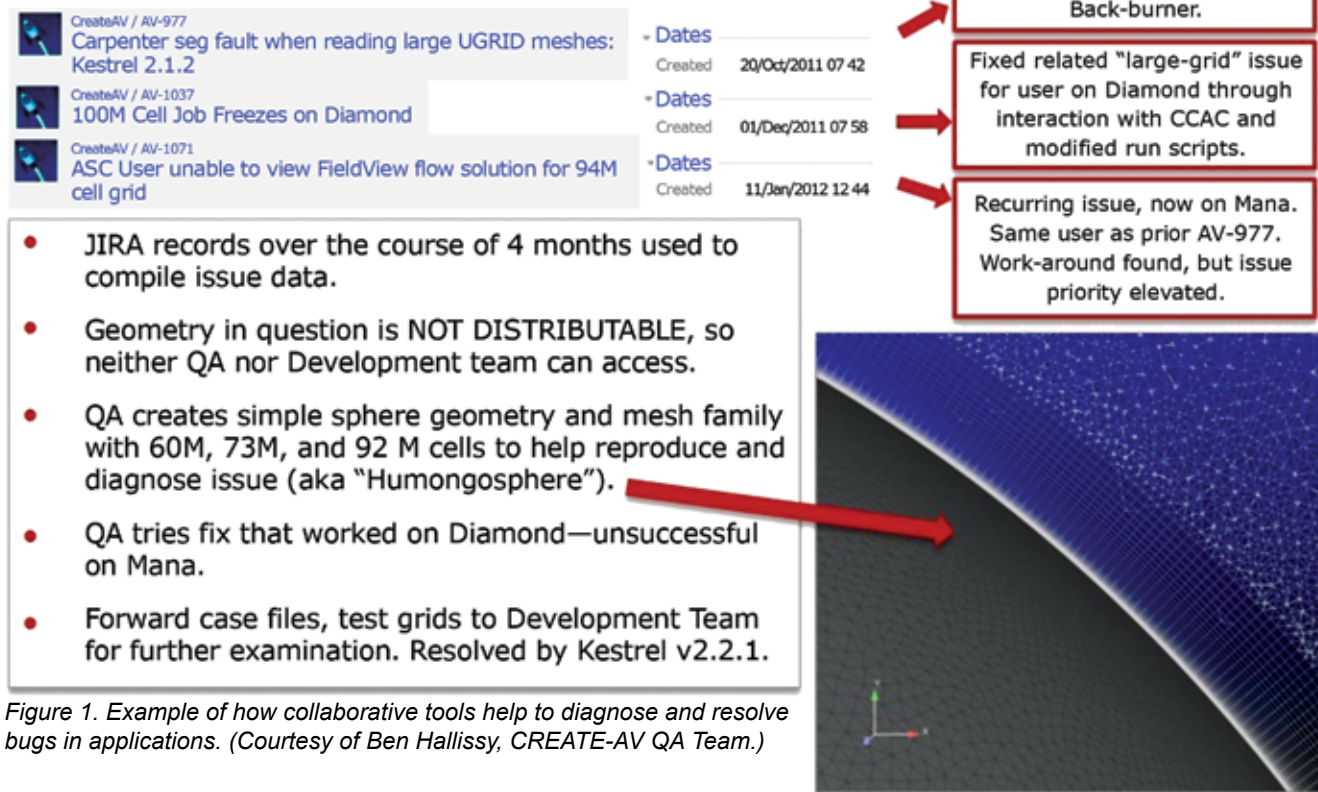
Application Development Lifecycle Services

Code Repository

One key difference between a research project and a sustainable application is the ability of multiple developers to edit, build, test, and debug via a shared code repository. Software version control enables collaborating on a specific version of the code or tracking the revision history for debugging and code release purposes. The key requirements for the configuration management software include (1) developer ease-of-use with multiple languages, (2) wide use in commercial or open-source projects to enable fast support as well as longevity of the software, and (3) integration with issue-tracking software for CREATE product sustaining use. For the current iteration, the development teams chose Subversion (SVN) because of its widespread use in both commercial and open-source projects, enabling rapid support and relative assurance of continued development with limited cost-impact. In addition, SVN is extensible via plug-ins available for many developer environments and languages and is accessible with GIT clients as well as with continuous integration and issue tracking systems.

Issue Tracking

A second key element of project management is tracking issue status in a way that incurs minimal interruption in the primary deliverables of the team. Issue-tracking systems are crucial for recording defects, tasks, requirements, and feature requests in a manner that sets expectations for due dates and action item ownership. Tracking systems also passively build a knowledge repository for the project, providing role-based context for project issues, in contrast to mail, which is associated with individual personnel at that particular time. In addition, issue tracking can be used to establish a feature list for future releases, as well as link to dependent issues, a code repository, and the user community forum. With such systems, metrics, such as the amount of code addition, defect time-to-fix, and feature time-to-implement, can be queried without interrupting the development, testing, or customer teams. Hence, this tool is useful for evaluating the health of the product from a full-lifecycle perspective, from concept, through development, to customer feedback.

Atlassian JIRA was chosen as the issue tracker based on its SVN integration, source-code availability, visibility scoping, robust database for backup and queries, packaging for hosting in a secured domain,

Figure 1. Example of how collaborative tools help to diagnose and resolve bugs in applications. (Courtesy of Ben Hallissy, CREATE-AV QA Team.)

and issue submission and modification capabilities via web or email. Current CREATE use of JIRA contains approximately 6000 assigned issues across 12 projects, which are in states ranging from open to resolved to closed. Figure 1 depicts a scenario in which the issue tracker was used to compile data on a recurring problem involving a large grid. Successive iterations between the CREATE-AV Quality Assurance (QA) and development teams and the users helped to enable a bug fix in the next release.

Discussion Forum and Wiki

In modern software development efforts, it is common for users to form a Community of Practice (COP) for self-help and for engagement with the development team in a scalable manner. Our DoD community is no different in this respect, with a goal of the CREATE project to enable DoD and defense industry engineers access to the integrated capability of HPC software and hardware, allowing them to solve problems that cannot be solved by other means. The purpose of the CREATE web-based user community discussion forum is to scale the knowledge base of

the CREATE Community via a searchable archive, which organically grows with each version of the codes released. The forum mechanism offers the advantage of passively capturing project institutional knowledge, information which might otherwise be lost in person-to-person mail and, hence, inaccessible to benefit future community members. In addition, development team experts moderate the forum, providing responses that are generally useful to the entire community and reducing the burden of duplicate responses to similar issues that users experience. As the COP grows, the forum enables the user community to help each other so that the development team does not have to respond individually to each question that arises in the use of the code. Although there are many forum software applications available, the forum software phpBB was selected based on the requirements of configurability for separable access by users and developers, healthy community support, search functionality, availability of integration with issue tracker and wiki, daily digests, low initial and recurring costs, and server maintenance ease. The currently implemented discussion forum
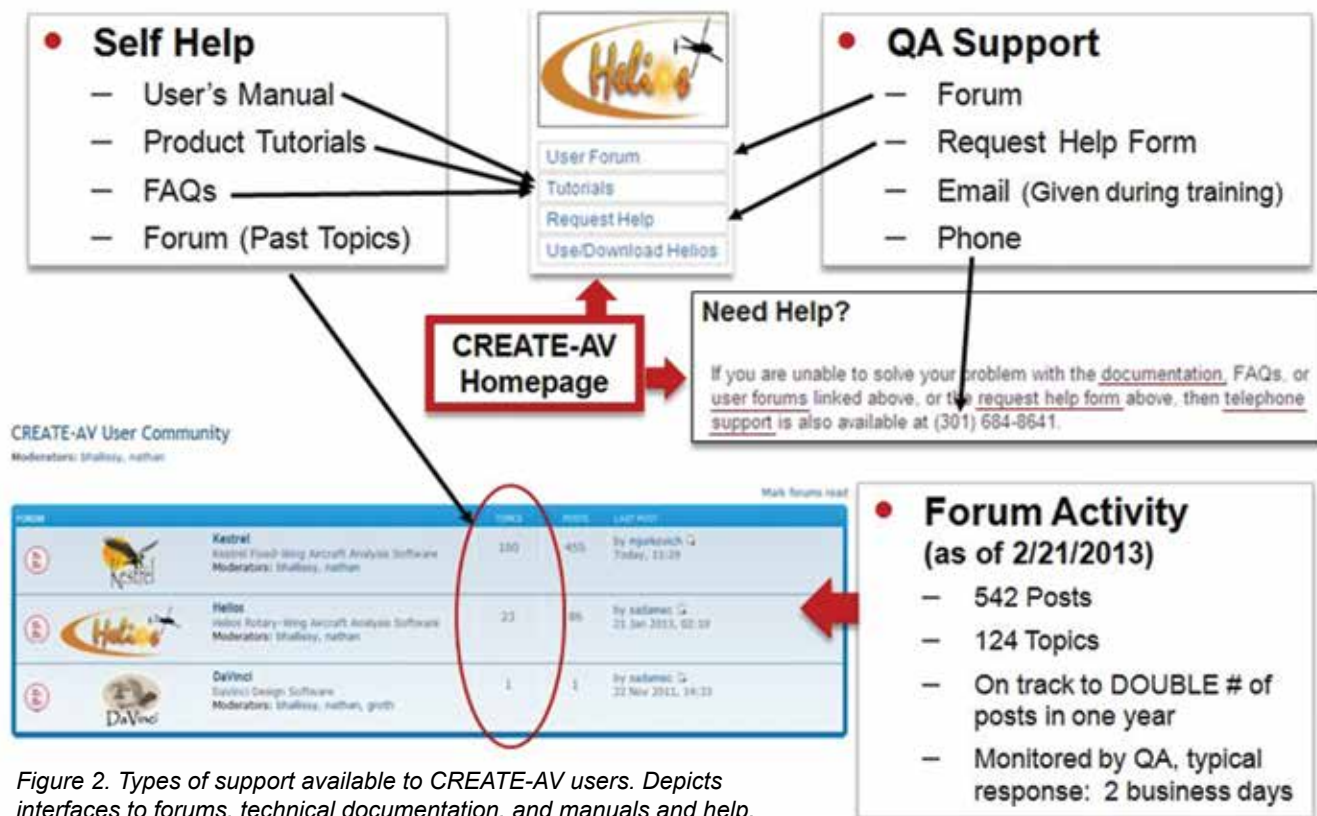

Figure 2. Types of support available to CREATE-AV users. Depicts interfaces to forums, technical documentation, and manuals and help.
(Courtesy of Ben Hallissy, CREATE-AV QA Team.)

contains approximately 3200 posts across 650 threads, primarily oriented at troubleshooting discussions within the development team members. The threads of discussions may culminate in an action item, for which a cross-posting link is generally provided to the issue tracker or wiki.

Separately from the forum, a wiki functions as a structured space for users to access tutorials, to download software releases, and to submit formal help requests. Confluence offers delegated access control mechanism and versioning features, which allow authors to manage distribution of their content (e.g., meeting minutes, project documents, and technical references) as well as to maintain a history of modifications. Figure 2 summarizes the types of support available from the forum and wiki collaboration services to users of one of the CREATE software products.

Server and Network Architecture

The CREATE Community server applications, consisting of issue tracking, code management, document repositories, and forum and wiki services, are hosted individually on separate web servers. This configuration minimizes interruptions to all other services should one server require maintenance or upgrades, and provides protection against application dependency and port conflicts. Application servers are deployed on virtual machines residing on two physical servers and are managed by KVM software. All systems run Red Hat Enterprise Linux 5 64-bit distribution.

A majority of the CREATE Community web server applications consists of components commonly found in the LAMP (Linux, Apache, MySQL, PHP/ Perl) architecture. All Apache web applications are configured securely and provide HTTPS encryption. The CREATE servers exist on a separate subnet within DREN, and a combination of information security technologies protect them. User management is based on a system that can quickly grant or revoke specific privileges, with service access based on roles assigned to an individual user. In this manner, CREATE can accommodate various access levels required for both developers and end-users without incurring significant administrative overhead.

All services are accessed by the user via the CREATE Community site, http://create.hpc.mil, which serves as the mechanism for licensing, single sign-on (SSO) authentication, and access control to available applications. The portal interface provides a consistent and intuitive method for accessing common services, as well as a customized content area, for each design project.

DSRC Application Deploy and Buildmastering

The AP subject matter expert community is generally quite busy. Therefore, it may benefit from a quick means to evaluate the CREATE tools and to then begin using them if the tools are deemed useful. Hence, it is beneficial to have CREATE software pre-installed, pre-configured, and verified on their systems. The CREATE Community team has engaged with the HPCMP Baseline Configuration Team to enable this capability on the DSRC systems (Graham and Kevorkian 2013). This enables users to avoid the tedious and error-prone process of building, installing, configuring, testing, and administering the software prior to their productive duties of using the software for their AP task. Once the engineer is authorized to access the ITAR material, use of the software is straightforward, typically starting with the use of the tutorial sample test cases provided at the CREATE Home locations on the DSRC systems.

Although the CREATE teams generally provide their software on the DSRC systems, they must build, test, and execute the CREATE software on a range of systems to satisfy the user requirements for access to the software on local workgroup clusters as well as on the DSRC systems. Since the CREATE-AV project will require that released versions of the code are supported on the then-current interpreter, development compilers, and runtime libraries, builds of the software are required for the targeted user-base. The Buildmaster performs these builds as part of the release engineering process of porting the CREATE-AV codes from development platform to the release-supported platforms, including DSRC systems, workgroup clusters, client workstations, and connecting networks. The Buildmaster collaborates closely with development teams on identification, scoping, and tracking of issues found during the build process and also implements automated build reporting, including performance measures across the supported platforms and identifying opportunities for refactoring common runtime libraries for CREATE-AV codes. In addition, the Buildmaster influences the DSRCs on implementation of common software configurations with the goal of improving the acquisition engineering usage of the software.

Software Integration and Test

To integrate components from disparate teams into one application, for example CREATE-RF SENTRi integration of CREATE-MG Capstone SDK, the combinatorics of components, build tools, operating systems, and libraries become difficult to manage. To enable this type of integration, test systems for CREATE software and hardware must support multiple

product platforms (Windows, OSX, Linux) and their variants (e.g., Vista and 7 for Windows, CentOS and other Linux distributions). Providing permanent and continuous access to each supported OS platform is critical for product testing, robustness, and stability. As the combination of supported platforms increases, the ability to purchase sufficient hardware becomes challenging. To fulfill these requirements, the most cost-effective method of addressing platform availability is virtualization. A single virtual machine could host each required platform and handle the build, test, and integration activities for that operating system. Multiple projects could use any virtual machine (VM), which reduces duplication of commonly supported platforms. The Common Build Server system will deploy upwards of 12 platforms running on individual VMs.

HPC Portal

The user-friendly HPC Portal, http://www.portal.hpc.mil, which allows users with varying levels of HPC experience to process jobs using nothing more than a web browser, provides access to the CREATE community tools from its interface. This integration offers users the convenience of access to tutorials, examples, documentation, and help without workflow interruptions. For example, a user who has submitted a job via the Portal has an issue with divergence because of the input mesh files. Without switching context, he can refer to the CREATE forum to research the problem or to submit a formal help request via the wiki.

Software Adoption and User Analytics

To successfully transition the CREATE technology to the stakeholder community, establishing familiarity and confidence in the successful use of the software tools by the users is key. A measure of success for the CREATE software is the adoption and use of the application. Ultimately, an acquisition program that successfully leverages the software is a good measure of value, with the AP engaging CREATE via its engineering members in the community. This engagement can be measured via forum activity, support requests, or software use. To gain better quantification of the CREATE user base, the team has begun gathering user analytics. Because of security concerns, accumulating site data statistics requires a self-hosted analytics solution. Piwik is an open-source alternative to software-as-a-service (SaaS) solutions, such as Google Analytics, and provides complete privacy and control of the data that is collected. Piwik has been integrated with CREATE wiki and forums to compile web traffic data (such as visitors' geolocation, sources, activity, and technical capabilities) and to generate user engagement reports.

Conclusion

The collaborative community software services that support the CREATE efforts have, in part, led to successful adoption of HPC computational physics software in several Air Force, Navy, and Army Programs. The community services have enabled feedback directly from the stakeholder and users to the CREATE software teams, creating quick response capability. Over the past few years, the community site has grown to 447 authorized users, 160 of whom have accessed the site in February 2013. Ultimately, the value of these services will be measured by impact to DoD programs and their engineering users, users whom we are engaging with the CREATE Community site.

References

Atwood, C. A., S. A. Adamec, M.D Murphy, D. E. Post, and L. Handley-Blair. 2010. Collaborative Software Development of Scalable DoD Computational Engineering. In *2010 DoD High Performance Computing Modernization Program Users Group Conference (HPCMP UGC), 14–17 June, Schaumburg, IL.*

Graham, J. and A. K. Kevorkian. 2013. *HPCMP DSRC Baseline Configuration*. http://www.ccac.hpc.mil/consolidated/bc/team.php.

Post, D. E. for the CREATE Team. 2008. A new DoD initiative: the Computational Research and Engineering Acquisition Tools and Environments (CREATE) Program. *Journal of Physics: Conference Series* 125(1).

Numerical Flow Analysis (NFA) Simulation of a Planing Boat in Waves

By Thomas O'Shea, Dr. Douglas Dommermuth, Dr. Kyle Brucker, Lucas Rhymes, and Donald Wyatt, NFA Simulations, NFA Simulations, SAIC; Dr. Michael Stephens, Chris Lewis, Richard Walters, Miguel Valenciano, and Michael Wissmann, Flow Visualization, Data Analysis and Assessment Center (DAAC); and John Levesque, Parallel Computing Guidance, SAIC

HPC Resources: Cray XE6 (Raptor), AFRL DSRC; Cray XE6 (Garnet), ERDC DSRC

A recent collaboration between DAAC; US Army Engineer and Research and Development Center; and the NFA simulation team, SAIC, has allowed high-fidelity numerical simulations of a planing craft in a seaway to be realistically rendered. High-quality visualizations allow researchers to gain insight into the complex multiscale physics involved.

Motivation

Personnel aboard high-speed planing boats can experience substantial accelerations in a seaway. They can be injured when the shock loads that cause these accelerations are extreme. These shock loads also wreak havoc on the boat's structural members. Time spent on heavy seas can significantly shorten a planing boat's overall operational lifetime.

Quantifying the extreme pressures and forces on a planing hull can help guide the design process but can also be difficult to accomplish. Computational Fluid Dynamics (CFD) offers one path, but the problem is especially challenging. Pressures on a planing boat hull are concentrated at the spray root and can move significantly with only slight changes in vessel orientation. This makes the force-mass balance sensitive to any errors in the simulation. High spatial resolution is needed to capture the short extent of the spray root and also any small features on the hull (steps, chines, etc.). The addition of a seaway necessitates high temporal resolution to capture wave slamming events, which can happen over extremely short periods of time.

Leveraging Department of Defense (DoD) high performance computing (HPC) resources, simulations were completed with the high spatial and temporal resolution necessary for accurate results. These resources are used efficiently because of recent advances in NFA input, output, and core algorithms. Guiding these improvements were scaling studies performed across multiple HPC resources. Insights gained should make even larger, more sophisticated calculations possible.

Approach

NFA solves the Navier-Stokes equations utilizing a Cartesian grid cut-cell formulation with a second-order accurate, volume-of-fluid (VOF) interface-capturing technique to model the unsteady flow of air and water around moving bodies. The Cartesian grid cut-cell method avoids a complex body-fitted grid generation process, which enables the rapid analysis of new hull geometries. Turbulence is modeled with an implicit subgrid-scale model that is built into the treatment of the convective terms in the momentum equations (Rottman *et al.* 2010). A domain decomposition method is used to distribute portions of the solution domain over a large number of processors. Detailed descriptions of the numerical algorithm and of its implementation, using Fortran 2003 and the Message Passing Interface (MPI), on distributed memory HPC platforms are provided in Dommermuth et al. (2007), O'Shea et al. (2008), and Brucker et al. (2011).

Planing Craft Simulation

Simulations of a high-speed planing boat accelerating from rest to a constant forward speed in head seas were performed using NFA. The boat is permitted to have four degrees-of-freedom, including sway, heave, roll, and pitch. Integration of the pressure on the wetted surface provides a means to understand the shock loads the crew will experience. Analysis of these loads will facilitate the design of vessels that will provide a smoother ride without adversely affecting craft performance. Visualization of the locations and spatial extent of the high pressure regions on the hull aids in the analysis of the fatigue life of structural members.

The planing boat in this simulation has a length, *L*, of 62.3 ft (19 m) and was traveling at a velocity, *U*, of 40 knots (20.6 m/s), equating to a Froude number of 1.5, defined as $Fr = U/\sqrt{gL}$ where *g* is the acceleration of gravity. The domain has a length, width, and depth of 4.3 (81.7 m), 2.0 (38 m), and 1.8 (34.75 m) boat lengths, respectively. The number of cells in x, y, and z was 1536, 1024, and 512, respectively, resulting in 805 million cells in the total simulation. Spacing near the body was 0.0018 *L*, or 1.3 in. (3.4 cm), necessitating a nondimensional time-step of 6.24×10^{-4} (6.77×10^{-4} seconds.) The simulation was run for 70,000 time-steps, or 44 body lengths, which corresponds to about 30 wave impacts.

The research was completed on the Cray® (Cray Inc.) XE6 located at the Air Force Research Laboratory (AFRL) DoD Supercomputing Resource Center
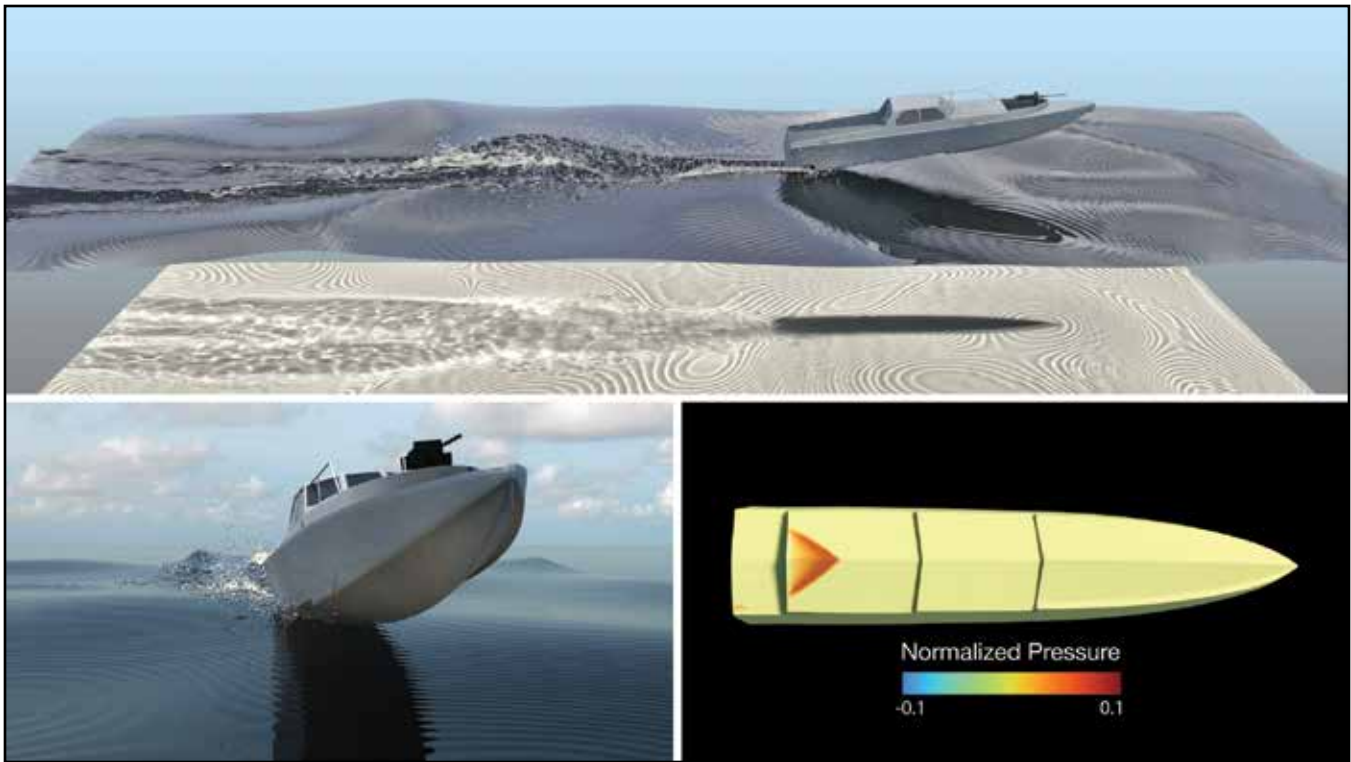
Figure 1. Planing boat motion off crest of a wave and resultant surface pressures on hull.
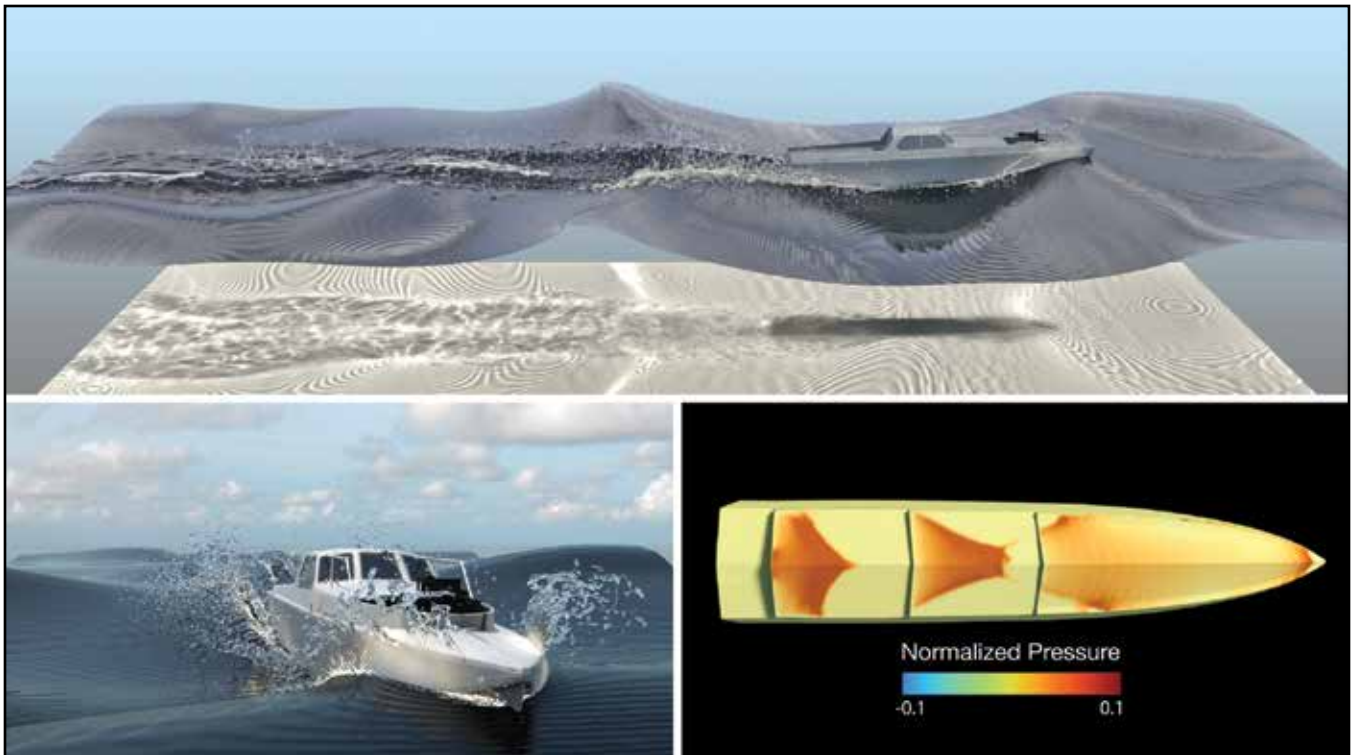
Figure 2. Planing boat motion slamming into waves and resultant surface pressures on hull.

(DSRC). The 805 million grid cells of this simulation were distributed over 3072 cores and run for 190 wall-clock hours.

Figure 1 illustrates the planing vessel launching off a wave. The highest pressures occur along the spray root, evident as the sharp interface between minimal pressure, yellow on the color bar, and high pressure, highlighted in red. At this instant, the entire craft is supported by the small wetted surface area with high pressure. Within a few seconds in a seaway, the portion

of the boat in contact with the free surface can either be a small area at the stern, as in Figure 1, or can be distributed over the extent of the craft when the boat slams back down between waves, as shown in Figure 2, or various alternate configurations. Slamming of the craft back down into a wave produces the greatest amount of wave breaking and spray generation. Throughout the simulation, waves break at the rooster tail behind the boat. Wave breaking also occurs along the edges of the Kelvin wake. Ambient waves form spilling breakers toward the back of the computational domain. Videos of this simulation, as well as those of other NFA simulations, can be seen at http://www.youtube.com/WaveAnimations.

Visualization

Isosurfaces output during the simulation were provided to the Data Analysis and Assessment Center (DAAC) in two sets of Visualization Toolkit (VTK) geometry files. The first set represented the free water surface and had on the order of 5.5 million triangles per isosurface. The second set represented the boat's geometry and was provided with a scalar pressure value per node.

From these two sets of data files, the DAAC produced three types of animations: several realistic perspective views, a bottom view showing normalized pressures experienced by the boat hull, and a caustics study of light through the turbulent wake.

The realistic animations were accomplished by ray tracing the scene with reflection and transparency attributes applied to the water using Adobe 3D Studio Max commercial rendering software. A special rendering model of the boat was developed by applying rich texture maps to one of the simulation boat isosurfaces. Each frame was rendered via 3D Studio Max scripts that performed the following sequence:

```
for each frame {
        import texture mapped boat file;
        import next simulation boat file;
        import next simulation water surface file;
        orient texture mapped boat to the
           simulation boat position;
        delete simulation boat from scene, (leaving
           only the correctly oriented textured
           boat);
        render the scene;
}
```

The bottom view was shown simultaneously with the realistic view. Dramatic slamming of the boat from the realistic views could be quantified by the color mapped pressure displayed on the boat hull, as shown in Figures 1 and 2.

Finally, a visualization employing caustics was produced by the DAAC team. Caustics are the collection of light rays that are reflected or refracted through a curved surface. Light rays were ray cast through the wake and projected onto a surface located some distance under the bottom of the boat. This visualization technique results in light and dark bands on the bottom surface that provides more information about the structure of the turbulence of the wake.

Assessment of Numerical Predictions

Past efforts to verify the use of NFA for the prediction of the hydrodynamic forces and moments associated with planing craft, through detailed comparison with available experimental data, are described in O'Shea et al. (2012), Fu et al. (2012), and Brucker et al. (2012).

Under the direction of Dr. Craig Merrill at the Naval Surface Warfare Center, Carderock Division (NSWCCD), NFA was delivered and installed at the Ship Engineering and Analysis Technology (SEATech) Center in January 2012, where it is currently being used to research stepped planing boat design. As part of that process, predictions made using NFA are being compared with ongoing towing tank experiments.

One example of this continual experimental verification includes comparison with a forced-roll experiment of a prismatic planing hull conducted by the United States Naval Academy (USNA), described in Judge (2012). This assessment illustrates the capability of NFA to accurately model the complex multiphase flow associated with high Froude number flows associated with planing craft. Experimental underwater photographs were compared with NFA simulations of a planing hull, with steady forward speed at fixed sinkage and trim, at a series of roll angles.

This simulation was run with a width of two boat lengths, or 10 ft (3 m), and a depth of one boat length, or 5 ft (1.5 m), corresponding to a smaller domain than the towing tank dimensions in order to cluster cells near the body without stretching the grid too heavily. The number of cells in x, y, and z was 1280, 896, and 448, respectively, resulting in 514 million cells in the total simulation. Spacing near the body was 0.0009 L or 0.054 in., necessitating a nondimensional time-step of 0.00023. Simulations were run for 13,000 time-steps, or three body lengths that were sufficient to reach steady state, on *Garnet*, a Cray® (Cray Inc.) XE6 at the U.S. Army Engineer Research and Development Center (ERDC) DSRC, and took approximately 24 hours to complete using 576 cores.

The generation of the spray sheet is correctly predicted; the spray edge matches well to experiments.

As the model rolls over, the flow over the chine initially separates cleanly; but at 20 degrees and higher, the flow wraps around the chine and wets the side of the model. The NFA simulations exhibit this behavior. Figure 3 shows comparisons with experimental photos taken at 10 degrees of roll. The spray root line, and thus wetted surface area, are demonstrated to agree. NFA appears to model the instabilities in the spray sheet breakup reasonably well. The bottom of the model was painted with a ruler, and the length along the keel (L_k) and along both the port and starboard chines (L_c) from the transom to the spray root was estimated from the underwater pictures by the experimentalist. Figure 4 presents the comparison of the same values extracted from the NFA results. The resulting agreement is excellent throughout the range of roll angles, indicating that the wetted surface area is being accurately predicted.

The parallel embedded isosurface output capability allows data to be extracted efficiently, facilitating much higher temporal resolution for quantities of interest while also massively reducing the amount of data output and postprocessing time. Strain on the mass storage and archival storage systems is therefore reduced while the quality of both the visualizations and scientific/engineering data is improved.

Isosurfaces for the planing boat simulations were output every 5 time-steps, corresponding to a dimensional sampling frequency of 295 Hz; and the entire data field (three velocity components, pressure, and volume fraction) was output every 20 time-steps (75 Hz) to permit analysis of turbulence, air entrainment, and spray formation. Each save of the entire data field requires 6.2 GB of storage, totaling 21.7 TB for the entire run. Extracting isosurfaces
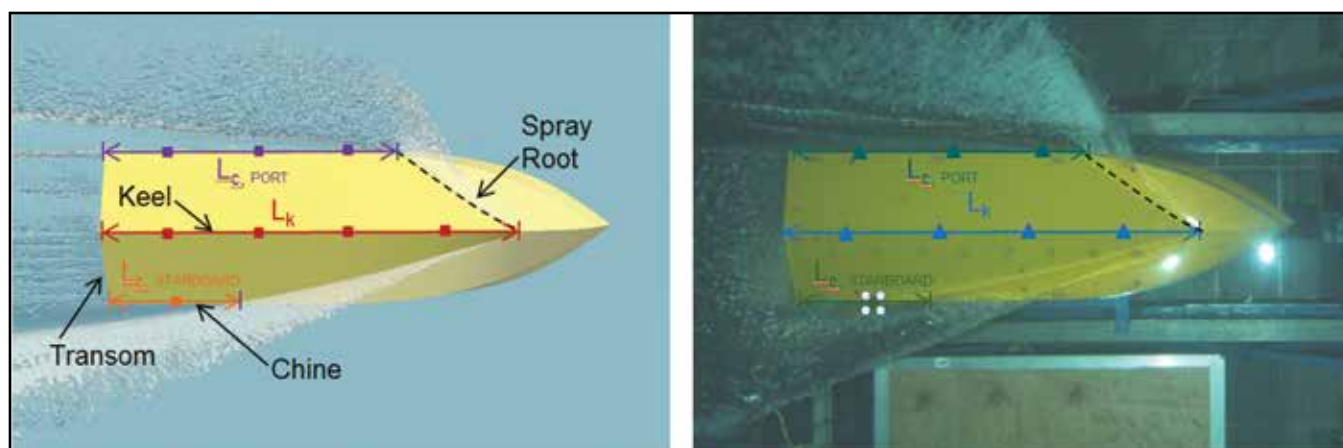

Figure 3. Planing boat with 10 degrees of roll, with NFA simulation results on left and experimental photograph on right. View is from underneath craft. Line colors and marker shapes correspond to quantitative values shown in Figure 4.

HPC Considerations

Dynamic Output Control

The flow visualization of large datasets, including rendering of isosurfaces and volumetric data, is difficult. NFA has developed many output tools that have proven to be invaluable in the analysis of large CFD datasets. Variables in the data can be extracted in the form of lines, planes, isosurfaces, or full fields. The ability to change these output types, values, and locations on the fly, as a simulation is running, is a valuable innovation that allows rapid isolation of important features in the data, allowing more efficient research. Data storage requirements, queue time, computer hours, and calendar time are all saved by being able to redirect focus and to change data analysis and visualization techniques dynamically as desired.
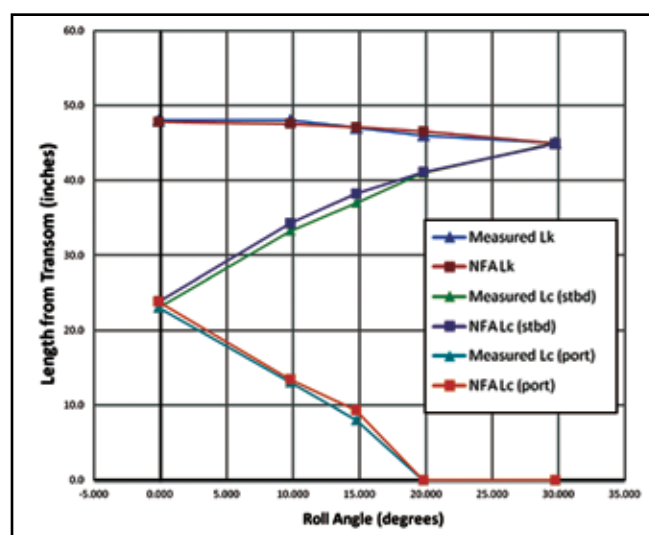

Figure 4. Comparison of spray sheet formation and contact line measurements.

reduces this by a factor of 92, allowing for the sampling frequency of the extraction to be increased by a factor of 4 while only requiring 3.8 TB of storage for the run. All output files are written in binary format, to reduce size, in a format requiring no postprocessing. Output files can be directly opened with visualization software based on the Visualization Tool Kit (VTK), such as the freely available Department of Energy (DOE) funded ParaViewTM (Kit, Inc.) software.

User Efficiency

Identical NFA code, analysis tools, and infrastructure can be run on desktops running Windows, Linux, or Mac-OSX and on many HPC platforms including Cray® (Cray, Inc.) XT3, XT4, XT5, and XE6, SGI Altix, or IBM P5, and P6, and iDataPlex. Small runs (10+ million grid points) are completed in minutes, medium runs (100+ million grid points) in hours, and large runs (1+ billion grid points) in a day. Portability, coupled with an intuitive GUI and a dynamic output capability, makes NFA an easy-to-use, robust, and efficient platform from the end-user perspective.

Solver Efficiency

Efficiently solving a complex elliptic problem, like the pressure field in an incompressible flow simulation, with billions of unknowns is a difficult problem. Poor conditioning because of a 1000:1 density jump across the air/water interface and an embedded irregular boundary condition on the ship hull make the problem even more difficult. NFA uses a multigrid solver, without which small simulations would be difficult and large parallel simulations intractable. Much effort has been put into the NFA's multigrid algorithm, as it accounts for 60 to 80 percent of the simulation time.

Under the 2010 Capability Application Project (CAP) program, scaling studies including iterations per hour and weak scaling efficiency of NFA were completed. A significant improvement to the scaling, at large core counts (>1024), was made as a result of replacing the one-sided remote memory access (RMA) MPI communication model, which required barriers, with a two-sided nonblocking send/receive communication model completed by John Levesque and the Cray Supercomputing Center of Excellence. Figure 5 shows the scaling results of simulations with fixed blocks of 256 × 128 × 128 grid cells per core distributed over numbers of cores ranging between 8, resulting in 33.6 million grid points, and 41,472, resulting in 174 billion grid points. These simulations require 1.9 GB of memory per core, which is near the 2-GB-per-core hardware limit. The largest simulation demonstrates
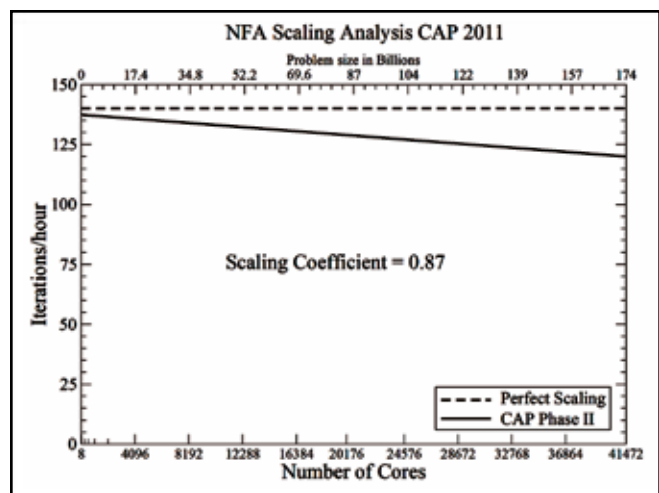

Figure 5. Iterations per hour achieved with 256 × 128 × 128 grid points per core (maximum based on 2 GB/core available memory).

the full ability of the AFRL DSRC Cray® (Cray, Inc.) XE6, *Raptor*, using nearly all of the available memory and all of the available cores concurrently to simulate a complex free-surface flow with 174 billion grid points. Notably, NFA only requires 2 GB of memory per core, or less, and has no increased memory requirements when expanding the problem size by adding additional compute cores.

The weak scaling coefficient of 0.87 bodes well for NFA to scale out to hundreds of thousands of cores and to one trillion or more grid points. Testing on such large problems should become possible once the DoD completes consolidation of the three Cray® (Cray, Inc.) XE6 machines (*Chugach*, *Raptor*, and *Garnet*), as 151,040 cores will become available on a single machine.

Conclusions

Advanced visualization techniques have brought an NFA simulation of a planing boat, with a simulated wavefield, to life by realistically rendering the hydrodynamic output. The detailed multiphase flow associated with the formation of the spray sheet and free-surface turbulence of planing boat hydrodynamics at high Froude number has been verified through comparison with experimental data. This demonstrates that NFA can successfully model these complex fluid flows and the resulting impact on vessel motion and stability.

Efforts to reduce memory usage and improve scaling performance allow runs to be spread across greater numbers of processors to drastically reduce run time. Because of efficient, dynamic output, NFA is able to reduce disk utilization to manageable levels even

when simulations involve billions of grid cells. NFA is now well positioned to serve as a useful tool for evaluating novel planing hull-forms and to solve even larger, more sophisticated hydrodynamic problems.

Acknowledgments

References

Brucker, K. A., T. T. O'Shea, and D. G. Dommermuth, J. Levesque. 2011. Numerical Flow Analysis Capability Application Project 2010. In *Proceedings of the 21st Department of Defense High Performance Computing Modernization Program's User Group Conference, Portland, OR*.

Brucker, K. A., T. T. O'Shea, and D. G. Dommermuth. 2012. Numerical Flow Analysis of Planing Boats. In *Proceedings of the 22nd Department of Defense High Performance Computing Modernization Program's User Group Conference*. IEEE.

Dommermuth, D. G., T. T. O'Shea, D. C. Wyatt, T. Ratcliffe, G. D. Weymouth, K. L. Hendrikson, D. K. Yue, M. Sussman, P. Adams, and M. Valenciano. 2007. An Application of Cartesian-Grid and Volume-of-Fluid Methods to Numerical Ship Hydrodynamics. In *Proceedings of the 9th International Conference on Numerical Ship Hydrodynamics, Ann Arbor, MI*.

Fu, T. C., T. T. O'Shea, C. Q. Judge, D. G. Dommermuth, K. A. Brucker, and D. C. Wyatt. 2012. A Detailed Assessment of Numerical Flow Analysis (NFA) to Predict the Hydrodynamics of a Deep-V Planing Hull. In *Proceedings of the 29th Symposium on Naval Hydrodynamics, Gothenburg, Sweden*.

Judge, C. Q. 2012. Static and Dynamic Forces and Wetted Lengths for a Planing Hull Model Forced in Roll. In *Proceedings of the 3rd Chesapeake Power Boat Symposium, Annapolis, MD*.

O'Shea, T. T., K. A. Brucker, D. C. Wyatt, D. G. Dommermuth, and T. C. Fu. 2012. A Detailed Validation of Numerical Flow Analysis (NFA) to Predict the Hydrodynamics of a Deep-V Planing Hull. In *Proceedings of the 3rd Chesapeake Power Boat Symposium, Annapolis, MD*.

O'Shea, T. T., K. A. Brucker, D. G. Dommermuth, and D. C. Wyatt. 2008. A Numerical Formulation for Simulating Free-Surface Hydrodynamics. In *Proceedings of the 27th Symposium on Naval Hydrodynamics, Seoul, Korea*.

Rottman, J. W., K. A. Brucker, D. G. Dommermuth, and D. Broutman. 2010. Parameterization of the Internal Wave Field Generated by a Submarine and its Turbulent Wake in a Uniformly Stratified Fluid. In *Proceedings of the 28th Symposium on Naval Hydrodynamics, Pasadena, CA*.

Real-Time Prediction of RF Signal Attenuation Using the CPU/GPU Hybrid Dedicated HPC Project Initiative at ARL

By Brian J. Henz, U.S. Army Research Laboratory (ARL), Aberdeen Proving Ground, Maryland; Evens Jean, Secure Mission Solutions, North Charleston, South Carolina; David Richie, Brown Deer Technology, Forest Hill, Maryland; Song Jun Park, ARL; James A. Ross, Dynamics Research Corporation, Reston, Virginia; and Dale R. Shiers, ARL

The scale and complexity of mobile ad-hoc networks (MANETs) used by the Department of Defense (DoD) are unique and continuously increasing, particularly within the Army, creating a mobile fighting force. The military is rapidly becoming a network-centric force, with substantial access to sensor derived surveillance information, as well as to an increasingly complicated application layer running over many different devices. Each layer introduces significant advantages to the Warfighter but also brings in new dependencies and new risks from the rapid change in configurations of the MANETs that provide network access across the battlefield. Large-scale testing, evaluation, and analysis of MANET platforms is an expensive yet necessary undertaking requiring extensive study of the quality of mobile radio transmissions that are affected by many environmental factors such as terrain, foliage, and weather conditions, to cite a few. Simulation and emulation tools have been developed to support such studies of MANET, with emulation holding great promise in limiting the amount of live experimentation required for platform development. Emulation provides for hardware-in-the-loop (HIL) testing and analysis where the physical medium is replaced by a virtual environment, and a physical or simulated radio can be used with real applications. At the physical layer, the interaction between devices is governed by the RF propagation characteristics of the environment.

Independent of the approach, the accurate simulation or emulation of mobile radios requires the computation of RF propagation path loss in order to predict connectivity and signal interference. RF propagation models play an essential role in the planning, analysis, and optimization of radio networks (Schmitz and Weng 2006; Rick and Mathar 2007; Michea and Komatitsch 2010). For instance, coverage and interference estimates of network configurations are based on field strength predictions; routing is also highly dependent upon computed path loss data. RF propagation path loss predictions for MANET emulation has traditionally relied on either off-line link analysis using various models including high-fidelity finite difference time domain (FDTD) and ray tracing methods or real-time calculations with stochastic models (Kaplan et al. 2009; Nitsche and Fuhrmann 2011). MANET emulation with HIL capabilities requires that the path loss data be computed and provided to the emulation environment in real-time despite the computational complexity of proposed path loss algorithms. A more realistic physical layer for MANET emulations and simulations can be achieved through the use of high performance computing (HPC) resources and graphics processing units (GPUs) to provide the floating point performance required to compute the RF propagation path loss algorithms in real-time. To that end, the Mobile Network Modeling Institute (MNMI), headed by the U.S. Army Research Laboratory (ARL), was established in 2007 to exploit HPC resources through the development of computational software for MANET analysis. The goal of the MNMI is to enable the DoD to design, test, and optimize networks at sufficient levels of fidelity and with sufficient speed to understand the behavior of network technologies in the full range of conditions under which they will be deployed.

Operational goals of the MNMI include the development of scalable computational modeling tools for simulations and emulations, the ability to understand *a priori* the performance of proposed radio waveforms in the field, and the optimization of the network for U.S. Army warfighters. In this article, we present an overview of our work in computing the RF path loss of mobile radios utilizing a dedicated HPC system and the ray tracing algorithm (Nitsche and Fuhrmann 2011; Bertoni 1999) in urban environments, such as depicted in Figure 1.

The results presented here stem from an ARL effort to develop a framework for large-scale MANET emulations (e.g., up to 5000 emulated devices) through the use of the Extendable Mobile Ad-hoc Network Emulator (EMANE) from DRS (formerly Cengen Labs) (Galgano et al. 2012). The ARL framework will serve as a test bed for research, development, and evaluation of network algorithms, applications, and devices. The framework makes extensive use of HPC resources and GPUs to mitigate the computational complexity of RF path loss algorithms that are typically $O(n^2)$ where n is the number of transmitter/receiver device pairs.

The available methods for computing RF path loss require a large number of floating point operations (FLOPs), necessitating a high FLOP rate for real-time path loss predictions. We have investigated the development of three algorithms for use on GPUs: the irregular terrain model, the transmission line matrix method, and the ray tracing method. The Irregular Terrain Model (ITM) is well suited for use with large-scale emulations of 1000s of devices but fares poorly in urban environments (Longley and Rice 1968). The Transmission Line Matrix (TLM)

Figure 1. Polygon representation of Tonsberg, Norway.

algorithm is targeted towards pico-cell scenarios within buildings or in relatively localized urban environments (Christopoulos 1995). TLM relies on the relationship between electromagnetic field quantities and voltage and current on transmission lines. The algorithm is computationally expensive, $O(n^3)$. The ray method is used primarily for small-scale to large-scale urban environments and is at the center of our current discussion.

The use of GPUs has been identified as a solution to provide raw FLOP performance (Song and Akoglu 2011) for real-time path loss computations. Initially developed for rapid rasterization, GPUs tend to exceed the performance of CPU architectures for raw FLOP rates, hence their use in our work. We further relied on Standard CL (STDCL) (Brown Deer Technology 2012), which leverages OpenCLTM (Khronos Group 2012) to present a more efficient and simplified interface designed for the complexities of heterogeneous computing platforms (GPUs/CPUs) and HPC applications.

Our efforts have recently focused on developing a ray tracing algorithm optimized for GPUs and HPCs for use in real-time MANET emulations. Previous work in the field has focused on efficient methods for computing the propagation path loss using ray tracing (Liang and Bertoni 1998) and also porting of the ray tracing algorithm to GPUs. Minimizing the path loss computation time through algorithm development and GPU optimizations is of prime importance. Within the study of RF propagation, ray tracing is a technique by which a large set of rays emanating from a transmitter are launched in all directions of interest. The rays are then traced to study the attenuation of the signal as they undergo reflection, diffraction, scattering, and refraction through their interactions with objects in the environment.

The virtual environment used in this study is a polygon-based 3D representation of the town of Tonsberg, Norway (Figure 1). The model consists of 68,356 triangles stored on the GPU. For each transmitter in the system, we rely on the GPGPU cores to trace the paths of individual rays that have been generated through initialization, reflection, and diffraction. The rays are generated based on user-specified values n_θ and n_φ that represents the number of θ and φ angle partitions of a unit sphere surrounding the transmitter. The product $n_\theta \cdot n_\varphi$ is the number of initially generated rays. The path of the reflected rays from interaction with the model is computed based on the laws of reflection in light propagation. Computing the RF path loss seen at a particular point in the environment in the presence of reflected rays takes into account the unfolded path length of the rays, as well as the number of reflections that the rays have undergone. As ray tracing exhibits a high degree of parallelism through the uncoupled transmission of individual rays, by increasing the number of processing cores available, the computation time is reduced proportionally. The path of a ray depends upon the interaction of the ray with the model's polygons. Figure 2 showcases the path of individual rays from two transmitters to one receiver through the environment.

Performance of the emulation test bed system is tightly coupled with the ability to efficiently compute the signal attenuation. With the ITM algorithm, we investigated the effect of changing the amount of work performed per kernel execution. We varied the stripe size from 1024 to 2048 point-to-point calculations and noticed an improvement of roughly 10 percent across different GPUs as shown in Table 1. When the block size is increased to 4096, the improvement is not as drastic, leading us to conclude that increasing the block size further would yield diminishing returns.

Table 1. Performance for AMD and NVIDIA GPUs as a function of block size in terms of rthe number of point-to-point paths evaluated per kernel execution

Processor	Block Size	Time (s)
ATI Radeon HD 5870	1024	83
ATI Radeon HD 5870	2048	75
ATI Radeon HD 5870	4096	72
AMD Radeon HD 6970	1024	65
AMD Radeon HD 6970	2048	58
AMD Radeon HD 6970	4096	55
NVIDIA Tesla C2070	1024	42
NVIDIA Tesla C2070	2048	40
NVIDIA Tesla C2070	4096	39

For comparison, the TLM is well suited for computation on GPU architectures. In this case, the biggest bottleneck is expected to be data transfer across the PCIe bus, which is also known to be a bottleneck for many applications executing on GPUs. By limiting the number of times results are transported across the PCIe bus in the TLM algorithm, we were able to optimize the calculation time by an order of magnitude (Figure 3). Notice in Figure 3 that the time per step for CPUs remains fairly constant from 10 to 1000 steps, whereas the GPU results show an order of magnitude decrease in time per step. This illustrates the importance of increasing the computation to communication ratio when using GPUs as a co-processor. In Figure 3, the cpu-opt and gpu-opt lines refer to the use of the shuffled grid optimization method that combines four single precision floating point operations into a single float4 SIMD operation. The gpu-opt time per step line shows a nearly ideal 4x speedup over the unoptimized version, whereas the cpu-opt line shows about a 1.2x speedup over the unoptimized CPU version. The nearly 4x speedup indicates that the algorithm is able to take advantage of the GPU ability to perform MADD functions on four 32-bit floating point values simultaneously.

The primary factor in determining execution time for the ray tracing algorithm is the number of rays generated and computed. The total number of rays in the system depends upon the number of rays emitted by individual transmitters; the number and degree of reflections, scattering, diffraction and refraction allowed; as well as the number of planar surfaces with which the rays interact. Our model contains 68,356 triangles, and the benchmark scenario contains two transmitters and one receiver. We used an NVIDIA Quadro FX4800 GPGPU with 1.5-GB GDDR3 of GPU memory and varied the maximum number of reflections that rays are permitted to undergo from 1 to 6. We also varied the number of rays initially generated from 642 to 2562. The results show a linear relationship between the number of rays and the run time with offsets depending on the initial angular partitioning used (Figure 4). The offsets are approximately equivalent to the difference between the square of the number of angular partitions.

The hardware architecture that supports the MANET emulation and simulation efforts at ARL (computing
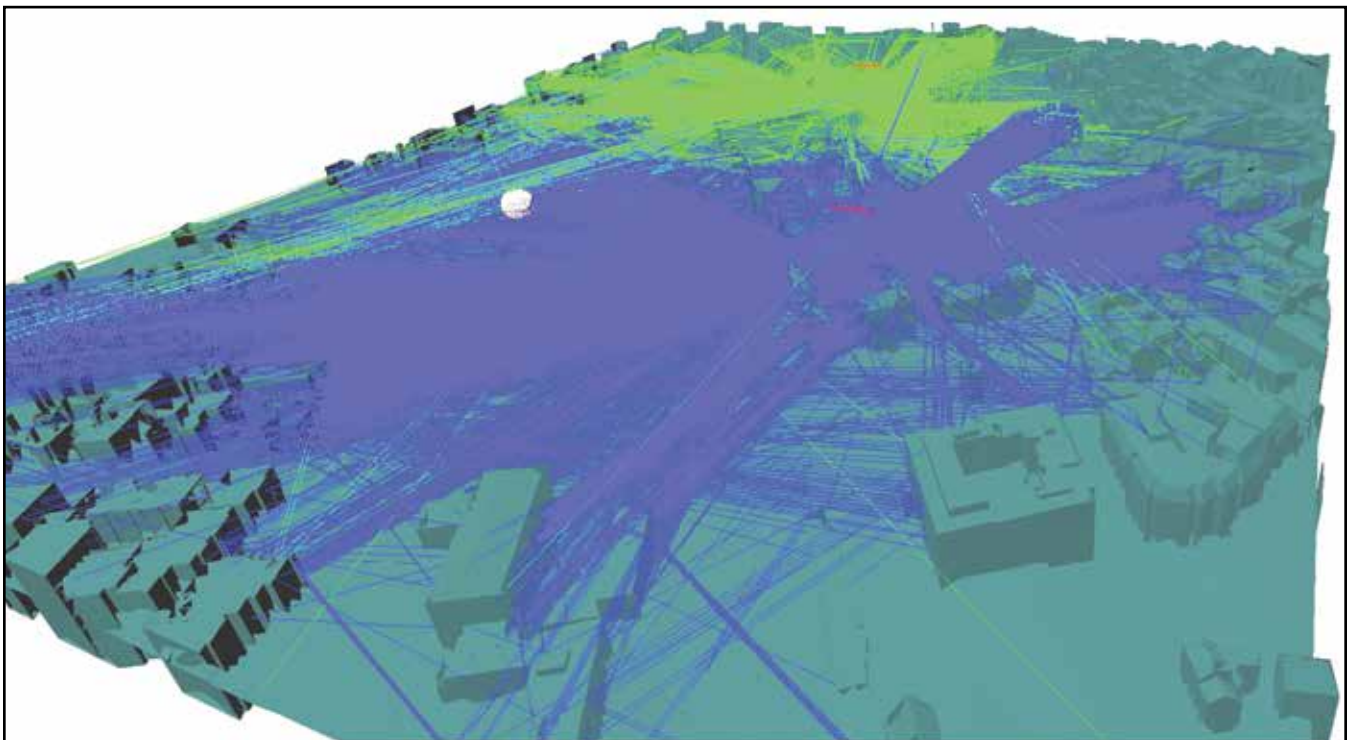
Figure 2. Transmitted signals from two transmitters represented by blue and green lines; receiver represented by white sphere.
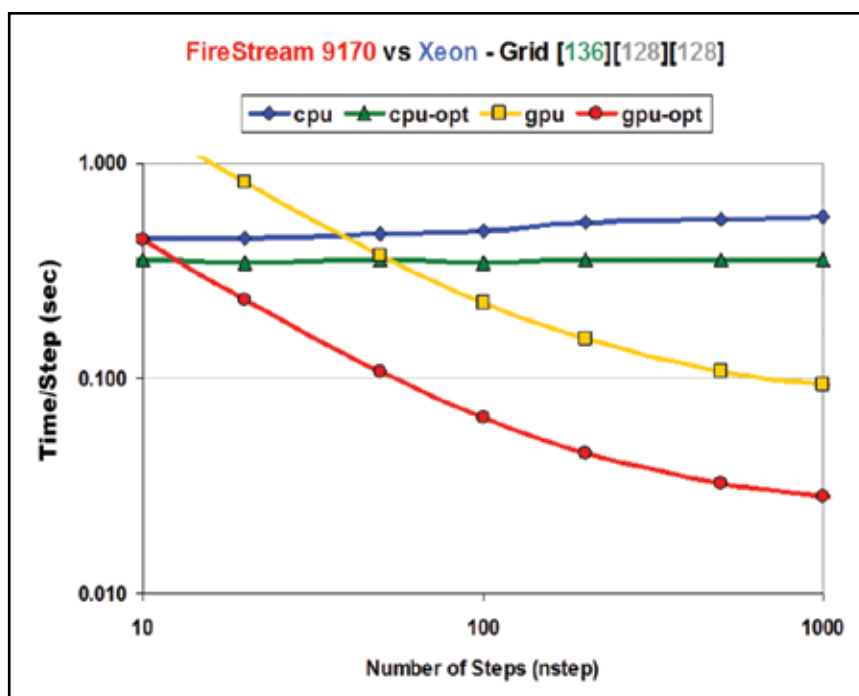
Figure 3. TLM performance improvements obtained by limiting memory I/O across PCIe interface.
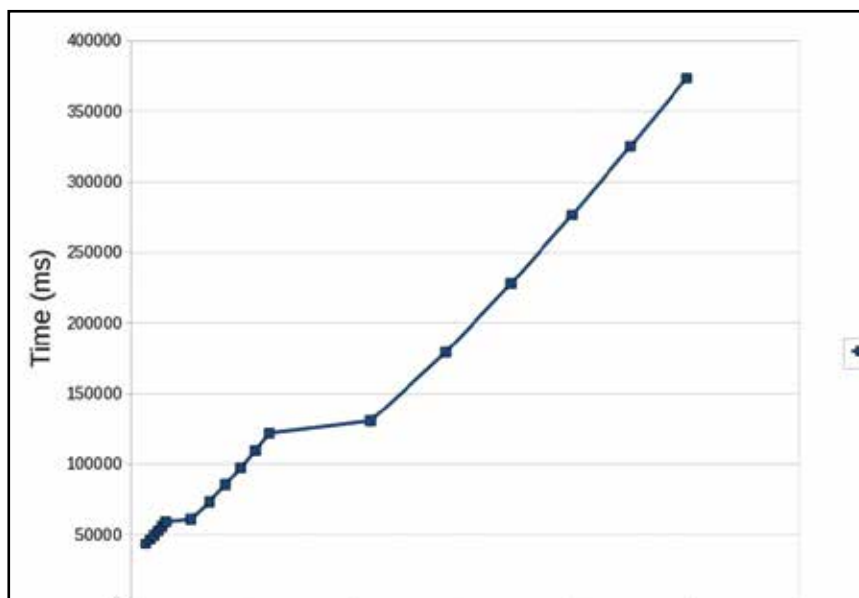
Figure 4. Ray tracing algorithm run time versus number of generated rays. The three slopes correspond to number of initial angular partitioning of 64,128, and 256 partitions.

facilities provided by the ARL DSRC [DoD Supercomputing Resource Center] at the Aberdeen Proving Ground, MD) was acquired through the HPCMP dedicated high performance computing project initiative (DHPI). Since the system is dedicated to MANET emulation, we were able to configure the system to efficiently perform large-scale calculations in a balanced approach. The specifications of the machine are as follows:

- 2.6 GHz 12c Interlagos—303 TFlops Peak DP (700 SP)
  - 160 compute 2S nodes, 2.67 GB Mem/core
  - 114 GPU 2S nodes, 2.67 GB Mem/core
  - 6576 total CPU cores
  - 456 Nvidia M2070 GPUs, 4 GPUs per node, 3 PCIe x16 Bus (222,528 compute cores)

- 10 GbE cluster interconnect
  - 10 Gb/s and 40Gb/s with Gnodal switches
  - 2-stage Fat Tree
  - Latency ~282ns min, ~546ns max
  - Non-oversubscribed
- (2) 40TB Panasas PAS12 Shelves, dedicated 1 GbE Interconnect
- 2 Login and 2 Management nodes, management network

The specifications for Thufir include a large number of physical CPU cores (6576). These cores are used for virtual machine (VM) hosting and allow process pinning so that VMs are provided real-time access to the underlying CPU core. Each CPU core also has 2.67 GB of physical memory available for VM support or combined on a single node. The 64 GB of memory supports large-scale computations. The GPUs are used primarily for RF propagation calculations but can also be used to off-load expensive calculations that support MANET emulation and simulation. The high-speed/low-latency Ethernet interconnect supports an industry standard interface while providing high performance. The low-latency and non-oversubscribed requirement is particularly important for wireless emulations where real signals travel at the speed of light and the order in which packets arrive must be as repeatable and accurate as possible. The system also includes about 60 TB of reliable storage for collecting large datasets during MANET emulations and subsequent analysis. As the reader can see, the architecture of this system is geared for large-scale MANET emulation and simulation.

Without realistic RF propagation, MANET emulation of large-scale networks is a useful tool for network analysts, but the accuracy of the results is questionable. Using GPUs, we have explored three RF propagation path loss methods, although we focused primarily on ray tracing in this article. The explored algorithms cover a broad range of typical scenarios encountered in MANETs. Real-time path loss calculations enable additional capabilities, such as the use of MANET emulation in live exercises and the driving of programmable attenuators for laboratory experimentation with physical devices. Dedicating large numbers of CPU cores to path loss computations is no longer necessary, and such cores can be dedicated to hosting the virtual machines required for large-scale MANET emulation. Use of GPU co-processors plays an important role in efficient large-scale MANET emulations. Our efforts, while ongoing, are enabling the emulation framework at ARL to provide real-time situational awareness data to live field exercises and will have applicability to the testing and evaluation of new radio technologies prior to fielding.

Acknowledgments

References

Bertoni, Henry L. 1999. *Radio Propagation for Modern Wireless Systems*. Upper Saddle River, NJ: Prentice Hall Professional Technical Reference.

Brown Deer Technology.2012. COPRTH. http://www.brown-deertechnology.com/coprth_r_stdcl.htm. (accessed 14 November 2012).

Christopoulos, C. 1995. *The Transmission Line Modeling Method: TLM*. Piscataway, NJ: IEEE Press.

Galgano, S. M., K. B. Patel, and E. Schreiber. 2012. *EMANE Developer Manual 0.7.3*. Bridgewater, NJ: DRS CenGen, LLC.

Kaplan, Michael A., Ta Chen, Mariusz A. Fecko, Provin Gurung, Ibrahim Hokelek, Sunil Samtani, Larry Wong, Mitesh Patel, Aristides Staikos, and Ben Greear. 2009. Realistic wireless emulation for performance evaluation of tactical manet protocols. In *IEEE Military Communications Conference (MILCOM), 18–21 October*.

Khronos Group. 2012. OpenCL–The open standard for parallel programming of heterogeneous systems. http://www.khronos.org/opencl/. (accessed 14 November 2012).

Liang, G., and H. L. Bertoni. 1998. A new approach to 3-D ray tracing for propagation prediction in cities. *IEEE Transactions on Antennas and Propagation* 46(6):853–863.

Longley, A. G., and P. L. Rice. 1968. *Prediction of Tropospheric Radio Transmission Loss over Irregular Terrain: A Computer Method—1968*. ESSA Technical report ERL 79-ITS67. Boulder, CO: Environmental Sciences Services Administration.

Michea, David, and Dimitri Komatitsch. 2010. Accelerating a three-dimensional finite-difference wave propagation code using GPU graphics cards. *Geophysical Journal International* 182(1):389–402.

Rick, T. and R. Mathar. 2007. Fast edge-diffraction-based radio wave propagation model for graphics hardware. In *INICA '07. 2nd International ITG Conference on Antennas, 28–30 March, Munich, Germany*.

Schmitz, Arne, and Martin Wenig. 2006. The effect of the radio wave propagation model in mobile ad hoc networks. In *The 9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2006), 2–6 October, Torremolinos, Spain*.

Song, Yang and Ali Akoglu. 2011. Parallel implementation of the irregular terrain model (ITM) for radio transmission loss prediction using GPU and Cell BE processors. *IEEE Transactions on Parallel and Distributed Syst*ems 22(8):1276–1283.

Thomas Nitsche and Thomas Fuhrmann. 2011. A tool for raytracing based radio channel simulation. In *SIMUTools '11, 21–25 March, Barcelona, Spain*.

Code Development at the ARL DSRC – Yes you can!

By George Petit, Customer Support Team Lead, ARL DSRC

For the majority of users at the HPCMP DSRCs, the production systems are used for running vendor-developed engineering/scientific analysis codes to generate the required computational analysis for their projects. However, there is a subset of users who run their own home-grown codes on these systems. Many of these users have special requirements in order to effectively develop and maintain their codes for our production systems. These include a large dedicated disk area to store the source code and to build and store libraries and executables; special groups to limit access to only the developer(s) and selected users; version control software for version maintenance; compilers/ MPI suites and other high-productivity languages, general purpose math and data modeling libraries; software for debugging and performance analysis; and finally, dedicated compute node access for iterative code verification, debugging, and rebuilding. The good news is all these capabilities are already here! The remainder of this article will describe how you, as a developer, can access and use these capabilities and resources to develop and maintain your code on the ARL DSRC systems.

Defining a Space for Code Development and Storage

The first decision a developer will have to make is what system to use to store and to build his code. For code development and modification, the Utility Server (abutil) is an ideal choice for several reasons:

• It is designed for interactive use and thus has a low load average allowing for efficient source code editing and executable building.

• It also supports batch processing with PBS but without the job backlog on the production machines. In most cases, test jobs will start immediately.

• There are identical Utility Servers set up at all the DSRCs, making porting to other sites as simple as a remote copy command.

• All the necessary software requirements to support code development are available on the system.

For developers who are simply porting a code developed and maintained at their local site, defining an area on one or more of the production machines will be sufficient. Developers using the Utility Server for development will need to do this as well in order to port the final production code to the system for production runs.

Once the developer decides on a platform, an area has to be set up to store the source code, libraries, executables, and other required files. ARL and all DSRC sites provide this storage capability. The base directory for this disk storage area is defined by the PROJECTS_HOME environment variable automatically defined in a user's default login environment. A user simply needs to put in a request to the ARL helpdesk to have a directory created with a requested directory name. This directory will be created for the user in the $PROJECTS_HOME directory, with the requester or whoever the requester designates as the owner.

The next step is to have a special group defined to control access to this directory. Users are now able to directly create and designate membership in groups through the pIE website (https://ieapp.erdc.hpc.mil/info/kerberosValidate.jsp). Some developers find it useful to set up two groups: one for access to the executables and libraries (*<code_name>*) for all users of the code and one that allows access to the source directories (*<code_name-dev>*) for developer access only.

Building and Managing Your Code

Once the source code has been ported into the development area, you are ready to begin editing, compiling, and linking your code. For editing, the following editors are available:

• vi
• vim
• emacs
• jove
• ruby

For compiling, the tables below provide a list of the compilers and MPI suites available on the ARL unclassified systems.

ARL DSRC System Compilers			
Compiler	Utility Server	*Harold*	*Pershing*
GNU	Yes	Yes	Yes
Intel	Yes	Yes	Yes
PGI	Yes	No	No
Java	Yes	Yes	Yes

ARL DSRC MPI Suites			
MPI Suites	Utility Server	*Harold*	*Pershing*
Intel MPI	No	Yes	Yes
OpenMPI	Yes	Yes	Yes
MPICH	No	No	No
IBM POE	No	No	Yes

In addition to the standard compilers, the following high-productivity languages and libraries are available through the Computational Science Environment (CSE) software suite:

- Python
- PyMPI
- SciPY
- OCTAVE
- MATPLOTLIB
- Boost libraries

The CSE also provides access to these data modeling packages:

- HDF5
- NetCDF
- XDMF

For a full listing of all the CSE software, execute the following on any of the systems:

```
>module load cseinit
>module load cse-tools
>module avail
```

There is also a full suite of math libraries available provided by PETTT on all systems. These include the following:

- ARPACK: Solution of Eigenvalues/Eigenvectors
- FFTW: Library for computing Discrete Fourier Transforms (Both MPI and non-MPI versions are maintained)
- PETSc: Suite of scientific computing routines
- SuperLU: Direct solution of sparse linear systems of equations
- LAPACK: Software library of numerical linear algebra routines
- ScaLAPACK: Subset of parallelized LAPACK routines
- BLAS: Basic linear algebra routines (ATLAS and GotoBLAS)
- GSL: C/C++ library with a wide variety of mathematical routines

These libraries and others can be found under $PET_HOME/pkgs on any of the ARL systems.

Another important part of code development is version management. This process is necessary to allow multiple developers to collaborate successfully in building and maintaining various versions of source code files. There are two popular version management packages available on the ARL systems: CVS and Subversion. Both are accessible within a user's default path.

Debugging and Analyzing Your Code

An integral part of any code development effort is debugging and analyzing code performance. There is a full suite of software packages on the ARL systems to assist developers in performing these tasks.

For debugging, both the TotalView and DDT debuggers are installed on our systems. Both are GUI-based, very capable, and feature-rich packages. Most users are familiar with TotalView; however, only a limited number of licenses are available for its use (50). For most debugging requirements, small test jobs are sufficient; so TotalView or DDT would be a good choice for these types of debugging efforts. However, occasionally code problems only manifest themselves when running on a large number of cores. In these cases, DDT is available for use on large-scale debugging jobs. See the system user guides for more information on the uses of TotalView and DDT.

There are many important aspects of code performance that impact high performance system code development. To address these requirements, several code performance analysis packages have been installed on the ARL systems, either as part of vendor-supplied compiler suites or by PETTT. The following tools are available:

- GNU gprof: Application execution profiling tool. Access to this software is defined in the user's default environment.

- Intel CODECOV: Serial code profile-guided optimization tool. CODECOV is provided by Intel as part of the compiler suite.
- PAPI library: Suite of routines to access hardware performance counters for instrumenting and analyzing code performance. The PAPI library can be found in the PETTT area under $PET_HOME/pkgs.
- TAU: Provides a GUI-based suite of tools for parallel performance analysis of FORTRAN 95, C, and C++ applications. The TAU software package can be found in the PETTT area under $PET_HOME/pkgs.
- Valgrind: An instrumentation framework for building dynamic analysis tools for debugging, profiling, and similar tasks that help improve code performance. The Valgrind software package can be found in the PETTT area under $PET_HOME/pkgs.

Dedicated Access

Most development efforts require dedicated access to a subset of the compute nodes, at least occasionally, for extended code debugging, analysis and verification, and demonstration purposes. In the case of intermittent requirements, the Advanced Reservation Service (ARS) is available to reserve up to 12 percent of the available compute nodes on a system for up to 1 week per reservation. Reservations should be made at least 24 hours in advance. The ARS can be accessed through the CCAC web page, www.ccac.hpc.mil. If a development effort requires full-time dedicated access to a subset of the available compute nodes on a system, the project may apply for a Dedicated System Partition (DSP) through the HPC Modernization Program Office (HPCMPO). Applications for DSPs are solicited by the HPCMPO on a fiscal year basis through a memo sent to all project leads.

Portal Development and Operations for High Performance Computing (HPC) at the Maui High Performance Computing Center

By Dave Morton, AFRL; Peter Colvin, Lance Terada, Betty Duncan, and Kelly Suzuki Payba, MHPCC DSRC Staff; Jeff Houchard, Eric Therkildsen, and Randy Goebbert, Pacific Defense Solutions; and MHPCC DSRC

Background

The High Performance Computing Modernization Office (HPCMO) is increasingly interested in ways to promote the access and use of High Performance Computing (HPC). Enhancing effective access and use of HPC resources includes (1) improving agile access to HPC computational and data storage resources by current users, (2) increasing transparent and user-friendly access to these resources by client-based "communities of practice" that have not been historical HPC users, and (3) establishing new ways that HPC can increase project cycle effectiveness and efficiencies in these user communities. The benefits of this program to the end application user include (a) eliminating traditional HPC application stove-pipes, making application and data sharable world-wide and on-demand; (b) requiring no installed software, eliminating user configuration and maintenance; (c) simplifying and enhancing the security model; (d) requiring no HPC knowledge or traditional hurdles to fully exploit HPC; and (e) providing immediate access to HPC resources using fieldable, low-power, portable devices, such as the iPad.

At the direction of the HPCMO, the MHPCC DSRC has taken a leadership role in expanding HPC support to DoD science and engineering organizations through the use of a web-enabled portal. MHPCC is working with the Army Research Laboratory (ARL) DSRC, the USACE Engineer Research and Development Center (ERDC) DSRC, the Computational Research and Engineering Acquisition Tools and Environments (CREATE) Team, and HPCMP leadership to integrate ongoing activities. An initial production version of the Portal is operational on the MHPCC *"Mana"* System, and a development path has been approved within AFRL and the HPCMP communities to provide portal capabilities across all DSRCs to users in a timely and cost-effective manner.

HPC resources demonstrate great value in supporting science, engineering, and business enterprises. However, historically, HPC use has been confined to specialized groups and has not expanded into other sectors where the derived value of HPC could be significant. This has been attributable, in part, to stringent access requirements to HPC resources, lack of application software, limited access to specialized talent, and cost constraints. Activities such as the CREATE initiative have focused on overcoming these constraints and "mainstreaming" HPC capabilities to support clients such as DoD acquisition programs. Increasing access and ease of use of HPC applications has the potential to dramatically expand the customer base both for CREATE applications as well as for third party applications such as MATLAB.

Access and Operations Support

The MHPCC Portal provides a single sign-on Software-as-a-Service (SaaS) working environment for users from their client web browsers via secure CAC or HPCMP YubiKey. After a user is authenticated to the Portal, an integrated framework with access to decentralized components allows users to run HPC "jobs" on available HPC resources in response to and in support of applications being served through the Portal.

Through the use of web-technologies, the Portal is a "zero footprint" install, meaning users will not require any additional software installs on their workstations to use the Portal. As all DoD sites already have a web browser client installed or approved for use on their workstations, it will not be a requirement to go through any separate software approval process to use the MHPCC Portal. The HPCMP highly desires this "zero footprint" install—which now opens up opportunities to reach out to DoD agencies and users who, in the past, could not access the HPCMP DSRCs because of the requirement to install Kerberos software to access the HPC resources. For the DoD acquisition community, this new approach also provides the opportunity to bring HPC to the otherwise isolated desktop user who may typically not be a command-line advanced user. During 2013, the Portal will switch over from local authentication to an OpenID authentication server installed by AFRL. MHPCC has also posted an educational Website that provides an introduction to the Portal at http://www.mhpcc.hpc.mil/portal/.

On-boarding: To those familiar with the process, obtaining an MHPCC Portal account is no different than requesting access to a system at any one of the five HPCMP DSRCs. For those new to the process, first and

foremost, a user must obtain an account on the Portal to the Information Environment, commonly referred to as a pIE account. Users can find information on obtaining a pIE account on the HPCMP Consolidated Customer Assistance Center (CCAC) web pages located at http://www.ccac.hpc.mil/accounts/index.html. This web page details the four actions you will need to complete to obtain a pIE account and to start your process of acquiring a system account at the MHPCC DSRC.

Description and Feature Updates

Initial Production Portal: Four stand-alone software product areas were specified for initial deployment and are available for use on the production Portal:
• Computational Research and Engineering Acquisition Tools and Environments (CREATE) Kestrel
• Distributed MATLAB
• Distributed TASAT
• Virtual Applications, including CREATE Capstone, Fieldview, Pointwise, and MATLAB IDE

An important new feature of our HPC Portal is the ability to access a traditional shell capability, as shown in Figure 1. For legacy or power users, this provides the ability to directly access HPC resources from within a browser without requiring installation of a Kerberos client kit.
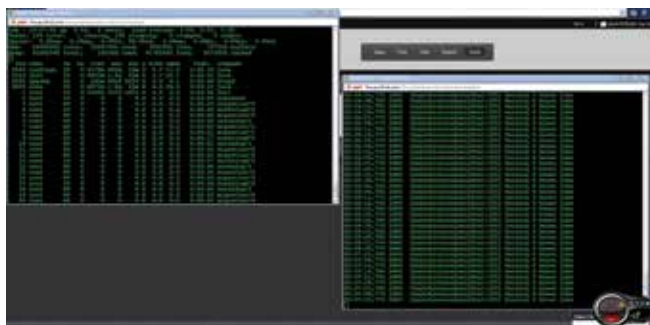

Figure 1. Multiple web-shells available from HPC Portal.

For power users, the shell can provide scripting or development of custom workflows in addition to native HPC Portal applications. An additional benefit is the ability to access shell resources from non-traditional devices, such as a CAC- or YubiKey-enabled iPad. Delivery to these types of devices was previously not possible as they are not compatible with the Kerberos client kit.

Branching Out: The PMO (AD-Centers) has directed the Portal Team to deliver HPC Portal access at all DSRCs in CY2013. Working towards this goal, the team has investigated and completed a design that leverages the capabilities of the Utility Server clusters common to each DSRC. The Utility Server was developed for interactive job debugging, preprocessing,
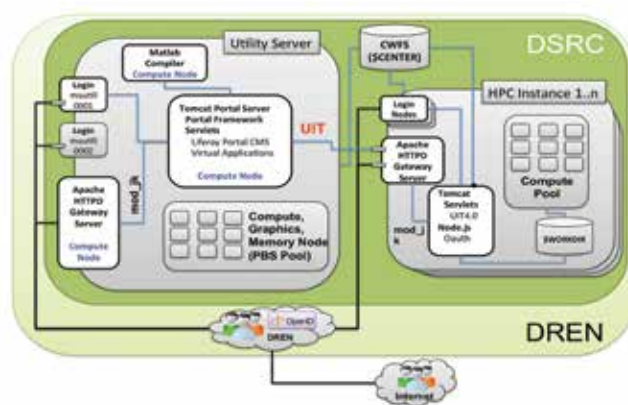
caption?

and post-processing. The Utility Server also includes two types of nodes, large memory nodes and graphic nodes, that are not typically found on HPC resources. These features map well to the Web Services and Virtual Application that are key to delivering the Portal functionality. Integrating the Utility Server resources into the Portal framework will be beneficial to the workflow of the user, providing accelerated graphics and high availability. The inherent commonality within the Utility Servers across DSRCs will enable the deployment of the Portal to all HPCMP DSRCs. The technology to support the Utility Server architecture will be included in the next software release of the Portal (July CY2013), installed at ERDC (August CY2013), and deployed at all DSRCs by the end of CY13.

New Applications: The planned July 2013 release includes the following additional software products:
• CREATE DaVinci
• CREATE Helios
• CREATE Kestrel V3/V4
• ARCADE JSpOC Mission System (JMS) HPC development and test platform
• OneSAF One Semi-Automated Forces application delivery
• TASAT BEEP (Batch Enabler Enhancement Program)
• SDK Release, which includes ReSTful services for language independence, OpenID/Oauth via UIT 4.0

Updated SDK: The portal SDK (Software Development Kit) is being developed for the DoD application developers. Developers share a common repository, and the MHPCC portal team operates as source code change moderators. In this way, new versions of the SDK are compatible with prior versions but may also include enhancements and bug fixes supplied by outside teams using the SDK to get their applications incorporated into the portal. The SDK is still in development for the July 2013 release and will include a ReSTful (Representational State Transfer) service. API REST facilitates transactions between web

servers by allowing loose coupling between different services and is the predominant Web API.

SDK Release 1.0, which provides a Java API, is currently being used to deploy OneSAF using an outside development team in conjunction with the HPCMO PETTT team. As shown in Figure 2, the SDK supplies a Portal FrameWork (backend) for HPC resource access and a user interface appUI (frontend) to provide a common look and feel interface to the portal application while AppTop supports legacy and COTS applications. When developers use the SDK, end users benefit from the commonality that the SDK provides in the user interface and in the resource management, including files, job setup and status, and quick-look capabilities, etc.
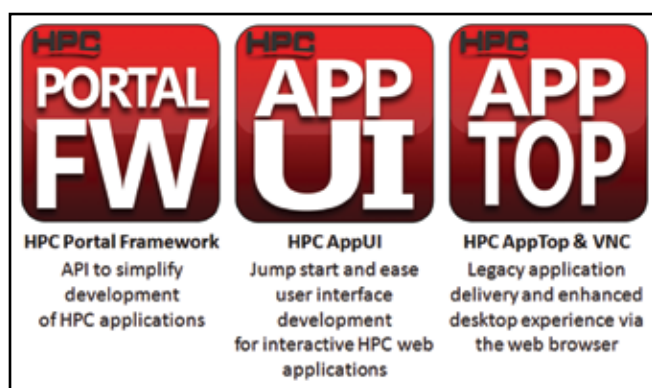

Figure 2. Portal developer tools.

Time-domain Analysis Simulation for Advanced Tracking (TASAT): TASAT is a legacy desktop application that often involves long-running, CPU-intensive, parametric studies. TASAT is a mainstay in the arsenal of modeling resources used to support space situational awareness (SSA). AFRL and other research and intelligence organizations have used it to provide a high-resolution, operational understanding of vehicles in orbit using image and non-image based simulations. Historically, effective use of TASAT has been limited to those organizations that have the local computational expertise and resources needed to support its use. TASAT was developed in the mid-1980s to support the Air Force Research Lab (AFRL) Directed Energy (DE) division and has been continually updated and upgraded over the years to support more capabilities and higher fidelity simulations. TASAT has acted as a test-bed for many programs and has been used to develop requirements and to predict mission performance. The software has also been the benefactor of many extensive validation efforts from empirical field data measurements and is now considered a standard by the Space Situational Awareness (SSA) community for accurate illumination, reflectance calculation, and imaging for space objects.

The distributed TASAT program is one of the core programs provided for the initial production Portal. Along with user training, the next release provides updates to TASAT.
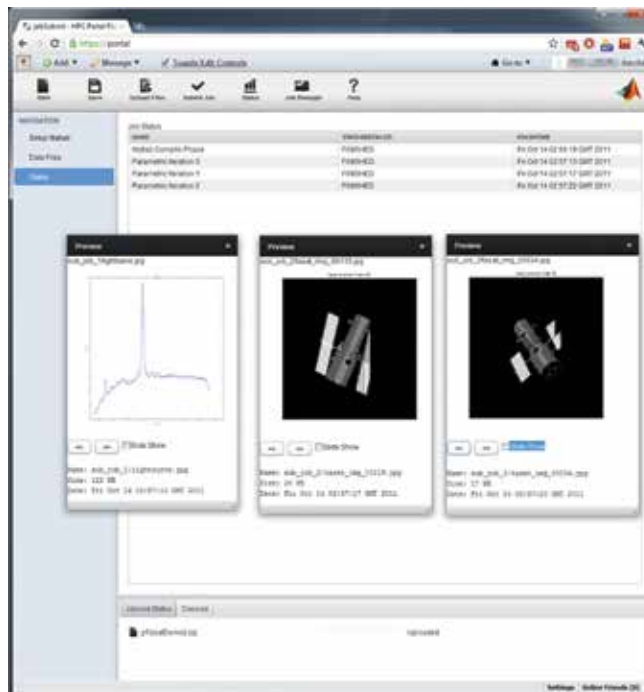
Figure 3. TASAT output after computations are performed on HPC Portal.

To this day, TASAT remains a powerful tool for supporting Air Force missions. However, like many other legacy codes, it has issues in the areas of installation interfacing, maintainability, portability, and error handling. A MATLAB based BEEP (Batch Enabler and Enhancement Program) tool has been developed for TASAT (officially called TASAT BEEP) to enhance the user experience, to develop powerful interfacing between other codes, and at the same time to provide a batch capability to take advantage of HPC.

TASAT Training and Outreach: Outreach and promoting user engagement is a key part of the HPCMP Portal initiative. The Portal as a resource is designed to support training and collaboration as well as the access and use of HPCMP resources. The TASAT BEEP HPC Portal capability was recently demonstrated during a virtual training session including staff from AFRL and MHPCC, as well as other contractors. Using a combination of Defense Connect Online (DCO) and HPCMP Portal capabilities, the training included TASAT BEEP tutorials across a range of applications that could quickly and easily benefit from the high-fidelity, validated, physics-based modeling engine within TASAT, particularly when combined with integrated HPC resources.

Figure 4 depicts preprocessing verification output directly from MATLAB, showing a TASAT BEEP ground sensor view of the Hubble Space Telescope prior to creating the final TASAT parameter file for HPC. TASAT BEEP, and the preprocessing checks, can be used on the client machine using local MATLAB or through the browser, connected to the Portal, using Virtual Application (Vapps) VNC technology.

Leadership within AFRL and HPCMP have endorsed the approach of using the HPCMP Portal to support training sessions, workshops, and the development of collaboration networks focused on improving the use of HPCMP resources.

Summary

The HPC Portal is on track to conclusively demonstrate the high utility a unified distribution mechanism provides. The Portal program has moved beyond the early pioneer and initial production at MHPCC towards a phased rollout to all DSRCs in CY2013. A number of collaboration services, along with the ability to seamlessly access HPC resources within a browser, provide substantial ease-of-use and workflow improvements for these application areas. CREATE, distributed MATLAB, JMS, OneSAF, and TASAT HPCMP applications are important to the DoD acquisition and engineering communities. The ability to quickly integrate legacy applications using outside teams is being tested under OneSAF; it is expected that the Portal SDK should be a force-multiplier across the DoD. For TASAT, enhancements and tutorials with the AF Space Vehicles Directorate were initiated to demonstrate TASAT capabilities within the Portal environment. The goal of this effort is to provide to the larger TASAT community seamless access to HPC capability. The HPCMP Portal can play a significant role in enhancing collaboration within DoD engineering and science and technology communities of practice. Besides empowering distributed, virtual training activities, the HPCMP Portal will provide a means for subject matter expert collaboration, for methodology best practices, and for the promotion of HPCMP resources in support of DoD missions. The goal of our longer term overall effort is to work closely with other Government sponsors and targeted HPC user communities to identify additional applications that should migrate to the portal environment.

We would like to thank our partners for their collaborations with the development of this new technology.

Regarding a recent TASAT training session, "…the session was promulgated and was effective. Personally, I think this was a very good start…Of course, many programs can be run on MHPCC, and eventually we should be able to run programs from our desktops or laptops. This should be a great enabler. I would like to raise this capability at the next SSA IPT meeting and increase participation across directorates."

— Michael J. Duggin, Ph.D., ST, Senior Scientist for Spacecraft Technology, Principal Science Advisor to SSA IPT, AFRL
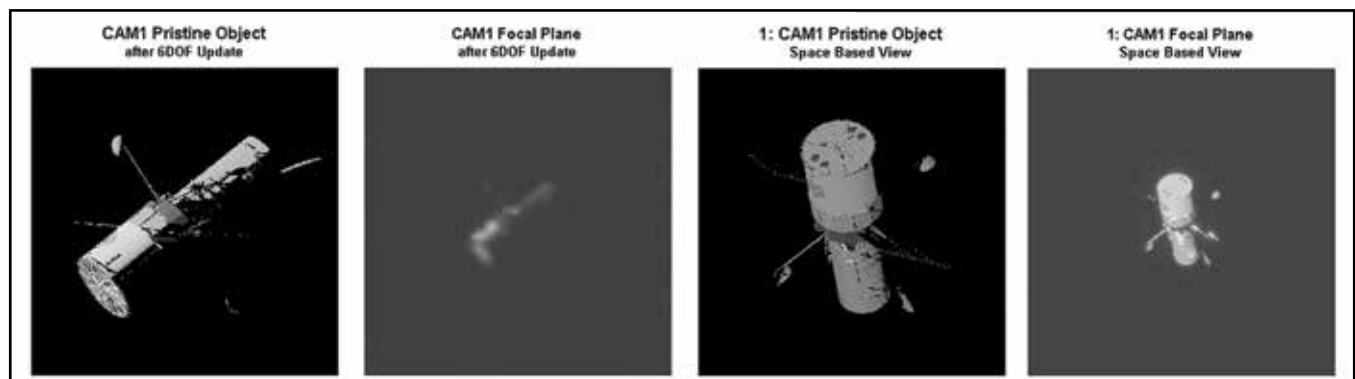

Figure 4. Output from MATLAB IDE using Virtual Apps (VNC) prior to parameter file creation.

From the Director's Desk – Dr. Raju Namburu

The ARL DSRC team has been hard at work as we bring our Technology Insertion 2011/2012 (TI-11/12) systems into full production status in our newly renovated supercomputing facility. Both the unclassified system, Pershing, and the classified system, Hercules, are in production. In addition to the deployment of the TI-12 systems, we have also moved and integrated the DoD HPCMP Dedicated HPC Project Investment (Mobile Network Modeling Institute) system and the Arctic Region Supercomputing Center's (ARSC) HPC Enhanced User Environment (HEUE) test lab equipment to include a utility server, storage silo, and front-end server into our DSRC infrastructure. The level of effort and teamwork by the ARL DSRC systems engineering team and the ARL DSRC facilities engineering team has been nothing short of remarkable since we completed an incredible amount of work in less than six months.

As I mentioned in the last edition of *HPC Insights*, we are partnering with the Army Test and Evaluation community by addressing their computational and I/O intensive requirements and the demand for "big data" analytics for the Army's Network Integration Evaluation (NIE) Process. As a result of this effort, we have established HPC accounts for seven new T&E users who are now using the ARL DSRC resources. Our DSRC team continues to work closely with this new T&E community to further expand their ability to leverage the Program's HPC resources to address their computational and advanced networking requirements.

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

New DSRC Facility Up and Running

By Michael Knowles, Site Engineer, Lockheed Martin/ARL DSRC

As of March 1, 2013, it has been 11 months since ARL (Army Research Laboratory) was awarded a new building on Aberdeen Proving Ground (APG). Within that time frame, ARL and the ARL/DSRC (DoD Supercomputer Resource Center) have moved extremely efficiently to transform the building into a fully functional High Performance Computing center and computing research facility. The ARL/DSRC leadership, in conjunction with the Lockheed Martin (LM), has

• Modified the building to support two new large IBM IdataPlex systems (Pershing: 20K cores; Hercules: 18K cores) delivered in September 2012. Two separate computer rooms were built because of different machine classifications.

• Added 2 MW of utility power with a refurbished generator backup capacity.

• Modified another refurbished 2-MW generator for future installation.

• Refurbished another existing 2-MW generator and power distribution system.

• Added 225 tons of chiller capacity with extensive piping loop extensions.

• Refurbished an existing 400-ton chiller plant.

• Relocated a tape library and HEUE Test Lab equipment (MCAT-ARSC, US TDS, and CWFS TDS) from the University of Alaska Fairbanks.

• Created another computer room ("White Cell") to house DSRC and ARL computational and mass storage assets.

Some of the details of this huge undertaking are discussed below.

The building enhancements that were embarked on in April 2012 included the design and build out of two separate areas of the new facility: the "Blue Cell" (BC) and the "Server Room" (SR). The BC was an area that housed a handful of computer racks and was a laboratory area for the buildings' previous occupants. To support TI-12 BC requirements, the power distribution system had to be built by tying in two 400-kVA UPS systems (repurposed from another building) to the existing building switchboard, adding power panels, and adding a Bus Bar (repurposed from another area) system to deliver power to the new TI-12 racks. Additionally, power distribution was required for the cooling system specifically designed for the room. The cooling system was composed of two 30-ton CRAH units for overall air cooling and four Cooling Distributed Units (CDU) that

Blue Cell with Pershing system installed. (Inset) Blue Cell before conversion to computer room.

provide cooling directly to the rack doors. All power distribution and cooling had to be run overhead as there was no raised floor in the room and little ceiling height to create one. All these modification were in place for the delivery of the SR system on September 7, 2012. The existing building cooling system was not fully functional until February 13 due to extended maintenance issues with two existing 200-ton building chillers.

The SR was a computer area already in April 2012. However, the TI-12 system requirements for that area significantly exceed the room's capabilities and capacities. To meet these new requirements, a new utility transformer and a refurbished backup 2-MW generator were installed along with all required switching equipment. Power distribution for the SR continued with a new high–efficiency, flywheel-based UPS; new PDU installations; reuse of existing PDU and UPS systems; and overhead Bus Bar systems to provide power to the racks. An additional 225-ton chiller system was added to the building chiller loop, as well as extensive piping modifications, CRAH, and CDU installations. A new raised floor was also required in the SR because of the weight of the TI12 system. All these modification were in place for the delivery of the SR system on September 27, 2012.

The generator work associated with this new facility capability expansion centered on three generators: an existing 2-MW generator that was already on site and two 2-MW generators that were excessed by Edgewood Arsenal and refurbished to provide backup power for a TI-12 system and future requirements. The existing generator was not actually attached to the building last

March. Now it is fully functional with the addition of a new 3000-gal. fuel tank, enclosure, and fully-working power distribution system. This generator now supports one of the new TI-12 assets, its cooling requirements, and the overall building systems. One of the excessed generators also is currently providing backup power for one of the TI-12 assets, its cooling requirements, and some building functions. The remaining excessed generator has been refurbished with a new alternator and radiator and is targeted for installation this calendar year.

In the fall of 2012, another project was initiated in the new facility to house equipment that was coming from the Arctic Region Supercomputing Center (ARSC). The "White Cell" (WC) area of the new facility was chosen to house those significant resources as well as other DSRC and ARL specific equipment. Additionally, an adjacent meeting room, the "Red Cell" (RC) was required to provide a working and display area for many of the DSRC and ARL technologies that are in production and are being developed. Another adjacent area, the "Gray Cell" (GC), was chosen to provide classified access space for researchers and staff. These requirements and the design for these areas were developed through the fall and into December 2012. Modification of these areas began in earnest in January 2013 and is ongoing as of March 1. The WC area is approximately 75 percent complete and already houses the ARSC equipment. The RC area is approximately 60 percent complete while the GC area is approximately 95 percent complete. By the end of March 2013, all of these areas will be fully functioning for their intended purposes.

Thanks to the diligent efforts of the ARL and DSRC staff, a huge undertaking is almost completed. Each of these individual efforts required significant resources in the design, implementation, procurement, and coordination. The overall body of work that has been completed towards making the new facility a premier HPCMP and ARL research facility within a year is a testament to the capabilities of the staff and to their willingness to meet demanding requirements and schedules.

White Cell area concept rendering.

From the Director's Desk – Jeff Graham

Who Am I and How Did I Get Here?

Well, I am the new kid on the block as far as Center Directors go, the new AFRL DSRC Director. I am honored to get the opportunity to lead a great bunch of dedicated individuals here in the heartland—all working to bring the power of HPC to researchers across the DoD. But—why me? How did I get here?

One big step on my career path was answering customer phone calls for the Dayton Power and Light Company 32 hours a week while at the University of Dayton. While earning a BS in Mathematics, I learned about life in a service organization.

After graduating, I became a "user first" by accepting a job at Wright-Patterson AFB in the Computational Aerodynamics Group (CAG) in what would become the Flight Dynamics Laboratory. I learned about how you can solve the Navier-Stokes equations that define the airflow around airfoils (which was all we could do in 1979) using supercomputers. How cool is that!

We were vagabond users, begging, borrowing, and stealing time on systems like the CDC 7600, the STAR-100, and eventually the Cray-2 at NASA Ames. My one thought was—why can't this be easier? My boss—Dr. Joe Shang—had a little black book highlighting the specific commands that would drive each computer. I learned that there was a real opportunity to demystify the use of supercomputers—so that a scientist could work on science as opposed to computer science.

After getting my next big job in 1987 in the Aeronautical Systems Center's Computer Center, I was back in the service business—starting a newsletter to provide information to our customers and hosting quarterly scientific user forums. Fast forward to 1992, and I was fortunate enough to become part of the Program that was destined to help my friends in the CAG with their growing and expensive supercomputing requirement. It was the DoD HPC Modernization Program—and it was all about corporately funding HPC resources so customers across the DoD could access supercomputers along with expert support from the top National Science Foundation centers in the Nation—all for absolutely nothing…free. How cool is that?

I was all in to help build this Program in any way I could—and contributed in as many ways across the years as I could. First I was on the original Requirements Gathering team, then on to the Shared Resource Center Advisory Panel (precursor to the UAG), next helped write the anti-Fee-for-Service report with HPCMP Director-to-be Kay Howell, and finally got involved with PET management (when there was only one "T"). Not long after, I helped create the Baseline Configuration Team in 2005—the ultimate opportunity to make using different systems and Centers better for our customers. And I learned how to bring some of the great talent across the Program together to focus on ease of use across our enterprise. How cool is that?

So our Program was like landing in Oz after the vagabond wasteland of Kansas computing we had found ourselves in 20 years earlier. (We actually bought cycles from a company in Kansas City in the 1980s). And now—here we all are—in the middle of a whirlwind not unlike the storm that brought Dorothy to Oz. I take over the Director's job in a time of big change—our Program is now run by the Army; the threats of sequester and big budget cuts are not just looming anymore; we are at the end of our current Technical Services contract with Lockheed Martin; we are developing a new Technology Insertion process; and locally, we are integrating the new Spirit into our new Information Technology Complex while we reorganize our parent organization after five years of being part of AFRL. And—I have learned that there are ways to cope with uncertainty and change. How cool is that?

How do we cope with all the change? All the fear? All the rumors? I think there are ways that will make sure the whirlwind lands us right back in the Land of Oz and not in dreary, black and white Kansas.

1. Work together—and hang together. We have come a long way, and we know our mission is important and necessary for the viability of our Nation's defense. Believe in each other!

2. Get back to the "User First" mentality. Without our users, we have no purpose. Take every opportunity to talk with our customers. Find ways—even without the Users Group Conference. Where there's a will, there's a way.

3. Roll with the punches. Don't let a couple of shots knock you out. Get off the mat, and refocus on the end goal.

4. Communicate, communicate, communicate. Rely on the trust relationships we have developed over the years across our Program. Share your thoughts with your colleagues. Heck—call me on the Tandberg!

5. Have fun! We still have a huge user base and a super-talented staff across the Program. Come to work with a smile, and seize the day! (Carpe diem.) Why not?

Let's all get after it and make sure those nasty witches (furloughs, cutbacks, political uncertainty, etc.) are all killed off. Drop a house on one, and hose down another.

Lancer – AFRL Shared-Memory Test Cluster

Lloyd Slonaker, Technical Director, Air Forse Research Laboratory, DoD Supercomputing Research Center, Wright Patterson Air Force Base

As HPC systems evolved into the now ubiquitous cluster, there's a segment of HPC users who need large amounts of memory that these clusters haven't been able to provide. Traditionally, the large-memory needs have been solved with hardware-based, shared-memory architectures. Those architectures tend to be much more expensive to not only design and manufacture but also, in the end, to purchase. With this in mind, the AFRL DSRC researched solutions to provide large shared-memory capabilities using traditional HPC cluster architectures and software solutions. One possible solution was found with an innovative product from ScaleMP called vSMP Foundations.

ScaleMP's vSMP Foundation software displays the potential to provide a production-quality globally shared-memory architecture at a much lower cost point than the traditional hardware-based globally shared-memory solutions. The Versatile SMP (vSMP) architecture aggregates multiple x86 systems into a single virtual x86 system, delivering a high-end symmetric multiprocessor (SMP) computer. ScaleMP uses software to replace custom hardware and components to generate a virtualized environment. Through vSMP's aggregation, multiple physical systems appear to function as a single logical system.

The AFRL DSRC directed Lockheed Martin to procure a test cluster for the express purpose of researching the performance characteristics of a globally shared-memory architecture provided by ScaleMP's vSMP Foundation software. During Lockheed's research, it was discovered that while there were several existing vSMP solutions, only one vendor had a proven history in providing ScaleMP clusters at the size of the required test system. Therefore, a new cluster was purchased, called *Lancer,* from Appro. This system is similar to the *Gordon* cluster built by Appro at the San Diego Supercomputing Center. However, unlike *Gordon*, the *Lancer* system does not have flash memory. Appro has established past performance and support capabilities through the Utility Server procurement for the HPCMP program.

Configuration Overview

Lancer has 2560 Intel Sandy Bridge-EP compute cores running at 2.60 GHz with 8 GB of memory per core, providing 20 TB of total computing memory. The system consists of 160 Compute Nodes, 2 Login Nodes, an LDAP server, and 2 Management Nodes

Lancer

connected via FDR Infiniband. All nodes in *Lancer* use RedHat Enterprise Linux version 6. *Lancer* has a parallel storage array, which is a Xyratex-based Lustre file system providing 192 TB of usable storage via a 10 GbE storage network. The key component in Lancer is ScaleMP's vSMP Foundation software. vSMP allows the ability to link multiple systems into different Single System Image sizes by allowing the vSMP software to "take over" and "front" most of the hardware between the BIOS and operating system on the individual compute nodes. vSMP uses the FDR InfiniBand interconnect as the primary communications channel in order to link multiple systems into a single system image, and thus this interconnect is invisible to the operating system and is used exclusively by the vSMP hypervisor.

Access

So how do you get an allocation? Since *Lancer* is a test system to investigate the potential for software-based shared-memory architecture, it will not be treated like a traditional allocated system. Access will be granted based upon shared-memory or special test requirements. Authorized DoD and contractor personnel may request an account on *Lancer* by submitting a proposal to the AFRL DSRC. Principal Investigators (PI) will need to submit a short one to two page proposal outlining their HPC experience, level of required support, project suitability for a shared-memory system, how the project supports the DoD mission and HPC technical advancement, and what will be the proposed workload. Any questions regarding the *Lancer* proposals can be directed to sp-proposal@ccac.hpc.mil.

From the Director's Desk – Dr. Robert S. Maier

Preparing for Cray XE6 Integration

This year, the ERDC DSRC is integrating three Cray XE6 systems into a single capability system to support larger core-count jobs on a routine basis. The integrated system will be called *Garnet*, after one of the three XE6 systems. It will include 4,720 32-core nodes, or about 150,000 CPU cores.

Garnet is one of several larger Cray systems supporting unclassified research in the U.S. The XE/XK system *Titan* at Oak Ridge National Laboratory has 300,000 cores. The *Hopper* system at the National Energy Research Scientific Computing Center has over 150,000 cores. The new XE/XK hybrid, *Blue Waters,* at the National Center for Supercomputing Applications has over 360,000 XE cores.

Each of these systems has slightly different scheduling policies to support rapid turnaround for large jobs. *Titan* has a maximum run time of 24 hours for jobs requesting at least 60,000 cores and 12 hours for those requesting at least 5,000 cores. *Hopper* system has a more complex scheduling policy, but basically limits large jobs to 36 hours, while smaller jobs are allowed 48 hours at a lower priority. *Blue Waters*, which is not yet in full production, limits larger jobs to 12 hours and less for smaller jobs.

There are at least two reasons for the scheduling policies adopted at the larger Cray sites. Achieving rapid throughput for large jobs requires both limits on the amount of time any job can run and increased priority for larger jobs. The mean time between failures of compute nodes also suggests a limit on maximum run time. According to ORNL staff *Titan* experiences an average of one node failure per day. To help address the MTBF, Cray systems administrators have adopted *warm boot* procedures, which place the downed nodes back in service without taking the system down.

Several initiatives will help us transition to a capability system. Many users have already integrated *restart* capabilities into their codes, enabling them to drop restart files periodically and resume calculations in the same or a subsequent job. We are providing a clearinghouse of examples and pointers to restart methods, as well as guidance for choosing the best restart I/O techniques. Many users have also learned to combine Monte Carlo simulations into a single job, allowing multiple instances of the same code to run simultaneously under a single job. However, the most important initiative is to increase the number of cores used in a simulation. For the great majority of codes, this will result in a faster time to solution.

How to Save a Cool Million Dollars

By Greg Rottman, ERDC DSRC

In the spring 2012 issue of *HPC Insights*, Paula Lindsay described how the ERDC DSRC was committed to installing a Daikin-McQuay chiller and a Baltimore Air Coil (BAC) evaporative tower to become more efficient in cooling the HPCMP supercomputers hosted at ERDC. What Paula was not at liberty to discuss at that time was the unique way the two pieces of equipment were to be teamed together to go beyond the sum of the efficiencies of the two pieces of equipment.

After searching for the most effective commercially available mechanical equipment, the following were found. In January of 2011, the 700-ton Daikin McQuay Magnitude™ Magnetic Bearing Centrifugal Chiller was introduced to the marketplace. This chiller with its patented magnetic bearing compressor is advertised as 40 percent more energy-efficient than a standard centrifugal chiller. This particular chiller is a water-cooled chiller.

Let's start by saying that Greg Rottman and Charles Rush had discussed a need to reduce mechanical operations cost at the ERDC. Knowing that Charles Rush, one of the owners of Hermetic Rush Services, Inc., had a deep background in innovative ways of providing cooling to unique processes, ERDC teamed up to discuss the goals that ERDC would like to achieve on a mechanical process rather than just meeting specifications of different pieces of mechanical gear. Mr. Rush indicated that he had a process he had been working on that used new and innovative ways to provide efficient chilled water to the ERDC high performance computing systems. The Controlled Water Results, CWRTM, design would meet the requirement for this high-efficiency cooling process. (The U.S. Patent Office has granted Patent Pending Status for the CWRTM design to Charles W. Rush.)

Knowing that BAC was a leader in efficient cooling towers, Mr. Rush contacted them and asked if their towers could be modified to meet the design criteria. He asked if the BAC cooling towers could be modified to include dry cooler coils to provide cooling in cold-climate areas. The answer was yes, and the path to greater efficiency took one more step toward reality.

The main goal of the chilled water system is to take full advantage of the environmental conditions that Mother Nature has to offer. Interconnecting the Daikin-McQuay chiller with the modified BAC cooling tower, along with the proper monitoring and control capabilities, allows this system to operate in four distinct modes. These modes allow the system to leverage the local environmental conditions, thus using the least power necessary to provide the proper chilled water for the ERDC DSRC supercomputers. In addition, as high performance computing (HPC) systems become capable of operating with higher water temperatures, they will increase reduction in power consumption.

On January 23, 2013, we performed the initial test of the system in comprehensive mode transitioning into normal mode. We were in a scheduled outage and started the new chiller system when the three major ERDC computers were partially shut down for maintenance. During the 36-hour period starting at 19:30 on January 22 and ending at 07:30 January 24, we performed various tests on the new mechanical system. We witnessed a power utilization effectiveness (PUE) ranging from 1.45 using our traditional air-cooled chillers to a PUE of 1.13 using the new chiller system in comprehensive cooling mode with an average PUE of 1.25 across the 24-hour period.

Current projections predict that we will be able to save in excess of $360,000 per year just implanting this system. Next year, our energy demand is projected to grow by 2.0 megawatts. Applying the savings to the projected growth the ERDC DSRC will save in excess of $650,000 next year. In addition, we plan to raise the operating temperature of the process chilled water system from 46 °F to 55 °F.

By combining the change in chilled water temperature with the improvements in the mechanical system operation, we anticipate an annual savings in excess of $1.2 million.

The four modes of operation are normal, precooling, comprehensive cooling, and load balance:

Normal Mode—In this mode, the return water is taken through the magnetic chiller; and the chilled water is sent to the HPC compute systems. The heat from the chiller is taken from the condenser barrel and dissipated to the outdoors via the water cooling towers. Use of the magnetic chiller will result in an approximate 35 percent reduction on power draw over a traditional air-cooled chiller.
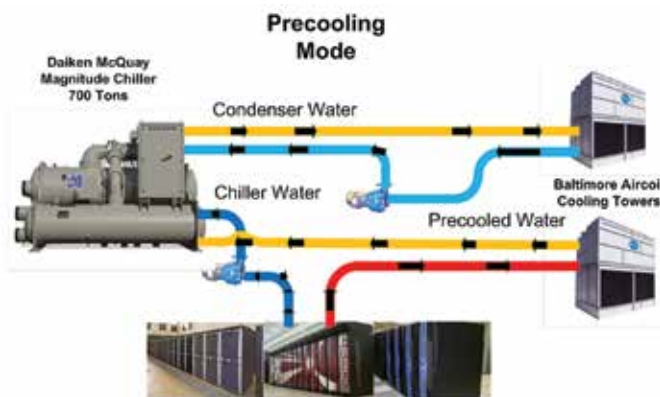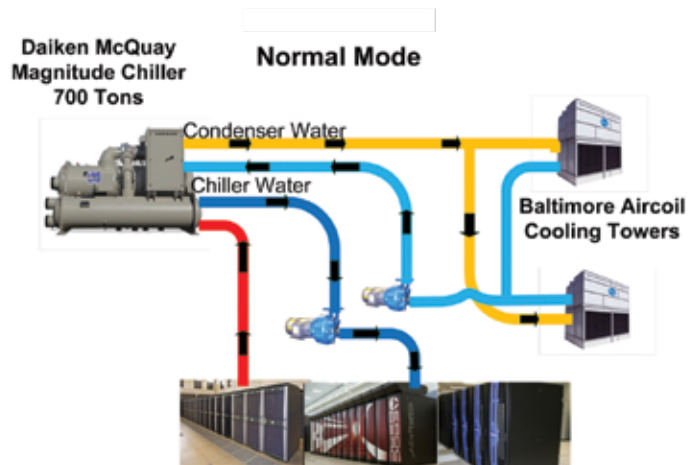
Precooling Mode—In the precooling mode, the controls will measure the ambient wet-bulb temperature along with the supply and return water temps from the computer systems. The controller will allow the return water from the computer systems to pass through the cooling towers, first for heat removal (precool), then through the chiller for the final cooling. This mode of operation is based on the wet-bulb temperature and provides the proper chilled water to the HPC system load by removing the process heat by both ambient conditions and mechanical cooling. When environmental conditions support this function, an additional 10 to 15 percent reduction in operating costs will be realized in relation to normal mode.

Comprehensive Cooling Mode—The wet-bulb temperature of the ambient air conditions will be the deciding factor for entering into the comprehensive cooling mode of operation. The chill water supply pipe allows the load return to pass through the cooling tower and then into the supply line back to the load with only the water pumps and fans to provide cooling to the load based on actual load demand. No mechanical cooling is used in this mode. The chilled water pumps are operated by the controls to meet the volume flow demand of the computer systems. When environmental conditions support this mode, an additional 25 to 30 percent reduction in operating costs will be realized in relation to normal mode.
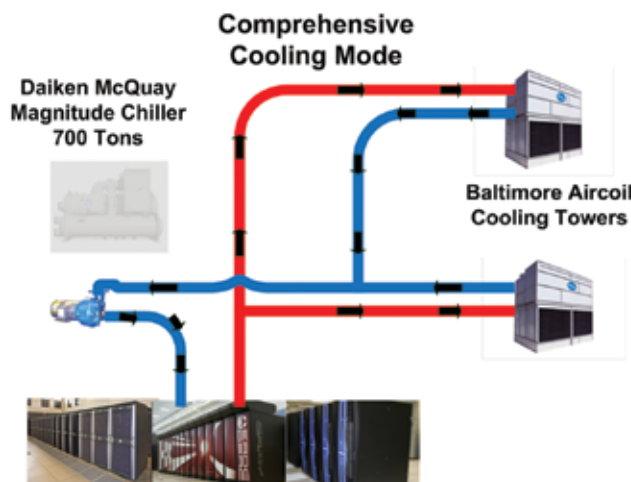
Load Balance Mode—The load balance mode is used to maximize the performance of the magnetic lift chiller under low load conditions. When the magnetic chiller has low heat demand, the cost per BTU rises because of the optimization of the magnetic chiller for large heat loads. Optimization under low heat load conditions requires warm water to be inserted into the evaporator barrel to raise the water temperature in the condenser barrel, thus allowing the chiller to operate at a lower KW/BTU heat rejection. This mode can reduce power consumption by up to 5 percent when the proper conditions are present.

Update on *Garnet*

By Paul Adams, ERDC DSRC

In the summer of 2012, the ERDC DSRC in Vicksburg, Mississippi, received upgrades to two of its Cray XE6 high performance computers (HPC). The first system, called *Garnet*, received an upgrade to its CPUs that doubled the number of cores from 20,160 to 40,320. The memory available to the CPUs on *Garnet* was doubled as well, increasing to 78.75 TB.

The second system at the ERDC DSRC, called *Chugach*, is an Open Research System (ORS). *Chugach* also received an upgrade to its CPUs that doubled the number of cores from 11,648 to 23,296. The memory available to the CPUs on *Chugach* was doubled as well, increasing to 45.5 TB. Both *Garnet* and *Chugach* use 2.5-GHz AMD 6200 series processors on its compute nodes. There are two processors per node, each with 16 cores, for a total of 32 cores per node. By design, the AMD 6200 series processor pairs cores into modules that allow them to share a "Flex FP" FPU (Floating Point Unit).

Cray also upgraded the CPUs on *Raptor*, which is located at the AFRL DSRC. *Raptor* received an upgrade to its CPUs that doubled the number of cores from 43,712 to 87,424 and its memory up to 174 TB. Just like *Garnet* and *Chugach*, *Raptor* is using the new AMD 6200 series of processors.

These three systems will be combined into one system in 2013 at the ERDC DSRC. The combined system will still be called *Garnet* and will feature 151,040 cores running at 2.5 GHz. The combined *Garnet* will also have over 298 terabytes of memory. This is not the first time that the ERDC DSRC has had a leadership-class Cray supercomputer. In 1997, the ERDC DSRC integrated a 544 CPU Cray T3E. By 2002, that system had increased in size to be a 1902-CPU Cray T3E, the largest in the world.

With a leadership-class Cray XE6 supercomputer, this allows the ERDC DSRC to focus on capability computing. Capability computing can be simply

defined as giving priority to larger jobs that compose a large fraction of the new *Garnet*. In contrast, capacity computing would be defined as giving priority to smaller jobs on the HPC system but running more of them. Capability computing allows computational scientists to attempt to solve complex problems that require tens of thousands of cores.

Changes will need to be made to the scheduling policy and file system structure in order to focus on capability computing. It is expected that the priority for large jobs will be increased. However, even with increased priority, it will still take a long time for large jobs to start running if the maximum run time remained at seven days. Therefore it is expected that the maximum

Garnet

run time will be decreased. This will lead to a higher queue turnover, which will give the capability jobs the ability to run in a timely manner.

Likewise, the 2.6 PB of disks will probably not be allocated to a single file system. Instead, the disks will likely be subdivided into smaller file systems from which the users may run their simulations. Whereas the crash of a single file system would bring down the entire machine, the potential crash one of several smaller file systems would allow the machine to stay operational.

From the Director's Desk – David Morton

M HPCC is pleased that this issue of *HPC Insights* is focusing on the CREATE project. It is my personal opinion that the efforts of CREATE have the potential to fundamentally change the way that the DoD develops, procures, and tests weapon systems. It won't be easy, and all organizations are resistant to change; but the current DoD budget shortfalls and exceedingly long procurement cycles all point to a necessity to update our acquisition activities to a more efficient model.

The ability to use supercomputers and to model complex systems, to understand design limitations, to identify potential mistakes, to optimize design, to optimize and reduce required testing, and to reduce the time to market have been proven time and time again in the commercial world. I've spent a considerable amount of my professional career at supercomputer vendors like Cray Research, SGI (I liked it so much I worked there twice!), and Linux Networx. During this period, I saw many commercial companies that turned to supercomputers to give them a competitive edge. They were able to produce better products faster and cheaper by optimizing designs early in the development cycle.

MHPCC really believes in the CREATE mission, and it has been our pleasure to support the talented CREATE team in many projects. Those efforts include

• *CREATE Software Development Environment:* MHPCC has helped to develop a software development environment that most of the development team uses. Features include code repository and checkout, wiki and other information exchanges, and bug tracking and resolution. Hundreds of CREATE developers and users across many organizations and locations have used this system.

• *CREATE Software Testing:* Through the use of Dedicated Support Partitions (DSPs) access at MHPCC, many of the CREATE software teams have dedicated access to a "virtual" cluster. Software development activities are particularly suited to dedicated use as many times jobs must be started, observed, and then stopped, changed, and resubmitted rapidly to optimize code or to correct errors. Regression testing prior to release is another activity that is well suited to the DSP environment. The access to these supercomputer resources can result in higher performance, higher quality software with a faster time to market.

• *HPCMP Portal:* The DoD HPCMP has directed the MHPCC DSRC to create and deploy a web-based Portal environment. This is a "Software-as-a-Service" (SaaS) model that helps to remove the barriers to accessing HPC resources and data. This is especially important to the CREATE effort—if efforts to access the software and HPC resources are perceived as too onerous, CREATE's goal of expanding the use and the impact of HPCMP assets could be limited. The SaaS model is also particularly advantageous to nimble software development teams such as CREATE. Instead of having to support versions of software on many platforms and locations, as one would historically have done, they can focus on having their software at just a few locations with much easier maintenance philosophy. Current plans are to make the Portal environment available at all DSRCs over the next nine months thus enabling Portal driven CREATE software at all DSRCs.

As detailed above, MHPCC has teamed with the CREATE team on many mutually beneficial projects. We look forward to supporting these leading-edge software efforts in the future.

From the Director's Desk – Tom Dunn

In this issue of *HPC Insights*, you'll find three articles related to the Navy DSRC's recent computational upgrades.

The first article pertains to the facility upgrades required to support three high performance computing systems with an aggregate computational capacity of 954 teraflops. It also details how the Center made a conscious decision to become "greener" and incorporate water cooling at the rack level into our data center.

The next article will hopefully help users get up to speed on the new IBM iDataPlex platforms, *Haise, Kilrain,* and *Cernan*, currently in production at the Navy DSRC. In this article, users will find general system information, tips and tricks, as well as links to Navy DSRC web documentation on the systems.

The final article covers a special ceremony the Navy DSRC hosted recently. We were pleased to be able to honor an American hero, former astronaut and naval aviator, Fred Haise, with a dedication ceremony. One of the new IBM iDataPlex systems is named after Haise, a native of Biloxi, Mississippi, who is a key contributor to the success of the NASA Outreach program locally.

As the Navy DSRC welcomes the new IBM iDataPlex systems into the fold, the Center also says farewell to two systems that have been highly popular with the user community. The 12,736-core Cray XT5 Einstein and 4800-core IBM Power6 Davinci arrived in late 2008 as part of the TI-08 acquisition. The systems have been two of the most stable the Center has had and have been tremendous computational platforms for the DoD HPCMP.

(Left to right) John West, HPCMP Director; Dave Cole, Navy DSRC Associate Director; Jerry Cook, NASA SSC Deputy Director; Dr. Rich Gilbrech, SSC Director; and Honoree Fred Haise listen as Navy DSRC Director Tom Dunn describes the new supercomputing capabilities at the Center.

Fred Haise is joined by special guests Steven Harrison, Acting Technical Director at the Naval Oceanographic Office; Ken Human, Associate Director of NASA Stennis Space Center (SSC); Dr. Richard Gilbrech, NASA SSC Director; John West, Director of the DoD High Performance Computing Modernization Program (HPCMP); Jerry Cook, NASA SSC Deputy Director; Myron Webb, Stennis Legislative Affairs Officer; Dr. Bill Burnett, Technical Director at the Naval Meteorology and Oceanography Command; Dr. Herbert Eppert, Superintendent of the Marine Geosciences Division and Head of Office Research Support at the Naval Research Lab at SSC; and Tom Dunn, Navy DSRC Director.

Improving Productivity on IBM iDataPlex Systems

By Bryan Comstock, Navy DSRC Outreach Staff, and Sean Ziegeler, Navy PETTT Site Lead

The Navy DSRC is pleased to offer three new IBM iDataPlex 1350 systems to the HPCMP user community. Each system is based on the Intel Sandy Bridge E2670 processor, running at 2.6 GHz and offering features like TurboMode and Hyperthreading, RedHat Enterprise Linux, the IBM General Parallel Filesystem (GPFS), and a Mellanox FDR-10 Infiniband interconnect. The two larger systems, *Haise* and *Kilrain*, are composed of 1236 total compute nodes—1220 contain 32 GB of memory and 4 are configured as large memory nodes with 256 GB of memory each. The remaining 12 compute nodes have 64 GB of memory available to them and are attached to an accelerator blade containing two Xeon Phi 5110P co-processors. Each Xeon Phi co-processor has 1 teraflop of computational capacity. The smaller system, *Cernan*, contains 252 compute nodes and 32 GB of memory per node as well. The three systems are named in honor of former naval aviators who have served as NASA astronauts.

General Tips

Now that the systems are in production, users are beginning to take full advantage of them. Here are some suggestions for getting the most out of *Haise*, *Kilrain*, or *Cernan*:

• The systems have three MPI libraries available: IBM PE, Intel MPI, and OpenMPI. Users are encouraged to try each. Some codes may perform better with one MPI while others may do better with another.

• When changing compiler options related to speed (both optimizations and debugging), be sure to recompile all components of the code (i.e., source files, modules, shared objects, etc.). If a component of the code does not get recompiled with the new options, the code may run at the rate of the slowest component.

• Keep an eye on memory usage, especially arrays and structures that get replicated in every MPI task. If the code is running out of memory, try running fewer MPI tasks per node. If this solves the problem, there is a good chance that the code is using a significant amount of memory in replicated data in every task.

Code Compilation Tips

The systems are equipped with multiple versions of the Intel Compiler Suite and the system default GNU Compiler Suite. As mentioned above, there are three MPI libraries available and multiple versions of each. The different versions of software can be accessed via the Modules software and the "module" command. For more information on the modules environment, please see http://www.navo.hpc.mil/docs/modulesUserGuide.html.

For best performance, users should compile source code with the Intel compiler. The Intel compiler will be able to best optimize executables for the Intel Sandy Bridge cores that drive the systems. In order to get the most optimized performance, users should implement the "*-fast -O3 -xHost -ipo*" flags. Please note that including optimization flags will increase compilation times. Users should also watch numerical operations carefully for accuracy when attempting to optimize. The "*-O3*" and "*-fast*" flags will substitute certain operations with faster but less precise routines. If the precision provided under those optimizations is unacceptable, the "*-prec-div*" and "*-prec-sqrt*" flags will enforce fully precise divisions and square-root routines even at the highest optimization levels. If further numerical precision is required, "*-fp-model precise*" will maintain only value-safe floating-point optimizations.

It is also highly recommended that users turn off all debugging, checking, and tracing options once code stabilization and debugging are complete. These options can negatively impact performance significantly when using the Intel compiler.

Below is a table containing pertinent Intel optimization options and their descriptions:

Option	Description
-O0	No Optimization.
-O1	Scheduling within extended basic blocks is performed. Some register allocation is performed. No global optimization.
-O2	Level 1 plus traditional scalar optimizations such as induction recognition and loop invariant motion are performed by the global optimizer. Generally safe and beneficial. (default in Intel).
-O3	Levels 1 and 2 plus more aggressive code hoisting and scalar replacement optimizations that may or may not be profitable. Generally beneficial.
-xHost	Compiler generates code with the highest instruction set available on the processor.
-fast	Chooses generally optimal flags for the target platform. Includes -ipo -O3 -static.
-ipo	Interprocedural analysis.
-parallel	Enables autoparallelization (serial codes).

Batch Configuration

Once code is ready to be executed, users should submit batch jobs to the PBSPro batch scheduler found on *Haise, Kilrain, and Cernan*. Below is a table of the queue configurations for *Haise* and *Kilrain*, as they are identical.

Queue	Max Wallclock Time	Max Cores Per Job
urgent	24 Hours	4096
high	168 Hours	6144
challenge	168 Hours	6144
special	24 Hours	4096
debug	30 Minutes	1024
standard	168 Hours	4096
transfer	12 Hours	1
background	4 Hours	512

For more information on submitting jobs, please see the PBS Guide available at http://www.navo.hpc.mil/docs/pbsUserGuide.html.

Additional system documentation may be found at http://www.navo.hpc.mil/docs/. For additional specific information on compiling code or submitting batch jobs on the new IBM iDataPlex systems, please see the *Haise and Kilrain* User Guides available at http://www.navo.hpc.mil/docs/haiseUserGuide.html and http://www.navo.hpc.mil/docs/kilrainUserGuide.html, respectively.

IBM iDataPlex

Facility Upgrades in Support of the IBM iDataPlex Systems

By Rob Thornhill, Navy DSRC Facilities Lead

During the fourth quarter of 2012, the Navy DSRC received, installed, and accepted three new IBM iDataPlex 1350 supercomputers. Many months in advance of the selection of the systems, the Navy DSRC made a conscious effort to plan for water-cooled high performance computers (HPCs). As a major chiller plant upgrade was already in progress, design changes were made to create a new, separate chilled water loop so that the chilled water plant could provide water at two temperatures simultaneously to help maximize the energy efficiency of the chilled water system. The existing chilled water plant was converted to provide relatively warmer water to the water-cooled HPC equipment. A second smaller chiller was then installed to provide cold water to computer room units (CRUs) supporting the Center's remaining air-cooled equipment.

A new 300 ton magnetic chiller being lowered next to a 20,000 gallon cool water storage tank.

The IBM iDataPlex systems offered in the Technology Insertion (TI) process had an option for rear-door heat exchangers (RDHX), which the Navy DSRC elected to accept. Subsequently, the warm-water distribution piping was modified to provide the 65 °F "warm" chilled water directly to the rear doors. By having a separate higher temperature cooling loop, the Center was able to forego the installation of multiple cooling distribution units (CDUs) that are generally required to support water-cooled systems. This saved not only cost but significant floor space that will be used to house future HPC equipment. Also, since approximately 80 percent of the cooling for the large HPC systems, *Haise* and *Kilrain*, was to be supported directly by their own RDHX, three CRUs were not needed, further freeing up floor space. When the cost of installing and supporting six additional CRUs is taken into consideration, the Center will ultimately save around $50K per year. Savings for the anticipated 4-year life cycle of the systems will amount to $200K.

Having the mechanical portion largely taken care of, the Center proceeded to modernize the electrical infrastructure. Power distribution units (PDUs) from multiple vendors were evaluated for physical size, ease of access, cost, quality, and efficiency. As part of site prep, six Powersmiths PDUs, with extremely energy-efficient transformers, were purchased and installed. Also installed were all of the necessary electrical whips needed to provide power to all of the new IBM racks.

It should also be noted that the Center worked extensively to do computational fluid dynamics (CFD) modeling of the space. The CFD simulations allowed small changes to be made easily. Sufficient and efficient airflow to each area of the room was achieved as a result of these changes. Without the modeling effort, it would likely have taken weeks of trial and error in order to balance the airflow throughout the space.

The 300 ton chiller's eventual location is just inside the double doors on the left of the picture. The chiller had about 1/2" of clearance going through the doors.

Navy DSRC Establishes HPC System Names and Honors NASA Astronaut and Former Naval Aviator Fred Haise at Ceremony

Supercomputers require distinctive names for identification to their users and the people who maintain them; and with the arrival of three new supercomputing systems and the decommissioning of three outgoing supercomputers, the Navy DSRC went in search of a new naming scheme that would present a plentiful list of names which the Center could draw from well into the future.

To recognize the unique fact that the Navy DSRC's mission is carried out on the grounds of NASA's Stennis Space Center, the Center chose a naming scheme that honors NASA astronauts who have served in the Navy.

At a ceremony in February, the Center dedicated one of its IBM iDataPlexes in honor of naval aviator and Apollo 13 astronaut Fred Haise, who attended the event. The two other IBM systems are named for retired Navy Cmdr. Susan Still Kilrain, a naval aviator and space shuttle pilot, and retired Navy Capt. Eugene Cernan, a naval aviator and the last man to step foot on the moon.

Installation of the new systems expanded the installed supercomputing capability of the Navy DSRC, which now peaks at 954 trillion floating point operations (teraflops) per second.

The chosen list of names includes some 28 men and women. The Navy DSRC is proud to recognize three iconic men and women who have served NASA and the Navy, and to continue naming its supercomputers after such an acclaimed and distinguished group.

"Today, we are proud to recognize the contributions of an iconic American and native Mississippian," said Dr. William H. Burnett, Deputy Commander and Technical Director of the Commander Naval Meteorology and Oceanography Command at Stennis Space Center. "Just as Fred Haise has made a great impact on the state, the Navy, and the Nation, so will the supercomputer named after him."

Apollo 13 astronaut and former naval aviator Fred Haise speaks to the Navy DSRC staff and special guests at the dedication ceremony.

"The Navy DSRC provides unique value within our supercomputing system," observed John West, Director of the HPCMP. "In addition to serving the users from the research, development, test, and evaluation communities of the Department served by all of our Centers, the Navy DSRC has a unique mission to assist the Navy in delivering wind, wave, and other oceanographic forecasts to the fleet on a 24/7 basis. We are proud of the work of our partners, and the men and women of the Navy DSRC, who have brought this added capability online for the Department."

Fred Haise stands with DoD High Performance Computing Modernization Program Director John West, NASA Stennis Space Center (SSC) Director Dr. Richard Gilbrech, Naval Meteorology and Oceanography Command (NMOC) Technical Director Dr. Bill Burnett, special guests, and the staff of the Navy DSRC at the conclusion of the ceremony dedicating one of the Center's supercomputers in Mr. Haise's name.
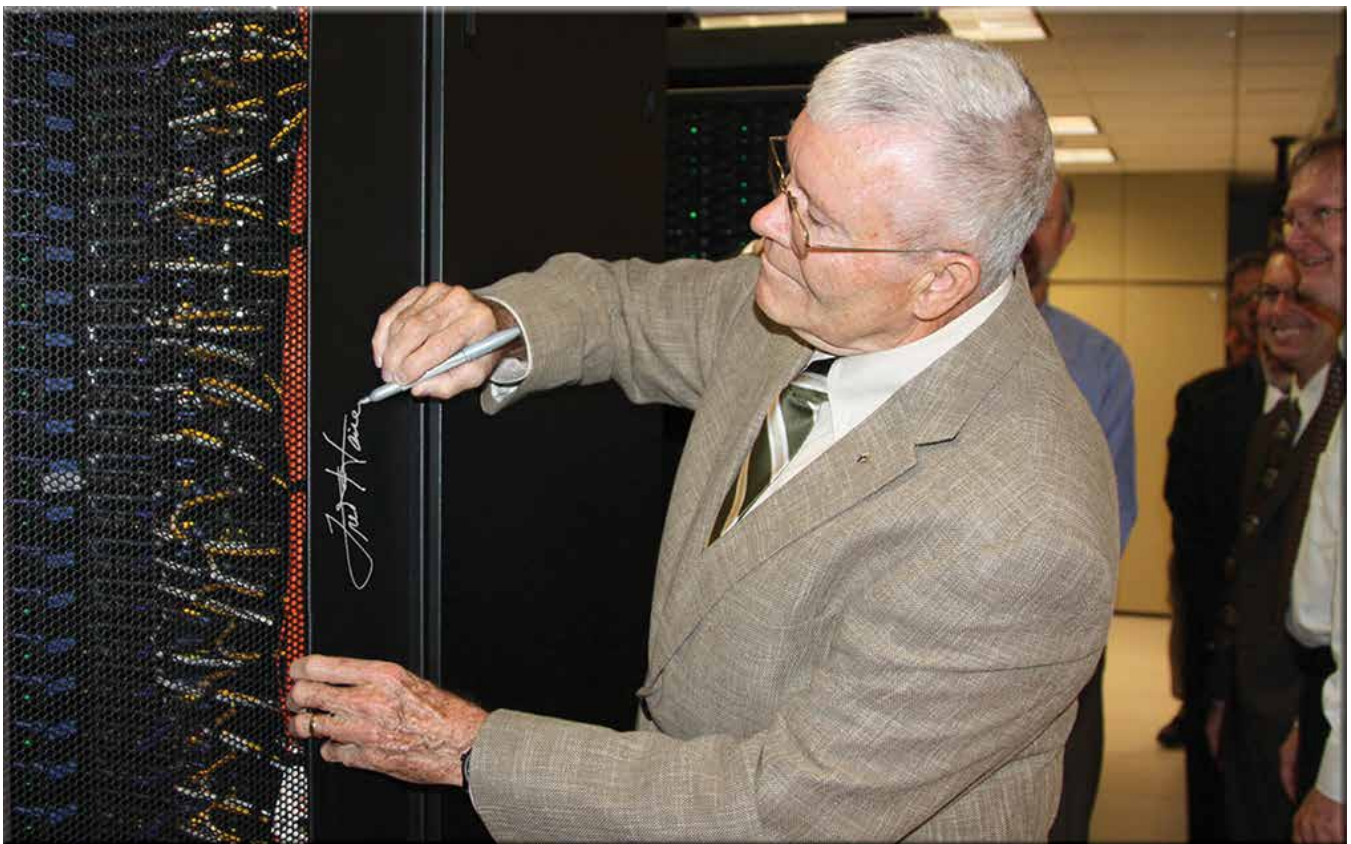
Astronaut and former naval aviator Fred Haise signs his namesake supercomputer, a 21,216-core IBM iDataPlex installed at the Navy DSRC.

Who's Your UAG Rep?

The User Advocacy Group (UAG) was established to provide a forum for users of DoD HPCMP resources to influence policies and practices of the Program, to facilitate the exchange of information between the user community and the HPCMP, to serve as an advocacy group for all HPCMP users, and to advise the DoD HPCMPO on policy and operational matters related to the HPCMP. Each Service has four representatives in the group providing a balanced representation on both its S&T and T&E communities. The DoD Agencies may each have one representative.

Users are encouraged to contact their UAG members with any issues or concerns they might have that fall within the scope of the group as outlined above. Along with regularly scheduled meetings, the group now plans to meet via Defense Connect Online several times throughout the year to discuss user concerns.

There is a complete list of UAG representatives on the HPCMP Bridge (https://www.hpcmo.hpc.mil/bridge/), along with their contact information. We have also created new exploder lists for the UAG representatives within each Service/Agency for your convenience. The email addresses are as follows:

Air Force:	AFUAGrep@hpc.mil
Army:	ArmyUAGrep@hpc.mil
Navy:	NavyUAGrep@hpc.mil
DoD Agencies:	AgencyUAGrep@hpc.mil.

Your UAG representatives will receive the email and be able to help bring your issues/comments/recommendations to HPCMP Leadership.

FY 2012 User Satisfaction Survey

Throughout its history, the High Performance Computing Modernization Program (HPCMP) has employed user satisfaction surveys. The purpose of the surveys is to determine performance through tangible metrics, which in turn provides the opportunity for continual Program improvement. At the annual Users Group Conference (UGC), the HPCMP Director relays the results of the annual User Satisfaction Survey to the community. As this year's UGC was cancelled, here are some of the highlights of the Survey.

In FY12, the survey was executed on a quarterly schedule, providing four snapshots of the Program over the year. The response rate was fairly consistent for all quarters, in the 28 percent range, providing a margin of error of +/- 2.59 percent.

The majority of the questions were on a 10-point Likert scale to evaluate effectiveness, usability, satisfaction, and usefulness. In addition, each section provided a section for comments, allowing the user community to provide additional feedback to assist the HPCMP in improving our service. The provision for comments has proved invaluable to the HPCMP. A case in point, in 2012 the Program Office engaged the "Last Mile Team" to assist users who indicated they were having connectivity issues. This information would have been unavailable had users not had a forum to provide comments. The Last Mile Team has since engaged directly with the user community, HPCMP DREN and Centers, along with local site network and security personnel to help resolve the connectivity issues.

The graph below shows a sampling of the areas covered in the survey. In rating their satisfaction, respondents gave an average 8.9 (on a 10-point Likert scale) value of using HPCMP hardware to achieve goals. Overall, respondents were satisfied with CCAC services and the HPCMP supporting their goals. They were also satisfied with the ease of use of queuing and scheduling and valued the simplification for information retrieval on Center websites.
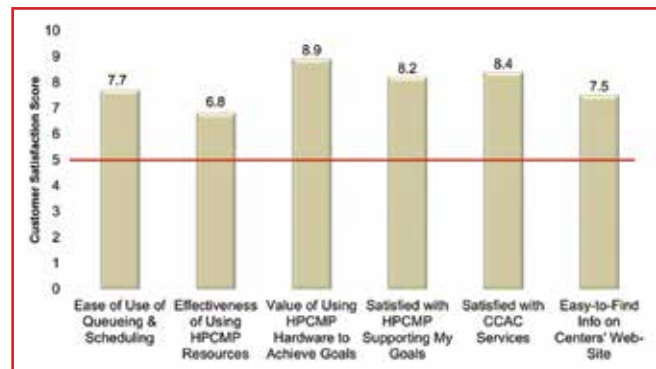
When questioned about the impact, value, and satisfaction of using HPCMP resources on projects or programs, the results are noteworthy. Of the 876 that responded to the question on improving time to solution, 70 percent of users indicated an improvement by 51 percent or more.

The information obtained through the User Satisfaction Survey (average score of 58 rating questions) concluded that the user community is satisfied with the services and products provided by the HPCMP, resulting in a score of 7.64 overall.
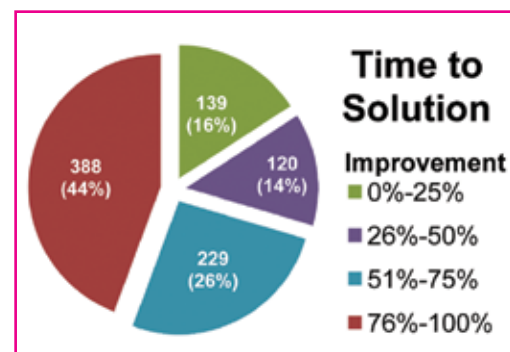
Users are encouraged to participate in the FY 2013 survey. We make changes based on your feedback.
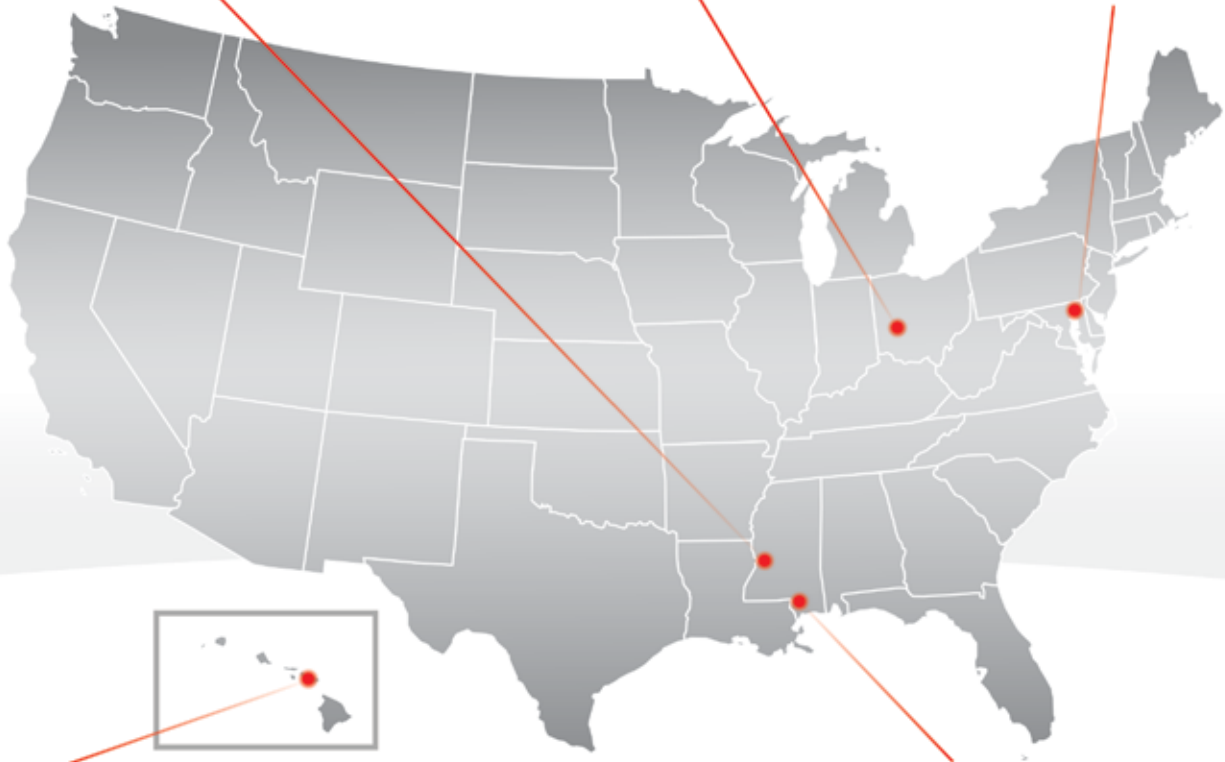
DoD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM

(ERDC DSRC)
US Army Engineer Research and Development Center
Vicksburg, MS,
http://www.erdc.hpc.mil/

(AFRL DSRC)
US Air Force Research Laboratory
Wright-Patterson AFB, OH,
http://www.afrl.hpc.mil/

(ARL DSRC)
US Army Research Laboratory
Aberdeen Proving Ground, MD,
http://www.arl.hpc.mil/

(MHPCC DSRC)
Maui High Performance Computing Center
Kihei, Maui, HI,
http://www.mhpcc.hpc.mil/

(NAVY DSRC)
Navy DoD Supercomputing Resource Center
Stennis Space Center, MS,
http://www.navo.hpc.mil/

Scientists and engineers throughout the U.S. leverage the capabilities of the High Performance Computing Modernization Program (HPCMP) to solve the most time-consuming computational problems. They know that the HPCMP's supercomputing centers continually architect, deploy, and sustain their equipment to deliver world-class network and supercomputing capabilities, resulting in simulations with greater resolution and faster results than achievable on conventional workstations. You too can gain access to these powerful resources by calling to register for an account!

CCAC Accounts Center
1-877-222-2039

http://www.hpc.mil